

Mesh Saliency and Human Eye Fixations

YOUNGMIN KIM, AMITABH VARSHNEY, DAVID W. JACOBS, and FRANÇOIS GUIMBRETIERE
University of Maryland

Mesh saliency has been proposed as a computational model of perceptual importance for meshes, and it has been used in graphics for abstraction, simplification, segmentation, illumination, rendering, and illustration. Even though this technique is inspired by models of low-level human vision, it has not yet been validated with respect to human performance. Here, we present a user study that compares the previous mesh saliency approaches with human eye movements. To quantify the correlation between mesh saliency and fixation locations for 3D rendered images, we introduce the normalized chance-adjusted saliency by improving the previous chance-adjusted saliency measure. Our results show that the current computational model of mesh saliency can model human eye movements significantly better than a purely random model or a curvature-based model.

Categories and Subject Descriptors: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling; I.3.m [Computer Graphics]: Perception

General Terms: Algorithms, Human Factors, Verification

Additional Key Words and Phrases: Visual perception, mesh saliency, eye-tracker

ACM Reference Format:

Kim, Y., Varshney, A., Jacobs, D. W., and Guimbretière, F. 2010. Mesh saliency and human eye fixations. *ACM Trans. Appl. Percept.* 7, 2, Article 12 (February 2010), 13 pages.
DOI = 10.1145/1670671.1670676 <http://doi.acm.org/10.1145/1670671.1670676>

1. INTRODUCTION

With the increase in size, number, and complexity of 3D graphics datasets, it will become increasingly important to integrate principles of saliency with geometric processing of meshes. Lee et al. [2005] have proposed a model of mesh saliency as a measure of regional importance. Their method for computing mesh saliency uses a center-surround mechanism that is inspired by the human visual system. Similar mechanisms have been widely used in models of 2D image saliency [Itti et al. 1998; Koch and Ullman 1985]. Previous research in vision has assumed that visual search simply relies on 2D aspects of an image. However, Enns and Rensink [1990] have shown that 3D information can play an important role in making salient objects pop-out in a cluttered image. As far as we know, there has been no work comparing models of 3D saliency to eye movements, although many experiments have measured eye

This work has been supported in part by the NSF grants: IIS 04-14699, CCF 04-29753, CNS 04-03313, CCF 05-41120, CMMI 08-35572 and ARO Grant #W911NF-08-1-0466. Any opinions, findings, conclusions, or recommendations expressed in this article are those of the authors and do not necessarily reflect the views of the research sponsors.

Authors' addresses: Y. Kim, A. Varshney, D. W. Jacobs, and F. Guimbretière, Department of Computer Science, University of Maryland; email: {ymkim,varshney,djacobs,francois}@cs.umd.edu.

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 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 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2010 ACM 1544-3558/2010/02-ART12 \$10.00

DOI 10.1145/1670671.1670676 <http://doi.acm.org/10.1145/1670671.1670676>

movements as participants examine 3D objects [Howlett and O’Sullivan 2005; Kim and Varshney 2006; Cole et al. 2006; Lu et al. 2006].

In this article, we present a user study that compares the previous mesh saliency approaches with human eye fixations. Having a validated model of mesh saliency will be extremely useful in several contexts. For example, it could be helpful for identifying the role of 3D information in visual search task as Enns and Rensink [1990] have explored in their work. Our user study and data analysis can be also helpful for designing a better visual saliency model, which is closer to human eye movements. The main contributions of this article are the following.

- (1) We compare models of 3D mesh saliency to eye movements through an eye-tracking–based user study.
- (2) We introduce the normalized chance-adjusted saliency to quantify the correlation between mesh saliency and fixations for 3D rendered images. This is more appropriate for 3D rendered images than the previous chance-adjusted saliency measure.
- (3) We show that the current computational model of mesh saliency models human eye movements significantly better than what can be expected by chance or due to curvature alone.

The rest of the article is organized as follows. A review of related work is provided in Section 2. We present our experimental design in Section 3. In Section 4, we introduce the normalized chance-adjusted saliency and present the results. Section 5 discusses limitations of our approach and future work, and Section 6 concludes this article.

2. RELATED WORK

When people examine an image, their eyes tend to fixate on certain points, then jump quickly, with saccades, to new points. Although viewers may attend to portions of an image on which they do not fixate, a good deal of evidence suggests that viewers tend to move their eyes to parts of an image that have attracted their attention (see Palmer [1999], Chapter 11, for a brief review). For this reason, many models of visual attention and saliency have been evaluated by their ability to predict eye movements. It is not realistic to expect any model to perfectly predict eye movements because of the variability between human participants and even for the same participant at different times. However, recent research demonstrates that there is a significant correlation between existing models and human eye fixations. For example, Privitera and Stark [2000] compare points of fixation by participants to clusters formed by the most salient regions predicted by a large number of simple models of 2D image saliency. They compare this with the degree to which fixations agree between participants. Of the three classes of images they have looked at, they have found that for one class of images (paintings), algorithms based on simple operators including symmetry, center-surround, and discrete wavelet transform cohere very well with human data and approach the coherence among fixations across participants. Parkhurst et al. [2002] measure the saliency at points of fixation and show that the model of 2D image saliency of Itti et al. [1998] is more predictive of fixation points than a random model. Previous research also makes the useful methodological points that bottom-up models can better predict the first fixations, which are less influenced by top-down knowledge [Parkhurst et al. 2002], and that the exact order of fixations is highly variable and difficult to predict [Privitera and Stark 2000].

There have been many experiments measuring eye movements as participants examine 3D objects [Howlett and O’Sullivan 2005; Kim and Varshney 2006; Cole et al. 2006; Lu et al. 2006]. For instance, Cole et al. [2006] emphasize a region of interest by rendering it in a different style from other regions. They validate the effects of their rendering styles by an eye-tracking experiment. Howlett and O’Sullivan [2005] capture saliency values of 3D models by recording where participants look in 3D rendered images. They use this saliency information to simplify 3D models.

The computational model of mesh saliency [2005] uses a center-surround mechanism that is inspired by the human visual system. There have been other approaches to identifying salient regions on a mesh. Watanabe and Belyaev [2001] have identified salient regions on meshes by estimating curvature values and guided the simplification process to preserve them better. Gal and Cohen-Or [2006] have constructed salient geometric features by clustering a set of local descriptors that have a high curvature and a high variance of curvature values for partial shape matching of meshes. Shilane and Funkhouser [2007] have presented a method to select distinctive regions of a 3D model that are consistent with the same type of objects and are different from other types of objects. They compute how distinctive the regions are with respect to multiple classes of objects in a database, and use this information for local shape matching, simplification, and icon generation. Feixas et al. [2008] have defined an information channel between the viewpoints and an object. They use this channel to compute the viewpoint mutual information, which is further used for viewpoint selection and mesh saliency computation. While the notion of saliency in Lee et al. [2005] and Feixas et al. [2008] is based on perceptual and visual importance, the notion of saliency in Gal and Cohen-Or [2006] and Shilane and Funkhouser [2007] is specific to the tasks of shape matching, shape similarity, and shape uniqueness.

In Kim et al. [2008], we describe a user study examining the correlation between human eye fixation and a mesh saliency model by varying 3D models, view angles, and orientation (right-side-up vs. upside-down views). This study has shown that there may be certain view-dependent effects, such as view angles and orientation, on human eye movements. Although studying view-dependence of mesh saliency is an important area, a view-independent model of mesh saliency is desirable for several applications. For instance, offline mesh simplification, mesh segmentation, view selection, lighting design, as well as allocation of computational resources in the rendering pipeline all benefit from a view-independent model of mesh saliency. For such applications, it is crucial to establish the correlation between human eye movements and view-independent models of mesh saliency. In this article, we compare a model of mesh saliency with a purely random model and a curvature-based model. There are a number of excellent approaches that generalize differential-geometry-based definition of curvatures to discrete meshes [Goldfeather and Interrante 2004; Meyer et al. 2003; Taubin 1995; Yang et al. 2006]; we use Taubin's method [Taubin 1995] for computing mean curvature in this article.

3. EXPERIMENTAL DESIGN

To gather objective evidence of the correlation between saliency models and human eye fixations, we have carried out an eye-tracking-based user study and have quantified the similarity between the models and human eye fixations.

3.1 Hypothesis

Our hypothesis is that the computational model of mesh saliency has better correlation with human eye fixations than a random model and a curvature-based model for the first few seconds after stimulus onset.

3.2 Physical Setup

We used the ISCAN ETL-500 monocular eye-tracker, which can record eye movements continuously at 60Hz. The study was carried out on a 17-inch LCD display with a resolution of 1280×1024 , placed at a distance of 24 inches, subtending a visual angle of approximately 31.4 degrees horizontally. The participants had a chin rest to minimize head movements and to maintain calibration. Our experimental setup is shown in Figure 1.



Fig. 1. Our experimental setup for the user study with the ISCAN ETL-500 eye-tracker.

3.3 Eye-Tracker Calibration and Participant Selection

The standard calibration of ETL-500 eye-tracker was performed with four corner points and one center point. However, this was not sufficiently accurate for our purposes due to nonlinearities in the eye-tracker-calibrated screen space. Therefore, we used the second calibration step, which involves a more densely-sampled calibration phase similar to [Parkhurst et al. 2002], with 13 additional points. For this, we asked the participants to successively look at and click on 13 points presented on the screen. This gave us an accurate correspondence between the eye-tracker space and the monitor space for that participant. After this, we tested the accuracy of the calibration by asking the participants to look at 16 randomly selected points on the screen. Of the 25 participants participating for pay, 18 were able to successfully calibrate to within an accuracy of 30 pixels (about .75 degree) for each of the 16 points. We proceeded with our study using these 18 participants with normal or corrected-to-normal vision. Our participants were not familiar with the goals of this study. The participants were told to freely view the images with no assigned goal.

3.4 Stimuli

We have illustrated five natural scanned models used for our study in Figure 3. Each user saw these five images, and we randomized the order of images to counterbalance overall effects.

The user study had seven trials (images). The data captured for the first two images was discarded, as they were intended to give the participants a sense of the duration of viewing. Each trial started with the participants seeing a blank screen with a cross at the center of the screen. The participant was asked to look at the cross before clicking the mouse to bring up the next image. This ensured that each trial started with the participant's eyes fixated at the center of the image. Each image was shown for 5 seconds.

4. DATA ANALYSIS

4.1 Fixation Points

We divide the raw data points from the eye-tracker into two groups—fixation points, which correspond to a user looking at a single location and saccade points, which correspond to fast eye movements from one fixation point to the next. We followed an approach similar to the one suggested by Stampe [1993] to identify fixations and saccades. Figure 2 shows a two-step process to extract fixation points from the raw data points. We considered data points that had a velocity greater than 15°/sec as saccade points and removed them. We then averaged consecutive eye locations that were within 15 pixels and classified them as a single fixation point. Some researchers have advocated discarding short (exploratory) fixations




Fig. 2. Image (a) shows all the raw data points from the eye-tracking device. Image (b) shows the points remaining after removing saccade points. Image (c) shows final fixation points after removing brief fixations and combining consecutive points if they are spatially close.




Fig. 3. The Dinosaur, Isis, Male, Armadillo, and Igea models used in our study.

in measuring the attention of the viewer [Henderson and Hollingworth 1998] to discriminate between distraction and attention. We ignored brief fixations below the threshold of 133ms. This corresponds to eight consecutive points in the ISCAN ETL-500 eye-tracking device.

4.2 Normalized Chance-Adjusted Saliency

4.2.1 Chance-Adjusted Saliency. Parkhurst et al. [2002] introduced the notion of chance-adjusted saliency to quantify the correlation between stimulus saliency and fixation locations for an image. They compute the chance-adjusted saliency as follows. Consider a collection of images I_i , $1 \leq i \leq N$. A participant is asked to look at each image in turn. This generates a set of fixation points f_{ij} , $1 \leq j \leq F_i$ for each image I_i , where F_i is the number of fixation points. Let us consider the k -th fixation points f_{ik} across all the images I_i . Let s_{ik} be the saliency value at the k -th fixation point f_{ik} in the image I_i . They compute the mean fixation saliency for the k -th fixation points as $\bar{s}_k^f = \frac{1}{N} \sum_{i=1}^N s_{ik}$. To compute the mean random saliency, they first generate F_i random points r_{ij} over each image I_i , where $1 \leq i \leq N$ and $1 \leq j \leq F_i$. Then, the mean random saliency \bar{s}_k^r is computed as the average saliency over the k -th random point r_{ik} across all the images I_i , $1 \leq i \leq N$. Finally, they define the chance-adjusted saliency (s_k^c) for the k -th fixation points as the difference between the mean fixation saliency (\bar{s}_k^f) and the mean random saliency (\bar{s}_k^r): $s_k^c = \bar{s}_k^f - \bar{s}_k^r$.




Fig. 4. The left image shows the random points in chance-adjusted saliency computation. These points are scattered all over the image. The right image shows the points that we consider in normalized chance-adjusted saliency computation. We only include the foreground pixels that are covered by projected triangles of the mesh. For each fixation point represented as a cross, we also take into account the eye-tracker accuracy of 20 pixels, which is represented as a circle.

4.2.2 Normalized Chance-Adjusted Saliency. We observed three shortcomings in using the previously defined chance-adjusted saliency to quantify the correlation between human eye fixations and the model of mesh saliency.

- (1) The chance-adjusted saliency was developed for images in which there is a well-defined saliency at every pixel. We are trying to measure the correlation between a mesh saliency approach and the fixation points on the mesh but not the fixations on the entire rendered image. Therefore, we should only consider the foreground pixels that are covered by projected triangles of the mesh. This ensures fairer comparisons between a random model and the saliency model for 3D rendered images because excluding the background pixels would prevent lowering the average saliency values in a random model. Figure 4 shows the points considered in chance-adjusted saliency and normalized chance-adjusted saliency.
- (2) The chance-adjusted saliency does not consider eye-tracker accuracy. Since the fixation point acquired from the eye-tracker can differ from the actual pixel that a user looked at, we have to consider the eye-tracker accuracy as shown in Figure 4(b) when we assign the mesh saliency value to the fixation point.
- (3) The chance-adjusted saliency is defined over a collection of images. This restricts the analysis of the effect of different models. We need a method that normalizes saliency on a per-image basis.

To address these problems, we define *normalized chance-adjusted saliency* in this section. First, we consider the eye-tracker accuracy ϵ , which depends on both the accuracy of the eye-tracking device and the calibration steps. We have used $\epsilon = 20$ pixels, subtending a visual angle of approximately 0.5 degree horizontally. Note that a fixation point and a pixel share the same coordinate system. Let us consider the pixel p_{ij} on which a fixation point f_{ij} falls. Instead of taking the saliency value on a fixation point f_{ij} , we compute the error-adjusted saliency s_{ij}^ϵ as the maximum of the saliency values within a radius of $\epsilon = 20$ pixels around p_{ij} in the image I_i , $1 \leq i \leq N$: $s_{ij}^\epsilon = \max_{k \in \mathcal{N}_j^\epsilon} s_{ik}$, where $\mathcal{N}_j^\epsilon = \{k | \text{dist}(p_{ij}, k) \leq \epsilon\}$. For each rendered image I_i , we compute the mean (\bar{s}_i^ϵ) of the saliency s_{ij}^ϵ over all the pixels j that are covered by the rendered mesh. We now define our normalized chance-adjusted saliency for the fixation point f_{ik} , as $s_{ik}^n = s_{ik}^\epsilon / \bar{s}_i^\epsilon$. Here we use the ratio instead of the difference to enable it to be used across different models and different view points; otherwise, we will need to normalize the 3D saliency values for each rendered image for fair comparisons with different models and different views of a




Fig. 5. The saliency models and human eye fixations. The first and second rows show the mean curvature and mesh saliency, respectively, on the models shown in Figure 3. Here warm colors indicate high saliency regions while cool colors indicate low saliency regions. The third row shows the human eye fixations from our eye-tracking-based user study with hot spot maps.

model. We note that we could have computed the mean of the error-adjusted saliency differently for each participant, as in chance-adjusted saliency: Generate random points on the pixels covered by the rendered mesh and compute the average saliency values on these random points. However, this causes high variance in the means among different participants, resulting in a high variance in normalized chance-adjusted saliency values even for the same image. To avoid this high variance, we have decided to use \bar{s}_i^e in computing normalized chance-adjusted saliency for all fixation points on the image I_i .

4.3 Results

Figure 5 shows the fixation points and computed mean curvature and mesh saliency for each model. Fixation points are illustrated with hot spot maps, where warm colors show highly fixated regions. We observe that most fixations are close to warm-colored salient regions computed by the model of mesh saliency.

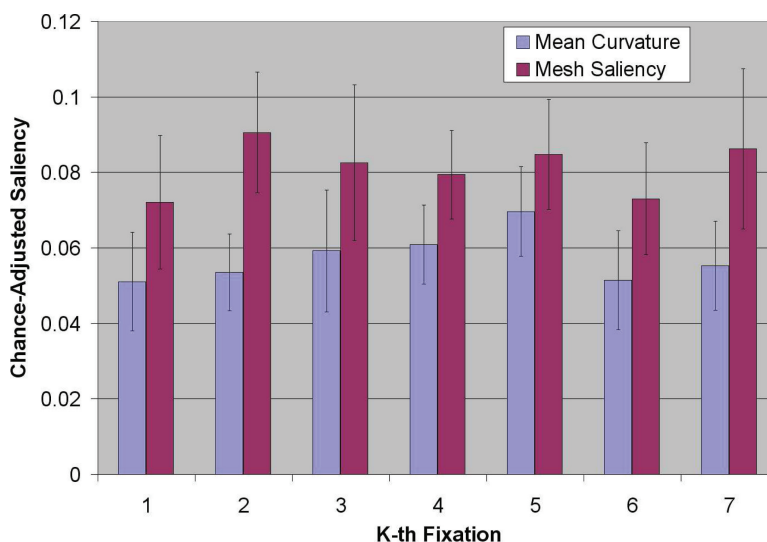


Fig. 6. Average chance-adjusted saliency values and 95% confidence interval using curvature and mesh saliency across all participants for each model. For all the cases, the values are higher than 0, which is the value that can be expected by chance. They exhibit high variances in computed chance-adjusted saliency values for all cases.

We first report the results of chance-adjusted saliency values and identify some shortcomings along the lines of what we have discussed in Section 4.2. We then present the results using our new normalized chance-adjusted saliency values. Figure 6 shows the average of the chance-adjusted saliency values across all participants. As in Manna et al. [1998] and Oliva et al. [2003], the study is restricted to the first seven fixations. In general, we observe that the curvature-based model and mesh saliency model have a higher correlation with human eye fixations than a random model as the chance-adjusted saliency values are higher than 0, the value that can be expected purely by chance. As noted in Section 4.2, one of the problems in chance-adjusted saliency is that it can be overvalued if random points on background pixels are included in the process of computing the mean random saliency. This overvaluation causes high variances in computed chance-adjusted saliency for all k -th fixation points ($1 \leq k \leq 7$) in Figure 6. We have performed a two-way ANOVA on the chance-adjusted saliency values with two variables: k -th fixation points ($1 \leq k \leq 7$) and different saliency models. For k -th fixation points, there is no significant difference ($F(6, 238) = 1.066, p = 0.3831, \eta_p^2 = 0.026$). However, for saliency models (curvature model and mesh saliency model), we observed significant differences ($F(1, 238) = 34.70, p < 0.001, \eta_p^2 = 0.127$). There was no interaction between two variables ($F(6, 238) = 0.494, p = 0.813, \eta_p^2 = 0.012$).

We further analyze the results of our normalized chance-adjusted saliency values. Figure 7 shows the average normalized chance-adjusted saliency values computed by the curvature model and the mesh saliency model across all participants. In general, we have observed that both computational models of saliency have higher correlation with human eye fixations than a random model as the normalized chance-adjusted saliency values are higher than 1, the value that can be expected purely by chance. A normalized chance-adjusted saliency score for a computational model of mesh saliency indicates how much better that model is compared to what we can expect purely by chance. For instance, a normalized chance-adjusted saliency score of 2 means that the model of saliency is twice as good as what one can expect purely by chance.

We observe lower variances in normalized chance-adjusted saliency than chance-adjusted saliency cases. To observe the effects of saliency models and 3D models, we carried out a two-way ANOVA on

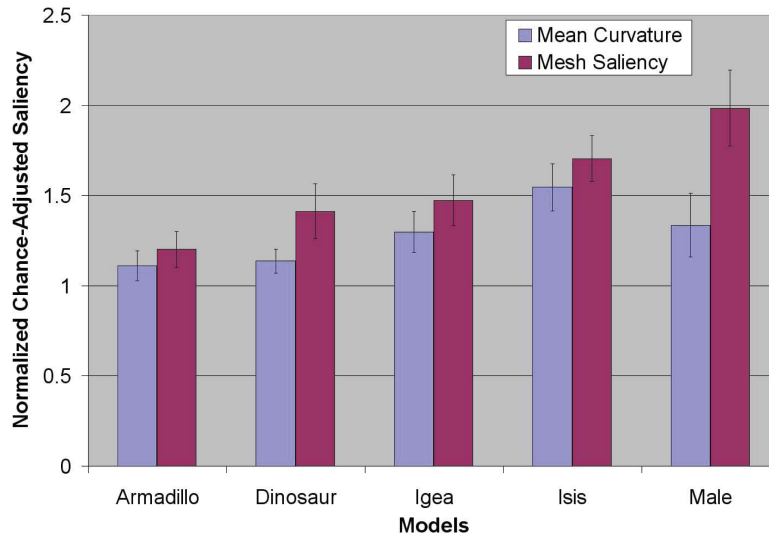


Fig. 7. Average normalized chance-adjusted saliency values and 95% confidence interval using curvature and mesh saliency across all participants for each model. For all the cases, the values are higher than 1, which is the value that can be expected by chance. A normalized chance-adjusted saliency score for a computational model of mesh saliency indicates how much better that model is compared to what we can expect purely by chance. For instance, a normalized chance-adjusted saliency score of 2 means that the model of saliency is twice as good as what one can expect purely by chance.

Table I. The Pairwise t -tests (two-tailed) between Curvature-Based and Mesh Saliency Models

Model	t -Value	p -Value
Armadillo	-2.643	0.017
Dinosaur	-4.122	< 0.001
Igea	-4.546	< 0.001
Isis	-3.786	0.001
Male	-7.489	< 0.001

the normalized chance-adjusted saliency values with two variables: different saliency models and different 3D models. We have found that there is a main effect of saliency models ($F(1, 170) = 37.75$, $p < 0.001$, $\eta_p^2 = 0.182$). This indicates that the mesh saliency model exhibits higher correlation with human eye fixations than the curvature-based model. A significant main effect was also obtained for 3D models ($F(4, 170) = 19.86$, $p < 0.001$, $\eta_p^2 = 0.319$). The result indicates that there are certain models that correlate better with human eye fixations than others. The possible cause is the difference in numbers of highly salient regions inherent in models. For instance, the Armadillo model has a large number of highly salient regions distributed on the model while the Male model has a small number of highly salient regions, as one can observe in Figure 5. We have also found there is a strong interaction between two variables ($F(4, 170) = 5.093$, $p = 0.01$, $\eta_p^2 = 0.107$), meaning that the effects of saliency models depend on 3D models under consideration. Next, we perform a pairwise t -test on the normalized chance-adjusted saliency values between two saliency models (curvature-based and mesh saliency). Table I shows the result for each of the models. We found a significant difference between two saliency models in the normalized chance-adjusted saliency values for each of the 3D models. Even




Fig. 8. The computed mesh saliency at scales of 2ϵ , 3ϵ , 4ϵ , 5ϵ , and 6ϵ . Here, ϵ is 0.3% of the length of the diagonal of the bounding box of the model.

with the Bonferroni correction, we found a borderline statistical significance for the Armadillo model and statistically significant results for all other models.

The results validate that the mesh saliency model has significantly higher correlation with human eye fixations than a random model and a curvature-based model.

5. DISCUSSION AND FUTURE WORK

We have used a few devices to reduce the effect of semantics in this article. (i) We did not give any tasks to the users when viewing images. (ii) We limited the time of a stimulus to the first 5 seconds. Others [Parkhurst et al. 2002; Santella and DeCarlo 2004] have also used similar durations. However, 5 seconds could be considered too long since semantic interpretation starts increasing right after the stimulus onset. We plan to further study the effect of semantics as we vary duration. Another thing we can do to reduce the effect of semantics is to experiment with objects that do not carry semantic information for most users (such as molecular models) or close-up views of scanned models.

Mesh saliency is sensitive to scale. Identifying the appropriate scales and their relative importance is another valuable area for future research. The previous mesh saliency approaches use five different scales (2ϵ , 3ϵ , 4ϵ , 5ϵ , and 6ϵ), where ϵ is 0.3% of the length of the diagonal of the bounding box of the model. Figure 8 shows the computed mesh saliency at these scales for the Dinosaur model. To understand the effect of scales on the normalized chance-adjusted saliency, we computed the average normalized chance-adjusted saliency values at five scales and at the nonlinear normalized sum across all participants for each model, as shown in Figure 9. We carried out two-way ANOVAs on the normalized chance-adjusted saliency values with two variables: different saliency models (mesh saliency model at a scale σ and curvature-based model) and different 3D models by varying σ from 2ϵ to 6ϵ . Table II shows that there were always significant differences found between mesh saliency model at each scale σ and curvature-based model when we varied σ from 2ϵ to 6ϵ . This may be because the center-surround operator plays an important role in modeling human eye fixations in the computational model of mesh saliency that the curvature-based model is lacking. It will be interesting to further study the effect of scales on the correlation between mesh saliency model and human eye movements in the future.

We currently compute mesh saliency in a view-independent way. However, there is some evidence in our study [Kim et al. 2008] that the correlation between eye fixations and mesh saliency is view dependent. We will further study the implications of view-dependent variables, such as illumination and viewing angles in the future. In this context, it will be interesting to compare and contrast the mesh saliency model to 2D image saliency models.

A natural question that arises when evaluating mesh saliency and human eye fixations is how well mesh saliency performs as compared to the image saliency approaches. To address this, we decided to examine the correlations of the mesh saliency and image saliency approaches with the human eye movements. We compute image saliency using Itti et al's method [1998]. Figure 10(b) shows the image saliency for the Armadillo model.

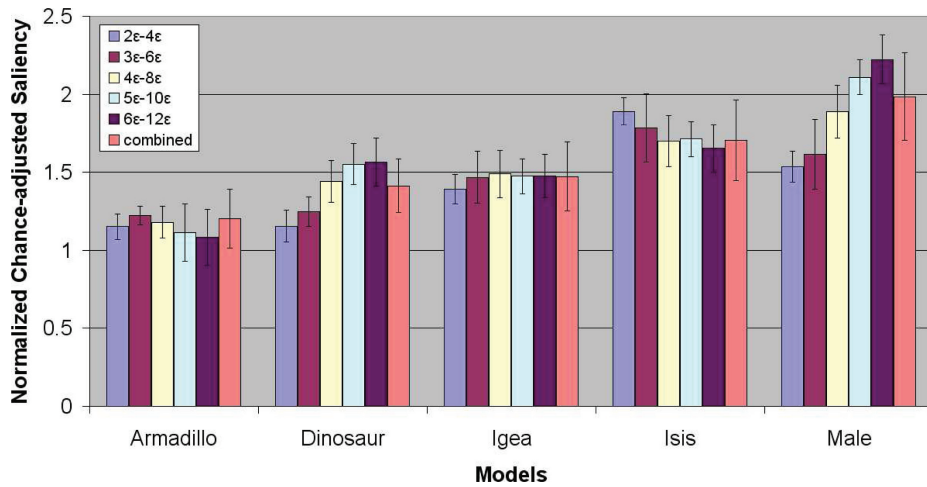


Fig. 9. Average normalized chance-adjusted saliency values and 95% confidence interval at five different scales and at the nonlinear normalized sum across all participants for each model.

Table II. The ANOVA Tests

Scales	<i>F</i> -Value	<i>p</i> -Value
2ε – 4ε	12.306	< 0.001
3ε – 6ε	21.247	< 0.001
4ε – 8ε	34.448	< 0.001
5ε – 10ε	42.686	< 0.001
6ε – 12ε	42.680	< 0.001

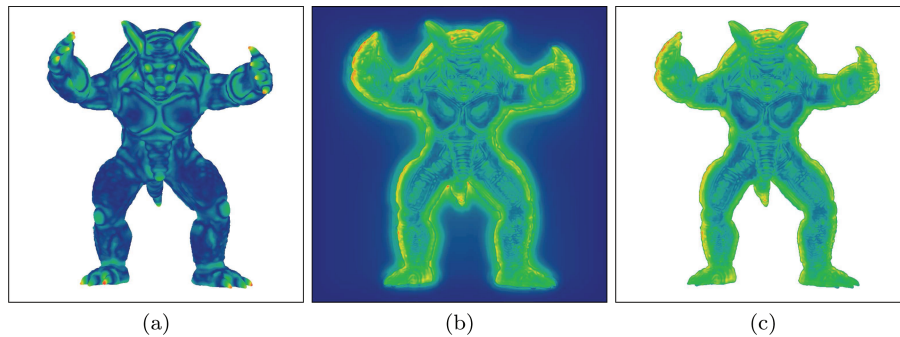


Fig. 10. Image (a) shows the computed mesh saliency by Lee et al.'s method, while Image (b) shows the computed image saliency by Itti et al.'s method. Image (c) only shows the image saliency values on the foreground pixels covered by the rendered Armadillo.

Following the discussion in Section 4.2, we consider image saliency values and fixation points only on foreground pixels, as shown in Figure 10(c). Figure 11 shows the results of the normalized chance-adjusted saliency values for mesh saliency and image saliency when we only consider the saliency values and the fixation points on foreground pixels. The results of normalized chance-adjusted saliency appear in Figure 11. Itti et al.'s model of image saliency was designed for human viewing of natural scenes. We wondered whether it would work well for viewing of isolated models. Our study indicates that in fact it is not. The reason is that Itti et al.'s approach assigns a high saliency value to the pixels at the mesh boundary due to a significant contrast with the background, but human observers do not attend

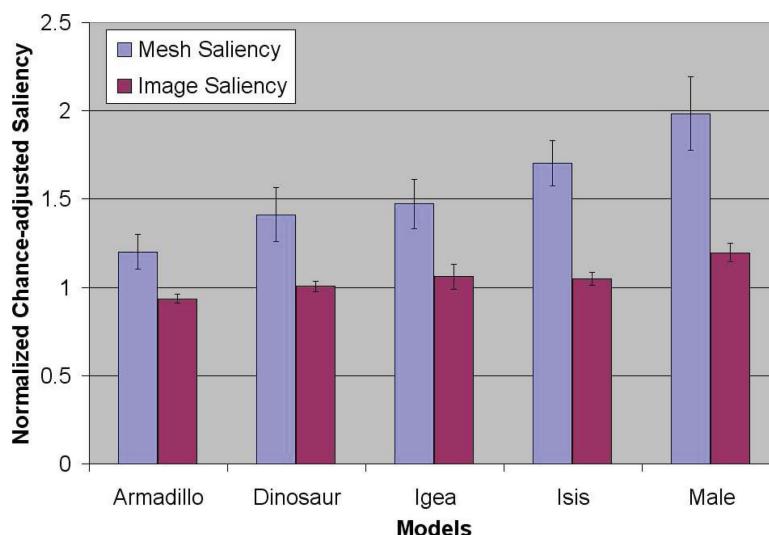


Fig. 11. Average normalized chance-adjusted saliency values and 95% confidence interval using mesh saliency and image saliency across all participants for each model. Due to the boundary effects and unfair removal of background pixels as shown in Figure 10, the model of image saliency makes poor performance. This suggests that a new measure should be introduced in the future for fair comparison between mesh saliency and image saliency.

to this particular case of contrast. It will be useful to design a more comprehensive study considering the broad aspects of image saliency and mesh saliency to gain a better understanding of how the human visual system directs attention based on 2D and 3D attributes of the scene.

6. CONCLUSIONS

In this article, we have taken the first steps towards validating an existing model of mesh saliency through an eye-tracking-based user study. We have introduced the notion of normalized chance-adjusted saliency, which is a robust measure of success of a mesh saliency model. We have observed significant correlations between the model of mesh saliency and human eye fixations. We believe that our carefully designed user study can be useful for designing a better visual saliency model, which is closer to human eye movements. This will also enable us to build further saliency-based systems for tasks, such as visual enhancement.

ACKNOWLEDGMENTS

We would like to thank Stanford 3D Scanning Repository and Cyberware for 3D models.

REFERENCES

- COLE, F., DECARLO, D., FINKELSTEIN, A., KIN, K., MORLEY, K., AND SANTELLA, A. 2006. Directing gaze in 3D models with stylized focus. In *Proceedings of the Eurographics Workshop/Symposium on Rendering*. Springer, Berlin, 377–387.
- ENNS, J. AND RENSINK, R. 1990. Sensitivity to 3D orientation in visual search. *Psychological Sci.* 1, 5, 323–326.
- FEIXAS, M., SBERT, M., AND GONZALEZ, F. 2008. A unified information-theoretic framework for viewpoint selection and mesh saliency. *ACM Trans. Appl. Percept.*
- GAL, R. AND COHEN-OR, D. 2006. Salient geometric features for partial shape matching and similarity. *ACM Trans. Graph.* 25, 1, 130–150.
- GOLDFEATHER, J. AND INTERRANTE, V. 2004. A novel cubic-order algorithm for approximating principal direction vectors. *ACM Trans. Graph.* 23, 1, 45–63.

- HENDERSON, J. M. AND HOLLINGWORTH, A. 1998. Eye movements during scene viewing: An overview. In *Eye Guidance in Reading and Scene Perception*. Elsevier Science Ltd., Oxford, UK.
- HOWLETT, S. AND O'SULLIVAN, C. 2005. Predicting and evaluating saliency for simplified polygonal models. *ACM Trans. Appl. Percept.* 2, 3, 286–308.
- ITTI, L., KOCH, C., AND NIEBUR, E. 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Patt. Anal. Mach. Intell.* 20, 11, 1254–1259.
- KIM, Y. AND VARSHNEY, A. 2006. Saliency-guided enhancement for volume visualization. *IEEE Trans. Visual. Comput. Graph.* 12, 5, 925–932.
- KIM, Y., VARSHNEY, A., JACOBS, D. W., AND GUIMBRETÈRE, F. 2008. A user study on relationship among mesh saliency, human eye fixations, and view dependence. Tech. rep. CAR-TR-1026, CS-TR-4913, and UMIACS-TR-2008-11. Center for Automation Research, Department of Computer Science, and UMIACS, University of Maryland, College Park, MD.
- KOCH, C. AND ULLMAN, S. 1985. Shifts in selective visual attention: towards the underlying neural circuitry. *Hum. Neurobiol.* 4, 219–227.
- LEE, C. H., VARSHNEY, A., AND JACOBS, D. 2005. Mesh saliency. *ACM Trans. Graph.* 24, 3, 659 – 666.
- LU, A., MACIEJEWSKI, R., AND EBERT, D. 2006. Volume composition using eye-tracking data. In *Proceedings of the Eurographics/IEEE VGTC Symposium on Visualization*. IEEE, Los Alamitos, CA, 147–154.
- MANNA, S., RUDDOCK, K. H., AND WOODING, D. S. 1998. Automatic control of saccadic eye movements made in visual inspection of briefly presented 2D images. *Spatial Vis.* 9, 3, 363.
- MEYER, M., DESBRUN, M., SCHR, P., AND BARR, A. H. 2003. Discrete differential-geometry operators for triangulated 2-manifolds. *Visual. Math.* III, 35–57.
- OLIVA, A., TORRALBA, A., CASTELHANO, M. S., AND HENDERSON, J. M. 2003. Top-down control of visual attention in object detection. In *Proceedings of the International Conference on Image Processing*. IEEE, Los Alamitos, CA, 253–256.
- PALMER, S. 1999. *Vision Science: Photons to Phenomenology*. MIT Press, Cambridge, MA.
- PARKHURST, D., LAW, K., AND NIEBUR, E. 2002. Modeling the role of salience in the allocation of overt visual attention. *Vision Res.* 42, 1, 107–123.
- PRIVITERA, C. M. AND STARK, L. W. 2000. Algorithms for defining visual regions-of-interest: Comparison with eye fixations. *IEEE Trans. Patt. Anal. Mach. Intell.* 22, 9, 970–982.
- SANTELLA, A. AND DECARLO, D. 2004. Visual interest and NPR: An evaluation and manifesto. In *Proceedings of the International Symposium on Non-Photorealistic Animation and Rendering*. ACM, New York, 71–150.
- SHILANE, P. AND FUNKHOUSER, T. A. 2007. Distinctive regions of 3D surfaces. *ACM Trans. Graph.* 26, 2, 7.
- STAMPE, D. 1993. Heuristic filtering and reliable calibration methods for video-based pupil tracking systems. *Behav. Res. Methods Instrum. Comput.* 25, 137–142.
- TAUBIN, G. 1995. Estimating the tensor of curvature of a surface from a polyhedral approximation. In *Proceedings of IEEE International Conference on Computer Vision*. IEEE, Los Alamitos, CA, 902–907.
- WATANABE, K. AND BELYAEV, A. G. 2001. Detection of salient curvature features on polygonal surfaces. *Comput. Graph. Forum* 20, 3, 385–392.
- YANG, Y.-L., LAI, Y.-K., HU, S.-M., AND POTTMANN, H. 2006. Robust principal curvatures on multiple scales. In *Proceedings of the Eurographics Symposium on Geometry Processing*. 223–226.

Received May 2007; revised August 2008; accepted January 2009