Viewpoint-Based Ambient Occlusion

Francisco González, Mateu Sbert, and Miquel Feixas ▪ *University of Girona*

Ambient occlusion is a powerful technique that mimics indirect global illumination at a fraction of its cost.[1] Researchers introduced obscurances, the first ambient-occlusion technique, in the computer-game context to allow fast editing, and later used it in production rendering.[2] To capture a good view of an object, we must first define what a "good view" means. Researchers have proposed several measures to this end, from heuristical measures to information-theoretic-based approaches. Computer graphics researchers have applied viewpoint quality measures in areas such as scene understanding,[3,4] scene exploration,[5] and volume visualization.[6,7] However, existing approaches don't exploit all the information that we can infer from captured data, such as how the object's polygon "sees" the viewpoints.

We've developed a new ambient occlusion technique based on an information-theoretic framework.[8] Essentially, our method computes a weighted visibility from each object polygon to all viewpoints; we then use these visibility values to obtain the information associated with each polygon. So, just as a viewpoint has information about the model's polygons, the polygons gather information on the viewpoints. We therefore have two measures associated with an information channel defined by the set of viewpoints as input and the object's polygons as output, or vice versa. From this polygonal information, we obtain an occlusion map that serves as a classic ambient occlusion technique. Our approach also offers additional applications, including an importance-based viewpoint-selection guide, and a means of enhancing object features and producing nonphotorealistic object visualizations.

A new ambient occlusion technique builds a channel between various viewpoints and an object's polygons, providing the information needed to create an occlusion map with multiple application possibilities.

Overview: Basic concepts

To set a context for our work, we first review some basic concepts of information theory, viewpoint selection, and ambient occlusion.

Information theory

To introduce the viewpoint channel and an object's associated information, we present here the two most basic measures of information theory: entropy and mutual information.

Let $\mathcal{X}$ be a finite set and a random variable, taking values x in $\mathcal{X}$ with distribution $p(x) = Pr[X = x]$. Likewise, let Y be a random variable taking values y in $\mathcal{Y}$. We define the Shannon entropy $H(X)$ of a random variable X by

$$H(X) = -\sum_{x \in \mathcal{X}} p(x) \log p(x) \qquad (1)$$

Logarithms are base 2, and we express entropy in bits. We use the $0\log 0 = 0$ convention and define conditional entropy by

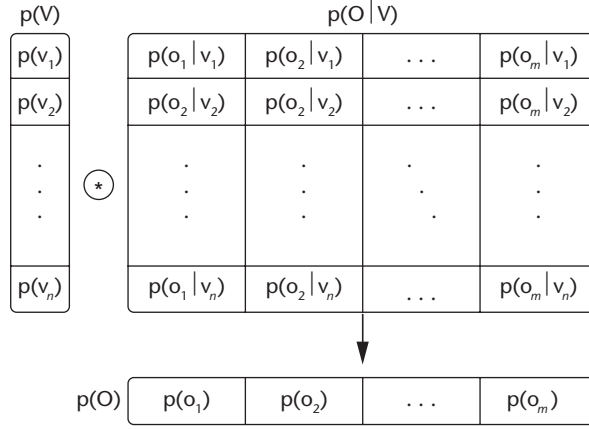$$H(Y \mid X) = -\sum_{x \in \mathcal{X}} p(x) \sum_{y \in \mathcal{Y}} p(y \mid x) \log p(y \mid x) \qquad (2)$$

where $p(y|x) = Pr[Y = y|X = x]$ is the conditional probability. The conditional entropy $H(Y|X)$ measures the average uncertainty associated with Y if we know X's outcome.

We define the *mutual information* (MI) between X and Y by

Published by the IEEE Computer Society

0272-1716/08/$25.00 © 2008 IEEE

Figure 1. Viewpoint information channel. (a) A sphere of viewpoints of an object. (b) Probability distributions of channel V → O.

The table in Figure 1b:

$p(V)$ column paired (via $*$) with $p(O|V)$ matrix, producing $p(O)$:

$p(V)$		$p(O\|V)$			
$p(v_1)$		$p(o_1\|v_1)$	$p(o_2\|v_1)$	$\ldots$	$p(o_m\|v_1)$
$p(v_2)$		$p(o_2\|v_2)$	$p(o_2\|v_2)$	$\ldots$	$p(o_m\|v_2)$
$\vdots$	$*$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$p(v_n)$		$p(o_1\|v_n)$	$p(o_2\|v_n)$	$\ldots$	$p(o_m\|v_n)$

$p(O)$	$p(o_1)$	$p(o_2)$	$\ldots$	$p(o_m)$

(a)　　　　　　　**(b)**

$$I(X,Y)=H(X)-H(X\,|\,Y)$$
$$=\sum_{x\in\mathcal{X}}p(x)\sum_{y\in\mathcal{Y}}p(y\,|\,x)\log\frac{p(y\,|\,x)}{p(y)} \tag{3}$$

The MI is a measure of the degree of dependence between X and Y, and $I(X,Y) = I(Y, X) \geq 0$.

Viewpoint channel

Researchers apply best-view selection algorithms to computer graphics domains such as scene understanding and virtual exploration[3-5] and volume visualization.[6,7]

To select the most representative or relevant object views, we define a viewpoint quality measure—the viewpoint MI (VMI)—from an information channel, $V \rightarrow O$.[9] The channel runs between the random variables V (input) and O (output), which represent an object's set of viewpoints $\mathcal{V}$ and polygons $\mathcal{O}$, respectively (see Figure 1a). We index these viewpoints by v and the polygons by o. Throughout this article, we use V and O as arguments of $p()$ to denote probability distributions. For example, while $p(v)$ denotes the probability of a single viewpoint v, $p(V)$ denotes the input distribution of the set of viewpoints. As Figure 1b shows, we characterize this information channel using both an input distribution and a probability transition matrix as follows:

- Input distribution $p(V)$ represents the probability of selecting each viewpoint; we obtain it by normalizing the object's projected area over each viewpoint. We can interpret a component of this distribution, $p(v)$, as viewpoint v's importance.
- We indicate conditional probability $p(o|v) = a_o/a_t$ by polygon o's normalized projected area over the directions sphere centered at viewpoint v, where a_o is the projected area of polygon o and a_t is the total area projected by the object. Conditional probabilities fulfill that $\Sigma_{o\in\mathcal{O}}\, p(o|v) = 1$.

- From $p(V)$ and $p(o|V)$, output distribution $p(O)$ is given by

$$p(o)=\sum_{v\in\mathcal{V}}p(v)p(o\,|\,v) \tag{4}$$

which is each polygon's average projected area.

From formula 3, the MI between V and O is given by

$$I(V,O)=\sum_{v\in\mathcal{V}}p(v)\sum_{o\in\mathcal{O}}p(o\,|\,v)\log\frac{p(o\,|\,v)}{p(o)}$$
$$=\sum_{v\in\mathcal{V}}p(v)I(v,O) \tag{5}$$

where

$$I(v,O)=\sum_{o\in\mathcal{O}}p(o\,|\,v)\log\frac{p(o\,|\,v)}{p(o)} \tag{6}$$

is defined as the VMI. This VMI gives us viewpoint v's quality and measures the dependence between viewpoint v and the polygons set. In this framework, we define the best viewpoint as the one with minimum VMI. High VMI values mean a high dependence between viewpoint v and the object, indicating a highly coupled view—such as that between the viewpoint and a few polygons with low average visibility. Low values correspond to more representative or relevant views, showing the maximum possible number of polygons in a balanced way. Research has shown that one of VMI's main properties is its robustness in dealing with any type of discretization or resolution of a volumetric data set.[7] We've observed the same advantage for polygonal data.[9]

Ambient occlusion

Ambient occlusion[1] is a simplified version of the obscurances illumination model.[2] Obscurances decouple direct and indirect illumination; developers
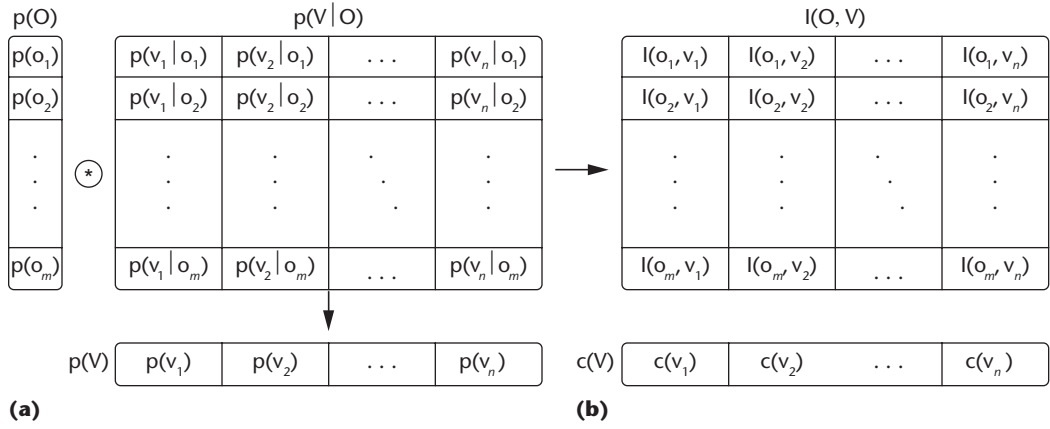
Figure 2. Defining a polygon's associated information. The channel is now inverted, so O is the input and V the output. The figure shows (a) the probability distributions of channel O → V and (b) the elements of matrix I(O,V) and the color distribution c(V) assigned to the viewpoint sphere.

introduced them to facilitate quick indirect illumination editing in the videogame context. Because the technique offers high-quality shadowing, it replaces radiosity in production. In the obscurances model, obscurance W is given by

$$W(x) = \frac{1}{\pi} \int_\Omega \rho(d(x,\omega)) \cos\theta d\omega \tag{7}$$

where ρ is a function of the distance $d(x, \omega)$ of the first intersection of a ray shot from x with direction ω. In this case, x is a surface point, θ is the angle between the normal vector at x and direction ω, and the integration is over the hemisphere, oriented according to the surface normal.

Ambient occlusion A is given by

$$A(x) = \frac{1}{\pi} \int_\Omega V(x,\omega) \cos\theta d\omega \tag{8}$$

where function ρ in the obscurances expression (formula 7) is substituted with the visibility function $V(x, \omega)$. When no geometry is visible in direction ω, the visibility function is 0; otherwise it has a value of 1.

To augment ambient occlusion, Méndez and colleagues introduced color bleeding, updated the obscurances dynamically in the presence of moving objects, and dealt with the problem of important secondary reflectors.[10] Sattler and colleagues use the graphics processing unit to compute the visibility from the object's vertices to the hemispherical mesh's vertices.[11] They also utilize the visibility function's coherence to achieve interactive frame rates with deformable objects, using point light sources at the hemisphere's vertices to provide illumination.

Information-theoretic ambient occlusion

As we described earlier, we obtain each VMI's as-

sociated information by defining a channel between the viewpoint sphere and the object's polygons. We now define a polygon's associated information from the inverted channel $O \rightarrow V$, so that O is now the input and V the output (see Figure 2a).

From the Bayes theorem, $p(v|o) = p(v)p(o|v) = p(o)p(v|o)$, we can rewrite the MI (formula 5) as

$$I(O,V) = \sum_{o\in\mathcal{O}} p(o) \sum_{v\in\mathcal{V}} p(v\,|\,o) \log \frac{p(v\,|\,o)}{p(v)}$$
$$= \sum_{o\in\mathcal{O}} p(o) \sum_{v\in\mathcal{V}} I(o,v)$$
$$= \sum_{o\in\mathcal{O}} p(o) I(o,V) \tag{9}$$

where $I(o,v) = p(v|o) \log p(v|o)/p(v)$ is a matrix element of MI, and we define the polygonal MI (PMI) as

$$I(o,V) = \sum_{v\in\mathcal{V}} p(v|o) \log \frac{p(v\,|\,o)}{p(v)} \tag{10}$$

which represents the degree of correlation between the polygon o and the viewpoints set, and can be interpreted as polygon o's associated information. As with VMI, low PMI values correspond to polygons that "see" the maximum number of viewpoints in a balanced way, while high values indicate low visibility.

To compute PMI, we estimate the projected area of the object's polygons at each viewpoint. From these data, we obtain the probabilities distribution $p(O|v)$ and $p(V)$. Before projection, we assign a different color to each polygon. The number of pixels of a given color, divided by the total number of pixels the object projects, gives us the relative polygon area that this color represents. In our experiments, all objects are centered in a viewpoint sphere, which we build from an icosahedron's recursive discretiza-
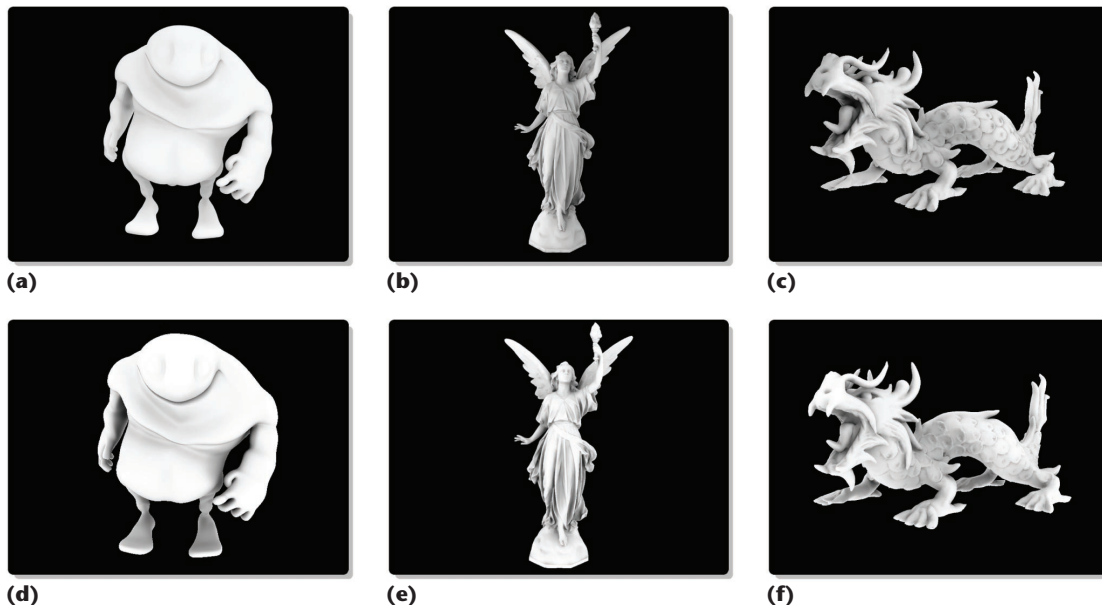
Figure 3. A comparison of viewpoint-based ambient occlusion maps (top row) with classic ambient occlusion maps (bottom row). The techniques were applied to three models: (a) the monster (46,400 polygons), (b) Lucy (65,000 polygons), and (c) the dragon (65,000 polygons).

tion, and the camera looks at the sphere's center. We obtain $p(O)$ from formula 4 and $p(V|o)$ from the Bayes theorem using $p(V)$ and $p(O|v)$.

Comparison with classical ambient occlusion maps

Figure 3 compares the results of applying our viewpoint-based ambient occlusion method (top row) with that of classic ambient occlusion maps (bottom row). We applied each technique over three models: the monster, which has 46,400 polygons (Figure 3a); Lucy, with 65,000 polygons (Figure 3b); and the dragon, which also has 65,000 polygons (Figure 3c). We used a 162-viewpoint sphere and a projection resolution of 1280×960. To obtain these images, we normalized the PMI of all polygons between 0 and 1. We subtracted from 1 because low PMI values (represented in the PMI map by values near 1) correspond to polygons that are nonoccluded (visible from many viewpoints), while high PMI values (represented in the PMI map by values near 0) correspond to occluded polygons. Once we've computed each polygon's associated information, we linearly interpolated the values at the model's vertexes.

Sattler and colleagues use a similar approach to compute ambient occlusion.[11] They compute a matrix, M_{ij}, as $n_i \cdot l_j$, where n_i is the normal to the object's vertex and l_i is the direction of a virtual light source placed at a bounding sphere. They use a number of virtual light sources to approximate ambient lighting. The final contribution to the vertex i is given by the sum for all visible $M_{ij}l_j$ light sources, where l_j is the source's intensity.

Applying classic ambient occlusion to the test models offers only a discrete set of possible occlusion values, because it's computed as a proportion of hits. As Figure 3 (bottom row) shows, this
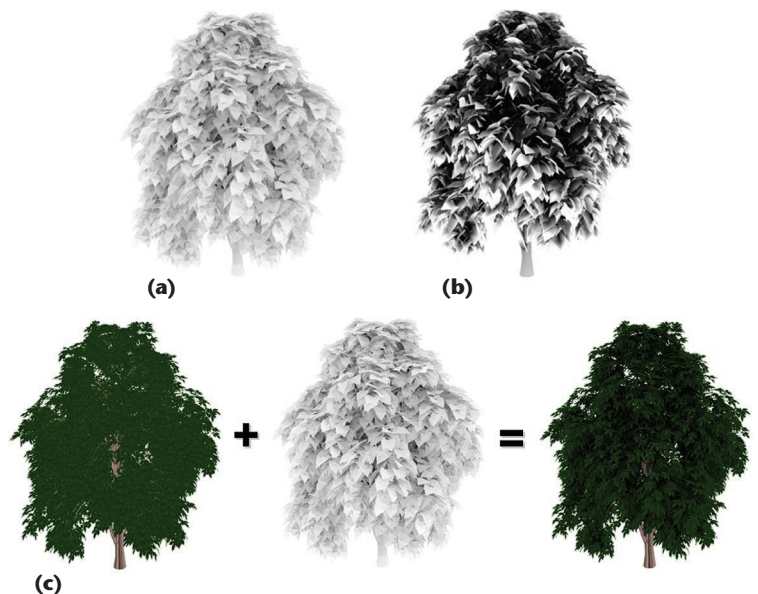

Figure 4. The Xfrog public plants tree model. (a) Applying viewpoint-based ambient occlusion offers smoother transitions than (b) classic ambient occlusion. (c) Composing the ambient occlusion map and the texture of the tree model.

results in transitions that are too sharp between the values. Comparing the tree images in Figures 4a and 4b offers an even more prominent illustration. Because our ambient occlusion technique accepts continuous values, it offers much smoother transitions. In Figure 4c, we compose our ambient occlusion with the tree model's texture to produce the final model rendering.

Calculation parameters

Our viewpoint-based ambient occlusion calculation depends on two parameters: the number of

Figure 5.
Viewpoint-
based ambient
occlusion maps.
We used three
projection
resolutions:
(top row)
640 × 480,
(middle row)
1280 × 960, and
(bottom row)
2560 × 1920.
We combined
these with
four viewpoint
levels: 12
(1st column),
42 (2nd
column), 162
(3rd column),
and 642 (4th
column).

Table 1. Computation time (in seconds) for Figure 5's polygonal mutual information maps.

Resolution	Number of viewpoints			
	12	42	162	642
640 × 480	6	8	16	53
1280 × 960	7	15	39	166
2560 × 1920	15	41	138	496

viewpoints and the resolution selected to project the object's polygons. Figure 5 shows how ambient occlusion quality changes when we modify the number of viewpoints and the projection resolution on the De Espona 3D encyclopedia's Lady of Elche model (51,980 polygons). We tested three resolutions (640 × 480, 1280 × 960, and 2560 × 1920) against various viewpoints (12, 42, 162, and 642). As the figure shows, the results get better as the resolution and number of viewpoints increase. Without a sufficient number of viewpoints, we can't properly capture visibility information, leading to sharper transitions between occlusion values (Figures 5a and 5b). However, projections with low resolution might miss small triangles and lose the corresponding visibility information; this makes triangles appear darker than they should. In Figure 5, this phenomenon appears as darker spherical bumps on the Lady of Elche's garments. As resolution increases, such artifacts disappear. Table 1 shows the time required to compute Figure 5's polygonal information maps. Using 162 viewpoints and 1280 × 960 resolution offers a good balance between computation time and quality.

Additional applications

In addition to using it as an ambient occlusion technique, we can use PMI to select viewpoints, assign viewpoint-importance weights, and relight nonphotorealistic renderings.

Viewpoint selection

Observing how much PMI is visible can help us select appropriate viewpoints. For each viewpoint, we add the PMI of all polygons weighted by how these polygons see the viewpoint—that is, the conditional probabilities $p(v|o)$ of the inverse channel. Thus, viewpoint v captures the following information:

$$PMI(v) = \sum_{o \in \mathcal{O}} p(v \mid o) I(o, V)$$

(11)

High $PMI(v)$ values mean that viewpoints can see a lot of PMI, while low values indicate low PMI visibility. Figure 6 shows the results for the Lady of Elche and Lucy models. Figures 6a and 6e and Figures 6b and 6f show the best and worst viewpoints, respectively. Both correspond to high and low PMIs, where high visibility shows significant detail and low visibility shows little detail. Figures 6c, 6d, 6g, and 6h show two views of the PMI sphere: warm colors represent high $PMI(v)$ values, while cold colors represent low values.

Viewpoint importance

If we modify the input distribution $p(V)$ according to each viewpoint's assigned importance, then (as defined by formula 10) PMI changes. Figure 7 shows the results on the Lady of Elche model, where
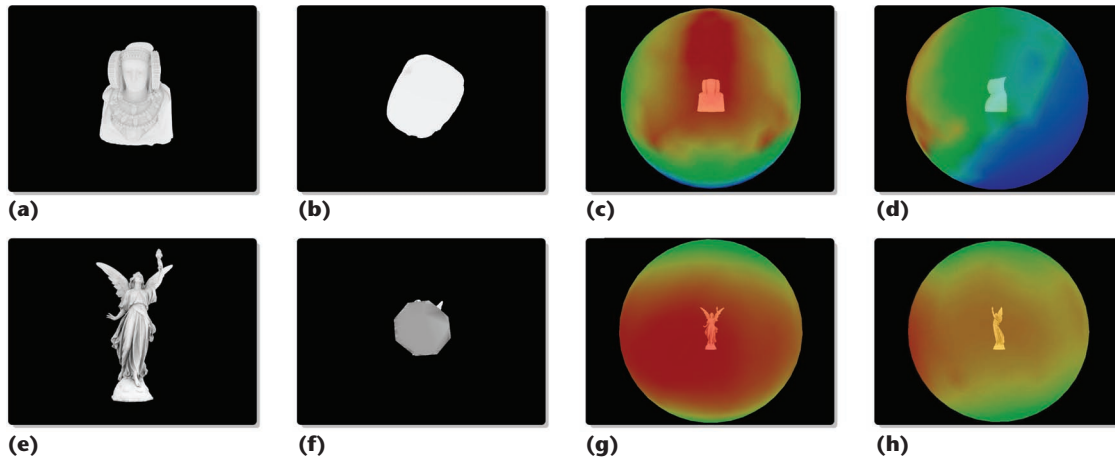
Figure 6. The effects of PMI visibility on two models: (top row) the Lady of Elche and (bottom row) Lucy. (a-b) The most and (e-f) least informative views of the models. Images (c-d) and (g-h) show two different projections of each model's PMI spheres.
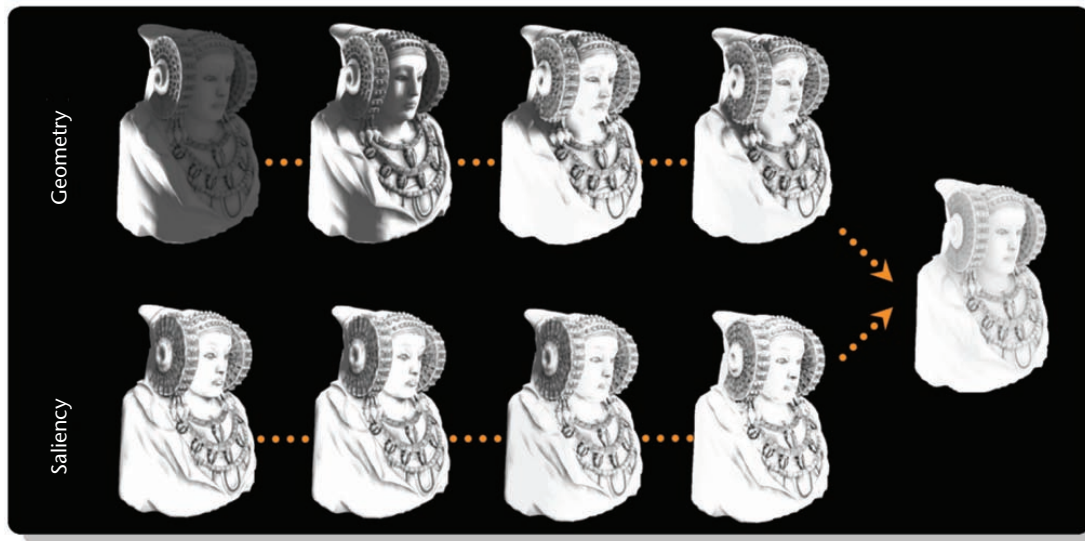

Figure 7. The affects of viewpoint importance on the Lady of Elche model. The left-most image in the series shows the effect of assigning importance to the best viewpoint. Moving to the right, the viewpoint's importance decreases in each image. The top row viewpoints are selected according to geometry; bottom row viewpoints are selected according to saliency. The final image assigns equal importance to all sphere viewpoints.

the most important viewpoints are the most brightly illuminated. The range of images includes

- assigning almost all importance to the best viewpoint in the first image,
- assigning equal importance to the two best viewpoints in the second image,
- assigning equal importance to the three best viewpoints in the third image, and
- assigning equal importance to the four best viewpoints in the fourth image.

We obtain the fifth and final image by assigning equal importance to all sphere viewpoints—the viewpoint-based ambient occlusion map. For each model, the top row shows images corresponding to viewpoints obtained from a best-view selection algorithm[9] and the bottom row shows images ob-

tained using the same algorithm driven by polygon saliency. As the figure shows, when the most important viewpoints are also the most salient ones, the images improve.
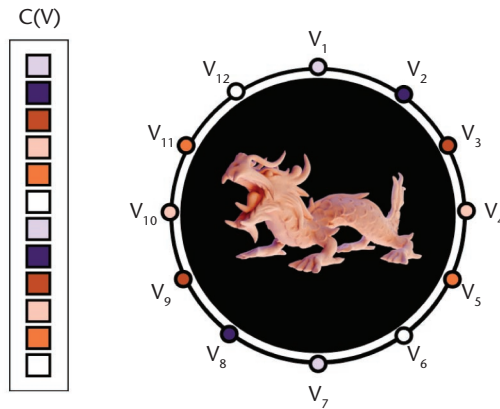
Relighting for nonphotorealistic rendering

In one sense, PMI using importance is a technique that places light sources at important points. Given this, we can investigate a PMI generalization that places colored light sources at key viewpoints. When all sources have the same intensity, the extension reverts to the original PMI.

We obtain color ambient occlusion from the scalar product of an $I(O,V)$ matrix row and a $c(V)$ color vector's complement:

$$I_\alpha(o,V) = \sum_{v \in V} I(o,v)(1 - c_\alpha(v))$$

$$(12)$$

Figure 8. A color ambient occlusion model. To compute the model, we warp a color texture to the viewpoint sphere.

where *a* stands for each color channel, $c_a(v)$ is channel *a*'s normalized vector, and $I(o,v)$ is $I(O,V)$'s matrix element (see Figure 2). After computing each channel's PMI, we derive the final color ambient occlusion by combining the channels.

To get a color vector, we can warp a color texture to the viewpoint sphere. We thereby assign a color to each viewpoint (see Figure 8). A comparison of Figure 8 and Figure 3c (top row) shows how color ambient occlusion improves model visualization. We can also combine this relighting technique using several Coloroid[12] palettes with a nonphotorealistic rendering technique[13] (see Figure 9). As the figure shows, combining these techniques enhances the NPR technique with the color ambient occlusion's indirect illumination.

Future work

In our future research, we plan to extend our ambient occlusion technique to deal with volumetric data sets to further investigate the relationship between ambient occlusion and saliency, and also to use generalized Tsallis-Havrda-Charvat mutual information. We're also planning a full GPU im-

plementation of our technique. A Web page with more information and additional examples of our work is available at www.gametools.org/iEEECG/index.html.

Acknowledgments

References

1. H. Landis, "Renderman in Production," *Course Notes of ACM Siggraph*, ACM Press, 2002, pp. 87-102.
2. S. Zhukov, A. Iones, and G. Kronin, "An Ambient Light Illumination Model," *Proc. Rendering Techniques Workshop*, Eurographics Assoc., 1998, pp. 45-56.
3. D. Plemenos and M. Benayada, "Intelligent Display Techniques in Scene Modelling (New Techniques to Automatically Compute Good Views)," *Proc. Int'l Conf. Computer Graphics and Vision* (Graphicon 96), 1996.
4. P. Vázquez et al., "Viewpoint Selection Using Viewpoint Entropy," *Proc. Vision, Modeling, and Visualization,* Aka GmbH, 2001, pp. 273-280.
5. D. Sokolov, D. Plemenos, and K. Tamine, "Methods and Data Structures for Virtual World Exploration," *The Visual Computer*, vol. 22, no. 7, 2006, pp. 506-516.
6. U.D. Bordoloi and H.W. Shen, "Viewpoint Evaluation for Volume Rendering," *Proc. IEEE Visualization*, IEEE CS Press, 2005, pp. 487-494.
7. I. Viola et al., "Importance-Driven Focus of Attention," *IEEE Trans. Visualization and Computer Graphics,* vol. 12, no. 5, 2006, pp. 933-940.
8. F. González, M. Sbert, and M. Feixas, "An Information-Theoretic Ambient Occlusion," *Proc. Int'l Symp. Computational Aesthetics in Graphics, Visualization, and Imaging*, Eurographics Assoc., 2007, pp. 29-36.

Figure 9. Combining viewpoint-based ambient occlusion with a nonphotorealistic technique. The examples use two Coloroid color palettes (left).

9. M. Feixas, M. Sbert, and F. González, "A Unified Information-Theoretic Framework for Viewpoint Selection and Mesh Saliency," to appear in *ACM Trans. Applied Perception*, 2008.
10. A. Méndez, M. Sbert, and J. Català, "Real-Time Obscurances with Color Bleeding," *Proc. 19th Spring Conf. Computer Graphics*, ACM Press, 2003, pp. 171-176.
11. M. Sattler et al., "Hardware-Accelerated Ambient Occlusion Computation," *Proc. Int'l Fall Workshop Vision, Modeling, and Visualization*, Aka GmbH, 2004, pp. 331-338.
12. A. Nemcsis, "Color Space of the Coloroid Color System," *Color Research and Applications,* vol. 12, no. 3, 1987, pp. 135-146.
13. A. Lake, "Cartoon Rendering Using Texture Mapping and Programmable Vertex Shaders," *Proc. Game Programming Gems 2*, Charles River Media, 2001, pp. 444-451.

Mateu Sbert *is a professor in computer science at the University of Girona, Spain. His research interests include application of Monte Carlo and information theory to computer graphics and image processing. He received a PhD in computer science from the Technical University of Catalonia, where he received the Best PhD award. Contact him at mateu@ima.udg.edu.*

Miquel Feixas *is associate professor in computer science at the University of Girona, Spain. His research interests include the application of information theory to computer graphics and image processing. Feixas received a PhD in computer science from the Technical University of Catalonia. Contact him at feixas@ima.udg.edu.*

Francisco González *is a PhD student at the University of Girona, Spain. His research interests include real-time computer graphics, realistic rendering, dynamic environments, and information theory applied to computer graphics. González received an MS in computer science and a Masters in Computing from the University of Girona. Contact him at gonzalez@ima.udg.edu.*

For further information on this or any other computing topic, please visit our Digital Library at http://www.computer.org/csdl.