

Viewpoint-based Ambient Occlusion

Francisco González, Mateu Sbert, and Miquel Feixas

Institut d'Informàtica i Aplicacions, Universitat de Girona, Spain

Abstract

Ambient occlusion is a powerful technique that mimics indirect global illumination at a fraction of the cost. In this paper, we introduce a new ambient occlusion technique based on information-theoretic concepts. A viewpoint quality measure is first defined using the concept of mutual information of the channel formed between a set of viewpoints and the polygons of an object. By reversing this channel we can speak of the mutual information of a polygon with respect to all viewpoints. From this polygonal information, we represent a kind of ambient occlusion which can be used to guide the viewpoint selection by preferring the most informative views of the object. The polygonal information is dependent on the importance assigned to each viewpoint and helps us to enhance features such as the most salient parts. Further, the assignation of color to each viewpoint combined with the polygonal information produces a nice visualization of the object. Examples are given with coloroid palettes and non-photorealistic rendering.

Paper Web page can be found at <http://www.gametools.org/iEEECG/index.html>

1. Introduction

Ambient occlusion [Lan02] is a family of techniques that allow to imitate the indirect illumination part of global illumination with a very cheap cost. Obscurances [ZIK98] was the first of such techniques, introduced in the context of computer games for fast editing purposes and then used in production rendering. Indirect illumination is decoupled from direct one, and some additional features include, for instance, color bleeding for further realism [MSC03]. In this paper we define a new ambient occlusion technique based on a viewpoint information-theoretic framework. In this approach, each polygon of an object shares information with the set of visible viewpoints. This shared (or mutual) information is a descriptor of the quality of that visibility and provides a natural ambient occlusion value.

In computer graphics, several *viewpoint quality* measures have been applied in areas such as scene understanding [PB96, VFSH01], scene exploration [SPT06], and volume visualization [BS05, VFSG06]. In [FSG07], a new viewpoint quality measure based on mutual information has been introduced from an information channel constructed between the set of viewpoints and the polygons of the object (Section 2). A non-polygonal version of this measure

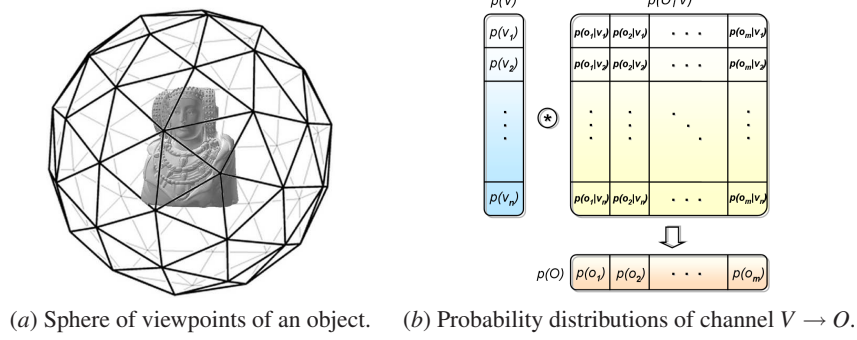
has also been applied to select the best views in volume visualization [VFSG06]. From the reversion of the viewpoint channel, the information associated with each polygon, which we call *polygonal mutual information*, is defined and used to obtain a new ambient occlusion (Section 3) technique. This information is used to select the most detailed and informative views, which correspond to the views seeing the maximum polygonal mutual information. Then, we show how the importance assigned to each viewpoint helps us to enhance characteristics such as the most salient parts and how the polygonal mutual information can be combined with the color assigned to each viewpoint to produce non-photorealistic visualizations (Section 4).

2. Background and Related Work

In this section we review some basic concepts of information theory, viewpoint selection and ambient occlusion.

2.1. Information-Theoretic Concepts

Let \mathcal{X} be a finite set, let X be a random variable taking values x in \mathcal{X} with distribution $p(x) = Pr[X = x]$. Likewise,


Figure 1: Viewpoint information channel.

let Y be a random variable taking values y in \mathcal{Y} . An information channel between two random variables (input X and output Y) is characterized by a probability transition matrix (conditional probabilities) which determines the output distribution given the input. The *Shannon entropy* $H(X)$ of a random variable X is defined by

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

All logarithms are base 2 and entropy is expressed in bits. The convention that $0 \log 0 = 0$ is used. The *conditional entropy* is defined by

$$H(Y|X) = - \sum_{x \in \mathcal{X}} p(x) \sum_{y \in \mathcal{Y}} p(y|x) \log p(y|x), \quad (2)$$

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 the outcome of X .

The *mutual information* (MI) between X and Y is defined by

$$\begin{aligned} I(X, Y) &= H(X) - H(X|Y) \\ &= \sum_{x \in \mathcal{X}} p(x) \sum_{y \in \mathcal{Y}} p(y|x) \log \frac{p(y|x)}{p(y)}. \end{aligned} \quad (3)$$

It is a measure of the shared information between X and Y . It can be seen that $I(X, Y) = I(Y, X) \geq 0$.

2.2. Viewpoint Channel

Best view selection algorithms have been applied to computer graphics domains, such as scene understanding and virtual exploration [PB96, VFSH01, SPT06] and volume visualization [BS05, VFSG06].

To select the most representative or relevant views of an object, a viewpoint quality measure, the *viewpoint mutual information*, was defined [FSG07] from an information channel $V \rightarrow O$ between the random variables V (input) and O (output), which represent, respectively, a set of viewpoints and the set of polygons of an object (see Figure 1(a)). Viewpoints are indexed by v and polygons by o . Throughout this

paper, the capital letters V and O as arguments of $p(\cdot)$ are used to denote probability distributions. For instance, while $p(v)$ denotes the probability of a single viewpoint v , $p(V)$ denotes the input distribution of the set of viewpoints.

This information channel is characterized by a probability transition matrix (conditional probability matrix) which determines the output distribution given the input (see Figure 1(b)):

- Input distribution $p(V)$ represents the probability of selecting a viewpoint. As we will mention below, $p(v)$ can be interpreted as the importance of viewpoint v . In our experiments, $p(V)$ is obtained from the normalization of the projected area of the object over each viewpoint.
- Conditional probability $p(o|v) = \frac{a_o}{a_t}$ is given by the normalized projected area of polygon o over the sphere of directions centered at viewpoint v . Conditional probabilities fulfil that $\sum_{o \in \mathcal{O}} p(o|v) = 1$.
- From $p(V)$ and $p(o|V)$, the output distribution $p(O)$ is given by

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

which is the average projected area of each polygon.

From (3), *mutual information* between V and O , that expresses the degree of *dependence* or *correlation* between a set of viewpoints and the object, is given by

$$\begin{aligned} I(V, O) &= \sum_{v \in \mathcal{V}} p(v) \sum_{o \in \mathcal{O}} p(o|v) \log \frac{p(o|v)}{p(o)} \\ &= \sum_{v \in \mathcal{V}} p(v) I(v, O), \end{aligned} \quad (5)$$

where we define

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

as the *viewpoint mutual information* (VMI), which gives us the *quality* of viewpoint v and measures the degree of dependence between the viewpoint v and the set of polygons. In our framework, the best viewpoint is defined as the one

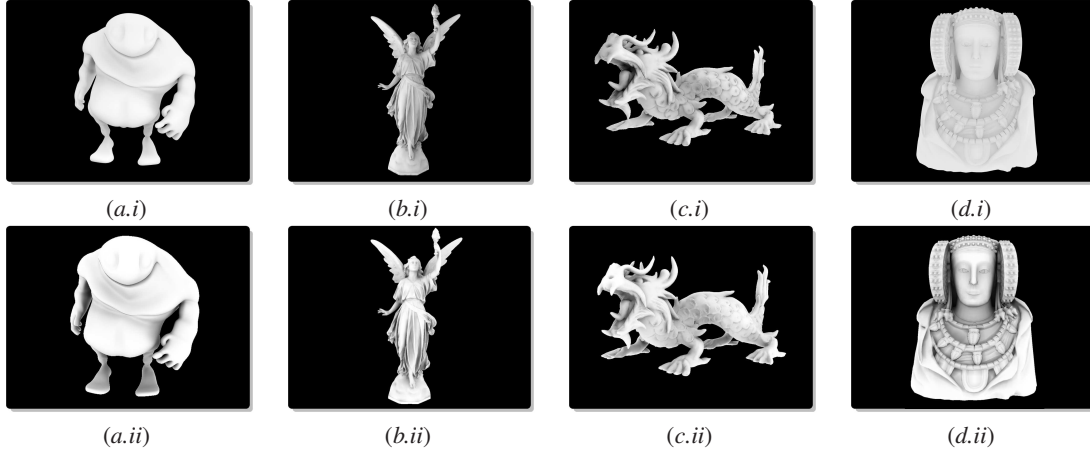


Figure 2: (i) Polygonal mutual information maps. (ii) Ambient occlusion maps.




Figure 3: Composition of our viewpoint-based ambient occlusion and the textures of the Ogre (left) and tree models (right).

that has *minimum* VMI. High values of the measure mean a high dependence between viewpoint v and the object, indicating a highly *coupled* view (for instance, between the viewpoint and a small number of polygons with low average visibility). On the other hand, low values correspond to more *representative* or *relevant* views, showing the maximum possible number of polygons in a balanced way. In [VFSG06], it has been shown that one of the main properties of VMI is its robustness to deal with any type of discretisation or resolution of a volume dataset. The same advantage can be observed for polygonal data [FSG07].

2.3. Ambient Occlusion

Ambient occlusion [Lan02] is a simplified version of the obscurances illumination model [ZIK98]. Obscurances decouple direct and indirect illumination and were first introduced in the videogame context as a technique to allow fast editing of indirect illumination. The high quality of shadowing obtained made them later to be included in production replacing radiosity. In the obscurances model, obscurance W is given by

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

where ρ is a function of the distance $d(x, \omega)$ of the first intersection of a ray shot from x with direction ω , 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

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

substitutes the ρ function in the obscurances (formula 7) by the visibility function $V(w, x)$ that has value zero when no geometry is visible in direction ω and one otherwise.

Méndez et al. [MSC03] introduced color bleeding, updated the obscurances dynamically in the presence of moving objects and dealt with the problem of important secondary reflectors. Sattler et al. [SSZK04] compute the visibility from the vertices of the object to the vertices of an hemispherical mesh using the GPU. They also utilize the coherence in the visibility function to achieve interactive frame rates with deformable objects with illumination coming from point light sources at the vertices of the hemisphere.

3. Information-Theoretic Ambient Occlusion

As we have seen in Section 2.2, the information associated with each viewpoint (VMI) is obtained from the definition of the channel between the sphere of viewpoints and the polygons of the object. Now, the *information associated with a polygon* will be defined from the reversed channel $O \rightarrow V$,



Figure 4: A panorama of a scene where our viewpoint-based ambient occlusion technique has been applied over the Lucy model.

so that O is now the input and V the output (see Figure 5). Note that MI is invariant to the reversion of the channel since $I(V, O) = I(O, V)$.

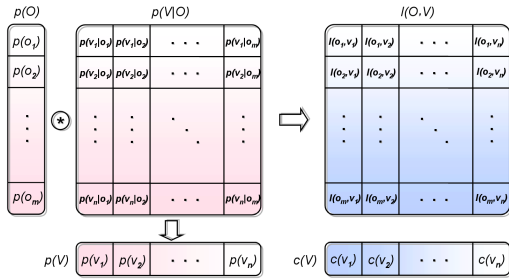


Figure 5: Probability distributions of channel $O \rightarrow V$, elements of matrix $I(O, V)$ and color distribution $c(V)$ assigned to the viewpoint sphere.

From the Bayes theorem $p(v, o) = p(v)p(o|v) = p(o)p(v|o)$, the mutual information (5) can be rewritten as

$$\begin{aligned} I(O, V) &= \sum_{o \in \mathcal{O}} p(o) \sum_{v \in \mathcal{V}} p(v|o) \log \frac{p(v|o)}{p(v)} \\ &= \sum_{o \in \mathcal{O}} p(o) \sum_{v \in \mathcal{V}} I(o, v) \\ &= \sum_{o \in \mathcal{O}} p(o) I(o, V), \end{aligned} \quad (9)$$

where $I(o, v) = p(v|o) \log \frac{p(v|o)}{p(v)}$ is a matrix element of MI, and the *polygonal mutual information* (PMI) is defined as

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

which represents the degree of correlation between the polygon o and the set of viewpoints, and can be interpreted as the *information* associated with polygon o . Analogous to VMI, low values of PMI correspond to polygons that 'see'

the maximum number of viewpoints in a balanced way. The opposite happens for high values.

To compute PMI, we estimate $p(o|v)$ from the projection of the visible polygons of the object on the screen. Before projection, a different color is assigned to each polygon. The number of pixels with a given color divided by the total number of pixels projected by the object or scene gives us the relative area of the polygon represented by this color. In our experiments, all the objects are centered in a sphere of 642 viewpoints[†], built from the recursive discretisation of an icosahedron, and the camera is looking at the center of this sphere. Using formula (4) and the Bayes theorem, $p(O)$ and $p(V|O)$ can be obtained from both $p(V)$ and $p(O|V)$.

In Figure 2(i) we show the information maps corresponding to the test models. To obtain these images, the PMI of all polygons has been normalized between 0 and 1 and subtracted from 1, because low values of PMI, represented in the grey-map by values near 1, correspond to non-occluded or visible (from many viewpoints) polygons, while high values of PMI, represented in the grey-map by values near 0, correspond to occluded polygons. Once we have computed the information associated to each polygon, these values have been linearly interpolated at the vertexes of the model. In Figure 2(ii) we show the results of applying classic ambient occlusion [Lan02]. Observe that these information maps look as an ambient occlusion or obscurity map (see Section 2.3). In [SSZK04], a similar approach was used to compute ambient occlusion. In this case a matrix M_{ij} is computed as $n_i \cdot l_j$, where n_i is the normal to the vertex of the object and l_j is the direction of a virtual light source placed at a bounding sphere. A number of virtual light sources is used to approximate ambient lighting. The final contribution to the vertex

[†] To obtain quality results would be in general enough to use a few hundred viewpoints.

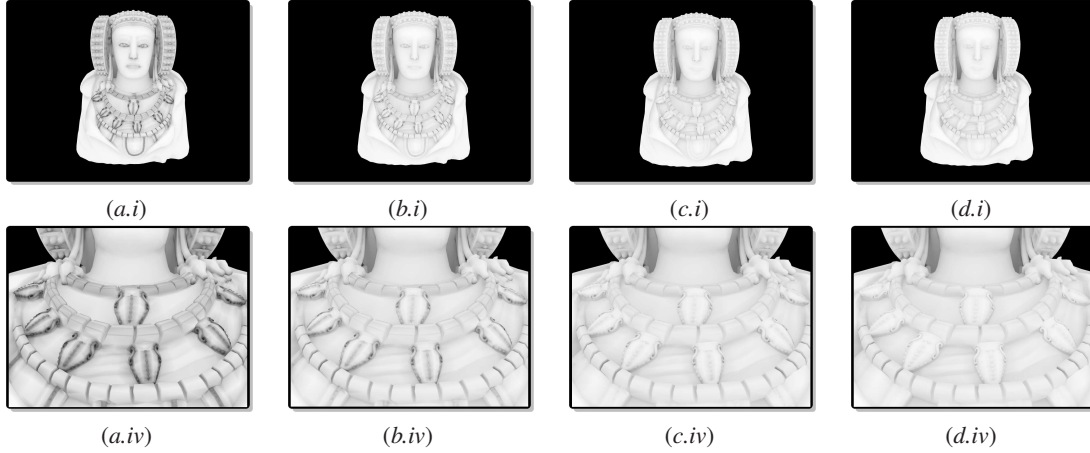


Figure 6: Several ambient occlusion models for the Lady of Elche. For each column we increase the projection resolution twice the row before, that is: (i) 640x480, (ii) 1280x960, (iii) 2560x1920 and (iv) 5120x3840. In the second row we show a close-up of each model of the first row.

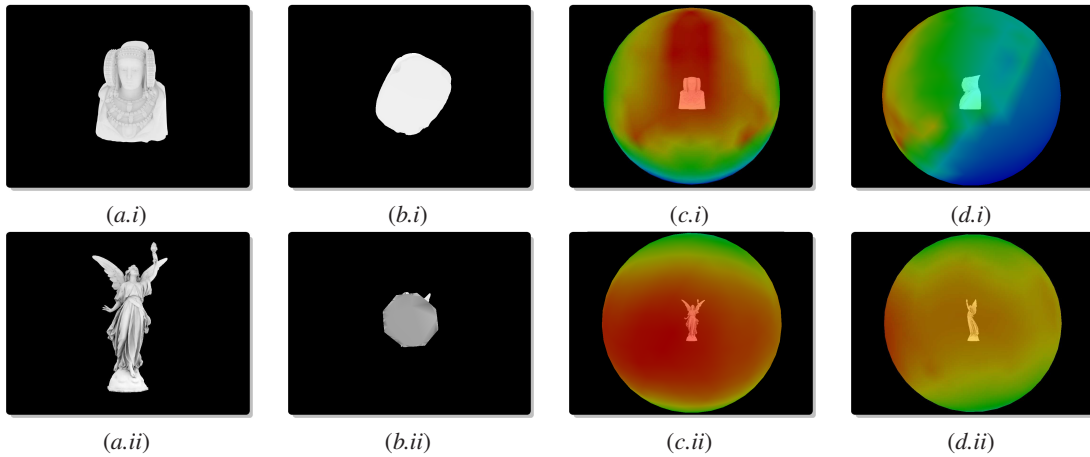


Figure 7: (a) The most and (b) least informative views of the Lady of Elche (i) and Lucy (ii) models. (c-d) Two different views of the PMI spheres.

i is given by the sum for all visible light sources of $M_{ij}I_j$, where I_j is the intensity of the source.

In Figures 2 we can observe our technique compared with classic ambient occlusion. In the latter case, there is only a discrete set of possible values, since it is computed as a proportion of hits. On the other hand, the quality of our ambient occlusion measure depends on the resolution used to project the polygons of the scene. If the resolution is not enough, the small triangles will not be projected correctly. This effect can be observed in the darker small spherical bumps of the Lady of Elche garments in Figure 6(a). From Figure 6, we see that these artifacts disappear when we increase the resolution of projection.

In Figure 3 we show several examples of the use of poly-

gonal information as ambient occlusion, where this is added to a textured model. In Figure 4 a frame of a video demo is shown, where our ambient occlusion technique has been computed over the Lucy model.

The models used in our examples come from Nvidia dynamic ambient occlusion demo (Figure 2(a)), Xfrog public plants (Figure 3) and De Espona 3D encyclopedia (Figure 2(d) and Figure 11).

4. Applications

As we have shown above, our polygonal mutual information can be used as an ambient occlusion technique. In this section, we show how this ambient occlusion can guide the viewpoint selection. PMI is also extended to enhance the




Figure 8: Effect of assigning importance to the best viewpoint (first left image), plus second best (second left image), plus third best (third left image), plus fourth best (fourth best image) for the coffee cup (i) and Lady of Elche (ii) models. Upper row viewpoints are selected according to geometry, in lower row according to saliency. Last image (both upper and lower row) corresponds to equal importance for all the viewpoints of the sphere.

most important viewpoints or to reflect the color of the environment, a sort of color bleeding. These extensions are explained below.

Viewpoint Selection

We can consider the selection of a viewpoint according to how much polygonal information is visible from it. For each viewpoint, we sum the PMI of all polygons weighted by how these polygons see the viewpoint, that is the conditional probabilities of the inverse channel. Thus, the information captured by viewpoint v is

$$PMI(v) = \sum_{o \in \mathcal{O}} p(v|o)I(o, V) \quad (11)$$

High values of $PMI(v)$ will correspond to viewpoints seeing much polygonal information, and viceversa. In Figure 7, we show the results for the lady of Elche and Lucy models. In the first and second columns of Figure 7, the best and worst viewpoints are shown, respectively. Observe how both of them correspond to the most and least informative viewpoints as they show the maximum and minimum of detail, given by the polygonal information. In columns three and four, we represent two views of the $PMI(V)$ sphere, i.e. for each viewpoint in the sphere we color code the $PMI(v)$ value. High $PMI(v)$ values are represented with warm colors and low values with cold ones.

Viewpoint Importance

From (10), importance can be introduced into the viewpoint space by modifying the input distribution $p(V)$ according to the importance we want to assign to each viewpoint. The polygonal information will be modified accordingly. The effect can be observed in Figure 8. For the two models shown, the range of images go from assigning almost all importance to the best viewpoint in the first image, to assign equal importance to the two best viewpoints in the second image, till assigning equal importance to the best 4 points in the fourth

image. Last image is obtained assigning equal importance to all viewpoints in the sphere. For each model, in the upper row we have considered the viewpoints obtained from the best view selection algorithm presented in [FSG07], while in the lower row the best viewpoints have been selected using the same algorithm driven by the saliency of polygons (see also [FSG07]). Observe the improvement of the images obtained when the most important viewpoints are the most salient ones.

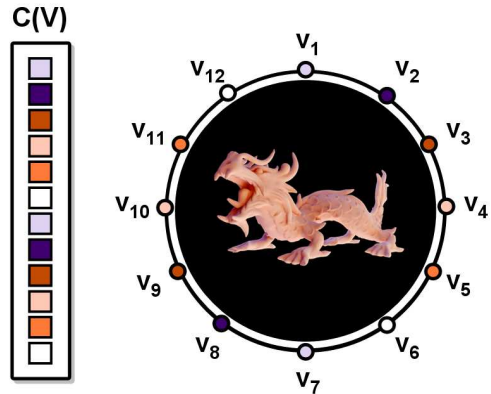


Figure 9: The result of warping a color texture to the viewpoint sphere for computing the color ambient occlusion model.

Relighting for Non-Photorealistic Rendering

Color ambient occlusion is obtained from the scalar product of a matrix row of $I(O, V)$ and the complementary of a color vector $c(V)$:

$$I_{\alpha}(o, V) = \sum_{v \in \mathcal{V}} I(o, v)(1 - c_{\alpha}(v)), \quad (12)$$



Figure 10: Combination of viewpoint-based ambient occlusion with a non-photorealistic technique using Coloroid color palettes (left).



Figure 11: Snapshots of an animation using our viewpoint-based ambient occlusion and the first row Coloroid palette shown in Figure 10, with a non-photorealistic technique on a boat model.

where α stands for each color channel, $c_\alpha(v)$ is the normalized vector for channel α and $I(o, v)$ is a matrix element of $I(O, V)$ (see Figure 5). After computing the polygonal mutual information for each channel, the final color ambient occlusion is given by the combination of the channels.

We can get a color vector by warping a color texture to the sphere of viewpoints. In this way, a color is assigned to each viewpoint (see Figure 9). In Figure 10 and 11 we show the combination of this kind of relighting technique with an NPR technique [Lak01], where the several color palettes used are Coloroid ones [Nem80]. Observe the nice effects obtained by this combination of techniques.

5. Conclusions and Future Research

In this paper we have presented a new view-point based approach to ambient occlusion. Our technique is based on the information channel between the polygons of an object and

the set of viewpoints, computing the mutual information associated with each polygon. This polygonal mutual information can guide the viewpoint selection in a sense that the best, and more informative, viewpoints will show the most details of the object. For model enhancement, the important viewpoints can modulate the obtained ambient occlusion values, and a relighting technique is shown in combination with an NPR algorithm. The effectiveness of our approach is demonstrated by the quality of the results obtained. Further research will be addressed to investigate the quality of the ambient occlusion obtained with generalized Tsallis-Havrda-Charvat mutual information and to obtain a full GPU implementation of our technique.

Acknowledgments

This project has been funded in part with grant numbers TIN2007-68066-C04-01 of the Spanish Government and IST-2-004363 (GameTools: Advanced Tools for Developing

Highly Realistic Computer Games) from the Vith European Framework.

References

- [BS05] BORDOLOI U. D., SHEN H.-W.: Viewpoint evaluation for volume rendering. In *Visualization, IEEE 2005* (May 2005), pp. 62–62.
- [FSG07] FEIXAS M., SBERT M., GONZÁLEZ F.: *A Unified Information-Theoretic Framework for Viewpoint Selection and Mesh Saliency*. Research Report IiA 07-03-RR, IiA - Institut d'Informàtica i Aplicacions, Universitat de Girona (Girona, Spain), 2007. Submitted to ACM Transactions on Applied Perception.
- [Lak01] LAKE A.: Cartoon rendering using texture mapping and programable vertex shaders. *Game Programming Gems 2* (June 2001), 444–451.
- [Lan02] LANDIS H.: Renderman in production, 2002.
- [MSC03] MÉNDEZ A., SBERT M., CATÀ J.: Real-time obscurances with color bleeding. In *Proceedings of the 19th spring conference on Computer graphics* (2003), ACM Press, pp. 171–176.
- [Nem80] NEMCSICS A.: The role of aesthetical uniformicity in the color harmony (in hungarian). In *In Proceeding of the XII. Color Symposium* (1980), pp. 24–29.
- [PB96] PLEMENOS D., BENEYADA M.: Intelligent display techniques in scene modelling. new techniques to automatically compute good views. In *International Conference GraphiCon'96* (July 1996). Held in St. Petersburg, Russia.
- [SPT06] SOKOLOV D., PLEMENOS D., TAMINE K.: Methods and data structures for virtual world exploration. *The Visual Computer* 22, 7 (2006), 506–516.
- [SSZK04] SATTLE M., SARLETTE R., ZACHMANN G., KLEIN R.: Hardware-accelerated ambient occlusion computation. In *Proceedings of International Fall Workshop on Vision, Modeling, and Visualization* (2004).
- [VFSG06] VIOLA I., FEIXAS M., SBERT M., GRÖLLER M. E.: Importance-driven focus of attention. *IEEE Trans. Vis. Comput. Graph.* 12, 5 (2006), 933–940.
- [VFSH01] VÁZQUEZ P. P., FEIXAS M., SBERT M., HEIDRICH W.: Viewpoint selection using viewpoint entropy. In *Proceedings of Vision, Modeling, and Visualization 2001* (Stuttgart, Germany, November 2001), Ertl T., Girod B., Greiner G., Niemann H., Seidel H.-P., (Eds.), pp. 273–280. Held in Stuttgart, Germany.
- [ZIK98] ZHUKOV S., IONES A., KRONIN G.: An ambient light illumination model. In *Rendering Techniques* (1998), pp. 45–56.