A Unified Information-Theoretic Framework for Viewpoint Selection and Mesh Saliency

Miquel Feixas, Mateu Sbert, Francisco González

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

Abstract

Viewpoint selection is an emerging area in computer graphics with applications in fields such as scene understanding, volume visualization, image-based modeling, and molecular visualization. In this paper, we present a unified framework for viewpoint selection and mesh saliency based on the definition of an information channel between a set of viewpoints and the set of polygons of an object. The mutual information of this channel is a powerful tool to deal with viewpoint selection and mesh visibility. In addition, the Jensen-Shannon divergence, closely related to mutual information, gives us a measure of viewpoint similarity and permits us to obtain the object saliency. Although we deal with the sphere of viewpoints around an object, our framework is also valid for any set of viewpoints in a closed scene. Several experiments show the robustness of the presented approach and the good behavior of the proposed measures.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computing Methodologies]: Computer Graphics

1. Introduction

In recent years, several methods that use the notion of *viewpoint quality* have been applied to computer graphics in fields such as image-based modeling [VFSH03], volume visualization [BS05, TFTN05], and mesh saliency [LVJ05]. The basic question underlying the viewpoint study and application is "what is a 'good' scene viewpoint?." Obviously, this question has not a unique answer. Depending on our objective, the best viewpoint can be, for instance, the most *representative* one or the most *unstable* one, i.e., the one that maximally changes when it is moved within its close neighborhood. Moreover, we probably need more than one view to capture the features of an object. The most representative views of an object can help us to understand it and the most unstable ones can enable us to obtain critical viewpoints to capture its structure. We will attempt here to answer the above question from the point of view of *information theory*, extending the work in [VFSH03, SPFG05].

In this paper, a unified and robust framework to deal with viewpoint selection and mesh saliency is presented. Given a set of viewpoints around an object, we define an *information channel* between the viewpoints and the polygons of the object. From this channel, we calculate the *viewpoint mutual information*, which will be used to obtain the best views of an object and to compute the degree of visibility of each polygon of the object. On the other hand, the *Jensen-Shannon divergence* will be used to select the most unstable viewpoints and to calculate the mesh saliency. The main advantages of our approach derive from the good properties of mutual information, such as its convergence to a finite value when the object mesh is infinitely refined. This framework is also applicable to any set of viewpoints in a closed scene and, although only the geometric properties of an object or scene are considered, other aspects such as lighting could be introduced.

This paper is organized as follows. In Section 2, we survey background and related work. In Section 3, we define an information channel between a set of viewpoints and the object. Section 4 and Section 5 are respectively devoted to viewpoint selection and mesh visibility and saliency. Finally, in Section 6, conclusions and future work are presented.

2. Background and Related Work

In this section we review some information-theoretic concepts [CT91] and related work.

2.1. Information-Theoretic Measures

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, let Y be a random variable taking values y in $\mathcal{Y}$. The *Shannon entropy* $H(X)$ of a random variable X is defined by

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

The Shannon entropy $H(X)$, also denoted by $H(p)$, measures the average uncertainty of random variable X. 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. In general, $H(Y|X) \neq H(X|Y)$, and $H(X) \geq H(X|Y) \geq 0$.

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

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

The mutual information $I(X,Y)$ is a measure of the shared information between X and Y. It can be seen that $I(X,Y) = I(Y,X) \geq 0$. A fundamental property of MI is the *data processing inequality* which can be expressed in the following way: if $X \rightarrow Y \rightarrow Z$ is a Markov chain, i.e., $p(x,y,z) = p(x)p(y|x)p(z|y)$, then

$$I(X,Y) \geq I(X,Z). \quad (4)$$

This result demonstrates that no processing of Y, deterministic or random, can increase the information that Y contains about X [CT91].

The *relative entropy* or *Kullback-Leibler distance* between two probability distributions p and q is defined as

$$KL(p|q) = \sum_{x\in\mathcal{X}} p(x)\log\frac{p(x)}{q(x)}, \quad (5)$$

where, from continuity, we use the convention that $0\log 0 = 0$, $p(x)\log\frac{p(x)}{0} = \infty$ if $p(x) > 0$ and $0\log\frac{0}{0} = 0$. The relative entropy $KL(p|q)$ is a measure of the inefficiency of assuming that the distribution is q when the true distribution is p [CT91].

A convex function f on the interval $[a,b]$ fulfils that $\sum_{i=1}^{n}\lambda_i f(x_i) - f(\sum_{i=1}^{n}\lambda_i x_i) \geq 0$, where $0 \leq \lambda \leq 1$, $\sum_{i=1}^{n}\lambda_i = 1$, and $x_i \in [a,b]$. For a concave function, the inequality is reversed. If f is substituted by the Shannon entropy, which is a concave function, we obtain the *Jensen-Shannon inequality* [BR82]:

$$JS(p_1,p_2,\ldots,p_N) = H(\sum_{i=1}^{N}\pi_i p_i) - \sum_{i=1}^{N}\pi_i H(p_i) \geq 0, \quad (6)$$

where $JS(p_1,p_2,\ldots,p_N)$ is the *Jensen-Shannon divergence* of probability distributions $p_1,p_2,\ldots,p_N$ with prior probabilities or weights $\pi_1,\pi_2,\ldots,\pi_N$, fulfilling $\sum_{i=1}^{N}\pi_i = 1$. The JS-divergence measures how 'far' are the probabilities p_i from their likely joint source $\sum_{i=1}^{N}\pi_i p_i$ and equals zero if and only if all the p_i are equal. It is important to note that the Jensen-Shannon divergence is identical to $I(X,Y)$ when $\pi_i = p(x_i)$ and $p_i = p(Y|x_i)$ for each $x_i \in \mathcal{X}$, where $p(x_i)$ is the marginal probability and $p(Y|x_i)$ represents the conditional probability distribution [ST00b]. Observe that capital letter Y is used to denote $p(Y|x_i) = \{p(y_1|x_i), p(y_2|x_i), \ldots, p(y_M|x_i)\}$.

2.2. Related Work

Viewpoint selection has been applied to several domains in computer graphics, such as scene understanding and virtual exploration [BDP00, Ple03, VFSH01, VS03, AVF04], molecular visualization [VFSL02], image-based modeling [VFSH03], volume visualization [BS05, TFTN05], and mesh saliency [LVJ05]. Different measures for viewpoint evaluation have been used in these fields. We review here the most relevant ones for our work.

In [BDP00, Ple03], the quality of a viewpoint v of a scene has been computed using the *heuristic* measure, given by

$$C(v) = \frac{\sum_{i=1}^{n}\lceil\frac{P_i(v)}{P_i(v)+1}\rceil}{n} + \frac{\sum_{i=1}^{n}P_i(v)}{r}, \quad (7)$$

where $P_i(v)$ is the number of pixels corresponding to the polygon i in the image obtained from the viewpoint v, r is the total number of pixels of the image (resolution of the image), and n is the total number of polygons of the scene. In this formula, $\lceil x\rceil$ denotes the smallest integer, greater than or equal to x. The first term in (7) gives the fraction of visible surfaces with respect to the total number of surfaces, while the second term is the ratio between the projected area of the scene (or object) and the screen area (thus, its value is 1 for a closed scene).

In [VFSH01], the *viewpoint entropy* has been defined from the relative area of the projected polygons over the sphere of directions centered at viewpoint v. Thus, the viewpoint entropy was defined by

$$H_v = -\sum_{i=0}^{N_f}\frac{a_i}{a_t}\log\frac{a_i}{a_t}, \quad (8)$$

where N_f is the number of polygons of the scene, a_i is the projected area of polygon i over the sphere, a_0 represents the projected area of background in open scenes, and $a_t = \sum_{i=0}^{N_f} a_i$ is the total area of the sphere. The maximum entropy is obtained when a certain viewpoint can see all

the polygons with the same projected area. The best viewpoint is defined as the one that has maximum entropy. The main drawback of viewpoint entropy is that it depends on the polygonal discretisation since a highly discretised region heavily attract the attention of the measure, penalizing big polygons in front of small ones. Viewpoint entropy has been recently extended to volume visualization [BS05] by substituting the area distribution by the voxel visibility distribution.

In [SPFG05], a new viewpoint quality measure based on the *Kullback-Leibler distance* (5) has been defined by

$$KL_v = \sum_{i=1}^{N_f} \frac{a_i}{a_t} \log \frac{\frac{a_i}{a_t}}{\frac{A_i}{A_T}}, \tag{9}$$

where a_i is the projected area of polygon i, $a_t = \sum_{i=1}^{N_f} a_i$, A_i is the actual area of polygon i and $A_T = \sum_{i=1}^{N_f} A_i$ is the total area of the scene or object. The viewpoint KL distance is interpreted as the distance between the normalized distribution of projected areas and the ideal projection, given by the normalized distribution of the actual areas. In this case, the background is not taken into account. The minimum value 0 is obtained when the normalized distribution of projected areas is equal to the normalized distribution of actual areas. Thus, to select views of high quality means to minimize KL_v. While the viewpoint entropy is very sensitive to both size and number of polygons, the KL measure only takes into account the proportion between the normalized projected area and the normalized actual area, trying to obtain a balanced vision of the object or scene. On the other hand, the fact that there exist many non visible or poorly visible polygons in a model can distort the quality of the measure.

3. Viewpoint Information Channel

In this section, we introduce an *information channel* between a set of viewpoints and the set of polygons of an object. We then define a *'goodness'* measure of a viewpoint and a *similarity* measure between two views, both based on the mutual information of this channel.

3.1. Viewpoint Mutual Information

We define an information channel $V \rightarrow O$ between the random variables V and O, which represent, respectively, a set of viewpoints and the set of polygons of an object. Viewpoints will be indexed by v and polygons by o. The marginal probability distribution of V is given by $p(v) = \frac{1}{N_v}$, where N_v is the number of viewpoints. That is, we assign the same probability to each viewpoint, although any other distribution could be used in our framework. The conditional probabilities $p(o|v)$ are given by the relative area of the projected polygons over the sphere of directions centered at viewpoint v. Finally, the marginal probability distribution of O is given by $p(o) = \sum_{x \in \mathcal{V}} p(v)p(o|v) = \frac{1}{N_v} \sum_{x \in \mathcal{V}} p(o|v)$.

From this channel, it follows that the *conditional entropy* (2) can be written as

$$\begin{aligned} H(O|V) &= -\sum_{v \in \mathcal{V}} p(v) \sum_{o \in \mathcal{O}} p(o|v) \log p(o|v) \\ &= \frac{1}{N_v} \sum_{v \in \mathcal{V}} H(v), \end{aligned} \tag{10}$$

where $H(v) = -\sum_{o \in \mathcal{O}} p(o|v) \log p(o|v)$ is the entropy of v. Observe that this coincides with the *viewpoint entropy* (8).

In this paper, we focus our attention on mutual information, that expresses the degree of *dependence* or *correlation* between a set of viewpoints and the object. The *mutual information* (3) is given by

$$\begin{aligned} I(V,O) &= \sum_{v \in \mathcal{V}} p(v) \sum_{o \in \mathcal{O}} p(o|v) \log \frac{p(o|v)}{p(o)} \\ &= \frac{1}{N_v} \sum_{v \in \mathcal{V}} I(v,O), \end{aligned} \tag{11}$$

where

$$I(v,O) = \sum_{o \in \mathcal{O}} p(o|v) \log \frac{p(o|v)}{p(o)} \tag{12}$$

represents the degree of correlation between the viewpoint v and the set of polygons. We propose to take $I(v,O)$ as our viewpoint goodness measure, called *viewpoint mutual information* (VMI). High values of the measure mean a high dependence between the viewpoint and a set of polygons, indicating a very *restrictive* view. On the other hand, low values correspond to low dependence, allowing for more *representative* views of the object. Observe that $I(v,O) = KL(p(O|v)|p(O))$, where majuscules indicate that $p(O|v)$ is the conditional probability distribution between v and the object and $p(O)$ is the marginal probability distribution of O. Note the difference with (9), where the distance is taken with respect to the actual areas, while in $I(v,O)$ the distance is taken with respect to the average projected area. This makes VMI sensitive to occlusions.

While the main drawback of *viewpoint entropy* is that it depends on the polygonal discretisation, VMI converges to a finite value when the mesh is infinitely refined [FdABS99]. As a consequence, VMI is more robust than viewpoint entropy when the object mesh is changed. This fact will be shown in the examples presented in the next section.

3.2. Viewpoint Clustering and Similarity

From the data processing inequality (4), we know that any clustering or quantization over V or O, respectively represented by $\widehat{V}$ and $\widehat{O}$, will reduce $I(V,O)$ [TPB99]. Thus, $I(V,O) \geq I(\widehat{V},O)$ and $I(V,O) \geq I(V,\widehat{O})$. For instance, merging neighbour viewpoints or polygons will reduce $I(V,O)$.

It can be demonstrated that the reduction of MI due to the merging of two viewpoints (or, in general, two view clusters)

v_i and v_j [ST00a] is given by

$$\delta I_{v_i,v_j} = (p(v_i)+p(v_j))JS(p(O|v_i),p(O|v_j)) \geq 0, \quad (13)$$

where $JS(p(O|v_i),p(O|v_j))$ is the Jensen-Shannon divergence (6), and $p(O|v_i)$ and $p(O|v_j)$ are the conditional probability distributions [ST00b]. After merging v_i and v_j, the probabilities corresponding to the resulting cluster $\widehat{v}_k$ are

$$p(\widehat{v}_k) = p(v_i)+p(v_j) = \frac{2}{N_v} \quad (14)$$

and

$$p(o|\widehat{v}_k) = \frac{p(v_i)p(o|v_i)+p(v_j)p(o|v_j)}{p(\widehat{v}_k)} = \frac{p(o|v_i)+p(o|v_j)}{2}. \quad (15)$$

Note that the loss of MI (3) is given by the multiplication of the 'weight' of elements we merge by the distance between them, given by the JS-divergence between the conditional distributions. The loss of MI can be interpreted as the *dissimilarity* between two views v_i and v_j. Thus, two views are similar when the JS-divergence value is small. The use of Jensen-Shannon as a measure of *view similarity* has been previously proposed in the volume rendering field by Bordoloi et al. [BS05].

From these concepts, a possible clustering algorithm appears naturally: we can successively merge two views into a new view in a way that locally minimizes the loss of MI. Using this algorithm, viewpoints (or polygons) can be clustered by preserving the maximum MI of the channel. It is important to stress that the variation of MI when two views are clustered does not depend on the level of clustering of the rest of viewpoints. That is, the convenience of clustering two views is a local decision.

4. Viewpoint Selection

The selection of *representative* viewpoints or *unstable* viewpoints can improve in an efficient way our understanding of the object or scene and can be used as an starting point for interactive scene exploration. In this section, VMI is used to find both informative views and a minimal set of representative views for an object, and JS-divergence is used to find unstable views.

4.1. Viewpoint Mutual Information Evaluation

We compare the behavior of VMI with respect to the following viewpoint quality measures: heuristic measure (7), viewpoint entropy (8), and Kullback-Leibler distance (9). All these measures are sensitive to the size of the view sphere with respect to the object. In this paper we have not taken into account this dependence.

In our experiments, all the objects are centered in a sphere of 642 viewpoints and the camera is looking at the center of this sphere. In all the figures, the view sphere is represented by a color map, where red and blue colors correspond respectively to high and low values of the measure utilized. Note that a high value for the heuristic measure (7) and the viewpoint entropy (8) indicates a good view, while a high value for KL (9) and VMI (12) corresponds to a restrictive view. Viewpoint entropy has been computed without taking into account the background.

To compute these viewpoint measures, we have estimated the projection of the visible parts of the scene on the screen. Before projection, a different color is assigned to each surface. 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 surface represented by this color.

To evaluate the behaviour of all the viewpoint quality measures, three objects are analyzed: a cow (Figure 1), a coffeecup (Figure 2), and a ship (Figure 3). Figure 1 has been organized as follows. Rows *i*, *ii* and *iii* show, respectively, the behaviour of the heuristic, entropy and VMI measures. Columns *a* and *b* show, respectively, the views with highest and lowest goodness, and columns *c* and *d* show two different viewpoint spheres. Figure 1 shows that the VMI measure performs better than the other ones, giving us the lateral views as the best.

Figure 2 shows the behaviour of the heuristic, entropy and VMI measures when the discretisation of the object varies outstandingly. Columns (a) and (b) show the viewpoint spheres computed on a uniform discretised coffeecup and columns (c) and (d) are computed on the same coffeecup with a more refined dish. We can clearly observe that heuristic and entropy measures change notably their behaviour with respect to the discretisation variation, while the VMI measure is insensitive to it. This is an important added value of the VMI measure.

In Figure 3, we show the difference between the KL measure (a-b) and the VMI measure (c-d). Due to the fact that the ship model has a lot of non visible polygons, the KL measure gives us a distorted vision of the object. Let us remember that the main difference between MI and KL is that MI computes the distance between the projected areas of the polygons and the area fraction *seen* by the set of polygons, instead of the *actual* areas of the object for the KL case.

4.2. Selection of n Best Views

In this section, we present a new viewpoint selection algorithm based on VMI. Its objective is to find the minimal representative set of views for a given object or scene in order to well understand it.

An ideal algorithm would be to select n views that maximize their JS-divergence, i.e., that maximize the loss of mutual information of the channel. The basic idea underlying this algorithm is to capture the maximum information of the object with the minimum number of viewpoints. Due to

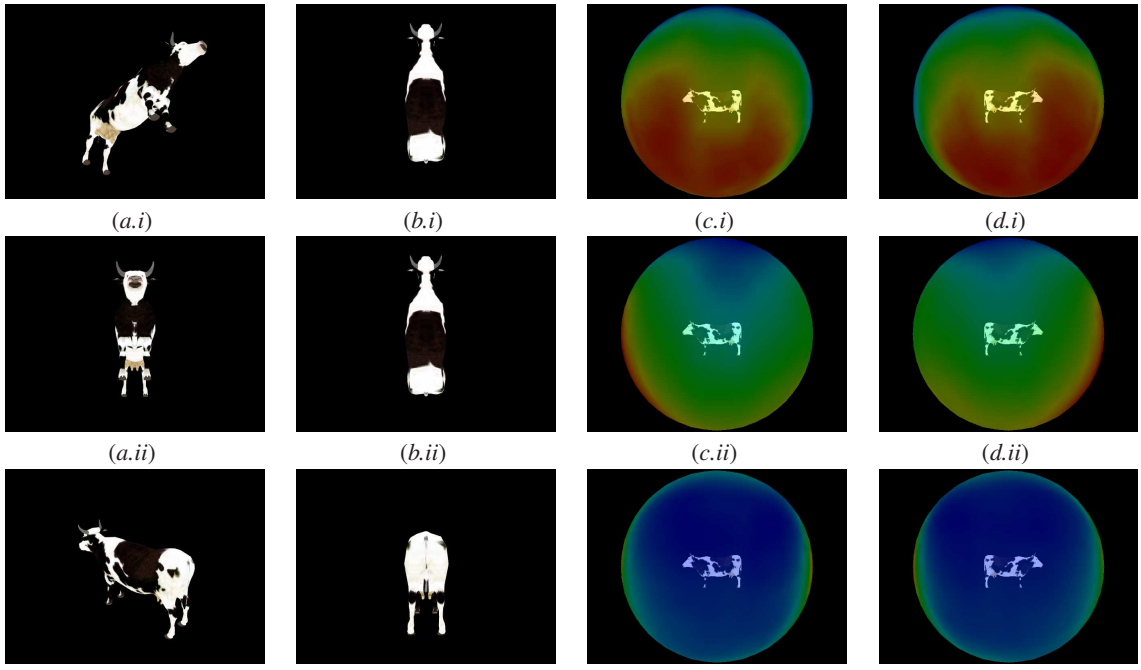

Figure 1: *The most representative (a) and the most restrictive (b) views and viewpoint spheres (c-d) obtained respectively from the heuristic (i), entropy (ii) and VMI (iii) measures. Red colors on the sphere represent high measure values (good viewpoints for the heuristic and entropy cases and restrictive ones for the VMI case), blue colors represent low measure values (restrictive viewpoints for the heuristic and entropy cases, good ones for the VMI case).*

the fact that this optimization algorithm is NP-complete, we adopt a greedy strategy by selecting successive viewpoints that maximize the JS-divergence. This algorithm permits us to find in an automated way the minimal set of views which represent the object or scene. A maximum value of the JS-divergence would be achieved when all the viewpoints are clustered. That is, the MI of the channel would be zero.

Taking into account that the most representative view corresponds to the minimum VMI, we propose a greedy algorithm consisting in finding a set of views where the mixed distribution of the projected areas has a minimum VMI with respect to the marginal probability distribution of O. The algorithm proceeds as follows. First, we select the view v_1 with distribution $p(O|v_1)$ corresponding to the minimum VMI (most informative). Next, we select $p(O|v_2)$ such that the mixed distribution $\frac{p(O|v_1)+p(O|v_2)}{2}$ also corresponds to the minimum VMI, i.e., $I(\widehat{v},O)$ is minimum, where $\widehat{v}$ represents the clustering of v_1 and v_2. At each step, a new mixed distribution $\frac{p(O|v_1)+p(O|v_2)+...+p(O|v_n)}{n}$ is produced until the decrease of the VMI distance is lower than a given threshold or a determined number of views is selected. Let us remember that the same weight has been assigned to all the views. Another alternative method could be to weight them by their projected area or other importance criteria.

Figures 4 and 5 present the six best views obtained with our selection algorithm. The algorithm stops when the VMI difference between two successive views is lower than a given threshold. Its behaviour is also shown in Figure 6, where we observe how the VMI values obtained for the successive mixed distributions for Figures 4 and 5 converge asymptotically to zero.

Another selection algorithm could be based on the clustering algorithm described in Section 3.2, selecting the most representative view for each cluster. However, due to its computational cost, we present here a greedy strategy to partition the sphere of viewpoints from the n best views selected using the previous VMI algorithm. We proceed by assigning each viewpoint to the 'nearest' representative view, i.e., the assignation is determined by the minimum JS-divergence between the viewpoint to be clustered and the selected views. In Figure 7, we show the behavior of this clustering algorithm for the coffeecup object (i) and for the cow (ii).

4.3. Unstable Viewpoints

In the volume rendering field, Bordoloi et al. [BS05] use the JS-divergence as a view similarity measure and define the view stability as "the maximum change in view that occurs

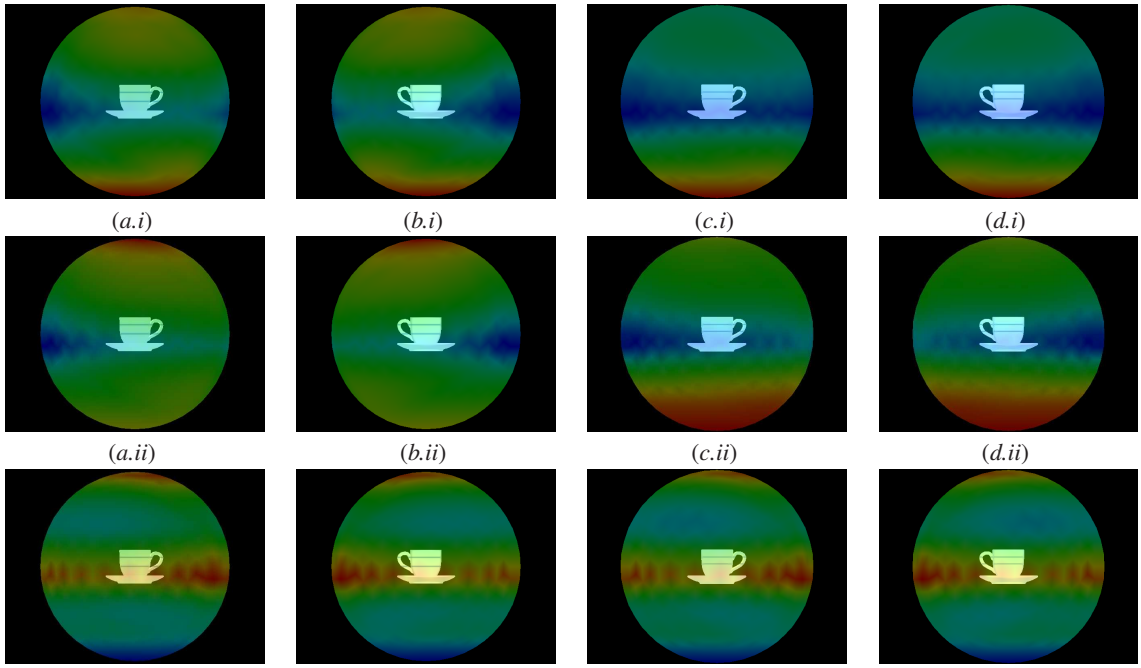

(*a.i*) (*b.i*) (*c.i*) (*d.i*)

(*a.ii*) (*b.ii*) (*c.ii*) (*d.ii*)

(*a.iii*) (*b.iii*) (*c.iii*) (*d.iii*)

Figure 2: *Viewpoint spheres obtained respectively from the heuristic (i), entropy (ii) and VMI (iii) measures. Viewpoint spheres (a-b) are computed on a uniform discretised coffeecup and (c-d) on a coffeecup with a more refined dish.*

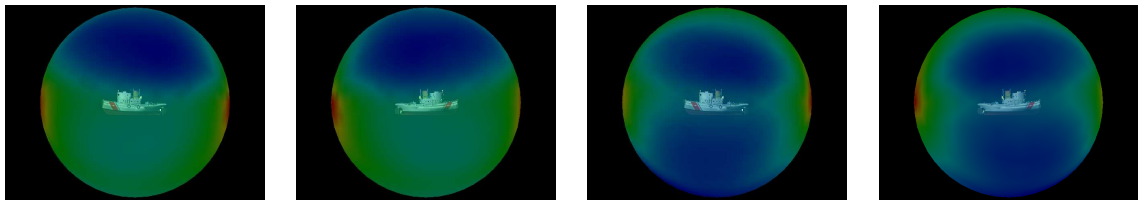

(*a*) (*b*) (*c*) (*d*)

Figure 3: *Viewpoint spheres obtained respectively from the KL (a-b) and VMI (c-d) measures.*

(*a*) (*b*) (*c*) (*d*) (*e*) (*f*)

Figure 4: *The six most representative views of the coffeecup object selected by the VMI algorithm.*

(*a*) (*b*) (*c*) (*d*) (*e*) (*f*)

Figure 5: *The six most representative views of the cow selected by the VMI algorithm.*

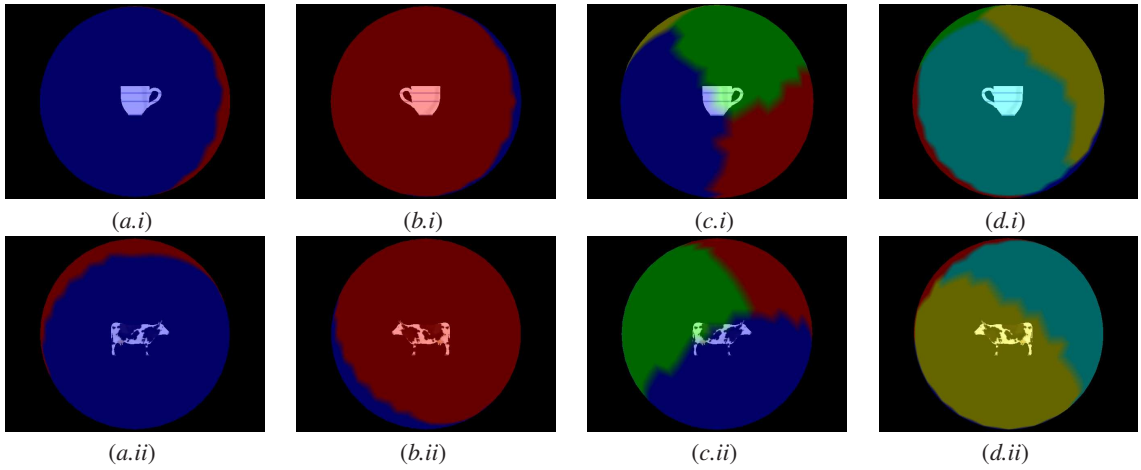

(a.i) (b.i) (c.i) (d.i)

(a.ii) (b.ii) (c.ii) (d.ii)

Figure 7: *Viewpoint clustering spheres obtained by the clustering algorithm for the coffeecup object (i) and for the cow (ii). Spheres (a-b) have two clusters and spheres (c-d) five.*

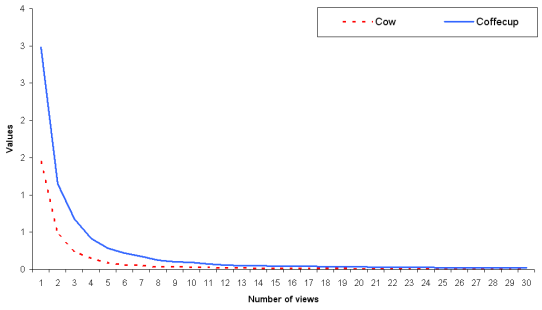

Figure 6: *VMI values obtained for the successive mixed distributions for Figures 4 and Figure 5.*

when the camera position is shifted within a small neighborhood." A small change corresponds to an stable viewpoint and a large change to an unstable view.

From our framework, the use of Jensen-Shannon to evaluate the stability of a viewpoint appears in a natural way. For each viewpoint, we propose to evaluate the average variation of MI when two neighbor views are clustered. Thus, the *viewpoint stability* is defined by

$$E(v_i) = \frac{1}{N_n} \sum_{j=1}^{N_n} JS(p(O|v_i), p(O|v_j)) \geq 0, \qquad (16)$$

where v_j is a neighbor viewpoint of v_i, N_n is the number of neighbors of v_i, and the conditional probabilities are respectively weighted by $\frac{p(v_i)}{p(v_i)+p(v_j)} = \frac{1}{2}$ and $\frac{p(v_j)}{p(v_i)+p(v_j)} = \frac{1}{2}$ in the JS-divergence.

Figure 8 shows the good behavior of the viewpoint stability measure for the coffeecup and ship models.

5. Mesh Visibility and Saliency

We now want to analyze how the object is *seen* or *perceived* from the set of viewpoints. Two aspects are considered: the *degree of visibility* of each polygon and its *saliency* with respect to its neighbors. We introduce two different perception-inspired measures of *regional importance* which can be utilized in mesh simplification. Moreover, both measures can be used to drive the viewpoint selection, giving more importance to the most visible or salient parts. Other tasks such as rendering, animation, and compression can benefit from mesh visibility and saliency computation.

5.1. Mesh Visibility

From the Bayes theorem $p(v,o) = p(v)p(o|v) = p(o)p(v|o)$, the mutual information (11) can be rewritten as

$$\begin{aligned} I(V,O) &= \sum_{o\in\mathcal{O}} p(o) \sum_{v\in\mathcal{V}} p(v|o) \log \frac{p(v|o)}{p(v)} \\ &= \sum_{o\in\mathcal{O}} p(o) I(V,o), \end{aligned} \qquad (17)$$

where $I(V,o)$, called as *polygonal mutual information*, represents the degree of correlation between the polygon o and the set of viewpoints, and can be interpreted as the *degree of visibility* of polygon o. The highest the value, the lowest the polygon visibility, and viceversa.

We show the visibility map of three figures: the coffeecup (9(i)), the angel (9(ii)), and the hebe (9(iii)). Observe how the interior of the coffeecup, the sides of the angel neck, and the hebe mouth and nose turn towards red (less visible parts). Meanwhile highly visible surfaces like the bottom of the dish or the back of the angel are blue indicating highly visible areas.

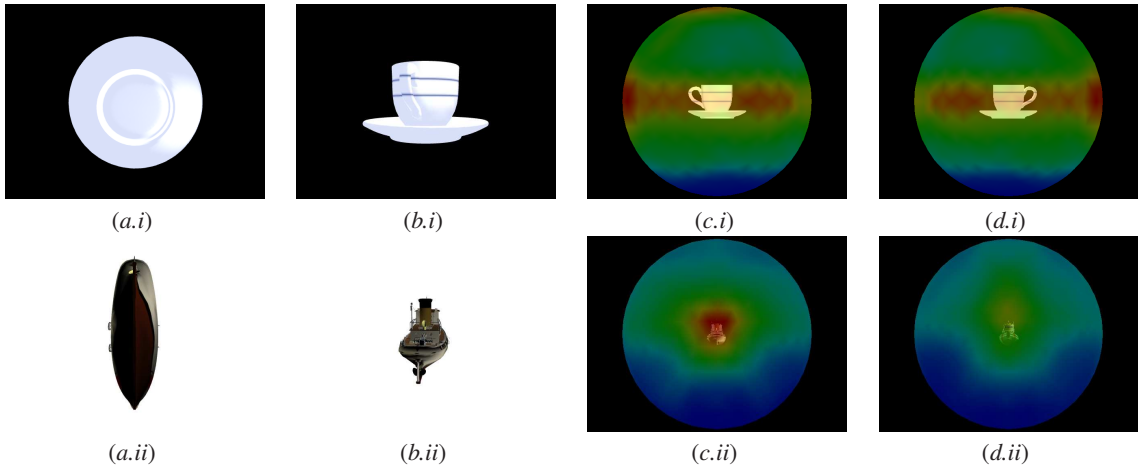

(*a.i*) (*b.i*) (*c.i*) (*d.i*)

(*a.ii*) (*b.ii*) (*c.ii*) (*d.ii*)

Figure 8: *More stable (a) and more unstable (b) views and unstability spheres (c-d) obtained for the coffeecup (i) and ship (ii) objects. Red colors on the sphere represent high unstable values, blue colors represent low unstable values.*

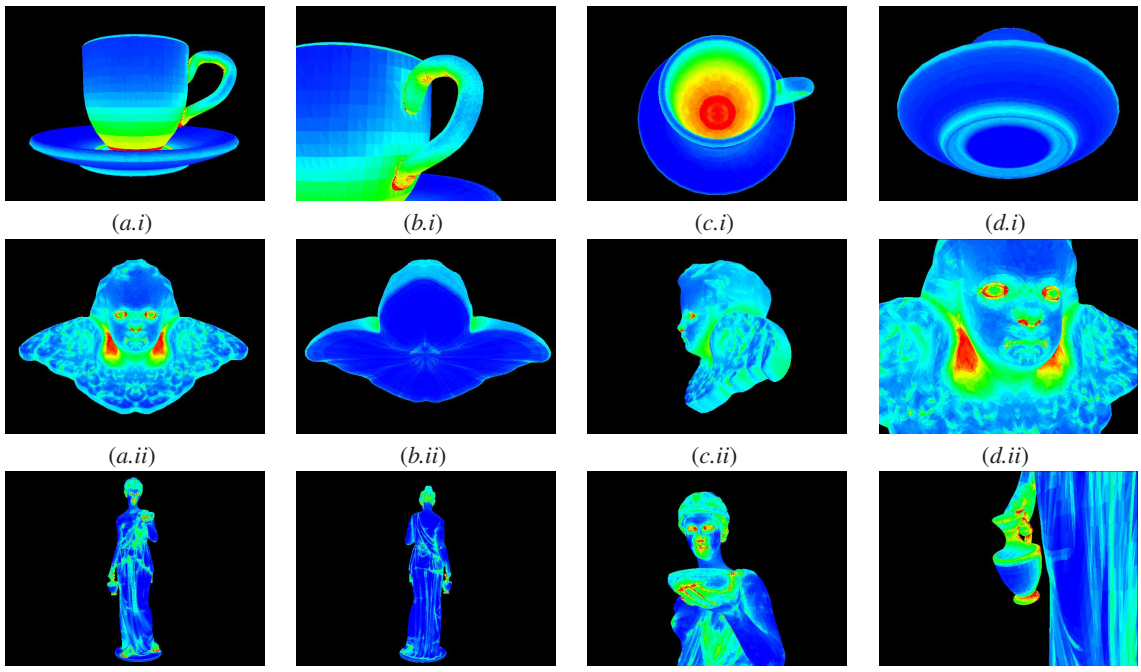

(*a.i*) (*b.i*) (*c.i*) (*d.i*)

(*a.ii*) (*b.ii*) (*c.ii*) (*d.ii*)

(*a.iii*) (*b.iii*) (*c.iii*) (*d.iii*)

Figure 9: *Mesh visibility for the coffeecup (i), angel (ii) and hebe (iii) models.*

5.2. Mesh Saliency

We introduce now a measure of mesh saliency based on the JS-divergence. Unlike [LVJ05], our approach formulates the mesh saliency in terms of how the polygons are seen by the set of viewpoints, i.e., according to the visual perception.

Similarly to the stability of a viewpoint (see Section 4.3), for each polygon we propose to evaluate the average variation of JS between two polygons. Thus, the *saliency of a polygon* is defined by

$$S(o_i) = \frac{1}{N_o} \sum_{j=1}^{N_o} JS(p(V|o_i), p(V|o_j)) \geq 0, \qquad (18)$$

where o_j is a neighbor polygon of o_i, N_o is the number of neighbor polygons of o_i, and the conditional probabilities are respectively weighted by $\frac{p(o_i)}{p(o_i)+p(o_j)}$ and $\frac{p(o_j)}{p(o_i)+p(o_j)}$ in the JS-divergence.

Figure 10 shows how the red parts as nose, mouth and fingers for the angel and hebe models and the handle of the coffeecup are showing the most salient surfaces. However, blue parts of the models as the bottom of the coffeecup dish and the back of the angel are showing the least salient values.

6. Conclusions and Future Work

In this paper, an information channel between a set of viewpoints and an object has been introduced. We have shown how viewpoint mutual information has led to several efficient algorithms in viewpoint selection and mesh visibility and saliency. This framework has integrated previous measures such as viewpoint entropy and Kullback-Leibler distance and has opened a wide range of applications. Experimental results have shown the robustness of the presented approach and the good behavior of the proposed measures. In our future research, the importance provided by mesh visibility and saliency could be incorporated to viewpoint selection by weighting adequately the conditional probabilities.

Aknowledgements

This project has been funded in part with grant numbers TIN2004-07451-C03-01 and FIT-350101-2004-15 of the Spanish Government and IST-2-004363 (GameTools: Advanced Tools for Developing Highly Realistic Computer Games) from the VIth European Framework.

References

[AVF04] ANDÚJAR C., VÁZQUEZ P. P., FAIRÉN M.: Way-finder: guided tours through complex walthrough models. *Computer Graphics Forum (Eurographics 2004)* (2004).

[BDP00] BARRAL P., DORME G., PLEMENOS D.: Scene understanding techniques using a virtual camera. In *Short papers proceedings, Eurographics 2000* (August 2000). Held in Interlagen, Switzerland.

[BR82] BURBEA J., RAO C. R.: On the convexity of some divergence measures based on entropy functions. *IEEE Transactions on Information Theory 28*, 3 (May 1982), 489–495.

[BS05] BORDOLOI U. D., SHEN H.-W.: Viewpoint evaluation for volume rendering. In *Visualization, IEEE 2005* (May 2005), pp. 62–62.

[CT91] COVER T. M., THOMAS J. A.: *Elements of Information Theory*. Wiley Series in Telecommunications, 1991.

[FdABS99] FEIXAS M., DEL ACEBO E., BEKAERT P., SBERT M.: An information theory framework for the analysis of scene complexity. *Computer Graphics Forum (Proceedings of Eurographics'99) 18*, 3 (September 1999), 95–106.

[LVJ05] LEE C. H., VARSHNEY A., JACOBS D. W.: Mesh saliency. *Computer Graphics (Proceedings of SIGGRAPH'05)* (July - August 2005). Held in Los Angeles, USA.

[Ple03] PLEMENOS D.: Exploring virtual worlds: Current techniques and future issues. In *International Conference GraphiCon'2003* (September 2003). Held in Moscow, Russia.

[SPFG05] SBERT M., PLEMENOS D., FEIXAS M., GONZÁLEZ F.: Viewpoint quality: Measures and applications. In *Proceedings of 1st Computational Aesthetics in Graphics, Visualization and Imaging* (May 2005). Held in Girona, Spain.

[ST00a] SLONIM N., TISHBY N.: Agglomerative information bottleneck. In *Proceedings of NIPS-12 (Neural Information Processing Systems)* (2000), MIT Press, pp. 617–623.

[ST00b] SLONIM N., TISHBY N.: Document clustering using word clusters via the information bottleneck method. In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (2000), ACM Press, pp. 208–215. Held in Athens, Greece.

[TFTN05] TAKAHASHI S., FUJISHIRO I., TAKESHIMA Y., NISHITA T.: Locating optimal viewpoints for volume visualization. In *Visualization, IEEE 2005* (May 2005).

[TPB99] TISHBY N., PEREIRA F. C., BIALEK W.: The information bottleneck method. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control and Computing* (1999), pp. 368–377.

[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

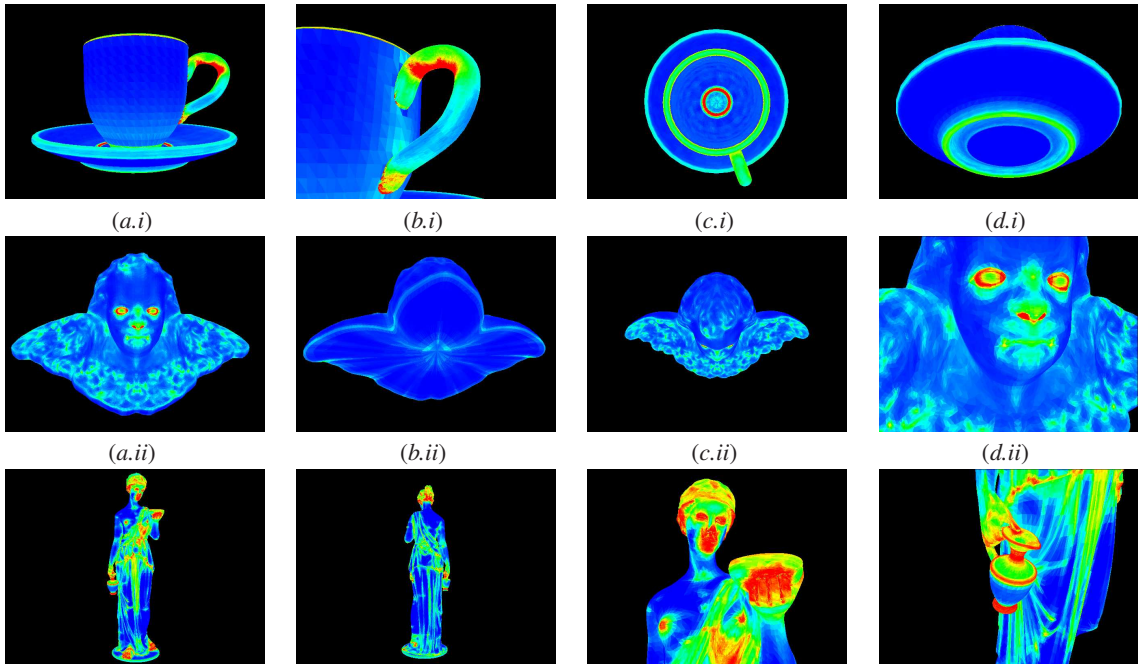

Figure 10: *Mesh saliency for the coffeecup (i), angel (ii) and hebe (iii) models.*

B., Greiner G., Niemann H.,, Seidel H.-P., (Eds.), pp. 273–280. Held in Stuttgart, Germany.

[VFSH03] VÁZQUEZ P. P., FEIXAS M., SBERT M., HEIDRICH W.: Automatic view selection using viewpoint entropy and its application to image-based modeling. *Computer Graphics Forum* (Desember 2003).

[VFSL02] VÁZQUEZ P. P., FEIXAS M., SBERT M., LLOBET A.: A new tool for obtaining good views for molecules. In *Proceedings of VisSym'02 (Eurographics-IEEE TCVG Symposium on Visualization)* (May 2002), Ebert D., Brunet P.,, Navazo I., (Eds.), pp. 0–1. Held in Barcelona, Spain.

[VS03] VÁZQUEZ P. P., SBERT M.: Automatic indoor scene exploration. In *Interantional Conference on Artificial Intelligence and Computer Graphics, 3IA* (Limoges, France, may 2003). Held in Limoges, France.