

Computational Aesthetics for Rendering Virtual Scenes on 3D Stereoscopic Displays

László SZIRMAY-KALOS, Pirkko OITTINEN, and Balázs TERÉKI

Introduction

Computer graphics builds virtual scenes that are rendered presenting images to the user. Both the virtual world and the image are represented by numbers in the computer, and there are many different possibilities to define a mapping between these numbers. What we usually need is a meaningful image that can be interpreted by the user without reading lengthy user manuals. To reach this goal, we use analogies like photography in photorealistic rendering¹, drawing or painting in illustrative rendering², X-ray, flows, etc. in abstract data visualization, because such images are natural and need no additional explanation (Figure 1). The analog process uses not only the primary data of the virtual scene but also additional parameters, called studio objects, defining how rendering should take place. Studio objects include camera and light parameters in photorealistic image synthesis, and drawing or hatching styles in illustrative rendering. Different studio object settings lead to different images that are better or poorer in characterizing the scene. To find good studio object settings that lead to expressive images, users usually initiate a long manual search based on trial and error.




Figure 1. Images obtained with photorealistic(left) and illustrative(right) rendering methods.

This long procedure can be replaced or helped by automatic algorithms that search for good rendering parameters. To support the search process, the quality of the images should be characterized by numeric values, and we should aim at maximizing these values. Measures expressing the quality of images are called *aesthetic measures*. Generally, aesthetics is related to order and complexity. One option to mathematically interpret such terms is to exploit the tools provided by *information theory*, which associate aesthetics (beauty) with objectively computable information content.

¹ L. Szirmay-Kalos, L. Szécsi, L. and M. Sbert, M., *GPU-Based Techniques for Global Illumination Effects*, San Rafael, Morgan and Claypool Publishers, 2008.

² T. Umenhoffer, L. Szécsi, L. Szirmay-Kalos, "Hatching for motion picture production", *Computer Graphics Forum* 30(2), (2011), pp. 533-542

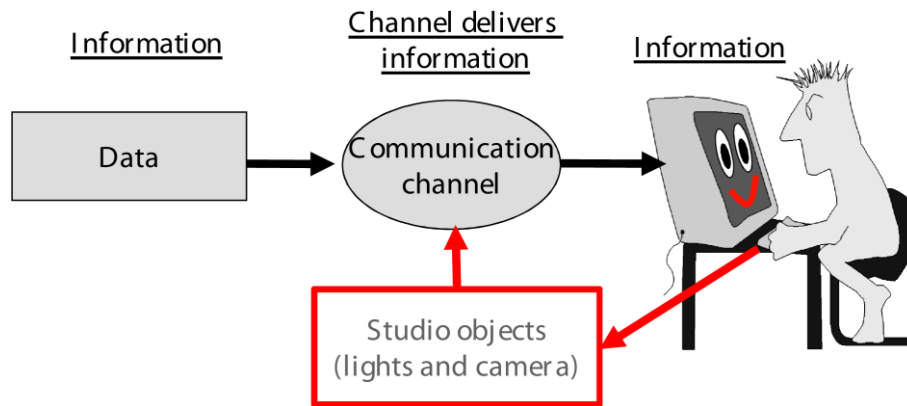


Figure 2. Information theoretic model of rendering: information of the original data goes through the information channel of the rendering process and is transformed to the image. A meaningful rendering process is expected to provide high information content, i.e. the entropy of the image must be high, and to maximize the correlation between the original data and the information that can be extracted from the image, i.e. the mutual information between the data and the image must be significant.

The application of information theory is justified by the observation that the rendering process can be interpreted as a channel communicating the information of the virtual world in form of images (Figure 2). An image is meaningful if it provides enough information and this information is strongly related to the data stored in the virtual world. In information theory, the amount of information is usually measured by the *entropy* and the strength of the relation between the data and its interpretation by the *mutual information*. In general, the model has focus data that should be communicated and data that can be controlled in order to improve the communication of the focus data. For example, the geometry and material properties of the objects may belong to the set of focus data while light sources and the camera can be set freely. In other scenarios, only the geometry or the density is important, and even the material properties can be subjects of user control. In abstract data visualization, the communication channel can be defined in a very flexible way.

In computer graphics, information theory has been used to describe complexity of 3D scenes³, to find the best views to render an object for triangle meshes⁴, volumes, and molecular models. Other applications include the optimal setting of light sources, registration of volumes, expressive shading, and automatic definition of transfer functions. For a review, refer to Sbert et al.⁵

In stereo imaging two images are rendered for the two eyes, which can reproduce *vergence*, i.e. the phenomenon that the same point is visible in two different directions from the two eyes. The vergence depends on the *interocular distance* and the distance of the point of interest. The vergence can be characterized by the angle difference in which a point is visible from the two eyes, or, expressing it from the point of view of the objects, the translation between the two projections of the same point, which is called the *disparity*. Disparity values of points visible in the pixels constitute the *disparity map*.

Significant research efforts have been devoted to the examination of human depth perception and also to the added problems of 3D displays that can reproduce just a limited set of depth cues. Due to the limitation of retinal

³ M. Feixas, E. Del Acebo, P. Bekaert, M. Sbert, “An information theory framework for the analysis of scene complexity”, *Computer Graphics Forum* 18(3), (1999), pp. 95-106.

⁴ P. Vázquez P., M. Feixas, M. Sbert, W. Heidrich, “Viewpoint selection using viewpoint entropy”, *Proceedings of the Vision Modeling and Visualization Conference*, (2001), pp. 21-23.

⁵ M. Sbert, M. Feixas, J. Rigau, I. Viola and M. Chover, *Information Theory Tools for Computer Graphics*, San Rafael, Morgan and Claypool Publishers, 2014.

disparity, when we focus on nearby objects, the depth perception of distant parts of the scene is lost. In stereo 3D displays, the vergence is reproduced, but other cues like accommodation or motion parallax are missing, which may cause discomfort or 3D fatigue for the viewer. To solve these problems, Lang et al.⁶ introduced the concept of non-linear transformation of the disparity map. Niu et al.⁷ examined the cropping and scaling of stereoscopic image pairs to avoid problems like the removal of a salient object from one of the images destroys its depth perception.

In this paper, we combine two lines of research for 3D stereo imaging: information theory based studio object setting and the perceptual issues of depth perception. First, we survey the most related previous work focusing on the application of entropy to the definition of aesthetics, then the new measure for the aesthetics of stereo images is defined and its computation presented.

Entropy based aesthetic measures

What a good camera or light setting means might be arguable, and there are several definitions and approaches focusing on different parameters, like the number of visible faces, the projected area, number of degenerate faces, saliency i.e. points that significantly differ from their environment, etc. However, it is intuitive that the purpose of tuning the camera and the lights is to help the user understand more of the scene, i.e. to increase the information arrived at the observer from the scene through the information channel associated with the rendering process.

The mathematical description of the information channel is given by the information theory. The information associated with a random variable X is described by its *Shannon entropy*:

$$H(X) = - \sum_x p(x) \log_2(p(x))$$

where $p(x) = P(x = X)$ is the probability density of random variable X .

This measure shows how much uncertainty is eliminated when the experiment is executed to get a value of the random variable. In order to apply this measure for images, a correspondence must be found between the important parameters of images and the definition of random variables. It means that the perception of the deterministic image should be interpreted as a random process.

For example, when quality focuses on the number and projected area of visible faces, the random variable can be defined as index i of the visible face and its probability as the ratio of its visible area A_i and the area of the screen S , which corresponds to the random model that the eye looks toward a point that is uniformly distributed on the camera window. With this model, the entropy of the random variable i is:

$$H = - \sum_i \frac{A_i}{S} \log_2 \frac{A_i}{S}$$

This definition is called the *viewpoint entropy*. Viewpoint entropy introduced this way has several problems. For example, it is based on the tessellation of surfaces, i.e. when every face is decomposed to smaller faces the image

⁶ M. Lang, A. Hornung, O. Wang, S. Poulakos, A. Smolic, M. Gross, “Nonlinear disparity mapping for stereoscopic 3D”, *ACM Transactions on Graphics* 29(3), (2010), 75:1-10.

⁷ Y. Niu, F. Liu, W.-H. Feng, H. Jin, “Aesthetics-based stereoscopic photo cropping for heterogeneous displays”, *IEEE Transactions on Multimedia* 14(3), (2012), pp. 783-796.

might not change at all, but the entropy H increases as stated by the Jensen inequality since the entropy is a concave function of face area A_i . From another point of view, this measure does not take into account perceptual properties. This problem is solved by re-interpreting A_i as the total area of the pixels that are perceptually different from other pixels in the neighborhood. In our work, we also use this approach. We propose a similar measure to characterize the quality of a stereo camera setting. As the objective of stereo imaging is to provide the illusion of depth, our model is based on the disparity map instead of the rendered image.

Depth information in a stereo image




Figure 3. Stereo imaging and the resulting image pair. Those objects are in the comfort zone where the accommodation cue caused by the distance of the display does not contradict the vergence cue caused by the binocular disparity. Objects showing up in only one image cause ghosting and have no correct depth perception.

In stereo imaging two pictures are rendered for the two eyes, which have slight differences, called binocular disparity, similarly to real-world images captured by the two eyes that have different positions. These differences are larger when the objects are closed, providing information to the brain to calculate depth in the scene (Figure 3).

The difference of the two images can be characterized by the *disparity map* that stores the translation vector in each pixel, that expresses the difference of the locations of this point on the two images. The disparity map is the potentially available maximum information about the depth, which can only be partially utilized by the brain since the human vision system's ability is limited in identifying corresponding point pairs in the two images. Correspondence can be easily established at object corners, where the point has a distinct color that is significantly different from its environment, or generally at salient parts, but is not possible inside faces where the color is constant or slowly changing.

Moreover, the human visual system has finite resolution in identifying the differences in depth and has a comfort zone where two points are believed to have a correspondence.

We propose an aesthetic measure that expresses the information content of the disparity map that can be utilized by the human visual system. The requirements for such a measure are the followings:

- It should be zero if two independent images are shown that do not depict the same part scene. Objects in the ghosting region should not have any contribution.
- It should increase when the scene has higher dynamic depth range, i.e. when more objects or objects parts of different perceived depths show up.
- It is expected to be higher when the object depths are more evenly distributed in the perceivable range.
- It should decrease when the disparities leave the control zone of the human eye.
- It should be identical for two different settings if the human eye has the same sensation, i.e. it must not be sensitive to the disparity of points for which stereo correspondence cannot be established, or to disparity differences that are too small to be noticeable.

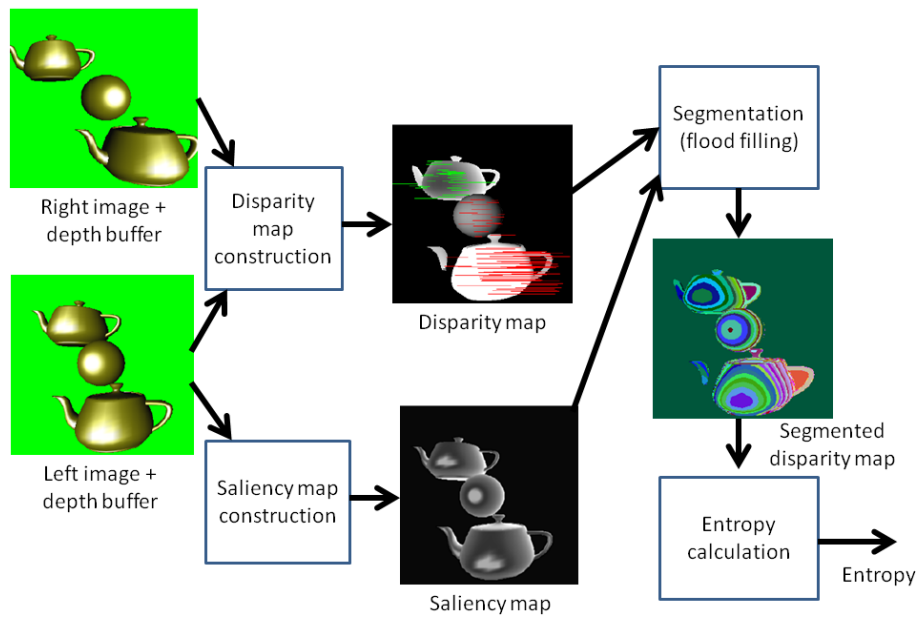


Figure 4. Steps of the calculation of the aesthetic measure.

The proposed measure is based on the entropy and the transformation of the information involved in the disparity map. The computation can be regarded as a sequence of the following steps (Figure 4):

1. *Disparity map computation*: We render the scene twice setting the camera according to the two eye positions. Reading back the depth buffer associated with the first image, for each pixel, surface points whose color are in the first image are unprojected and their 3D world positions are reconstructed. Projecting these points with the transformation associated with the second camera, we compare the depth of this projection to the depth stored in the depth buffer of the second camera. If the two depth values are similar, then the point is visible from both cameras, thus a stereo point pair is found. The image space translation is stored in the pixels of the disparity map. If the point is not simultaneously visible in both images, then this fact is indicated in the disparity map. Pixels that do not have a matching pair in the other image are given a *background disparity*, which are ignored in information calculation.
2. *Saliency map computation*: The human visual system may not be able to match corresponding points if their colors are not characteristically different from their environment. To incorporate this into our model, we also compute a saliency map from the color image. Saliency expresses how different the color of a point is from its neighborhood. Saliency map can be constructed using the difference of

Gaussians to express the change of color and intensity and on the Sobel filter to express the orientation dependence⁸ [Long.

3. *Perception based segmentation of the disparity map*: Using a *flood-fill* type algorithm, the disparity map is partitioned into a finite number of connected regions in a way that each region is a collection of pixels that are perceived by the human eye in the same depth. Without considering perceptual phenomena, partitioning would start with the quantization of the disparity map, then the flood-fill algorithm would identify the connected regions of the quantized map. To improve this approach, we should incorporate the features of human perception into the region filling process and classify two neighboring pixels similar if they cannot be distinguished by the human eye. We consider the ability of matching only salient points and that matching is more effective when different cues are not contradicting, i.e. when the object is in the comfort zone where accommodation and vergence are consistent. Current stereo displays force the eye to accommodate to the display distance, which is a different cue than vergence if the objects are far from this, i.e. when objects are out of the comfort zone. These far away objects cannot be fused by our visual system due to too large retinal disparities. To take both saliency and accommodation into account, we make the tolerance inversely proportional to the saliency. This way, image regions of low saliency and of uncomfortable depth are not decomposed by the segmentation algorithm since the human eye cannot perceive significant depth differences. Note that we are looking for *connected regions* formed by adjacent pixels since the information content of two smaller objects being in the same distance is higher than that of a single larger object.
4. *Entropy computation*: We use equation of the viewpoint entropy to compute the information content in the discrepancy map with the following substitutions: A_i is the screen space area (i.e. the number of pixels) of a connected region after segmentation and S is the total screen space area of valid pixels, i.e. where the stored value is different from the background disparity.

Results

The discussed method has been implemented in OpenGL environment and tested with different scenes. In these experiments we examined the effect of the interocular distance on the aesthetics of stereo imaging, but it would be also straightforward to modify other parameters, like the light source positions that affect the result via the saliency map, or the distance of the screen, which would modify the ghosting regions and also the comfort zone of the scene.

Figure 5 shows a Cornell box like scene where a dragon is placed close to the eye position while the Happy Buddha is in the background. Interactively modifying the interocular distance, we computed the entropy of the segmented disparity map. When the interocular distance is zero, then the disparity map is constant, so there is only one class in its segmentation, which results in zero entropy. Increasing the interocular distance, the dynamic range of the disparity map gets larger, which in itself would result in more segmented classes and thus would increase the information. However, ghosting, i.e. objects lacking stereo pairs show up, which reduce information content, and disparity values that are far from zero are quantized more radically, which results in further reduction. The different effects establish an optimum, and we observe an information increase when the interocular distance grows from zero, then after reaching its maximum it starts to decrease. The maximum can be used as an optimal setting.

Figure 6. shows a Room scene where a lot of objects can be found near to the edge of the view frustum. These objects quickly disappear from one of the images causing ghosting and a rapid information loss.

⁸ P. Longhurst, K. Debattista, A. Chalmers, "A GPU based saliency map for high-fidelit selective rendering.", In *Proceedings of the 4th international conference on Computer graphics, virtual reality, visualisation and interaction in Africa, AFRIGRAPH '06*, ACM, (2006), pp. 21-29.

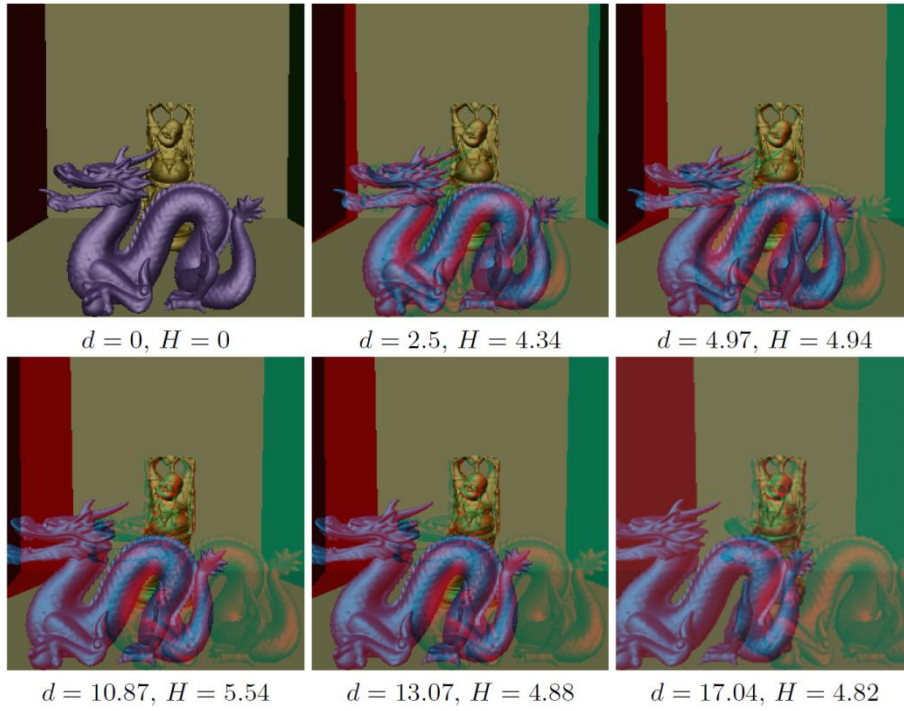


Figure 5. Cornell box with a dragon and a Buddha: Stereo images with interocular distances d and entropy values H .

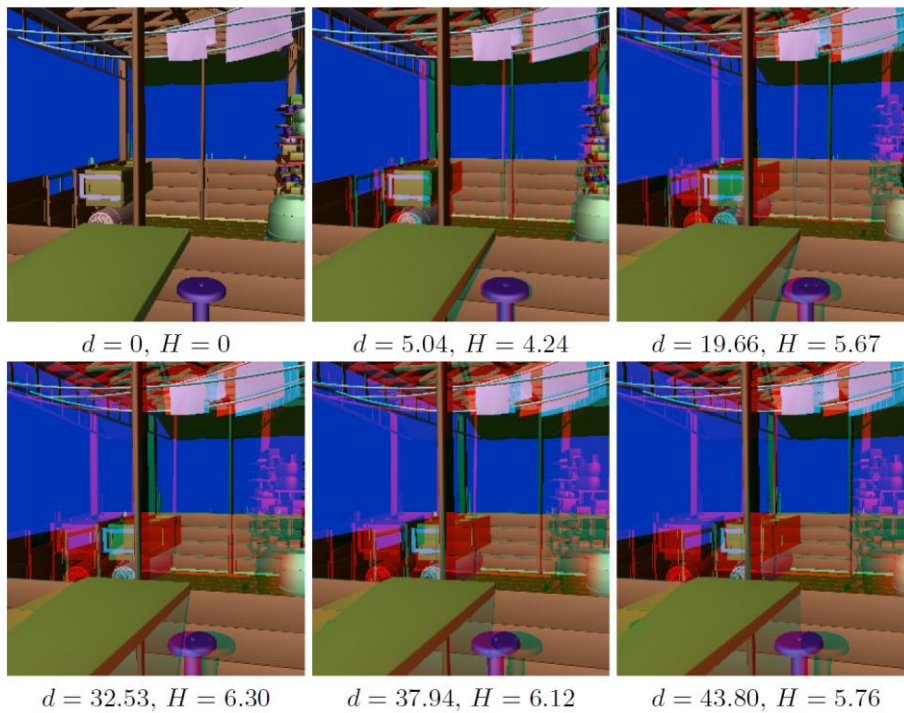


Figure 6. Room scene: Stereo images with interocular distances d and entropy values H .

Conclusions

This paper proposed the application of information theoretic measures for automatically finding stereo imaging parameters that provide aesthetic 3D images.

Based on the recognition that stereo imaging should be responsible for depth perception, we build our model on the information available in the disparity map. However, to consider also perceptual phenomena, like the ability of the human eye to match point pairs and the contradiction of accommodation according to the screen distance and vergence caused by objects at a different distances, we non-linearly transformed and quantized the disparity map before the entropy is calculated.

Acknowledgments

This work has been supported by OTKA K-104476.

Dr. László Szirmay-Kalos is the head of the Computer Graphics Lab and the Department of Control Engineering and Information Technology of the Budapest University of Technology and Economics. He received PhD in 1992 and DSc in 2001 from the Hungarian Academy of Sciences. His research interests include photorealistic and illustrative rendering, scientific visualization, visual aesthetics, games and medical systems. He is the author of more than 200 scientific publications in this field. He is a fellow of Eurographics.

Web: www.iit.bme.hu/en/munkatarsak/szirmay-kalos-laszlo
email: szirmay@iit.bme.hu




Dr. Pirkko Oittinen is full professor at the Department of Media Technology of Aalto University's School of Science where she is the head of the Visual Media research group. She received the degrees of MSc (Eng) in 1971 and the Doctor of Science in Technology in 1977 from Helsinki University of Technology (now part of Aalto University), was nominated associate professor in 1982 and full professor in 1992 in the same university. Her current research interests focus on computational and experiential understanding of the high level quality of images and video.

email: pirkko.oittinen@aalto.fi



Balázs Teréki is an MSc student of the Informatics Course of the Budapest University of Technology and Economics. His interests include computer graphics, realistic rendering using Monte Carlo methods, games programming and image processing.

email: terekib@gmail.com



Abstract

Computer graphics renders images about virtual scenes using some natural analogy, like photography, drawing, X-ray, flows, etc. To communicate the data of the virtual world, we need additional data like light sources, camera settings, drawing style, etc. The setting of such parameters should maximize the information received by the user. Informational aesthetic measures can quantify the quality of images in terms of information content. Using such measures we can automatically select the “better” solution from two alternatives, or optimizing with the objective of maximizing the aesthetic measure, the “best” option can be found. Such approaches can replace lengthy user interaction sequences tuning camera or light source parameters and automatically propose good camera and lighting settings. In stereoscopic imaging one of the goals is to provide depth perception, so the quality of such images can be characterized by how naturally the depth is perceived by human observers. In this paper, we propose an algorithm to measure the depth information of a stereo image pair, which is based on the entropy of the disparity map, and also takes into account perceptual metrics. With the proposed measures, the parameters of a stereoscopic virtual camera can be optimized, including, for example, the interocular distance.