Comparison and Evaluation of Viewpoint Quality Estimation Algorithms for Immersive Virtual Environments – Additional Material

Sebastian Freitag, Benjamin Weyers, Andrea Bönsch, Torsten W. Kuhlen

Virtual Reality Group – RWTH Aachen University
JARA – High Performance Computing

E-mail: {freitag|weyers|boensch|kuhlen}@vr.rwth-aachen.de.

The knowledge of which places in a virtual environment are interesting or informative can be used to improve user interfaces and to create virtual tours. Viewpoint Quality Estimation (VQE) algorithms approximate this information by calculating quality scores for viewpoints. In [FWBK15], we introduced three new VQE algorithms, and compared them against each other and six existing techniques, by applying them to two different virtual scenes. Furthermore, we conducted a user study to obtain a quantitative evaluation of viewpoint quality.

This document describes the additional material provided for [FWBK15], which includes additional scene information, and raw data from algorithm outputs and the user study.

1 License

2 Information about the Scenes

For the comparison, we used two different static scenes: a house and an office floor. In both of them, the ground is at the same height everywhere. Furthermore, only little information is contained in textures, and there are no (semi-)transparent objects. High-res perspectivic and orthographic top views of both scenes can be found in the subfolder `scenes` of the dataset. The perspectivic views are labeled, open doorways are marked with dotted lines. Additionally, the folder also contains an overview video (`overview.wmv`) showing a walkthrough through both scenes. For access to the original scene files, please contact the first author.

House This scene contains a small, detailed house, with furnished rooms and many different objects, especially in the kitchen (kitchen appliances and tools) and the living room (books and DVDs in cupboards, a laptop on a table, etc.). The house contains five rooms (bedroom, bathroom, hall, kitchen, living room), a patio, and an empty outside area without modeled details. All of them are connected by open doors. The total size of the scene is 20.35m$\times$15.45m, the total indoor area is $\approx$113m^2, the total walkable outdoor area is $\approx$110m^2.

Office This scene contains a large office floor, significantly less-detailed than the house scene, that we used as a representative for sparsely-modeled, repetitive indoor scenes. It contains similar interiors for many rooms, as well as mostly empty corridors. Most rooms are offices with one or several desks that are very similar to each other, except for some details like the presence/absence or the arrangement of certain office utensils. The offices also contain an (always identical) pinboard. In addition, there is a significantly larger and differently furnished boss office, a seminar room, a meeting room, and a kitchen. All these rooms are connected by open doors and corridors. Furthermore, there are some closed rooms without content, as is often the case for irrelevant areas to save modeling effort (marked “empty” in the perspectivic top view). The total size of the scene is 46.40m$\times$30.60m, the total indoor area excluding closed and inaccessible rooms is $\approx$1020m^2.

3 Algorithm Results

The algorithm scores were computed using a regular 5cm grid in the House scene and a 10cm grid in the Office scene, i.e., in the result images, one pixel corresponds to 5$\times$5 / 10$\times$10 cm. To account for different realistic eye heights, values were averaged over heights of 1.45m, 1.50m, 1.55m, 1.60m, 1.65m, 1.70m, 1.75m and 1.80m, resulting in 1,006,104 (house) or 1,135,872 (office) score computations for 125,763 or 141,984 resulting data values (one per x/z coordinate). Note that the scores are mostly very similar for the different eye heights, as almost all objects are similarly visible from all considered heights, varying more only in obstructed areas (e.g., in cupboards or potted plants).

The dataset includes both raw data (algorithm scores averaged over the given heights) and visualizations thereof (where one pixel corresponds to one data value).

3.1 Data

The data files contain the scores produced by all algorithms on both scenes in ASCII format. Each line in the files corresponds to one data value and contains the X and the Z coordinate in the scene (in meters) as well as the algorithm score, separated by tab characters. All data files of algorithm results for both the House and Office scene can be found in the subfolder `algorithm_results/data` of the dataset.

3.2 Visualizations

In all visualizations, the following color map was used:

All values were normalized before applying the color map, i.e., the lowest value in the scene was always mapped to black, and the highest value to white. For broadly logarithmically-scaled measures, exponentiated versions are included as well. For these, all values were first exponentiated ($newscore = 2^{score}$) and then normalized.

All visualizations of algorithm results for both the House and Office scene can be found in the subfolder `algorithm_results/visualizations` of the dataset.

3.3 Files

The following files of algorithm results are included in the dataset:

Surface area entropy (also called *viewpoint entropy*) [VFSH01, VFSH03]
`surfaceareaentropy_{house|office}.log` – Scores

`surfaceareaentropy_{house|office}.png` – Vis. of scores
`surfaceareaentropy_{house|office}_exp.png` – Vis. of exponentiated scores
`surfaceareaentropy_{house|office}_exp2.png` – Vis. of exponentiated scores, alternative color mapping [1]

Relative surface area entropy [SPFG05]
`relativesurfaceareaentropy_{house|office}.log` – Scores
`relativesurfaceareaentropy_{house|office}.png` – Vis. of scores
`relativesurfaceareaentropy_{house|office}_exp.png` – Vis. of exponentiated scores

Ratio of visible area [PPB+05]
`visibleararatio_{house|office}.log` – Scores
`visibleararatio_{house|office}.png` – Vis. of scores

Curvature entropy [PKS+03]
`curvatureentropy_{house|office}.log` – Scores
`curvatureentropy_{house|office}.png` – Vis. of scores
`curvatureentropy_{house|office}_exp.png` – Vis. of exponentiated scores

Mesh saliency [LVJ05]
`meshsaliency_{house|office}.log` – Scores
`meshsaliency_{house|office}.png` – Vis. of scores

Depth map stability (also called *depth-based view stability*) [Váz09]
`depthstability_{house|office}.log` – Scores
`depthstability_{house|office}.png` – Vis. of scores

Object area entropy [FWBK15]
`objectareaentropy_{house|office}.log` – Scores

[1] Some of the maxima produced by *surface area entropy* are extreme, such that the exponentiated scores often contain very large differences between most values and the highest values. The alternative color mapping images show the exponentiated scores with a different normalization (for illustration purposes only, due to the arbitrary normalization that cuts off very high values).

`objectareaentropy_{house|office}.png` – Vis. of scores
`objectareaentropy_{house|office}_exp.png` – Vis. of exponentiated scores

Relative object area entropy [FWBK15]
`relativeobjectareaentropy_{house|office}.log` – Scores
`relativeobjectareaentropy_{house|office}.png` – Vis. of scores
`relativeobjectareaentropy_{house|office}_exp.png` – Vis. of exponentiated scores

Object uniqueness [FWBK15]
`objectuniqueness_{house|office}.log` – Scores
`objectuniqueness_{house|office}.png` – Vis. of scores

4 User Study Data

In the user study, participants were teleported to a series of viewpoints in random order. They rated each one with a score between 0 and 4 using five buttons on an input device before they were moved to the next one. They were told to use the lowest score 0 for the most uninformative viewpoint(s) in the scene, and the highest score 4 for the most informative one(s). To include possibly different concepts of informativeness (of different people) in the rating, participants were not given concrete examples of what constituted an informative viewpoint. To avoid bias, the viewpoints were chosen by a regular sampling, using a 1m grid in the house scene, and 3m in the office. Furthermore, in the house scene, only viewpoints in the house and on the patio were considered to reduce their total number. In total, participants saw 159 viewpoints in the house and 176 in the office (which we found to be a high, but manageable number if evaluated in separate sessions), and took 16 minutes to complete each scene on average. They were allowed to physically turn around 360°, but not to move away or crouch down.

For the evaluation, all user scores were averaged for each viewpoint. The resulting mean scores as well as visualizations thereof can be found in the `study/user_evaluations` subfolder of the dataset, folder structure and file naming scheme are the same as used for the algorithm results in section 3. In the visualizations, the color map is also the same as used for the algorithm results (black corresponds to an average score of 0, white to the highest average score present), each data point is represented by a block of 20×20 pixels. Areas for which no user values were collected are colored gray (only applies to the house scene).

For each viewpoint quality estimation algorithm, we calculated the score for each viewpoint as the average over scores in a 0.30m×0.30m regular grid (0.05m increments) centered over the position, including heights from 1.45m to 1.80m in 0.05m increments, to account for small user movements. The raw scores as well as visualizations of the normalized results can be found in the `study/algorithm_results` subfolder of the dataset, folder structure, naming scheme and color map of the visualizations are the same as used for the algorithm results in section 3. For logarithmic measures, their exponentiated score is also included, as the user evaluation was linear in scale. In [FWBK15], these results were compared against the user evaluations using Bhattacharyya coefficients.

References

[FWBK15] S. Freitag, B. Weyers, A. Bönsch, and T. W. Kuhlen. Comparison and Evaluation of Viewpoint Quality Estimation Algorithms for Immersive Virtual Environments. In *Proceedings of the 25th International Conference on Artificial Reality and Telexistence and the 20th Eurographics Symposium on Virtual Environments*, ICAT-EGVE 2015, 2015.

[LVJ05] Chang Ha Lee, Amitabh Varshney, and David W Jacobs. Mesh Saliency. In *ACM Trans. on Graphics*, volume 24, pages 659–666, 2005.

[PKS+03] David L Page, AF Koschan, Sreenivas R Sukumar, Besma Roui-Abidi, and Mongi A Abidi. Shape Analysis Algorithm Based on Information Theory. In *Int. Conf. on Image Processing*, volume 1, pages 229–232, 2003.

[PPB+05] Oleg Polonsky, Giuseppe Patané, Silvia Biasotti, Craig Gotsman, and Michela Spagnuolo. What's in an image? *The Visual Computer*, 21(8-10):840–847, 2005.

[SPFG05] Mateu Sbert, Dimitri Plemenos, Miquel Feixas, and Francisco González. Viewpoint Quality: Measures and Applications. In *Proc. Eurographics Workshop on Computational Aesthetics in Graphics, Visualization and Imaging*, pages 185–192, 2005.

[Váz09] Pere-Pau Vázquez. Automatic View Selection through Depth-Based View Stability Analysis. *The Visual Computer*, 25(5-7):441–449, 2009.

[VFSH01] Pere-Pau Vázquez, Miquel Feixas, Mateu Sbert, and Wolfgang Heidrich. Viewpoint Selection using Viewpoint Entropy. In *VMV*, volume 1, pages 273–280, 2001.

[VFSH03] Pere-Pau Vázquez, Miquel Feixas, Mateu Sbert, and Wolfgang Heidrich. Automatic View Selection Using Viewpoint Entropy and its Application to Image-Based Modelling. In *Computer Graphics Forum*, volume 22, pages 689–700, 2003.