

Towards a quality metric for dense light fields

Vamsi Kiran Adhikarla¹ Marek Vinkler¹ Denis Sumin¹ Rafał K. Mantiuk³
 Karol Myszkowski¹ Hans-Peter Seidel¹ Piotr Didyk^{1,2}

¹MPI Informatik ²Saarland University, MMCI ³The Computer Laboratory, University of Cambridge

Abstract

Light fields become a popular representation of three-dimensional scenes, and there is interest in their processing, resampling, and compression. As those operations often result in loss of quality, there is a need to quantify it. In this work, we collect a new dataset of dense reference and distorted light fields as well as the corresponding quality scores which are scaled in perceptual units. The scores were acquired in a subjective experiment using an interactive light-field viewing setup. The dataset contains typical artifacts that occur in light-field processing chain due to light-field reconstruction, multi-view compression, and limitations of automultiscopic displays. We test a number of existing objective quality metrics to determine how well they can predict the quality of light fields. We find that the existing image quality metrics provide good measures of light-field quality, but require dense reference light-fields for optimal performance. For more complex tasks of comparing two distorted light fields, their performance drops significantly, which reveals the need for new, light-field-specific metrics.

1. Introduction

A light field can be seen as a generalization of a 2D image, which encodes most of the depth cues and allows to render a scene simulating arbitrary optics (e.g., defocus blur) [16]. It is a convenient representation for multiscopic and light-field displays [43], but also attractive format for capturing high-quality cinematographic content, offering new editing possibilities in post-production [19]. Due the enormous storage requirements, light fields are usually sparsely sampled in spatial and angular dimensions, stored using lossy compression, and reconstructed later. It is unclear how the distortions introduced on the way affect the perceived quality.

Similar problems have been addressed for 2D images, videos, and sparse multiview content. Many quality metrics have been designed to predict perceived differences between vari-

ous versions of the same content [1]. However, measuring quality for dense light fields still remains a complex task. While several works applied the existing metrics to such content [12, 8], their performance has never been systematically evaluated in this context. One of the challenges is acquiring dense light-field data to validate a metric. Wide baselines as in multi-camera rigs [44] need to be considered, and the reference light fields should be sufficiently dense to avoid uncontrolled visual artifacts. Obtaining human responses for light-field distortions is also difficult due to current display limitations. This work is an attempt to overcome these problems by first building a new dense light-field dataset which is suitable for testing quality metrics, and second, using a custom light-field viewing setup to obtain the quality judgments for this dataset. The collected subjective scores are used to evaluate the performance of existing metrics in the context of dense light fields.

We focus on light-field-specific angular effects akin to motion parallax, complex surface appearance, and binocular vision that arise in free viewing experience. To capture a rich variability over these effects and make quality scaling in our perceptual experiments tractable, we design fourteen real and synthetic scenes and introduce light-field distortions that are specific to light-field reconstruction, compression, and display. We then run a pair-wise comparison experiment over light-field pairs, and derive perceptual scaling of differences between original and distorted stimuli. This allows us to investigate the suitability of a broad spectrum of existing image, video, and multiview quality metrics to predict such perceptual scaling. We also propose simple extensions of selected metrics to capture the angular aspects of light-field perception. While the original metrics are not meant for light fields, our results show that they can be used in this context, given a dense light field as the reference. We also demonstrate that the robustness of such metrics predictions drops when evaluating the quality between two distorted light fields. The main contributions of this work are:

- a publicly available dense light-fields dataset that is designed for training and evaluating quality metrics;
- a perceptual experiment that provides human quality

judgments for several typical light-field distortions;

- an evaluation, analysis, and extensions of existing quality metrics in the context of light fields;
- identified challenges of quality assessment in light fields, such as the need for a high quality reference.

2. Previous works

In this section, we provide an overview of existing datasets for light fields as well as the experiments that measure the perceived distortions in various types of content.

Light-field datasets: There are several publicly available light-field datasets. The most popular ones are: 4D light-field dataset [42] containing seven synthetic scenes and five real-world scenes, Stanford archive [7] with twenty 4D light fields, and Disney 3D light-field dataset [14] containing five scenes. Although the first two datasets provide a good quality and reasonable number of light fields, they are captured over very narrow baselines that are insufficient for the new generation autostereoscopic displays. The Disney dataset provides high spatio-angular resolution light fields; however, they are few and do not have consistent spatial and angular resolution, which makes it difficult to use in quality evaluations. In the context of quality evaluation of 3D light fields, three real-world light fields are provided in the IRCCyN/IVC DIBR Images database [4]. These contain several scenes captured along a wide baseline at the cost of reduced angular resolution. Tamboli et al. [35] provided 360° round table shots of three scenes that are used for quality evaluation on a 3D light-field display. These are rather simple scenes with single objects and the images contain a lot of noise. In our work, we provide first consistent dataset of dense, complex-scene light fields with large appearance variation. We use the dataset for training and evaluating quality metrics. The database can also serve as a ground truth for automultiscopic displays.

Metrics and experiments: Because of their proven efficiency on 2D images, 2D objective metrics are viable candidates for evaluating light-field quality. Yasakethu et al. [46] tested the suitability of objective measures – Structural Similarity (SSIM) [40], Peak Signal-to-Noise Ratio (PSNR) and Video Quality Metric (VQM) [25] for quality assessment for stereoscopic and 2D+Depth videos that are compressed at different bitrates. They carried out subjective experiments on an autostereoscopic display and showed that 2D metrics can be used separately on each view to assess 3D video quality. They used few sequences and studied only compression artifacts. Several metrics have been proposed to determine the quality of synthesized views from multiview images. Bosc et al. [4] advocated two measures for assessing the quality of synthesized views. However, they did not conduct thorough subjective studies. Solh et al. [34] presented a metric

for quantifying the geometric and photometric distortions in multiview acquisition. Bosc et al. [3] suggested a method to assess the quality of virtual synthesized views in the context of multiview video. Battisti et al. [2] proposed more sophisticated framework for evaluating the quality of depth image based rendering techniques by comparing the statistical features of wavelet subbands and used image registration and skin detection steps for additional optimization. Sandic et al. [30] exploited multi-scale pyramid decompositions with morphological filters for obtaining the quality of intermediate views and showed that they achieve significantly higher correlation with subjective scores. These methods form a class of metrics specific to view-interpolation artifacts, and 2D stimuli containing the interpolated views are used for subjective experiments.

Vangorp et al. [38] ran a psychophysical study to account for the plausibility of visual artifacts associated with view interpolation methods. They considered such artifacts as a function of different number of input images; however, they limited their study to monocular viewing and Lambertian surfaces. An experiment was also performed for precomputed videos, so that the impact of user’s interaction and dynamic aspects of free viewing could be judged. More recently, this work was extended to transitions between videos [37]. Similar studies were also performed in the context of panoramas [23]. Tamboli et al. [35] conducted subjective studies on a 3D light-field display. Users were asked to judge the quality as perceived from different viewing locations in front of the display and the scores were averaged over all locations. The user could rate the quality only from a certain viewing position. Moreover, they only considered three distinct scenes. We believe that, for inferring a light-field quality, all the views should be taken into account at the same time.

Light-field displays: Our work focuses on wide-baseline 3D light fields which enable perfect simulation of stereoscopic viewing and continuous horizontal motion parallax crucial for new light-field displays. Although many light-field display designs exist [22], including more advanced ones that provide focus cues [20], they suffer from several drawbacks such as limited field of view, discontinuous motion parallax, visible crosstalk, and limited depth budget. Several strategies have been proposed to minimize these artifacts by filtering the content [47, 10] and manipulating depth [15, 9, 22]. However, display designs that enable displaying reference light fields for quality measurements are still unavailable.

3. Data collection

Our dataset consist of light fields which are parameterized using two parallel planes [16]. We consider only horizontal





Figure 1: Representative images of all light fields in our collection. Below each image representative EPIs are presented.

motion parallax that can be described using one plane and a line that is parallel to it.

More formally, we denote our light fields as $L(\omega, \mathbf{x}) \in (\mathbb{R} \times \mathbb{R}^2) \rightarrow \mathbb{R}^3$, where ω is a position on the line, and \mathbf{x} is a position on the plane. We refer to them as *angular* and *spatial* coordinates, respectively. In practice, ω describes a position of the viewer, and \mathbf{x} is a coordinate of the observed image. Below, we describe the acquisition of our light fields.



3.1. Scenes

We designed and rendered nine synthetic and captured five real-world scenes (Figure 1). They span a large variety of different conditions, *e.g.* outdoor/indoor, daylight/night etc. They also contain objects with large range of different appearance properties. The scene objects distribution in depth is widely varied to study the artifacts resulting from disocclusions and depth discontinuities. For capturing real-world scenes, we used a one-meter long motorized linear stage with Canon EOS 5D Mark II camera and 50 mm and 28 mm lenses. After capturing all views, we performed lens distortion correction using PTLens [26], estimated the camera poses using Voodoo camera tracker [39], and rectified all the images using the baseline drawn from the first to last camera using the approach in [11]. For rendered images, we used cameras with off-axis asymmetric frustums. For real-world scenes, the same effect was achieved by applying horizontal shift to the individual views. All the light fields are of identical spatial and angular resolution ($960 \times 720 \times 101$). The angular resolution was chosen high enough to avoid visible angular aliasing. This was achieved by assuring that the maximum on-screen disparities between consecutive views are around 1 pixel. To guarantee a comfortable viewing, the

total disparity range during the presentation was limited to 0.2 visual degree [31].

3.2. Distortions

We considered typical light-field distortions that are specific to transmission, reconstruction, and display. For each distortion, we generated multiple light fields by varying the distortion severity level. The exact levels were chosen to keep the differences between two consecutive levels small and similar. To this end, we conducted a small pilot study with 10 distortion levels, and then, selected the final levels manually.

Transmission: To transmit the light-field data, an efficient data compression algorithm is highly required. We consider well-known 3D extension of HEVC encoder [36]. The light-field views are encoded into a bit stream at various quantization steps, and then, decoded back from the bit stream using the 3D-HEVC coder. We chose the following quantization steps: {25, 29, 33, 37, 41, 45}.

Reconstruction: Light-field reconstruction techniques are used to recover a dense light field from sparse view samples. They interpolate the missing views using several techniques which alter the nature and appearance of the distortion. We chose the distortions resulting from linear (LINEAR) and nearest neighbor (NN) interpolation, as well as image warping using optical flow estimation (OPT). We also investigated the impact of using quantized depth maps (DQ). All the distortions are parametrized by the angular subsampling factor k (the distortion severity) that defines the angular resolution of the light field prior to applying the reconstruction technique. We considered $k \in \{2, 5, 8, 11, 18, 25\}$. The linear filter reconstructs dense light field by blending the reference views, and the NN method clones the closest reference view. For OPT method, we used the TV-L1 optical flow [29], and apply an image-warping technique [5] to synthesize in-between views. For DQ, we considered ground-truth depth map which is quantized using 8 discrete levels. Then, we

used the same image-warping technique [5] to reconstruct the light field. As this distortion requires ground-truth depth information, it is only applied to the synthetic scenes.

Display: As an example of multiview autostereoscopic display artifacts, we chose a crosstalk between adjacent views, which can be modeled using a Gaussian blur in angular domain [18]. Consequently, we include such artifacts into our dataset (GAUSS). In particular, we considered the same angular subsampling parameters used in light-field reconstruction distortions and created hypothetical displays with corresponding number of views. The upsampling to higher resolution light field was achieved by using the display crosstalk model.

Four different distortions with all severity levels were applied to every scene. To all synthetic scenes we apply NN, LINEAR, OPT, and DQ. For all real-world scenes, we used NN, OPT, GAUSS, and HEVC. Including original light fields, our database consists of 350 different light fields and it is available online [24]. The examples of resulting artifacts are presented in Figure 2. Please refer to supplemental materials for the whole light-field dataset.




Figure 2: Examples of distortions introduced to our light field for one of our scenes (BARCELONA). The images visualize central EPI of each of the distorted light fields and the enlarged portion of it is shown in the bottom row.

4. Experiment

To acquire subjective quality scores that enable both training and testing different quality metrics, we performed a large-scale subjective experiment.

Equipment: To simulate stereoscopic viewing with high-quality motion parallax, we used on our own setup (Figure 3) that consists of ASUS VG278 27" Full HD 120 Hz LCD desktop monitor and NVIDIA 3D Vision 2 Kit for displaying stereoscopic images. Motion parallax was reproduced using a custom head tracking in which a small LED headlamp was tracked using a Logitech HD C920 Pro webcam (refer to the supplemental video). The head tracking allowed the partici-

pants to view light fields in an unconstrained manner. The viewing distance was approximately 60 cm, and users could move their heads along a baseline of 20 cm in the direction parallel to the screen plane. The eye accommodation was fixed to the screen and did not change with eye vergence. The display was operated at the full brightness to minimize the effect of luminance on depth perception [9].




Figure 3: Experiment session: viewer's position is tracked using a head lamp and a webcam, a pair of NVIDIA 3D Vision 2 Kit active glasses provides stereoscopic viewing.

Stimuli: Each stimuli was a pair of light fields. As our scaling procedure used for obtaining quality scores (Section 5) can handle an incomplete set of comparisons and prefers when more comparisons are made for pairs of similar quality [33], each pair consisted of light fields with neighboring severity levels of the same distortion type. This results in 336 different stimuli which were presented stereoscopically.

Task: We experimented with direct rating methods, such as ACR [13], in order to measure Mean-Opinion-Scores of the distorted images. However, we found these methods to be insensitive to subtle but noticeable degradation of quality. Also participants found the direct rating task difficult. Therefore, we decided to use a more sensitive pair-wise comparison method with a two-alternative-forced-choice. In each trial, the participants were shown a pair of light fields side-by-side, and the task was to indicate the light field that a user "would prefer to see on a 3-dimensional display". Participants were given unlimited time to investigate the light fields, but they were allowed to give their response only after 80% of all perspective images were seen. The order of the light-fields pairs as well as their placement on the screen were randomized. Before each session, the participants were provided with a form summarizing the task, and a training session was conducted to familiarize participants with the experiment.

Participants: Forty participants took part in the test, including both male (20) and female (20) aged 24–40 with normal or corrected-to-normal vision. Each subject performed the test in three sessions within one week. In one session, the participants saw 120–180 light-field pairs consisting of all the test conditions, but for a subset of the scenes. For a given subject, two test sessions were allowed during a single day, and these were separated by at least an hour of break.

5. Analysis of subjective data

The results of pair-wise comparison experiment are usually scaled in just-noticeable-differences (JNDs). We observed that considering measured differences as “noticeable” leads to incorrect interpretation of the experimental results. Two stimuli are 1 JND apart if 75% of observers can see the difference between them. However, our experimental question was not whether observers can tell if the light fields are different, but rather which one has higher quality. As shown in Figure 4, a pair of stimuli could be noticeably different from each other ($\text{JND} > 1$), but they could appear to have the same quality. For that reason, we denote measured values as just-objectionable-differences (JODs). These units quantify the quality difference in relation to the perfect reference image. Note that the measure of JOD is more similar to visual equivalence [28] or to the quality expressed as a difference-mean-opinion-score (DMOS) rather than to JNDs.




Figure 4: Illustration of the difference between just-objectionable-differences (JODs) and just-noticeable-differences (JNDs). The image affected by blur and noise may appear to be similarly degraded in comparison to the reference image (the same JOD), but they are noticeably different and therefore several JNDs apart. The mapping between JODs and JNDs can be very complex and the relation shown in this plot using Cartesian and polar coordinates is just for illustration purposes.

To scale the results in JOD units we used a Bayesian method based on the method of Silverstein and Farrell [33]. It employs a maximum-likelihood-estimator to maximize the probability that the collected data explains JOD-scaled quality scores under the Thurstone Case V assumptions [27]. The optimization procedure finds a quality value for each pair of light fields that maximizes the likelihood modeled by the binomial distribution. Unlike standard scaling procedures, the Bayesian approach robustly scales pairs of conditions for which there is unanimous agreement. Such pairs are common when a large number of conditions are compared. It can also scale the result of an incomplete and imbalanced pair-wise design, when not all the pairs are compared and some are compared more often. As the pair-wise comparisons provide relative quality information, the JOD values are relative. To maintain consistency across the scenes, we fix the starting point of the JOD scale at 0 for different dis-

tortions and thus the quality degradation results in negative JOD values.

The results of the subjective quality assessment experiment are shown in Figure 5. The error bars represent 95% confidence intervals, relative to the reference light field, computed by bootstrapping by sampling with replacement. The results show interesting patterns in the objectionability of different distortions. OPT offers a consistent performance improvement over NN. The only exception is the *Furniture* scene featuring thin and irregularly shaped foreground objects, in which case all types of view interpolation are more objectionable than the selection of the nearest single view. The optical flow interpolation works better for real-world scenes as there are more features that can be detected. The LINEAR interpolation in most of the cases results in the worst performance, except for small distortion levels, which may indicate that visible blur due to this distortion is strongly objectionable. Similar findings have been reported by Van-gorp *et al.* [38] in their study on the visual performance of view interpolation methods in monocular vision. HEVC and GAUSS distortions are usually the easiest to detect as they induce significant amount of spatial distortion when compared to others. Overall the results show clearly that light-field quality is scene-dependent and successful quality metric must predict the effect of scene content on the visibility of light-field distortions.

6. Evaluation of quality metrics

We considered several popular image, video, stereo, and multiview quality metrics. We briefly describe the metrics and then show their individual performance on our dataset. For obtaining the quality of a light field using image quality metrics, we apply the metrics on individual light-field images and then average the scores over all images.

Quality metrics: Although studies show that perceptual metrics perform better than an absolute difference (AD) [17], because of its significant usage in image quality assessment, we considered peak signal-to-noise ratio (PSNR). We also investigated SSIM_{2D} [40], which is widely used on 2D images, and its extensions to angular domains – $\text{SSIM}_{2D \times 1D}$ and SSIM_{3D} . SSIM_{3D} computes the same statistics as standard SSIM_{2D} but on 3D patches extracted from the light-field volume. $\text{SSIM}_{2D \times 1D}$ uses $2D \times 1D$ patch which contains a 2D window extracted from a particular view and a 1D row of pixels that extends from the center of the 2D window in the angular domain (see Figure 6). We applied the metrics to all light-field images without resampling and averaged the scores over all images. Although we experimented with various pooling strategies, we found that the average value performs best. Due to better performance, we chose the angular window sizes of 32 and 64 pixels for $\text{SSIM}_{2D \times 1D}$ and

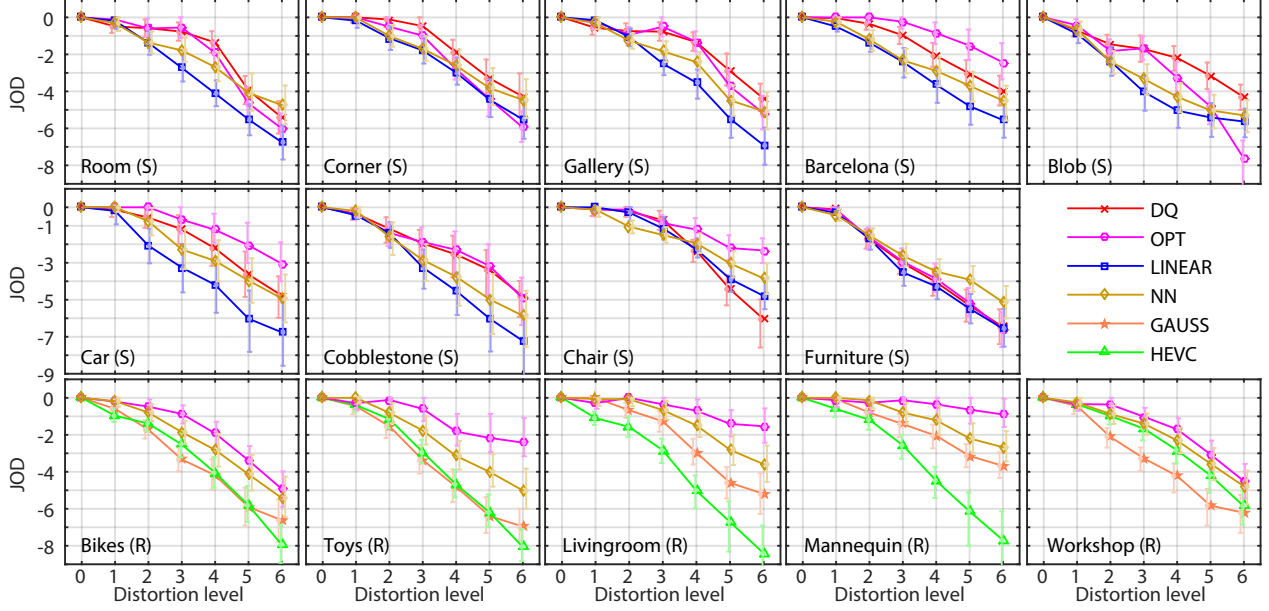


Figure 5: The results of the subjective quality assessment experiment. The distortion level indicates the distortion severity: 0–reference 6–severest distortion level. JOD is the scaled subjective quality value. The error bars denote 95% confidence interval. The bars are horizontally displaced to avoid overlapping. The scene names are indicated in the corner of each plot. The character in parenthesis after the scene name indicates whether the scene is synthetic (S) or real-world (R).

SSIM_{3D} respectively. We also considered a multi-scale version of SSIM_{2D} – MS-SSIM [41] which extends SSIM_{2D} to compute differences on multiple levels. We also used GMSD [45] which provides good performance over a rich collection of image datasets. The most advanced 2D metric considered in our experiments was HDR-VDP-2 [21] which stands out among perception-based quality metrics.




Figure 6: Patches used in our extensions of SSIM_{2D}.

We further considered the NTIA General Model – VQM [25] which was standardized for video-signals evaluation (ANSI T1.801.03-2003). For this metric, light-field images are input in a form of video panning from the leftmost view to the rightmost view and back. We also chose the stereoscopic image quality metric – SIQM [6] that is based on the concept of cyclopean image where, we averaged scores obtained from all stereo pairs shown in our experiment. To capture the full range of stereo quality metrics, we also included a stereoscopic video quality metric STSD_{LC} [32].

Finally, we chose metrics that address multiview data and account for interpolation artifacts. 3DSWIM proposed by

Battisti et al. [2] first shift-compensates blocks from the reference and distorted (interpolated) images. These matched blocks undergo the first level of Haar wavelet transform and histogram of the sub-band corresponding to horizontal details in the block is computed. Finally, the Kolmogorov-Smirnov distance of these histograms is taken as the metric prediction. Another metric for the multiview video is MP-PSNR [30]. It computes the multi-resolution morphological pyramid decomposition on the reference and test images. Detail images of the top levels of these pyramids are then compared through the mean squared error. The resulting per pixel errors maps are then pooled and converted to a peak signal-to-noise ratio measure.

6.1. Metric performance comparison

The quality values predicted by each metric are expected to be related to JOD values, but this relation can be complex and non-linear. To account for this relation, we follow a common practice and fit a logistic function:

$$q(o) = a_1 \left\{ \frac{1}{2} - \frac{1}{1 + \exp[a_2(o - a_3)]} \right\} + a_4 o + a_5 \quad (1)$$

where o is the output of a metric. The parameters $a_{1..5}$ are optimized to minimize a given goodness-of-fit measure. We computed several such measures, such as Spearman rank-order correlation, or MSE, which can be found in the supplementary materials. Here we report the reduced chi-squared

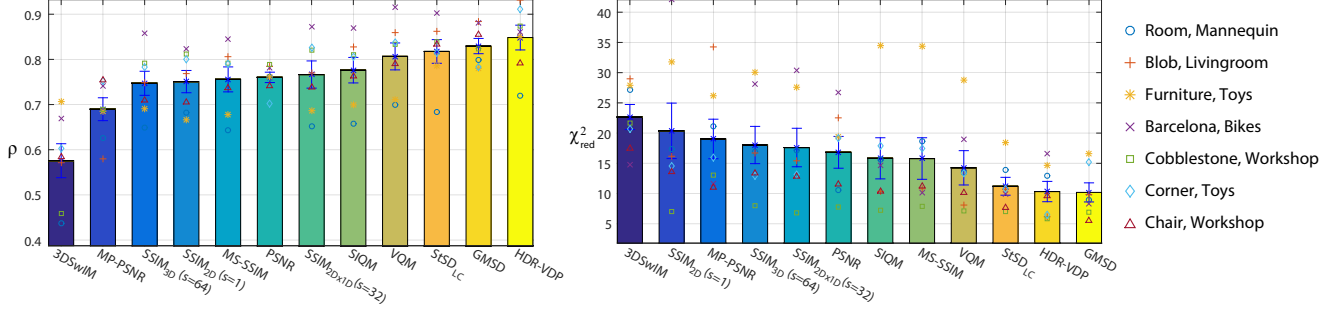


Figure 7: The goodness-of-fit scores for the metrics expressed as Pearson Correlation Coefficient (ρ) and reduced chi-square (χ^2_{red}) after cross-validation. The results for each cross-validation fold are shown. $\chi^2_{red} = 1$ indicates that the goodness of fit between the metric predictions and the subjective data is in perfect agreement with the measured subjective variance and $\rho = 1$ indicates perfect positive linear relation between objective scores and JODs. The error bars represent standard error.

statistic (χ^2_{red}) and Pearson correlation coefficient (ρ). χ^2_{red} is computed as a weighted average of the squared differences, in which weights are the inverse of sample variance. This accounts for the fact that larger JOD values are more uncertain (refer to Figure 5), and therefore, the accuracy of their prediction can be lower. For a fair comparison, we employed a seven fold cross-validation across different scenes. We measured the goodness-of-fit on two randomly chosen scenes in a cross-validation fold and averaged the results over all folds. The resulting Pearson correlation and χ^2_{red} values are shown in Figure 7. The performance of the metrics on individual distortions are shown in Figure 8. A more elaborate analysis including the evaluation on real-world and synthetic scenes separately is presented in the supplementary materials.

The results show good performance of 2D image and video quality metrics. This is unexpected as our dataset was meant to emphasize visibility of angular artefacts, which are not directly considered by these metrics. We observed, however, that angular distortions indirectly translate into the differences in spatial patterns, which could explain the good performance. We hypothesize that relatively better performance of HDR-VDP-2 and GMSD is achieved by detecting changes in contrast across multiple scales, which in case of HDR-VDP-2 is additionally backed by perceptual scaling of distortions and discarding of those that are invisible. A comparable performance of video (VQM) and stereoscopic (STSD_{LC}) metrics can be explained by their emphasis on the relation between neighboring views, which in some way captures angular aspects of light-field viewing. Figure 8 shows that some metrics are better at predicting some distortion types than the others. For example, HDR-VDP-2 consistently under-predicts quality for HEVC. Training such metrics for a particular distortion type could substantially boost their performance. Unexpectedly, our *ad hoc* attempts to extend the SSIM_{2D} metric by adding the angular dimension (SSIM_{3D}) or right away considering 3D patches (SSIM_{2Dx1D}) that should account for angular changes has

led to significantly worse results. Clearly, there is a room for improvements and a suitable dataset, such as the one provided in this work, should help to develop a better metric in future.

6.2. Sparse light-field reference case

In all our tests, we provided a high quality, 101-view light field as a reference for the quality metrics. In practice, in most applications only sparsely sampled light field is unavailable. When a sparse light field is used as a reference, a full-reference metric is given to compare two distorted light fields without a perfect reference. This is a task that such metrics were not designed for as they are intended to predict JODs relative to the perfect reference image, not JNDs relative to any other image (refer to Figure 4). This issue is potentially shared with other quality assessment tasks, for example when a metric is trained on 4K images, but it is used on much lower resolution images. However, this problem is exacerbated in case of light fields, where the reduction of angular resolution is often substantial.

To test whether the metrics can predict the quality of distorted light fields using sparse light fields as a reference, we measured the performance of the metrics on a subset of our dataset. As a reference, we chose light fields with distortion NN and severity level two, which correspond to original light fields subsampled to 21 angular views. For the testing light field, we considered all light fields with a higher distortion levels. For a fair comparison, we also ran the metrics on the same subset using full 101-view light fields as a reference. The results of these tests are shown as cyan and blue bars in Figure 9. The significant difference in goodness-of-fit scores (marked with dots) show that metrics predictions get worse if imperfect (sparse) reference is used. This suggests that the existing metrics must be provided with a high-quality reference light field to predict reliably the quality.




Figure 8: Left: The prediction accuracy per-distortion reported as reduced chi-squared goodness of fit score. Middle and right: χ^2_{red} -fit for the metrics HDR-VDP and GMSD over all scenes. The prediction accuracy for individual distortions are shown inside the plots and the overall accuracy is indicated on the top of the plots.




Figure 9: The goodness-of-fit scores for the subset of the dataset when a dense LF is used as a reference (blue), when nearest-neighbour at the 2nd distortion level is a reference (cyan), or when optical flow is used to up-sample the reference LFs. The dots at cyan bars mean that the value is statistically different from the dense LF case and the dots on the yellow bars that the values are statistically different from the NN case. The significance is computed by bootstrapping and running one-tailed test ($p = 0.05$).

But if such high quality reference is not available, can it be approximated? Our subjective data shows that optical-flow interpolation (OPT) produces the highest quality results. Therefore, we used OPT to produce reference 101-view light fields from sparse 21-view light fields and reran the metrics on the subset. The results indicate that the predictions improved as compared to using sparse light field (yellow vs. cyan bars in Figure 9). This suggests that a potential solution to the problem of imperfect reference is to use high-quality interpolation method in order to generate reference.

7. Conclusions and future work

We have established a new 3D dense light-field dataset together with the subjective quality scaling for various distortions that occur in light-field applications. Different methods in light-field processing lead to visual artifacts with quite different appearance, e.g., blur for LINEAR, ghosting for OPT,

image flickering and jumping for NN. Our experiments reveal how these different artifacts affect perceived quality. Our subjective scores are derived from an interactive 3D light-field viewing setup and correspond precisely to overall quality of light fields rather than individual views. We have evaluated the potential of existing image, video, stereo, and multiview quality metrics in predicting the subjective scores. Our observations show that the metrics – HDR-VDP-2, GMSD, STSD_{LC} and VQM perform reasonably well when comparing a distorted light field to a dense reference, and can be used in applications requiring such comparisons. When dense light field is not available, which is the case in some applications, the usage of these metrics for quality assessment is not justified. The perceptually scaled data that we provide can be used for training and validating new light-field quality metrics. Of practical interest for such development is the problem identified in this work, where incomplete, sparse light fields must serve as the reference. Our results also reveal the quality of different light-field reconstruction method, which can directly guide the choice of the light-field reconstruction technique. In the current work, we did not consider aspects such as masking properties of the human visual system. It could be interesting to investigate how much the metrics gain by considering this effect rather than simple averaging of scores over all views. When creating our dataset, we did not consider focus cue. We are, however, not aware of any display setup that could be used to evaluate both motion parallax and focus cue quality. We also believe that the problems revealed in this work should be addressed before including additional cues.

Acknowledgements: This project was supported by the Fraunhofer and Max Planck cooperation program within the German pact for research and innovation (PFI). Denis Sumin was supported by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 642841. The authors would like to thank Tobias Ritschel for the initial discussions and providing synthetic scenes.

References

- [1] T. O. Aydin, M. Čadík, K. Myszkowski, and H.-P. Seidel. Video quality assessment for computer graphics applications. In *ACM Trans. Graph.*, volume 29, page 161, 2010. 1
- [2] F. Battisti, E. Bosc, M. Carli, P. L. Callet, and S. Perugia. Objective image quality assessment of 3D synthesized views. *Signal Processing: Image Communication*, 30(C):78–88, 2015. 2, 6
- [3] E. Bosc, F. Battisti, M. Carli, and P. Le Callet. A wavelet-based image quality metric for the assessment of 3D synthesized views. In *Proc. SPIE*, volume 8648, pages 86481Z–86481Z–9, 2013. 2
- [4] E. Bosc, R. Pépion, P. Le Callet, M. Köppel, P. Ndjiki-Nya, M. Pressigout, and L. Morin. Towards a new quality metric for 3-D synthesized view assessment. *IEEE Journal on Selected Topics in Signal Processing*, pages J–STSP–ETVC–00048–2011, Nov. 2011. 2
- [5] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *Comput. Vis.-ECCV 2004*, pages 25–36. Springer, 2004. 3, 4
- [6] M.-J. Chen, C.-C. Su, D.-K. Kwon, L. K. Cormack, and A. C. Bovik. Full-reference quality assessment of stereopairs accounting for rivalry. *Signal Processing: Image Communication*, 28(9):1143–1155, 2013. 6
- [7] Computer Graphics Laboratory, Stanford University. The (new) Stanford light field archive. <http://lightfield.stanford.edu/acq.html>, 2008. Accessed: 2016-04-23. 2
- [8] D. G. Dansereau, D. L. Bongiorno, O. Pizarro, and S. B. Williams. Light field image denoising using a linear 4d frequency-hyperfan all-in-focus filter. In *Proceedings of the SPIE Conference on Computational Imaging (SPIE'13)*, volume 8657, 2013. 1
- [9] P. Didyk, T. Ritschel, E. Eisemann, K. Myszkowski, H.-P. Seidel, and W. Matusik. A luminance-contrast-aware disparity model and applications. *ACM Trans. Graph.*, 31(6):184, 2012. 2, 4
- [10] S.-P. Du, P. Didyk, F. Durand, S.-M. Hu, and W. Matusik. Improving visual quality of view transitions in automultiscopic displays. *ACM Trans. Graph.*, 33(6):192, 2014. 2
- [11] A. Fusiello, E. Trucco, and A. Verri. A compact algorithm for rectification of stereo pairs. *Mach. Vision Appl.*, 12(1):16–22, July 2000. 3
- [12] R. S. Higa, R. F. L. Chavez, R. B. Leite, R. Arthur, and Y. Iano. Plenoptic image compression comparison between jpeg, jpeg2000 and spith. *Cyber Journals: JSAT*, 3(6), 2013. 1
- [13] ITU-T-P.910. Subjective audiovisual quality assessment methods for multimedia applications. Technical report, 2008. 4
- [14] C. Kim, H. Zimmer, Y. Pritch, A. Sorkine-Hornung, and M. Gross. Scene reconstruction from high spatio-angular resolution light fields. *ACM Trans. Graph.*, 32(4):73:1–73:12, July 2013. 2
- [15] M. Lang, A. Hornung, O. Wang, S. Poulakos, A. Smolic, and M. Gross. Nonlinear disparity mapping for stereoscopic 3D. *ACM Trans. Graph.*, 29(4):75:1–75:10, July 2010. 2
- [16] M. Levoy and P. Hanrahan. Light field rendering. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, pages 31–42, New York, NY, USA, 1996. ACM. 1, 2
- [17] W. Lin and C.-C. J. Kuo. Perceptual visual quality metrics: A survey. *J. Vis. Commun. Image Represent.*, 22(4):297–312, 2011. 5
- [18] J. Liu, T. Malzbender, S. Qin, B. Zhang, C.-A. Wu, and J. Davis. Dynamic mapping for multiview autostereoscopic displays. In *Proc. SPIE*, vol. 9391, pages 11:1–11:8, 2015. 4
- [19] Lytro. Lytro cinema. <https://www.lytro.com/cinema>, 2016. Accessed: 2016-15-11. 1
- [20] A. Maimone, G. Wetzstein, D. Lanman, M. Hirsch, R. Raskar, and H. Fuchs. Focus 3D: Compressive accommodation display. *ACM Trans. Graph.*, 32(5):1–13, 2013. 2
- [21] R. Mantiuk, K. J. Kim, A. G. Rempel, and W. Heidrich. HDR-VDP-2: a calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Trans. Graph.*, 30(4):40:1–40:12, 2011. 6
- [22] B. Masia, G. Wetzstein, P. Didyk, and D. Gutierrez. A survey on computational displays: Pushing the boundaries of optics, computation, and perception. *Comput. Graph.*, 37(8):1012–1038, 2013. 2
- [23] Y. Morvan and C. O’Sullivan. Handling occluders in transitions from panoramic images: A perceptual study. *ACM Trans. Appl. Percept.*, 6(4):1–15, 2009. 2
- [24] MPI. Light-field archive. <http://lightfields.mpi-inf.mpg.de/Dataset.html>, 2017. Accessed: 2017-07-04. 4
- [25] M. H. Pinson and S. Wolf. A new standardized method for objectively measuring video quality. *IEEE Transactions on Broadcasting*, 50(3):312–322, 2004. 2, 6
- [26] PTLens. Lens distortion correction software. <http://www.epaperpress.com/ptlens/>, 2016. Accessed: 2016-15-11. 3
- [27] Rafał K. Mantiuk. Thurstonian scaling for pair-wise comparison experiments. <https://github.com/mantiuk/pwcmp>, 2016. Accessed: 2016-15-11. 5
- [28] G. Ramanarayanan, J. Ferwerda, and B. Walter. Visual equivalence: towards a new standard for image fidelity. *ACM Transactions on Graphics (TOG)*, 26(3):76, 2007. 5
- [29] J. Sánchez Pérez, E. Meinhardt-Llopis, and G. Facciolo. TV-L1 optical flow estimation. *Image Processing On Line*, 3:137–150, 2013. 3
- [30] D. Sandic-Stankovic, D. Kukolj, and P. Le Callet. Multi-scale synthesized view assessment based on morphological pyramids. *Journal of Electrical Engineering*, 67(1):3–11, 2016. 2, 6
- [31] T. Shibata, J. Kim, D. M. Hoffman, and M. S. Banks. The zone of comfort: Predicting visual discomfort with stereo displays. *J. Vis.*, 11(8):11:1–11:29, 2011. 3
- [32] V. D. Silva, H. K. Arachchi, E. Ekmekcioglu, and A. Kondoz. Toward an impairment metric for stereoscopic video: A full-reference video quality metric to assess compressed stereoscopic video. *IEEE Transactions on Image Processing*, 22(9):3392–3404, Sept 2013. 6
- [33] D. Silverstein and J. Farrell. Efficient method for paired comparison. *J. Electron. Imaging*, 10(2):394–398, 2001. 4, 5

- [34] M. Solh and G. AlRegib. MIQM: A novel multi-view images quality measure. In *Quality of Multimedia Experience, 2009. QoMEX 2009. International Workshop on*, pages 186–191, July 2009. 2
- [35] R. R. Tamboli, B. Appina, S. Channappayya, and S. Jana. Super-multiview content with high angular resolution: 3D quality assessment on horizontal-parallax lightfield display. *Signal Processing: Image Communication*, 47:42 – 55, 2016. 2
- [36] G. Tech, Y. Chen, K. Müller, J. R. Ohm, A. Vetro, and Y. K. Wang. Overview of the multiview and 3D extensions of high efficiency video coding. *IEEE Transactions on Circuits and Systems for Video Technology*, 26(1):35–49, Jan 2016. 3
- [37] J. Tompkin, M. H. Kim, K. I. Kim, J. Kautz, and C. Theobalt. Preference and artifact analysis for video transitions of places. *ACM Trans. Appl. Percept.*, 10(3):13:1–19, 2013. 2
- [38] P. Vangorp, G. Chaurasia, P.-Y. Laffont, R. W. Fleming, and G. Drettakis. Perception of visual artifacts in image-based rendering of façades. *Comput. Graph. Forum*, 30(4):1241–1250, 2011. 2, 5
- [39] Viscoda. Voodoo Camera Tracker. <http://www.viscoda.com/en/voodoo-download>, 2016. Accessed: 2016-15-11. 3
- [40] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004. 2, 5
- [41] Z. Wang, E. P. Simoncelli, and A. C. Bovik. Multiscale structural similarity for image quality assessment. In *Signals, Systems and Computers, 2004. Conference Record of the Thirty-Seventh Asilomar Conference on*, volume 2, pages 1398–1402 Vol.2, Nov 2003. 6
- [42] S. Wanner, S. Meister, and B. Goldluecke. Datasets and benchmarks for densely sampled 4d light fields. 2013. 2
- [43] G. Wetzstein, D. Lanman, M. Hirsch, and R. Raskar. Tensor displays: Compressive light field synthesis using multilayer displays with directional backlighting. *ACM Trans. Graph.*, 31(4):1–11, 2012. 1
- [44] B. Wilburn, N. Joshi, V. Vaish, E.-V. Talvala, E. Antunez, A. Barth, A. Adams, M. Horowitz, and M. Levoy. High performance imaging using large camera arrays. *ACM Trans. Graph.*, 24(3):765–776, 2005. 1
- [45] W. Xue, L. Zhang, X. Mou, and A. C. Bovik. Gradient magnitude similarity deviation: A highly efficient perceptual image quality index. *IEEE Transactions on Image Processing*, 23(2):684–695, Feb 2014. 6
- [46] S. L. P. Yasakethu, C. T. E. R. Hewage, W. A. C. Fernando, and A. M. Kondo. Quality analysis for 3D video using 2d video quality models. *IEEE Transactions on Consumer Electronics*, 54(4):1969–1976, November 2008. 2
- [47] M. Zwicker, W. Matusik, F. Durand, and H. Pfister. Antialiasing for automultiscopic 3D displays. In *Proc. of EGSR*, pages 73–82, 2006. 2