DOI: 10.2352/ISSN.2470-1173.2016.5.SDA-440

Evaluation of the Perception of Dynamic Horizontal Image Translation and a Gaze Adaptive Approach

Stefan Eickelberg; Communication Technology Institue / TU Dortmund University; Dortmund, Germany

Abstract

In order to reduce visual fatigue induced by viewing stereo 3d content, a horizontal image translation (HIT) is commonly applied to the stereo views such that the presented disparities are contained within a range that can be viewed comfortably. The technique is also frequently applied in a temporally dynamic manner. The perception of this dynamic HIT (DHIT) by a human observer needs to be studied in order to be able to parametrize it properly in stereo 3d production or automated systems. One such automated system is our previously proposed gaze adaptive approach, where the current point of gaze of the subject is brought into the zero parallax setting by applying DHIT. In this paper, the perceptual properties of dynamic HIT and our gaze adaptive approach are investigated in three subjective experiments.

1 Introduction

There is an inherent problem with all stereo 3d displays: Whenever an observer looks at a stereoscopic stimulus behind or in front of the display, his eyes have to *converge* to the depth of that stimulus, while retaining the focal distance of the eyes, i.e. *accommodation*, on the depth of the display, in order to achieve a sharp perception. The decoupling of these two usually cross-connected processes [1] is unnatural and known as the *accommodation-vergence-conflict* (AV conflict) [2] or *accommodation-vergence-discrepancy*. It is tolerable to small extends. However, excessive violations are known to be one of the main sources for *visual discomfort* and *visual fatigue* as induced by watching stereo 3d content [2]. Lambooij et al. defined visual discomfort as the subjective sensation of discomfort during exposure, while visual fatigue represents the objectively measurable "decrease in performance of the human vision system" [2]. This distinction will be adopted throughout this paper. In order to limit the AV conflicts one has to limit the presented disparities[1] to a range that can be viewed comfortably, the *zone of comfort* [1, 2]. In order to achieve that, two things have to be done:

1. Adjust the distance between the two stereo cameras, i.e. the *camera baseline*, to the depth of the scene, such that the recorded *disparity budget* is small enough to fit into the zone of comfort. The disparity budget is defined as the difference between the maximum and the minimum disparity of a recorded scene.
2. Shift the whole disparity range up or down on the disparity axis, such that the maximum and minimum disparity don't exceed the zone of comfort. This can be done by simply shifting the recorded stereo views in opposite directions, effectively altering the disparity between corresponding stereo-points. This operation is called *horizontal image translation* (HIT) [3, 4] or *reconvergence* [3]. Objects of interest are also commonly placed in the display plane by zeroizing their disparity using HIT, which establishes the *zero parallax setting*[2] (ZPS) [3] for those objects. This way, the average AV conflict is reduced even further.

A distinction has to be made between the unshifted disparities D, as given by the underlying scene, and the shifted disparities

$$\tilde{D} = D - D_{\text{conv}} \tag{1}$$

after applying HIT. Here, D_{conv} is the chosen convergence disparity, i.e. the unshifted disparity D that is zeroized through the HIT. The HIT is carried out by shifting the left view by $-D_{\text{conv}}/2$ and the right view by $+D_{\text{conv}}/2$, where a positive shift value corresponds to a shift to the right. The tilde will also relate other variables to the shifted domain throughout this paper.

Another source for visual fatigue are high velocity convergence movements of the eyes [2]. These can be induced by object motion along the z-axis, that is orthogonal to the display plane. Limiting object speed during recording is an obvious solution to this problem. However, strong depth discontinuities at cuts between shots pose the same problem. To solve this issue, the HIT can be applied in a temporally dynamic manner. The ZPS is shifted shortly before and after the cut in such a way that objects of interest lie in the same depth when the cut occurs. This technique is also called *active depth cut* [3], and it is done so slowly that it usually cannot be perceived. Little is actually known about the perceptual properties of dynamic HIT (DHIT) with its most crucial parameter being the *DHIT speed*. Let the DHIT speed be defined as the speed at which each view is shifted[3]. A stereographer can design the DHIT heuristically just by checking if it looks right. But since there are individual differences in the perception of DHIT [5], it might just happen that the stereographer is less sensitive to it than others and tunes it wrongly. Chamaret et al. did some quick tests to find out how much HIT may change between two successive frames without the observer noticing [6]. Other than that, variations of disparity have only been investigated in the context of camera baseline variation for example by Ware [7], with the result that mild disparity variations are not perceivable. Because of this low number of publications on this topic, we carried out some experiments in order to investigate the DHIT in more detail [5].

---

[1] In the scope of this paper disparity is used as a synonym for parallax, as it is commonly done. The term sometimes also refers to the retinal or angular disparity of the human eye.

[2] The ZPS is also called *plane of convergence*.

[3] Since the views are shifted in opposite directions, the change in disparity is twice as high as the DHIT speed.

©2016 Society for Imaging Science and Technology
DOI: 10.2352/ISSN.2470-1173.2016.5.SDA-440

Knowledge about its properties is especially important in the context of an automated DHIT design. There are many applications where automation is useful or even mandatory. One approach is to analyze the visual saliency of a given scene and have the DHIT set in such a way that salient areas are in the ZPS [4, 6]. The topic of 3d visual attention has been investigated for example by Huynh-Thu et al. [8, 9]. A per-pixel disparity map is necessary in these approaches. Since temporal correlations are very important in visual saliency, the computation also involves evaluation of multiple frames. This renders saliency based approaches inappropriate for real-time applications. Instead of the visual saliency, the actual visual focus, i.e. the point currently gazed upon, can be used to control DHIT by utilizing an eye tracker. A system like that called "GACS3D"[4] was proposed by us in 2013 [10]: The visual focus is slowly brought into the ZPS, thereby reducing the AV conflict. Hanhart et al. compared a similar approach to a visual saliency based DHIT design and the unprocessed stereo views in a subjective experiment [4] and showed that the gaze adaptive approach yielded the best results in picture quality, depth quality and visual comfort. Bernhard et al. evaluated visual fatigue by linking it to the measured stereoscopic fusion times with and without gaze adaptive DHIT [11]. The validity of this link was not discussed, though. To the best of our knowledge, long fusion time and small fusional limits are merely an indicator of subjects prone to visual fatigue [12]. It is also known that fusion time is dependent on disparity [12], which is reduced in the case of gaze adaptive DHIT. Bernhard et al. also carried out a subjective evaluation that yielded a slight improvement of gaze adaptive DHIT over the unprocessed stereo views in depth quality and visual comfort [11].

Independent of the way the HIT is controlled, there is a certain side effect that needs to be dealt with. Due to the shifting operation, black borders appear on opposite sides in the stereo views, for example on the left side of the left view and on the right side of the right view. This means that the border now also has a certain disparity and in the case of DHIT that disparity actually changes over time. The stereo 3d display can be interpreted as a viewing window through which the stereoscopic scene is observed. Because of the aforementioned border disparity, this viewing window now floats in front of the display in the above example. This side effect by itself is not harmful. It is actually a well known technique called *dynamic floating window* (FW) [3, 13] that is commonly used in order to avoid *window violations*. A window violation is given when a stereoscopic object floats in front of the screen, but is partially cut off by the stereoscopic window behind the object. Here, the depth cue "occlusion" dominates the depth cue "disparity", which leads to a distorted depth perception and visual discomfort [13]. With the FW approach, the border disparity is simply increased, such that the stereoscopic window floats in front of the problematic object. This is done by rendering black borders on opposite sides of the views. However, new window violations can be generated through the DHIT when content is shifted out of the display area, so that the disparity of the FW may need to be adjusted continuously.

This paper is structured as follows. In **section 2** the perceptual properties of DHIT are analyzed in an in-depth review of our previous results [10]. The findings are used to parametrize our proposed approach GACS3D, which is described in **section 3**, where differences to the approach by Hanhart [4] are also pointed out. In another review [10], the perceptual effect of GACS3D is analyzed in **section 4** and, lastly, visual comfort of GACS3D is evaluated in **section 5**. The paper is finally concluded in **section 6**.

Table 1: Parameters of all test images of the first experiment, sorted according to disparity budget

Sequence	Frame #	Disparity (px) Budget	Max.[a]	Min.[a]
EBU "Lupo Hands"	945	46	26	-20
RMIT3DV 46	277	62	39	-23
RMIT3DV 29	956	73	28	-45
Industrial Pump	-	92	42	-50

[a] Shifted disparity values after setting a reference convergence disparity through HIT to ensure comfortable viewing. This convergence disparity was actually changed by plus/minus half the shift range during the test by the DHIT. This also induces a change in the maximum and minimum shifted disparity so that some sequences protruded the zone of comfort (±58 pixels in disparity in our configuration).

2 Experiment 1: Investigation of basic perceptual properties of DHIT

As mentioned in the introduction, there have been only few publications on the perception of DHIT. Because of that, we carried out an explorative experiment to find out when the DHIT is perceivable or annoying. There are two parameters controlling the DHIT: The amount of HIT applied to each view, i.e. the shift range, and the DHIT speed. As noted in the introduction, the DHIT speed is usually chosen heuristically by the stereographer and the shift range is dictated by the content. An application-oriented DHIT design like this prevents proper orthogonalization of the DHIT parameters, so an interactive test under artificial conditions was constructed. Along with the two DHIT parameters, we also hypothesize that the disparity budget of the underlying scene affects DHIT perception. This section is to be considered an in-depth review of our previous results [5] with some additions.

2.1 Experimental setup

The following experimental setup was designed in accordance with the respective ITU recommendations [14, 15] and has been used in all three experiments except for a few changes that will be pointed out. The stimulus display was a 47 inch Full-HD 3D-LCD with passive polarizer glasses and it was placed on a long table. All signal processing was disabled on the display and it was driven in Full-HD at a framerate of $f = 60$ Hz, so that any effects of low framerate can be neglected. The subject sat in an office chair in a distance of $B = 3.1H$ away from the screen, with H representing the height of the display. For the subject interaction, there was also a black keyboard with labeled keys on the table.

2.2 Stimuli

For every test scene, the views were cyclically shifted back and forth over one of three defined shift ranges and at a user configurable DHIT speed, which effectively orthogonalizes the two

---

[4] GACS3D stands for "Gaze Adaptive Convergence in Stereo 3D Applications".

DHIT parameters. The shift ranges s_{px} were 20, 30 and 40 pixels. Assuming square pixels, the horizontal pixel-pitch on a Full-HD display is $p = H/1080$, so that the angular shift range s_{deg} at a viewing distance $B = 3.1H$ is given by

$$s_{\text{deg}} = 2\cdot\tan^{-1}\left(\frac{0.5\cdot s_{\text{px}}\cdot p}{B}\right) = 2\cdot\tan^{-1}\left(\frac{s_{\text{px}}}{6696}\right), \tag{2}$$

which yields 0.34 °, 0.51 °and 0.68 °, respectively.

The third parameter is the disparity budget, which is dictated by the underlying scene and camera setup, specifically the camera baseline. High resolution multi-baseline material is rather rare. However, the main reason why we did not use such material is that we did not want to bore the subjects by showing only a single sequence in different conditions, thereby risking inaccurate results due to tiredness. Instead, four different video sequences were used: EBU 3DTV Test sequence "Lupo_Hands" [16], RMIT3DV no. 29 and 46 [17], and a custom synthetic test image showing an industrial pump has been created. Considering the three shift range settings, this makes a total of 12 test stimuli. We chose to show a still image of those videos, because it is the worst case scenario for DHIT detection, since there is nothing distracting the subjects. Some parameters of the used test images are summarized in **Table 1**.

2.3 Procedure

For each stimulus, the DHIT speed v_{px} is initialized to a very annoying setting, which was $v_{\text{px}} = 0.5$ px/frame or

$$v_{\text{deg}} = 2\cdot\tan^{-1}\left(\frac{v_{\text{px}}}{6696}\right)\cdot f = 0.513\ °/\text{s}. \tag{3}$$

The subject was asked to lower the speed in steps of $\Delta v_{\text{deg}} = 0.0257\ °/\text{s}$ using the labeled keyboard, until the DHIT was just not deemed annoying anymore and confirm that setting by pressing the respective button. This is the annoyance threshold. Furtherly, the subject would lower the speed even further, until the DHIT was just not perceivable anymore (perception threshold) and move on to the next sequence. The subject was also able to increase speed again and correct the vote, if necessary. The 12 stimuli were presented in a random permutation. In order to familiarize the subject with the DHIT, the controls and the general procedure, some anchor sequences were shown prior to the test, while the instructions were presented. The subject was asked to explore the screen freely except for the border regions on the left and right side of the display. These were excluded, because the visible frame of the display serves as a reference that makes the HIT almost always perceptible.

2.4 Subject screening and rejection criteria

Prior to the test, the visual performance of all subjects was examined. A good overview of stereo 3d evaluation related examination methods along with rejection criteria is given by Lambooij et al. [18]. In our experiments, we used Snellen charts, the Ishihara test, the randot butterfly stereogram and the circle test to assess binocular visual acuity, color perception, gross stereopsis and fine stereopsis, i.e. stereo acuity, respectively. Subjects exhibiting visual acuity worse than 80% were rejected. Failing the randot butterfly stereogram also leads to rejection. However, nobody failed in our experiments. Although it is common to reject subjects with more than 60" in stereo acuity [18], we allowed subjects with a stereo acuity of up to 140" to take the test. Those subjects did not vote significantly different than the others. They were not classified as outliers by the recommended methods [14], which legitimizes their inclusion. Some details on the subjects, that participated in this experiment, as well as their examination results are summarized in **Table 2** and **Table 3**.

Table 2: Subject overview for all three experiments

	Exp. 1	Exp. 2	Exp. 3
No. of subjects[a]	26	19	36
No. of rejections	2	2	6
No. of females	4	3	5
Experienced subjects[b]	7	2	18
Experts	4	0	2
Research Assistants	8	0	10
Students	17	19	21
Other occupation	1	0	5
Minimum Age	21	23	22
Maximum Age	31	30	36
Average Age	25	24.7	26.4

[a] Rejections are included.
[b] Subjects with prior experience in subjective image quality evaluation.

Table 3: Visual performance of subjects in all three experiments

	No. of subjects		
Visual acuity	Exp. 1	Exp. 2	Exp. 3
100%	17	7	20 (1)[c]
100%, corrected[a]	4	10	10 (4)[c]
80%	4	1	3
80%, corrected[a]	0	0	2
< 80% (reject)	1	1	1
Stereo acuity[b]			
40"	21	10	27 (5)[c]
50"	1	1	4
60"	1	2	1
80"	0	2	3
140"	2	3	1
> 140" (reject)	1	1	0
Color perception			
Unimpaired	23	18	34 (5)[c]
Mild Deuteranopia	3	1	2

[a] Subjects wearing either glasses or contact lenses.
[b] Values are seconds in angle of stereopsis.
[c] Number of subjects rejected due to bad eye tracker performance.

2.5 Results and discussion

The results are shown in **Figure 1**. The plots show the perception and annoyance thresholds for DHIT speed v_{deg} as a function of shift range and disparity budget. In other words, the abscissas represent the different sequences. On average, the DHIT

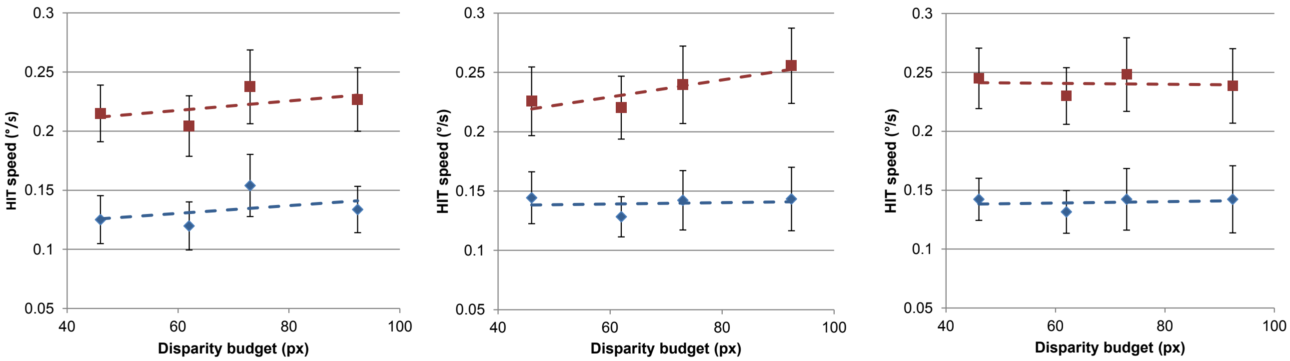

(a) Shift range $s_{\text{deg}} = 0.34°$ (b) Shift range $s_{\text{deg}} = 0.51°$ (c) Shift range $s_{\text{deg}} = 0.68°$

Figure 1: Perception (blue) and annoyance (red) thresholds of experiment 1 with confidence intervals and linear line fits

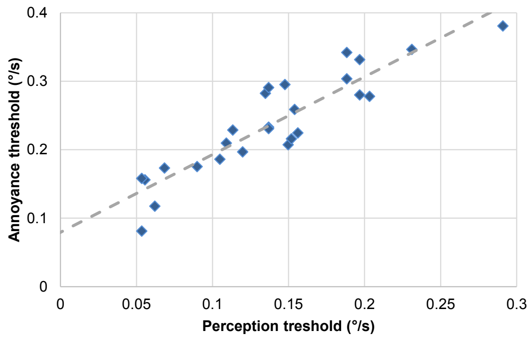

Figure 2: Annoyance threshold over perception threshold averaged over all sequences for each subject separately along with a linear line fit.

was not perceivable for $v_{\text{deg}} \leq 0.14$ °/s and annoyance did not set in until after $v_{\text{deg}} = 0.24$ °/s. The average annoyance threshold is not of interest, however, because it is important to ensure that the DHIT is never deemed annoying by anyone. This makes the lower boundaries of the confidence intervals a lot more interesting, which is as low as 0.18 °/s for the second sequence “RMIT3DV 46”.

The annoyance and perception thresholds are clearly correlated. The reason for this is rather simple: One cannot be annoyed by a visual artifact one cannot see. Neither the annoyance thresholds nor the perception thresholds are correlated with disparity budget or shift range, which renders two of our three initial hypotheses invalid. The thresholds do exhibit roughly the same value progression in all three figures, though, which indicates that there might still be an unidentified content dependency remaining. In interviews conducted directly after the test, some subjects said that it was easier to detect DHIT in image regions with big depth discontinuities. Consequently, the local distribution of depth discontinuities would be a good candidate for further testing. The reason for the heightened sensitivity towards DHIT in these regions could be the unnatural distortion of depth induced by the DHIT [19]. This distortion becomes more distinct the bigger the disparity gradient is. A distortion-free DHIT approach is currently being investigated.

The perception and annoyance thresholds were also averaged over all sequences for each subject separately and the results are shown in **Figure 2**. This graph shows once again the strong correlation between perception and annoyance. But what is more interesting is that there were strong individual differences in the DHIT sensitivity. One subject was annoyed by DHIT speeds as slow as 0.08 °/s, but one other extreme example only started to be annoyed by DHIT speeds as fast as 0.38 °/s.

3 Description of the gaze adaptive DHIT approach

The basic approach has already been described in one of our previous publications [10]. However, since that publication dates back three years, some new modifications have been implemented into the prototype. Because of that, a description of the current state of the system shall be given in this section. The general concept of GACS3D is to estimate the visual focus, retrieve the unshifted disparity at that point and use that as the new target convergence disparity in the DHIT process. Thereby, the ZPS is slowly established at the visual focus. The approach comprises five steps as depicted in **Figure 3**, that will be described in more detail in the following sections. Some benchmark results are given afterwards.

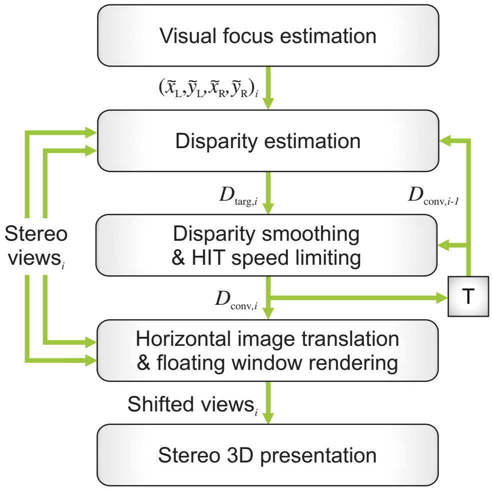

Figure 3: Processing steps in GACS3D

3.1 *Visual focus estimation*

A professional, research-grade, remote eye tracker of a renowned manufacturer is used. The eye tracker combines both bright and dark pupil tracking and allows for free head movement inside a virtual tracking box. The tracking box is big enough to ensure comfortable viewing conditions. The eye tracker delivers 60 binocular visual focus coordinates $(\tilde{x}_\mathrm{L}, \tilde{y}_\mathrm{L}, \tilde{x}_\mathrm{R}, \tilde{y}_\mathrm{R})$ per second, that correspond to pixel coordinates on the display. These coordinates are affected by noise due to technological limitations: According to the eye tracker specifications, the accuracy[5] is 0.4 ° and precision[6] is 0.34 °. Aside from these limitations, precision is also degraded by the nature of the eye movements. Only the voluntary eye movements are of interest, which are fixations, smooth pursuit and saccades, i.e. stepwise changes of visual focus. But even during fixations the eye never stays completely still, due to involuntary eye movements, namely microsaccades, tremor and drift [20]. Futhermore, saccades almost never end up perfectly in their final location, but are corrected afterwards by glissades [20]. These sources of noise prevent a direct computation of disparity

$$\tilde{D} = \tilde{x}_\mathrm{L} - \tilde{x}_\mathrm{R} \tag{4}$$

from the visual focus coordinates and call for a lookup in precomputed disparity maps, as was previously done [4, 10], or a sophisticated disparity estimation.

3.2 Disparity estimation

In the second step, the disparity currently looked at needs to be retrieved, yielding the new target convergence disparity D_targ. In our initial proposal, the visual focus coordinates were assumed noise-free and a lookup on precomputed, perfectly conditioned disparity maps was used, with some specific processing for occlusion regions. Since the views being looked at have been shifted by $\mp D_\mathrm{conv,i-1}/2$ prior to the current iteration i, the disparity maps would also have to be shifted by the same amount for the lookup. However, a much simpler solution, that also avoids interpolation problems, is to just shift the visual focus coordinates in the opposite direction and perform the lookup on the unshifted disparity maps for the left and right view:

$$(x_\mathrm{L}, y_\mathrm{L}) = (\tilde{x}_\mathrm{L} + D_\mathrm{conv,i-1}/2, \tilde{y}_\mathrm{L}) \tag{5}$$

$$(x_\mathrm{R}, y_\mathrm{R}) = (\tilde{x}_\mathrm{R} - D_\mathrm{conv,i-1}/2, \tilde{y}_\mathrm{R}). \tag{6}$$

In order to deal with the noisy visual focus coordinates, Hanhart et al. proposed to apply a 13-tap median filter to the coordinates and use a 15 by 15 maximum filter for the lookups on the disparity maps [4]. The median filter removes outliers, but does not necessarily improve accuracy, which is the reason for the increased lookup window. The maximum filter is applied to that window, because foreground objects are assumed to be more salient. While this approach may yield satisfactory results in some conditions, the actual visual focus will usually not be contained in the lookup window, because the visual focus samples of a remote eye tracker typically exhibit a spread exceeding that windows size by an order of magnitude. Since the actual visual focus might yield a completely different disparity, we believe that the lookup window size needs to be increased to accommodate to that sample spread, in order to achieve high reliability. Furthermore, in order to be able to use a gaze adaptive DHIT with regular stereo 3d content, one would need a real-time disparity map estimation instead of relying on unavailable precomputed disparity maps. Our approach is based on a similar ideas as the one by Hanhart et al., but also includes a real-time region-of-interest disparity map estimation in CUDA, which implicitly handles occlusions. However, a detailed description and evaluation of our complete disparity estimation algorithm is beyond the scope of this paper.

3.3 Disparity smoothing and DHIT speed limiting

In the last step, the new target convergence disparity $D_{\mathrm{targ},i}$ has been retrieved. As can be seen from the results of the first experiment in **section 2**, the DHIT speed needs to be limited to a certain value $v_\mathrm{px,max}$. Furthermore, we did not investigate the effect of sudden changes in DHIT speed, but increased it smoothly instead. So, in order to avoid such speed changes, a simple 1-tap recursive filter is used to smooth the target disparity

$$\overline{D}_{\mathrm{targ},i} = \alpha \cdot D_{\mathrm{targ},i} + (1-\alpha) \cdot D_{\mathrm{conv},i-1}, \tag{7}$$

where the smoothing factor was empirically chosen to $\alpha = 1/8$. The current unlimited speed is given by

$$v_{\mathrm{px},i} = \left| \overline{D}_{\mathrm{targ},i} - D_{\mathrm{conv},i-1} \right| \tag{8}$$

and the shift direction is

$$\beta_i = \mathrm{sign}\left(\overline{D}_{\mathrm{targ},i} - D_{\mathrm{conv},i-1} \right) . \tag{9}$$

After applying the speed limit, the new convergence disparity is

$$D_{\mathrm{conv},i} = \begin{cases} D_{\mathrm{conv},i-1} + \beta_i \cdot v_{\mathrm{px},i} = \overline{D}_{\mathrm{targ},i} & \text{if } v_{\mathrm{px},i} < v_\mathrm{px,max} \\ D_{\mathrm{conv},i-1} + \beta_i \cdot v_\mathrm{px,max} & \text{otherwise.} \end{cases} \tag{10}$$

3.4 Horizontal image translation and floating window rendering

As explained before in **section 1**, the stereo views are shifted by $\mp D_{\mathrm{conv},i}$ to the right. Since the convergence disparity is non-integer, a horizontal cubic interpolation is used to render the shifted views, which can be done very efficiently in parallel computing using CUDA. Hanhart et al. use a nearest neighbor interpolation for performance reasons [4]. This kind of interpolation leads to unsmooth and asymmetric motion of the views, because the DHIT speed can only discretely transition between $v_\mathrm{px} = 0$ px/frame and $v_\mathrm{px} = 1$ px/frame to achieve the desired non-integer DHIT speed on average. These sudden and asymmetric high speed variations are likely to affect DHIT sensitivity, but further investigations are necessary.

There is still the problem of possible window violations generated by image content shifted out of the display area and the temporal changes in FW disparity, as mentioned in **section 1**. Both factors might pose a distraction from the system under test in our experiments. Because of that, we chose to crop the views on both sides in such a way, that no content is shifted out of the display area, and manually set the FW to a fixed disparity. This decreases the width of the stereoscopic window, but completely removes window violations in a static fashion. Approaches for an automated dynamic FW design compatible with GACS3D are also available, though.

Afterwards, the shifted views are finally multiplexed into a stereo 3d frame format and passed to an OpenGL displaying routine.

---

[5] Accuracy represents the average divergence from the target [20].

[6] Precision evaluates the scattering of samples [20].

©2016 Society for Imaging Science and Technology
DOI: 10.2352/ISSN.2470-1173.2016.5.SDA-440

3.5 Benchmark results

The prototype has been implemented in a combination of MATLAB and CUDA C++ in order to utilize the parallel processing power of a graphics card. The prototype achieves 120 frames/s at Full-HD resolution on our test system, including video decoding, the OpenGl displaying routine and two disparity map estimations per frame for 201 by 101 pixels big regions of interest with a nine by nine pixel correlation window size. Without the displaying routine, the prototypes achieves 135 frames/s. Our test system consists of a 3.4 GHz quadcore with hyperthreading, 16 GB DDR3-1600 and a CUDA compute capability 3.5 graphics card that achieves 5.0 TFlops in single precision.

4 Experiment 2: Analysis of the perceptual effect of gaze adaptive DHIT

The DHIT was designed completely independent of the underlying scene structure in the first experiment. So, in order to find practical limits to DHIT speed, the DHIT was designed in a more application-oriented way in this experiment: Different interesting objects were slowly brought into the ZPS. Furthermore, the effect of GACS3D on the perception of DHIT was investigated. Some of these results have already been described in our previous publication [5].

4.1 Experimental setup

The experimental setup was the same as described in **subsection 2.1**, except for the computer keyboard, because no user interaction was required in this test.

4.2 Stimuli

The aim of this experiment is to evaluate how different DHIT speeds are perceived in a gaze adaptive DHIT design compared to the non gaze adaptive design. In order to truly isolate the effect of gaze adaptivity, the presented DHIT sequence must be the same in both conditions. Furthermore, the DHIT speed must be predetermined for each stimulus, so that results between subjects are comparable. These restrictions prevented the actual usage of GACS3D and called for an emulation of gaze adaptivity. Again, still images were used as the base material for the same reason as in the first experiment. Every test image was presented under three different conditions: GACS3D, movie watching experience and one control condition, to validate the results.

1. "GACS3D"
   A stereoscopic pointer was shown at a deterministic series of different interesting locations for 0.5 s to 1.5 s each and wherever the pointer was, the ZPS was slowly established through DHIT at a defined maximum speed. The subject was asked to always fixate on that pointer, which effectively emulates GACS3D by replacing the eye tracker with deterministic pointer locations. The pointer strongly attracts attention, because it is the only moving object on those still images. However, in order to truly synchronize eye movements and DHIT, one has to compensate for the human saccadic reaction time, which is 240 ms on average [21], but can vary under certain circumstances. Because of that, the DHIT was delayed by 150 ms, which was judged most natural by a small test group.

   The DHIT sequence generated by this test condition was used in all other test conditions as well. The only difference is the presentation of the pointer.
2. "Movie"
   Here the pointer was simply hidden and the subject was allowed to explore the screen freely, just like watching a movie. The results of this condition will serve as a reference point, because it is somewhat similar to experiment 1.
3. "Control"
   A comparison of the results for the first two conditions enables us to analyze the effect of gaze adaptivity, but one might argue that the pointer poses a distraction that alters the results. So, in order to check the validity of that comparison, a control condition was introduced: A pointer was shown at a different deterministic series of locations, that did not actually converge to the zero parallax setting. Hence, the DHIT was not gaze adaptive just like in the "Movie" condition, while a pointer was shown. If the control condition yields approximately the same results as the "Movie" condition, it is safe to say that the effect of the visible pointer is neglectable.

There were 25 still images accompanied by estimated disparity maps exhibiting budgets ranging from 14 to 92 pixels. The images were taken from the EBU [16] and RMIT3DV [17] test sequence libraries in addition to the industrial pump test image. The mean DHIT speed is dependent on the disparity budget of a scene, the pointer locations and the maximum DHIT speed setting. The series of pointer location was chosen in such a way that maximum and mean DHIT speed were approximately of the same order of magnitude. Each test image was assigned with one of seven maximum DHIT speed values equally distributed on the interval 0.125 to 0.5 px/frame, which is equivalent to 0.128 °/s and 0.513 °/s.

4.3 Procedure

The test images were rated under the three conditions independently in random order in a non-interactive single stimulus impairment scale test [15], where the subject was instructed to rate their perception of the DHIT. A discrete scale was used with labels translated to German: "5: imperceptible", "4: perceptible, but not annoying", "3: slightly annoying", "2: annoying", "1: very annoying". Again, the subject was instructed not to rate a sequence while watching at the border of the display, because the DHIT is always perceivable there. The first five ratings were neglected as recommended in [14]. In order to familiarize the subject with the procedure and the DHIT, some anchor sequences were shown before the test.

4.4 Results and discussion

There were 19 subjects taking part in this experiment. A summary on the subjects and their visual performance can be found in **Table 2** and **Table 3**. The mean opinion scores (MOS) were calculated for each sequence and plotted over the mean DHIT speed in **Figure 4a**. There is a clear correlation between mean DHIT speed and MOS. The linear line fits show that the "Control" and "Movie" condition are approximately the same, which renders the comparison between "Movie" and "GACS3D" legit. The "GACS3D" condition exhibits higher scores than "Movie". This suggests that gaze adaptivity has a beneficial effect on DHIT sensitivity since it is decreased. The spread of samples

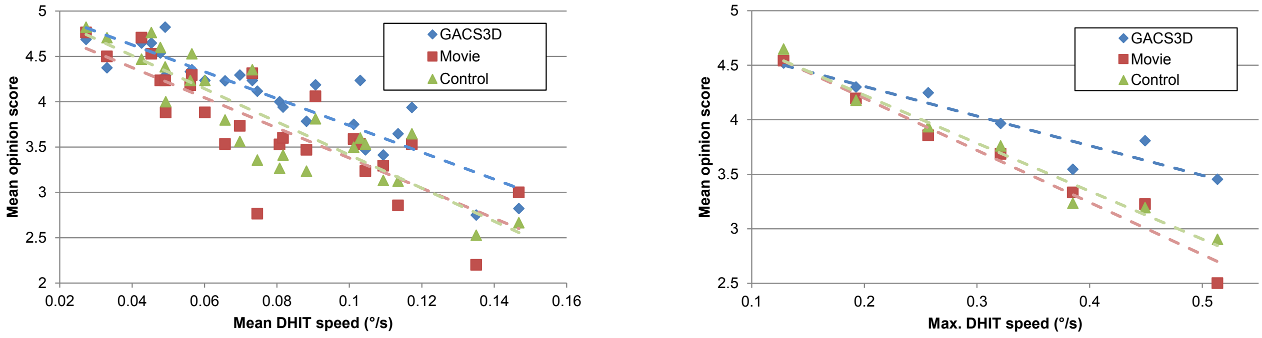

(a) MOS over mean DHIT speed.

(b) MOS over max DHIT speed.

Figure 4: Mean opinion scores in experiment 2 for all three conditions with linear line fits.

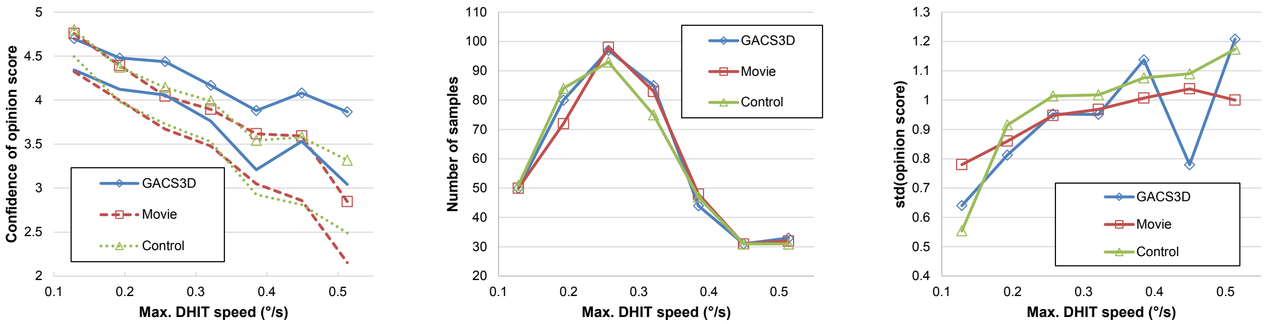

(a) Maximum and minimum (no symbols) grade series, i.e. 95% confidence intervals

(b) Number of samples per maximum DHIT speed score

(c) Standard deviation of MOS over maximum DHIT speed

Figure 5: Statistical analysis of the results of experiment 2.

is very high, though, which is understandable, because the mean DHIT speed is influenced by so many factors as mentioned earlier. The maximum DHIT speed seems like a much better candidate.

The MOS were therefore also calculated for each maximum DHIT speed value. This means that the scores of multiple different sequences go into a single MOS. This is legit since experiment 1 showed that the effect of content on the DHIT perception is insignificant. The result is shown in **Figure 4b** and exhibits a much better correlation with MOS. Again, “Movie” and “Control” have approximately the same linear line fits, while the line of “GACS3D” is a lot more flat-angled. It yields higher scores for big maximum DHIT speeds. The reason for the decreased HIT sensitivity in the “GACS3D” condition could be saccadic suppression: Shortly before and after the start of a saccade, visual processing is suppressed for a few 100 ms [21] and since the dynamic HIT was synchronized to saccades, HIT sensitivity might be affected. The confidence intervals do slightly overlap, as can be seen in **Figure 5a**, which renders the effect of gaze adaptivity rather insignificant. The intervals are generally dependent on the number of samples and the standard deviation. In this experiment, they grew bigger towards higher speed values for two reasons. On the one hand, the number of samples per maximum DHIT speed setting was lower for high speed values[7] (see **Figure 5b**). On the other hand, the standard deviation, plotted in **Figure 5c**, is mildly correlated to maximum DHIT speed. This is because subjects exhibited individual systematic differences in their sensitivity towards DHIT, as already mentioned in **subsection 2.5**.

It is also possible to derive annoyance and perception thresholds from this experiment and compare them to experiment 1. The transition from “imperceptible” to “perceptible, but not annoying” at score 4.5 represents the perception threshold and the next transition to “slightly annoying” at score 3.5 represents the annoyance threshold. Based on the linear line fits in **Figure 4b**, these average thresholds are 0.13 °/s and 0.35 °/s of maximum DHIT speed in the “Movie” condition. The perception threshold strongly agrees with the results from experiment 1, but the annoyance threshold is increased. This might be due to the different DHIT design. The DHIT occurred only sporadically and over comparatively short time periods in this experiment. The gaze adaptive condition “GACS3D” yields a perception threshold of 0.12 °/s on average and an annoyance threshold of 0.50 °/s. We conclude that gaze adaptive DHIT reduces annoyance.

5 Experiment 3: Visual discomfort evaluation of gaze adaptive DHIT

Our proposed approach GACS3D is supposed to reduce visual fatigue. Evaluating visual fatigue directly would be a very elaborate and ambitious undertaking. Instead, we chose to evaluate visual discomfort as an indicator for visual fatigue in a pair comparison test, similarly to what Hanhart et al. did [4].

---

[7] The variation of the number of samples is due to the restrictions during sequence generation and the rejection of the first five random scores.

Figure 6: Experimental setup in experiment 3.

Table 4: Parameters of all test video sequences of experiment 3.

		Disparity (px)			
Sequence	Frames	Max.	Min.	ZPS	FW
EBU "Lupo Hands"[a]	301[a]-600[a]	46	-36	-24	20
NAMA3DS1 "Umbrella"	76-325	19	-39	-23	20
RMIT3DV 02	1-300	31	0	15	22
RMIT3DV 29	951-1250	45	-28	19	23
RMIT3DV 43	1-300	0	-69	-30	35
RMIT3DV 46	251-550	-9	-72	-40	30

[a] This sequence is actually available in 50 frames/s, but was downsampled to 25 frames/s so that it has the same framerate as the other sequences.

5.1 Experimental setup

The basic experimental setup was the same as in **subsection 2.1**. However, some modifications were necessary because of the utilized eye tracker. The eye tracker was placed on the table on a customly made stand that enables a precise eye tracker calibration an this comparatively high $3.1H$ operating distance. Due to the passive polarized stereo 3d glasses, the eye tracker performance was drastically degraded. Because of that, the subject was illuminated by two spotlights positioned left and right of the table, as can be seen in **Figure 6**. In order not to blind the subject with those spotlights, the background illumination was increased by enabling the ceiling lighting of the laboratory. Because of that, the light density of the background was approximately 45 $\mathrm{cd/m^2}$ and its color was CIE $(x,y) = (0.39, 0.40)$. Hence, this experimental setup was not compliant with the respective ITU recommendations [14, 15]. Prior to the test, subjects were asked whether the spotlights blinded them and none affirmed. Furthermore, black blinders were added on both sides of the glasses to prevent visible reflections on the inner side of the polarization filters. The keyboard was removed from the table, because all interactions were handled by an operator.

5.2 Stimuli

Since the final prototype should be evaluated in this experiment, moving sequences from the EBU [16], NAMA3DS1 [22] and RMIT3DV [17] stereo 3d test sequence libraries were used, rather than still images. The six videos were presented under three different conditions:

1. "Raw": No HIT was applied to the stereo views. The scene was presented as is.
2. "ZPS": A static HIT was applied to establish a certain zero parallax setting. The convergence disparity was manually chosen by an expert with the aim to minimize visual discomfort while simultaneously generating a visually pleasing depth sensation.
3. "GACS3D": The full prototype was used to realize the gaze adaptive DHIT. The views were shifted at a maximum speed of 0.12 °/s, which was the perception threshold for GACS3D as of experiment 2. The DHIT was updated at a rate of 60 Hz, in accordance with our previous experiments.

The video sequences are available in 25 frames/s. In order to avoid any motion judder, the sequences were played at an increased speed of 30 frames/s, so that every frame was played twice on the 60 Hz display. The resulting test sequences still looked natural at this increased speed and were 10 s long. Details on the video sequences are summarized in **Table 4**[8]. Snapshots of the sequences are displayed in **Figure 11**.

As mentioned in **subsection 3.4**, the stereo views were cropped and a floating window was applied. This is done in all three conditions. The respective disparity can be found in the referenced table, alongside the disparity for the ZPS of the second condition. The listed FW disparities can be a lot bigger than the respective maximum disparities, because the FW must be able to eliminate window violations even in the most extreme settings, i.e. when the background is looked at with GACS3D, so that the whole scene is shifted in front of the display.

5.3 Procedure

The visual comfort of the conditions was evaluated using the pair comparison method [15] in a simple preference judgment ("A is better", "equal", "B is better"). A graded scale was not necessary, because the differences were hardly perceivable. In an effort to help the subject to see the subtle differences, each condition was repeated once per trial (A-B-A-B). In the case of GACS3D, this means that the DHIT applied to the sequence was not the same in both playthroughs due to the gaze adaptivity. However, this ensures that it is actually GACS3D being rated, and not some specific DHIT sequence. All condition combination pairs were tested (Raw vs. ZPS, Raw vs. GACS3D, ZPS vs. GACS3D), but not in both possible orders to keep the required time acceptable. Instead, the order of each combination was randomized per subject.

Prior to the test, the eye tracker was calibrated individually and some anchor sequences were shown to familiarize the subject with the concept of visual discomfort related to excessive AV conflicts. The Subject was instructed to sit still during a trial, so that proper eye tracker performance was ensured. Between trials, a window showing the position of the eyes was presented to the subject to ensure that the optimal tracking position was maintained throughout the whole test. In order to reduce subject move-

---

[8] Please note that some values may differ from those in **Table 1**, because an HIT was applied there and only a single frame was shown, which might exhibit a different disparity budget than the whole scene.

©2016 Society for Imaging Science and Technology
DOI: 10.2352/ISSN.2470-1173.2016.5.SDA-440

ment to a minimum, the scores were furthermore collected by an operator, who also triggered the start of the next trial.

5.4 Results and discussion

In this experiment, many subjects wearing optical aids had to be rejected due to bad eye tracking performance, as can be seen in **Table 2** and **Table 3**. Because of that, there was a certain preference for people without optical aids. There were 30 subjects, excluding the rejections. Each sequence, where eye tracking was involved, was inspected manually, because the eye tracker still occasionally yielded bad results for some subjects. In that case, individual scores of problematic sequences were rejected, which lead to a reduced number of 15 to 25 samples per MOS in condition combinations with “GACS3D”, as can be seen in **Figure 8**.

The MOS are shown in **Figure 7**. The comparison between “Raw” and “ZPS” exhibits no clear tendency, which also shows in the almost zero average over all sequences. In **Figure 7b** “GACS3D” is preferred slightly over “Raw” for nearly all sequences, but a dominantly negative rating for the Umbrella sequence pushes the average to almost zero again. Finally, the comparison between “GACS3D” and “ZPS” shows no clear tendency again and, on average, “GACS3D” is rated slightly negative.

The MOS magnitudes are generally very small, though, while their standard deviations in **Figure 8** are very big. Furthermore, the score histogram in **Figure 10** shows that “equally comfortable” was picked most often in 15 out of 18 trials and was a very close second place in the rest of the trials. There were also no subjects exhibiting a clear systematic preference for any condition. We had already expected results like that, judging from the subject interviews conducted after the test. Most subjects said that it would look all the same or that they tried to concentrate on details, because they could not tell the difference between the sequences. However, all of them affirmed a strong visual discomfort when an anchor sequence with big AV conflicts was shown prior to the test. A very small group of subjects also hinted that they disliked the DHIT, while the rest of the subjects did not detect it at all. This shows again that there are individual differences in the DHIT sensitivity, as mentioned in **subsection 2.5**. The results of the “Raw” vs. “ZPS” comparison should be most reliable, since no individual scores had to be rejected, but this comparison actually exhibits the highest standard deviations in **Figure 8**. Considering all the evidence in this paragraph, we conclude that no visual discomfort was perceived in any condition. In other words, HIT, DHIT and GACS3D do not affect visual comfort, as long as the resulting disparities don’t protrude the zone of comfort and the DHIT speed is kept below the annoyance threshold.

In contrary to ours, the results of Hanhart et al. [4] suggest that gaze adaptive DHIT improves visual comfort, despite their DHIT was designed in a more critical way. Their maximum DHIT speed was $v_{\text{px}} = 0.5$ px/frame, which using **Equation 3** corresponds to $v_{\text{deg}} = 0.21$ °/s at the framerate $f = 25$ Hz. This is almost as high as the average annoyance threshold from the first experiment and a lot higher than some of the individual annoyance thresholds. As already pointed out in **section 3**, the group furthermore uses a nearest neighbor interpolation for the HIT, which actually makes the DHIT speed discretely transition between $v_{\text{px}} = 0$ px/frame and $v_{\text{px}} = 1$ px/frame to achieve the aforementioned $v_{\text{px}} = 0.5$ px/frame on average. This interpolation also leads to asymmetric shift behavior of the views.

In **Figure 11**, snapshots of the six video sequences are shown along with heatmap overlays visualizing the actual areas of visual attention of all subjects. It becomes apparent, that scene elements like faces or occlusions are highly salient. In return, when these attractors are absent (RMIT3DV 02, 43, 46), the scene is explored a lot more freely. This also shows in the normalized histograms of watched disparities in **Figure 9**, where concentrations can be observed on the disparities of visual attractors. This strongly influences the behavior of GACS3D, because it almost degenerates to a static ZPS in presence of visual attractors. Considering these observations, a design recommendation for the ZPS can be easily derived. Whenever one or more visual attractors are present, the ZPS should be placed such that AV conflicts on those objects are minimized on average. Otherwise, the whole disparity budget should be fitted into the zone-of-comfort, which extends a lot further behind than in front of the display [2]. Ideally, the scenedepth statistics should be considered in that process. However, this design recommendation is not as easily implemented as it is described, as can be seen in **Figure 9**. For some sequences, the ZPS and the average watched disparity differ quite a bit due to wrongly identified visual attractors in the design process.

6 Conclusion

In this paper, we evaluated the perceptual properties of dynamic horizontal image translation (DHIT) in general and our previously proposed approach for gaze adaptive DHIT called “GACS3D”, which is supposed to reduce visual fatigue. The basic method and recent modifications of our prototype were described. In contrary to most other approaches, our prototype does not rely on precomputed disparity maps, but estimates them in a region of interest instead, which makes it applicable to regular stereo 3D content. The prototype achieves real-time performance at 120 frames/s. In order to parametrize it, two experiments were carried out to analyze the properties of DHIT. We found out that the DHIT was on average not perceivable for DHIT speeds, i.e. shift speeds of the views, below 0.14 °/s. For fast DHIT speeds of 0.24 °/s and more, the average observer got annoyed by the DHIT. Gaze adaptivity raised this threshold up to 0.5 °/s, which means that gaze adaptivity reduces annoyance. However, it is important to note that there were strong individual differences in the sensitivity towards DHIT, with one extremely sensitive subject already being annoyed by DHIT speeds as low as 0.08 °/s. Since it is mandatory to ensure that nobody gets annoyed by the DHIT, while also considering that DHIT sensitivity was decreased under application-oriented conditions, we recommend to keep DHIT speed in the range of 0.1 to 0.12 °/s.

Furthermore, we evaluated visual discomfort as an indicator for visual fatigue with and without our approach. Contrary to Hanhart et al. [4], we were not able to show that gaze adaptive DHIT has a beneficial effect on visual comfort, despite our efforts to design the DHIT in a more comfortable way. The referenced group used a comparatively fast DHIT speed of 0.21 °/s, which is almost as high as the average annoyance threshold of our first experiment. Our results showed that neither HIT nor GACS3D induce visual discomfort (or visual comfort), as long as the disparity range is kept inside the zone of comfort and the DHIT speed below the annoyance threshold.

This does not mean, however, that visual fatigue is also unaffected by GACS3D. Visual Fatigue is induced through hour-long

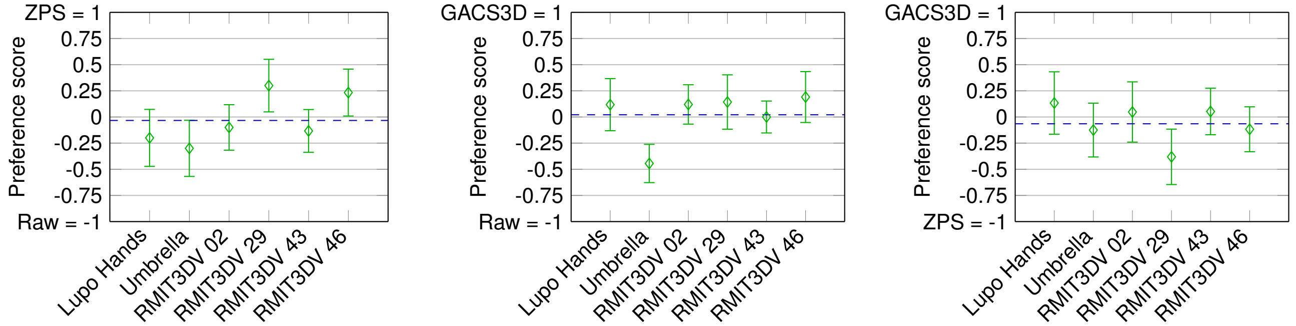

(a) “Raw” vs. “ZPS” (b) “Raw” vs. “GACS3D” (c) “ZPS” vs. “GACS3D”

Figure 7: Mean opinion scores (green) of all three condition combinations in experiment 3 with 95% confidence intervals and the mean value over all sequences (dashed blue).

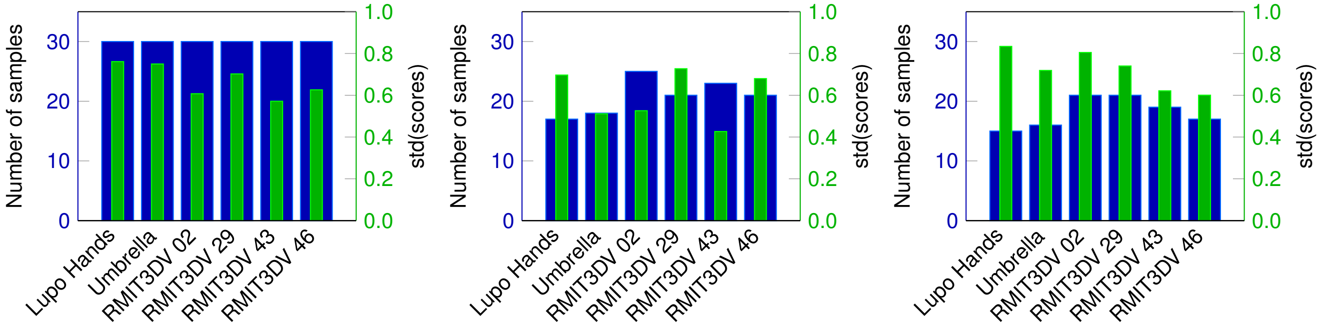

(a) “Raw” vs. “ZPS” (b) “Raw” vs. “GACS3D” (c) “ZPS” vs. “GACS3D”

Figure 8: Number of samples per MOS (blue) and standard deviation (green).

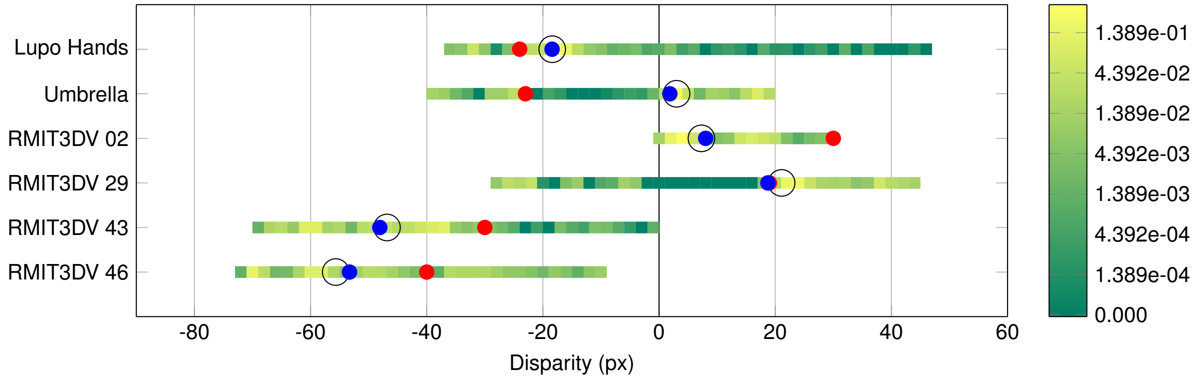

Figure 9: Relative frequencies of watched disparities (greenscale) in experiment 3 for each sequence. The plot also includes the “Raw”-Setting (black line), the “ZPS” (red circle), the average watched disparity (blue cicle) and the average converge disparity (black circle). This plot is displayed in the unshifted disparity domain. In the “ZPS” condition, the whole disparity budget would be shifted, so that the red dot is placed at zero disparity.

DOI: 10.2352/ISSN.2470-1173.2016.5.SDA-440

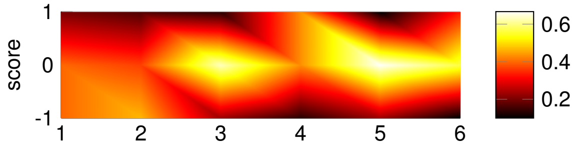

(a) “Raw” vs. “ZPS”

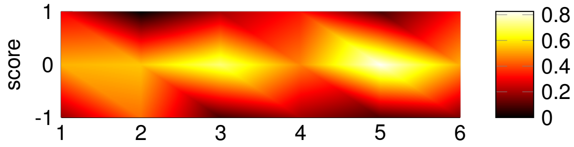

(b) “Raw” vs. “GACS3D”

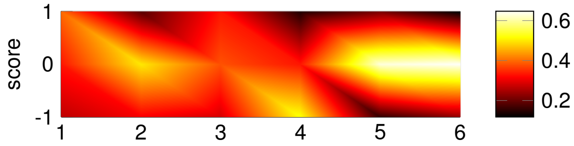

(c) “ZPS” vs. “GACS3D”

Figure 10: Relative frequencies of the scores in experiment 3.

exposures to harmful material and it takes a few hours for the human visual system to be restored afterwards. So, in order to evaluate visual fatigue, subjects would have to watch an appropriately long stereo 3D movie twice, once with and once without GACS3D and on separate days. The eye tracker would have to function properly the whole time, which proves to be the most challenging task at the moment. Using an autostereoscopic display to get rid of the stereo 3D glasses might help to improve eye tracker data quality, but has its own disadvantages.

We also found out that GACS3D almost degenerates to a static HIT in presence of very salient objects like faces or occlusions. Due to its fully automated and flexible nature, the approach still has many advantages over the conventional method, though, which is why it will be analyzed in even more detail in the future. As a next step, we will focus on a DHIT method free of any distortions of depth. The distortion is compensated by adjusting the camera baseline through depth image based rendering and has already been implemented in CUDA, yielding realtime performance.

References

[1] T. Shibata, J. Kim, D. Hoffman, and M. Banks, The zone of comfort: Predicting visual discomfort with stereo displays, J. Vis., 11, 8. (2011).

[2] M. Lambooij, W. IJsselsteijn, M. Fortuin, and I. Heynderickx, Visual Discomfort and Visual Fatigue of Stereoscopic Displays: A Review, J. Imaging Sci. Technol., 53, 3. (2009).

[3] B. Mendiburu, 3D TV and 3D Cinema, Elsevier, Oxford, 2011.

[4] P. Hanhart and T. Ebrahimi, Subjective evaluation of two stereoscopic imaging systems exploiting visual attention to improve 3D quality of experience, IS&T/SPIE Electronic Imaging, pg. 90110D. (2014).

[5] S. Eickelberg, Perception of Dynamic Horizontal Image Translation in Stereo 3D Content, Proc. ICCE-Berlin. (2015).

[6] C. Chamaret, S. Godeffroy, P. Lopez, and O. Le Meur, Adaptive

(a) EBU Lupo Hands

(b) NAMA3DS1 Umbrella

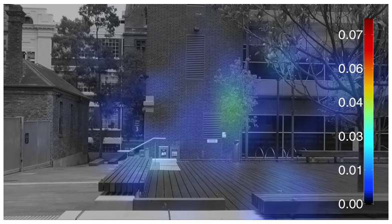

(c) RMIT3DV 02

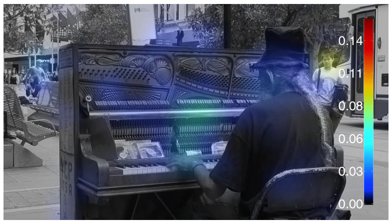

(d) RMIT3DV 29

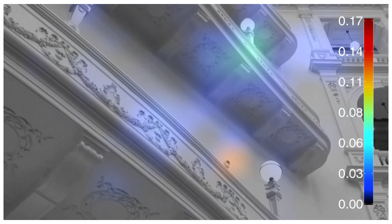

(e) RMIT3DV 43

(f) RMIT3DV 46

Figure 11: Exemplary frames of the test sequences with heatmap overlay, that shows the relative observation frequencies of certain image regions in experiment 3.

©2016 Society for Imaging Science and Technology
DOI: 10.2352/ISSN.2470-1173.2016.5.SDA-440

3D Rendering based on Region-of-Interest, in IS&T/SPIE Electronic Imaging, pg. 75240V. (2010).

[7] C. Ware, Dynamic stereo displays, Proc. SIGCHI, pg. 310. (1995).

[8] Q. Huynh-Thu and L. Schiatti, Examination of 3D visual attention in stereoscopic video content, Proc. SPIE/IS&T EI, pg. 78650J. (2011).

[9] Q. Huynh-Thu, M. Barkowsky, and P. Le Callet, The Importance of Visual Attention in Improving the 3D-TV Viewing Experience: Overview and New Perspectives, IEEE Trans. Broadcast., 57, 2, pg. 421. (2011).

[10] S. Eickelberg and R. Kays, Gaze adaptive convergence in stereo 3D applications, Proc. ICCE-Berlin. (2013).

[11] M. Bernhard, C. Dellmour, M. Hecher, E. Stavrakis, and M. Wimmer, The effects of fast disparity adjustment in gaze-controlled stereoscopic applications, Proc. ETRA, pg. 111. (2014).

[12] D. Kim, S. Choi, S. Park, and K. Sohn, Stereoscopic visual fatigue measurement based on fusional response curve and eye-blinks, Proc. DSP. (2011).

[13] F. Dvernay and P. Beardsley, Stereoscopic Cinema, in: Image and Geometry Processing for 3-D Cinematography, 5th ed., R. Ronfard and G. Taubin, Springer Verlag, Heidelberg, 2010, pg. 11.

[14] ITU, ITU-R BT.500: Methodology for the subjective assessment of the quality of television pictures. (2012).

[15] ITU, ITU-R BT.2021: Subjective methods for the assessment of stereoscopic 3DTV systems. (2012).

[16] EBU TECH, 3DTV Test Sequences [online], Available: `https://tech.ebu.ch/testsequences/3dtv_test`

[17] E. Cheng, P. Burton, J. Burton, A. Joseski, I. Burnett, RMIT3DV: Pre-Announcement of a Creative Commons Uncompressed HD 3D Video Database, Proc. QoMEX. (2012).

[18] M. Lambooij, M. Fortuin, W. a. Ijsselsteijn, and I. Heynderickx, Measuring Visual Discomfort associated with 3D Displays, Proc. SPIE/IS&T EI, pg. 72370K. (2009).

[19] D. K. Broberg, Guidance for horizontal image translation (HIT) of high definition stereoscopic video production, Proc. SD&A, pg. 11. (2011).

[20] K. Holmqvist, M. Nystrm, R. Andersson, R. Dewhurst, H. Jarodzka, and J. de Weijer, Eye tracking: A comprehensive guide to methods and measures, 1st ed., Oxford University Press, New York, 2011.

[21] M. Joos, M. Rtting, and B. M. Velichkowsky, Spezielle Verfahren I: Bewegungen des menschlichen Auges : Fakten, Methoden und innovative Anwendungen, in: Psycholinguistik. Psycholinguistics: Ein internationales Handbuch. An International Handbook, no. 1, Walter de Gruyter, Berlin, 2003, pg. 142.

[22] M. Urvoy, M. Barkowsky, R. Cousseau, Y. Koudota, V. Ricordel, P. Le Callet, J. Gutierrez, and N. Garcia, NAMA3DS1-COSPAD1: Subjective video quality assessment database on coding conditions introducing freely available high quality 3D stereoscopic sequences, Proc. QoMEX. 2012.

Author Biography

Stefan Eickelberg received his diploma in information technology from the TU Dortmund University (2011). Afterwards, he stayed at the university to start working at the Communication Technology Institute as a doctoral student in the field of electrical engineering. His research is focused on stereo 3d related visual comfort and image processing. He is a member of FKTG, IEEE, SID, SMPTE and VDE.