Adaptation Speed for Exposure Control in Virtual Reality

Claus B. Madsen[a] and Johan Winther Kristensen
Computer Graphics Group, Aalborg University, Aalborg, Denmark

Keywords: Virtual Reality, Computer Graphics, Exposure Control, Dynamic Range.

Abstract: We address the topic of real-time, view-dependent exposure control in Virtual Reality (VR). For VR to realistically recreate the dynamic range of luminance levels in natural scenes, it is necessary to address exposure control. In this paper we investigate user preference regarding the temporal aspects of exposure adaptation. We design and implement a VR experience that enables users to individually tune how fast they prefer the adaptive exposure control to respond to luminance changes. Our experiments show that 60% of users feel the adaptation significantly improves the experience. Approximately half the of users prefer a fast adaptation over about 1-2 seconds, and the other half prefer a more gradual adaptation over about 10 seconds.

1 INTRODUCTION

Virtual Reality (VR) technology has witnessed unprecedented growth in recent years, revolutionizing various fields including gaming, education, healthcare, and industry. The fundamental appeal of VR lies in its ability to transport users to synthetic environments, creating a sense of presence and immersion. However, achieving a seamless and comfortable user experience in VR necessitates overcoming several technical challenges. One critical aspect is the management of exposure levels within the virtual environment.

Exposure control encompasses the manipulation of visual parameters such as brightness, contrast, and dynamic range to ensure that the visual content presented to the user aligns with their physiological and perceptual capabilities. In VR, accurate exposure control is paramount for achieving a realistic impression of an environment in terms of its dynamic range of luminance values.

This paper investigates how VR users evaluate aspects of exposure adaptation in VR. Specifically, we design and implement an experiment which enables users to tune their personal preference regarding how quick and responsive the exposure adaptation should be in VR. The core contribution of this research lies in demonstrating that VR users prefer exposure adaptation and that they feel it increases the realism of the experience, in addition to giving specific guidelines for the temporal aspects of such adaptation.

[a] https://orcid.org/0000-0003-0762-3713

The paper is organized as follows. In section 2 we expand on the background for this work and describe related work. Section 3 then briefly describe the specific goals of this research and the approach taken. In section 4 we go through how we designed and implemented the VR experience used in the presented experiments. The design of the experiment is presented in section 5, followed by a presentatio of experimental results in section 6. Finally, section 7 offers a conclusion.

2 BACKGROUND AND RELATED WORK

The human eye can handle a vast range of luminance levels, where luminance here is the photometric concept of candela per square meter. In radiometry the corresponding concept would be radiance, but in this paper we shall stick to photometric terminology. With the rods and the cons in the retina, the human eye can *adapt* to widely varying light conditions, from the dim glow of starlight to the intense glow sunlight So, the real physical world has a huge dynamic range in terms of luminance levels, (Reinhard et al., 2010).

Unfortunately, cameras do not at all support the same dynamic range. Cameras have to *adjust the exposure* in order to find a sensitivity that is suitable for a given scene; and this exposure can be a compromise, for example when taking pictures indoor the windows may end up over-exposed. And probably everyone have experience how the camera on their smartphone

Madsen, C. and Kristensen, J.
Adaptation Speed for Exposure Control in Virtual Reality.
DOI: 10.5220/0012450300003660

In *Proceedings of the 19th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2024) - Volume 1: GRAPP, HUCAPP and IVAPP*, pages 307-312
ISBN: 978-989-758-679-8; ISSN: 2184-4321

will automatically adjust the exposure to ensure a decent image given the luminance levels of the scene. Humans also adapt their own sensitivity to luminance levels when, e.g., walking from an indoor to an outdoor environment.

Similarly, various types of displays, TVs, cinema screens, computer monitors, VR headsets, etc., do not support close to real world dynamic range, either. They simply cannot recreate luminance levels of the same magnitude as found in, say, and outdoor daylight scene.

Tone mapping and exposure adaptation are two terms that relate to the challenges of having to display high dynamic range content on low dynamic range displays. The goal of tone mapping is to, as closely as possible, on a given display, recreate the visual impression the viewer would have if they observed the real scene (Haines and Hoffman, 2018). In this paper we focus solely on exposure adaptation, i.e., adjusting the light sensitivity in response to changing luminance levels.

In film, exposure control can be used as a means for supporting viewer understanding of lighting conditions in a scene and how they change over time, (Bordwell et al., 2008). And in VR, we definitely also need to understand the pros and cons of exppsure control and tonemapping, if we want VR users to move around naturalistic looking scenes, and getting close to realistic perception of changes in luminance levels.

A method for view-dependent tone-mapping for in the case of 360-degree video was presented in (Najaf-Zadeh et al., 2017). This work concluded that users preferred a view-dependent over a global, fixed exposure for the whole sequence. Other research has focused on the challenges with adapting existing 2D single-image tone mapping approaches to view-dependent 360-degree and VR context. The challenges are related to achieving global consistency, while retaining the benefits of the view-dependent mapping, (Goudé et al., 2019; Goudé et al., 2020; Melo et al., 2018).

It would seem that there has been quite some research, although not a massive amount, into various technical aspects of employing exposure control and tone mapping in VR contexts. But, we have not really been able to find studies of whether users find such adaptive techniques appealing, or conducive to an enhanced sense of realism. Hence, the purpose of this paper was to investigate some of these aspects.

3 OVERVIEW OF APPROACH

The main research question behind this work was: If users experience a VR environment, where the rendering continuously performs view-dependent exposure adaptation such that center-view luminance levels are properly exposed, at what rate do these users then prefer the adaptation to happen? Should it happen quickly? Or in a more slow and subtle way? We were curious to find out, if we might experimentally determine a user concensus regarding this.

To experiment with this we designed a VR experience, where the user should be able to freely adjust the speed with which the exposure adaptation happens. The VR scenario should be familiar to the user, and it should entail some realistic luminance ranges. Our idea was then to subject a number of users to this environment and study if there was any systematic trends regarding which adaptation speed they preferred.

4 IMPLEMENTATION

For the purposes of conducting user experiments into the aspects of real-time dynamic tone mapping we designed and implemented a VR experience. The VR experience is single static space with no moving objects, but with illumination elements that offer various levels of luminance values, from a dark floor in shadow under a table, to ceiling mounted light sources, and a projector projecting a photo onto a wall. We used the Unity game engine for the implementation of the environment.

4.1 Scenario Design

We decided to let the VR environment for the tests be a recreation of the actual physical space where the experiments would take place. For two reasons. Firstly, because it was convenient to have a real physical version of what the VR environment should like in terms of dimensions, materials and illumination. And secondly, there is plenty of literature supporting the fact that sense of presence and ability to sense for example distances in VR is heightened by using transition environments, where the experience in VR is a recreation of the physical space the user ”comes from” when entering the VR, (Steinicke et al., 2009; Okeda et al., 2018; Soret et al., 2021).

The chosen space was an 8x8x4 meter room with no windows, so there is no natural daylight coming into the room. All elements in the VR scene had similar sizes, materials and placement as their real world

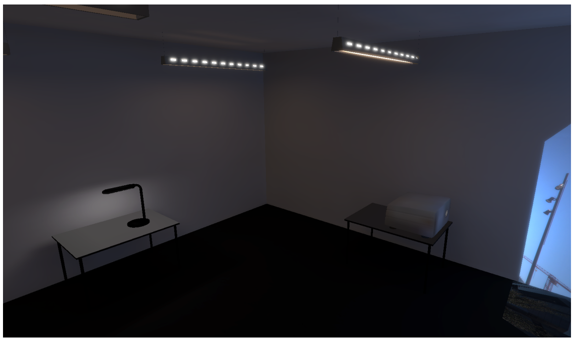

Figure 1: View of the VR environment developed for the experiment. The enviroment is a virtual recreation of the physical 8x8x4 meter room actually used for the experiments. Towards the right in this view it is possible to see how a project is projecting an color photograph onto a wall.

counterparts. Figure 1 shows a view of the room from one corner.

In the virtual scene there were three types light sources: 1) ceiling lights, 2) a table lamp aimed at a table, and 3) a projector light aimed at a wall.

The ceiling and the table lights were each implemented in Unity as two lights: One to light up the scene, and one with a very small range to light up fixture/light origin without spilling light into the scene. For the ceiling and table lights, the two light sources pointed in opposite directions: one pointed downwards to spread light into the scene, and one pointed upwards to light up the fixture. Ceiling lights were set to 650 lumen for the upwards light with a small range (1 meter as the holes in the fixture should allow for light on the ceiling) while the downwards light was set to 1900 lumen with a range of 10 meters. The table light upwards light was set to 1557 lumen at a range of 0.06 meters and a downwards light of 2850 lumen and a range of 10 meters.

The projector in the scenario was implemented in Unity as projector light at 40000 lumen and set up to use a High Dynamic Range image as a decal, while a point light at 295 lumen was placed at the front of the projector to light up the front and glass.

All lumen values were chosen based on subjective experience of the environment, not on actual measured values from real world lights. The lighting in all the scenes was baked to allow more natural lighting with global illumination.

4.2 Rendering and Exposure Control

The scenario was built using the High Definition Render Pipeline (HDRP) in Unity, as the HDRP has built-in functions for controlling exposure based on the light visible in the virtual camera. The HDRP has several modes of controlling the exposure control,

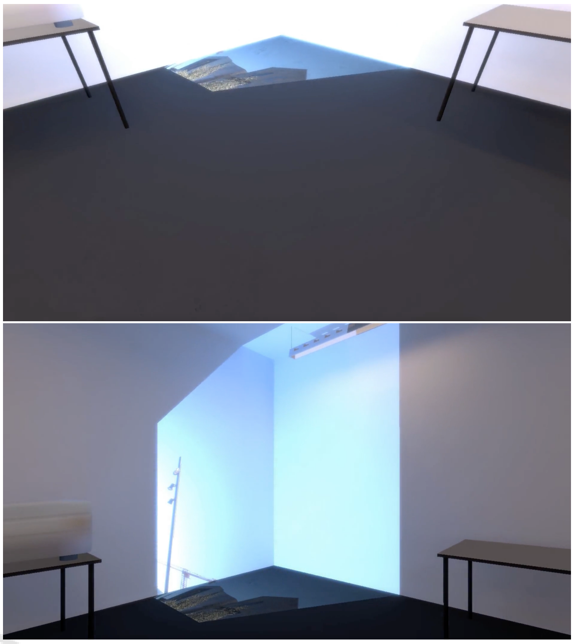

Figure 2: Top: the user is looking towards the dark floor in the VR room, and, with the center of the field of view being so dark, the dynamic adaptation adjust the exposure resulting parts of the wall becoming over-exposed and saturated. Bottom: the user is looking towards a projection on the wall, and the dynamic exposure control has adapted to the much higher luminance levels in the center of the view.

where the Automatic mode was used for this study. In Automatic mode, exposure control is done based on the within-frame content of the scene. In Automatic mode we opted to utilize the option of letting the exposure control be center-weighted such that whatever is in the center of the field-of-view at any given time is weighted higher when computing the proper exposure adjustment.

Unity generates the environment for VR by calculating each eye as its own camera, which means that in some cases one eye might have different light exposure to the other eye. To combat this we selected a softness of 5 so that pixels on the edge of the mask had less influence. This way, having a light at the edge of the camera would have less influence on the exposure which is similar to the eye's adaptation. The adaptation speed for the exposure was set to be between 0.01 seconds and 15 seconds, as any speed higher or lower had no discernible differences in adjustment. The test participants only controlled the speed from dark to light, with the speed from light to dark being set as one third of their chosen value. This was to recreate the standard human adaptation as our eyes more easily adapt to light than to dark environments.

4.3 Interaction and Adjustment Options

For the purposes of our experiments we wanted three types of interaction with the VR environment: 1) being able to freely look and move around the scenario, 2) being able to turn lights in the scene on and off to experiment with how the exposure control looked and felt, and 3) being able to adjust the speed with which the exposure control adapts to changes in luminance levels.

With regards to moving and looking around we opted to base all navigation on the tracking from the VR headset, so the participants had 6 degree of freedom movement, but no ability to do for example teleporting. So, all visual exploration of the scene was based on physically moving and looking around.

In terms of allowing participants to switch lights on or off, this was implemented as button presses on the controllers: using one face button on the right hand controller and two face buttons on the left hand controller, such that in total three lights could be toggled (desk lamp, ceiling lights, projector). Changing the status of a light was actually implemented as a load of a different scene, in order to have the lighting solution baked for increase performance and visual realism. So, we had scenes baked corresponding to all combinations of lights being on or off, respectively.

As mentioned in section 4 participants could adjust adaptation speed in the interval from 0.01 seconds (essentially instantaneous) to 15 seconds (almost too slow to be discernible). The current adaptation speed was visualized to test participants via an opaque sphere inside a large semi-transparent sphere locked to the position of their right controller. The inner sphere changed size depending on the current adaptation speed, with a higher speed meaning a larger sphere, and the outer sphere representing the fastest possible value, Figure 3. This way participants could see if they were close to the maximum or minimum speed available.

A tutorial scene was available to users with no prior VR experience, allowing them to get familiar with the VR controllers. This scene used the same setup as the other scenes, although the projector light was turned off. Here the user could press the buttons to squish and stretch a sphere or move a cube up and down using the joystick. A total of five participants tried the training scenario, each of which had never tried VR or only a short time before.

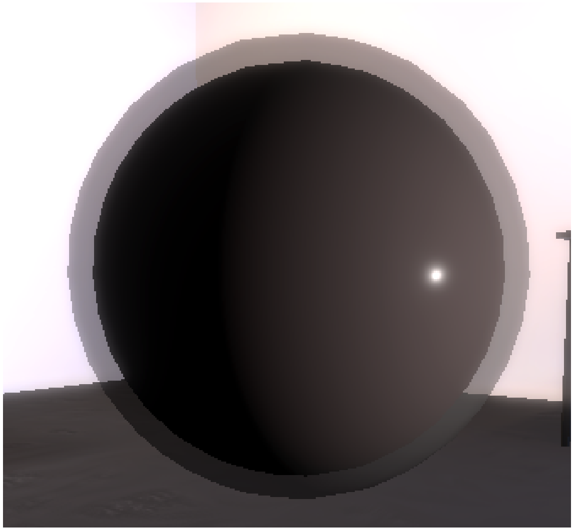

Figure 3: A view of the test participants adaptation speed as symbolised with an opaque sphere inside a transparent sphere. The opaque sphere would change size depending on the chosen speed, with a high speed resulting in a large sphere and vice versa. The maximum speed would set the opaque sphere at the same size as the transparent sphere.

5 EXPERIMENT DESIGN

The aim of the experiment was two-fold: 1) what adaptation speed does each test person prefer if there *is* dynamically adapting tone mapping, and 2) does the test person prefer dynamic adaptation, or not. The experiment was run with 21 participants of varying age (22 to 32, mean 27) and varying levels of VR experience (novice to experienced). The Meta Quest 2 VR headset was used for all experiments.

The room used for all experiments was identical to the scene that test participants were going to experience in VR, in terms of size, materials and illumination, etc.. So, once test participants put on the VR headset the whole VR scene and its appearance would be completely familiar to test participants.

Each test consisted of 2 phases. **Phase 1** started with a short introduction to the experiment, and gathering the age and previous VR experience of the participant. Each participant was then offered to try a training scene to familiarize themselves with using the controllers. Five participants with little to no prior VR experience opted to try the training scene.

Once in the actual **Phase 1** VR test scenario, test participants were instructed how to move around, how to turn individual lights in the scene on and off, and how to adjust the exposure adaptation speed, in the range described in section 4. Participants were not told about adaptation speed values or had access to any numerical read-out. When adjusting the adaptation speed they could just see the solid sphere inside the transparent sphere visually indicating the current value relative to the maximum, as described in section 4.3. Participants were then given the time they wanted to move and look around, explore, and experiment with setting the adaptation speed to whatever value they preferred. Once a participant had settled on a preferred adaption speed they were allowed to take off the VR headset, and the chosen adaptation speed would be logged.

After this, the participant would be instructed of the objective of **Phase 2**. Here the participants would be exposed once more to the exact same VR scene, but this time they could not adjust adaptation speed, they could only toggle dynamic adaptation on/off, where on would be using the adaptation speed the participant had elected as preferred in **Phase 1**, and off would be no dynamic adaptation at all, only a static generic exposure value chosen by experimenters to visually mimic an impression of the general luminance level of the real room. Participants were then given the time they wanted to move and look around, explore, turn lights on and off, and toggle dynamic adaptation. Once satisfied, participants would be allowed to take off the VR headset, asked if they preferred dynamic adaptation or not, and their answer to this would be logged. Subsequently, the participant be asked to rate the how much the adaptation improved they experience, and how natural the adaptation felt. In both cases using a Likert 1 - 7 scale where 1 would be not at all, and 7 would be significantly. And this would conclude the test session for the participant.

6 RESULTS AND DISCUSSION

Participants tested the adaptation speed in various ways. Most of the participants looked toward the projector in the virtual space to test the speed when going from dark to bright, as it was the brightest source in the scene. Some participants turned to look at the floor to get the dark adaptation, as the floor was dark. Other participants turned the lights on and off instead, although not so many utilized this option.

Figure 4 shows a histogram over what adaptation speeds test participants preferred in Phase 1 of the experiment. Participants tend to fall in two groups. Either they prefer a slow and subtle adaption, or the prefer a faster, clearly visible adaptation. The preferred values do not conform to a normal distribution in a Shapiro-Wilk test ($p = 0.07$). More test persons would be needed to establish whether this parameter is actually a bi-modal distribution.

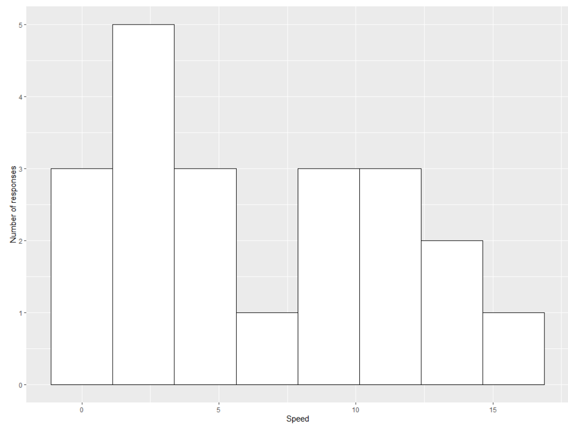

Figure 4: Histogram over preferred adaptation speeds logged during **Phase 1** of the experiment.

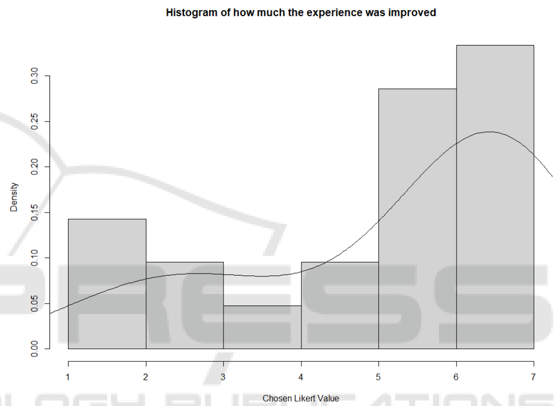

Figure 5: Histogram over test participant responses to the question "To what extent do you feel the adaptation improved the visual experience?", using a 1-7 Likert scale, with 1 being not at all, and 7 being significantly.

The clear results from **Phase 2** was that 17 out of 21 participants preferred dynamic adaptation over no adaptation. Three preferred without, and one was undecided. This particular participant set an adaptation speed at its lowest possible value, making the adaptation happen over several seconds, and therefore the participant had a hard time telling when the adaptation was on or off. The participant commented on this but chose to continue with the very low speed. Figure 5 shows how participants scored to what extent the adaptation improved the experience, with approx. 60% scoring significantly (6 and 7).

After the experiment, participants were also asked to rate how certain they felt in their choice of preferred adaptation speed, using a 1-7 Likert scale. Roughly one third scored their certainty around 3 to 4, and the remaining participants scored their certainty high (5, 6, or 7). There was no statistically significant correlation between time spent in the test environment, and self-reported certainty with preferred adaptation speed. There was also no statistically significant correlation between previous VR experience and self-reported certainty in preferred adaptation speed.

7 CONCLUSIONS

We designed a VR experience aimed at testing aspects of user preference regarding aspects of dynamic tone mapping in VR applications. The VR experience was a virtual recreation of the physical room used for user experiments. The VR experience allowed users to individually set how fast they preferred the dynamic luminance adaptation be, ie., how quickly they wanted adaptation to respond to drastic changed in luminance values between different parts of the VR scene, for example between looking at a dark floor and then shifting your viewing direction towards a light source or a bright wall.

The experiment clearly showed that test participants prefer dynamic adaptation, and that they feel this adaptation significantly improves the experience. Out of 21 test participants, 17 preferred the version of the VR experience that had dynamic luminance adaptation over the version that did not. And 60% reported that it significantly improved the experience.

The experiment seemed to indicate that participants fall in two groups in terms of preferred adaptation speed, i.e., how quickly they want the adaptation to respond to luminance changes. About half of the participants preferred a slow, relatively subtle adaptation, whereas the other half preferred a faster, clearly perceptible, and more immediate adaptation response.

For future work it would be extremely interesting to investigate whether does dynamic adaptation somehow tricks test participants into believing that the dynamic range of luminance in the scene is actually higher than what the VR headset can recreate. For example, is a light source in a VR scene perceived brighter when experienced with dynamic luminance adaptation than without.

ACKNOWLEDGEMENTS

This research was partially funded by the CityVR project, funded by the Department of Architecture, Design and Media Technology, Aalborg University, and partially by the BUMUS project, funded by the Danish Ministry of Food, Agriculture and Fishery. This funding is gratefully acknowledged.

REFERENCES

Bordwell, D., Thompson, K., and Smith, J. (2008). *Film art: An introduction*, volume 8. McGraw-Hill New York.

Goudé, I., Cozot, R., and Banterle, F. (2019). Hmd-tmo: a tone mapping operator for 360 hdr images visualization for head mounted displays. In *Computer Graphics International Conference*, pages 216–227. Springer.

Goudé, I., Lacoche, J., and Cozot, R. (2020). Tone mapping high dynamic 3d scenes with global lightness coherency. *Computers & Graphics*, 91:243–251. ID: 271576.

Haines, E. and Hoffman, N. (2018). *Real-time rendering*. Crc Press.

Melo, M., Bouatouch, K., Bessa, M., Coelho, H., Cozot, R., and Chalmers, A. (2018). Tone mapping hdr panoramas for viewing in head mounted displays. In *Proceedings of the 13th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, volume 1, pages 232–239. Scitepress.

Najaf-Zadeh, H., Budagavi, M., and Faramarzi, E. (Sep 2017). Vr+hdr: A system for view-dependent rendering of hdr video in virtual reality. pages 1032–1036. IEEE.

Okeda, S., Takehara, H., Kawai, N., Sakata, N., Sato, T., Tanaka, T., and Kiyokawa, K. (2018). Toward more believable vr by smooth transition between real and virtual environments via omnidirectional video. In *2018 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, pages 222–225. ID: 1.

Reinhard, E., Heidrich, W., Debevec, P., Pattanaik, S., Ward, G., and Myszkowski, K. (2010). *High dynamic range imaging: acquisition, display, and image-based lighting*. Morgan Kaufmann.

Soret, R., Montes-Solano, A.-M., Manzini, C., Peysakhovich, V., and Fabre, E. F. (2021). Pushing open the door to reality: On facilitating the transitions from virtual to real environments. *Applied Ergonomics*, 97:103535. ID: 271441.

Steinicke, F., Bruder, G., Hinrichs, K., Steed, A., and Gerlach, A. L. (2009). Does a gradual transition to the virtual world increase presence? In *2009 IEEE Virtual Reality Conference*, pages 203–210. IEEE.