Foveated Displays: Toward Classification of the Emerging Field

Josef Spjut
Ben Boudaoud
NVIDIA

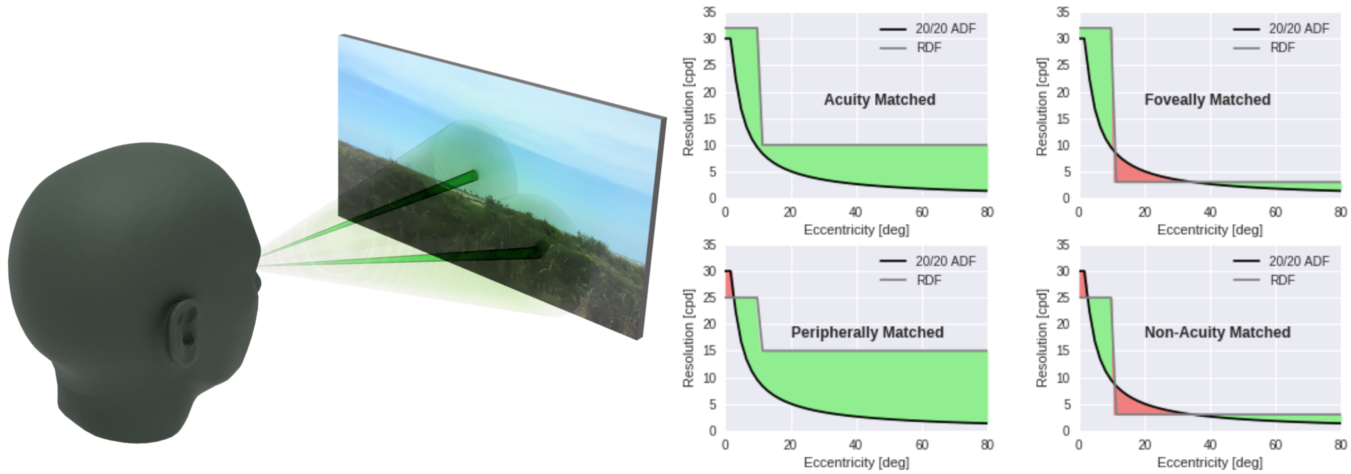

Figure 1: A "Foveated Display" is a display designed to function in the context of user gaze. This can mean it follows gaze direction or expects to be gazed upon in a certain region. Additionally, foveated displays have the potential to vary their resolution (in cycles per degree) using a perception-inspired resolution distribution function. The right image illustrates our 4 proposed classes for comparing the user acuity distribution function with imagined display resolution distribution functions.

ABSTRACT

There is not yet consensus in the field on what constitutes a “foveated display”. We propose a compromise between the perspectives of rendering, imaging, physiology and vision science that defines a foveated display as a display designed to function in the context of user gaze. This definition enables us to describe 2 axes of foveation, gaze interaction and resolution distribution, which we then subdivide to provide useful categories for classification. We view this proposal as the start of a discussion among the community rather than a final taxonomy.

CCS CONCEPTS

• **Computing methodologies** → **Mixed / augmented reality**; **Perception**; **Virtual reality**; • **Hardware** → **Displays and imagers**.

KEYWORDS

displays, foveated rendering, gaze detection, head mounted displays

SIGGRAPH '19 Talks, July 28 - August 01, 2019, Los Angeles, CA, USA

ACM ISBN 978-1-4503-6317-4/19/07.
https://doi.org/10.1145/3306307.3328145

ACM Reference Format:
Josef Spjut and Ben Boudaoud. 2019. Foveated Displays: Toward Classification of the Emerging Field. In *Proceedings of SIGGRAPH '19 Talks*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3306307.3328145

1 INTRODUCTION

Due to the diverse backgrounds of those building and advocating foveated displays, it has become apparent that some disagreement exists regarding what the term “foveated display” refers to. This disagreement can challenge communication and make integration of the components needed to create foveated displays more difficult. In defining foveation, we believe it is vital to build on common ground and find compromises that accelerate progress.

In computer graphics, well known proposals [Guenter et al. 2012; Patney et al. 2016] have led to a wide understanding of foveation as primarily matching computational load (resolution) to human visual acuity. As a result, someone from the field of rendering or imaging is likely to hear the term “foveated display” and immediately think of a display that has a distribution of pixels that is somehow reduced in the periphery relative to the center (fovea). However, in physiology and vision science, the verb “foveate” is used to describe the motion of angling ones eye in a particular direction, aligning the fovea with content of interest [Kaufman et al. 2011]. Thus the term “foveated display” is likely to invoke thoughts of a display that moves in response to such changes in gaze direction.

While clearly related, these definitions often lead to different design choices which, without clear communication, can become

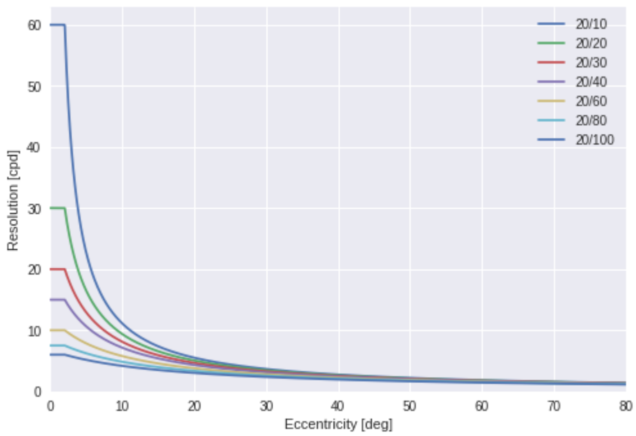

Figure 2: Acuity distribution function approximation over a range of common visual acuity

challenging to manage. Some have used terms like "fixed foveated display" or "dynamic foveated display" to describe subsets of the possible displays that might be considered "foveated", but depending on the background of the speaker and the listener, these terms may take on different meanings. As a result, we propose a compromise that is simple in purpose with a few concrete recommendations that we hope will be useful in bringing these distinct perspectives together. Our proposal hinges on the statement that a foveated display is designed to function in the context of user gaze. This statement includes both: (1) displays that have a resolution distribution that mimics that of cones in the retina and (2) displays that match this distribution across the gaze angle(s) of the user. We propose further classification along each of these 2 axes. This talk is a presentation of content from a to be published technical report [Spjut et al. 2019].

2 ACUITY MATCHING

The density of cones in the retina is much higher in the center of the field of view, or fovea, than in the periphery [Curcio and Allen 1990]. To help characterize this, we define a visual Acuity Distribution Function (ADF) as the angular resolution in cycles per degree (cpd) *as perceived* by a user as a function of *gaze eccentricity*. We approximate the ADF for various levels of visual acuity as seen in Figure 2. Similarly, we define the Resolution Distribution Function (RDF) of a display as its resolution as a function of eccentricity across its field of view.

We propose 4 classifications based on a comparison of RDF of the display to the ADF of a user as in Figure 1. These are classes: (A) Acuity matched, (B) Foveally matched, (C) Peripherally matched, and (D) Non-acuity matched. It is important to specify along with this letter classification which acuity is used for the classification (e.g. 20/20). When the RDF of a display meets/exceeds the ADF in both the fovea and periphery it is said to be *acuity matched*, while if the display is matched *only* in the fovea or periphery it is matched to this region. A display that meets neither foveal nor peripheral acuity is *non-acuity matched.*

3 GAZE CONTINGENCY

The other axis of foveation in our proposal is the range of gaze interaction supported by the display, in other words, the range of user eye motion over which the display RDF exceeds user ADF. Similar to the previous axis, we propose 4 classifications we believe to be useful for various displays, depending on the intended application space. These are classes: (1) Fully foveated, (2) Practically foveated, (3) Partially foveated and (4) Non-foveated or fixed foveated.

A *fully foveated* (Class 1) display is one in which acuity matching (at any given level) holds for every possible gaze direction. While this is ideal, it is unlikely necessary to achieve this level of motion for a wide range of applications, which is why *practically foveated* (Class 2) displays are useful to define. Based on personal experiences and rough surveys of gaze across a narrow range of applications, we believe a useful threshold to consider something practically foveated is $\pm 15°$ of eye motion. The *partially foveated* (Class 3) classification exists for displays supporting less eye motion than what is practical for a wide range of experiences, but more than a single gaze direction, that is $<15°$ ($\pm 7°$) but $>0°$ of gaze direction. Finally, displays which do not maintain an ADF vs RDF relationship for any eye motion, are referred to as *non-foveated* (Class 4).

4 CONCLUSION

We propose both a definition and a classification methodology for the emerging field of foveated displays, based on comparison between the RDF of the display and the ADF of the user, as well as the supported range of gaze directions. To our knowledge, few comprehensive hardware solutions to these problems have yet to be proposed, thus this represents an attempt to organize the efforts of the field around what matters most for foveated displays. We hope that by proposing this classification, we help establish common ground, make steps toward consensus, and provide strata that encourage hardware developers to move toward more human-centered display design.

ACKNOWLEDGMENTS

This work is the result of many heated discussions among members of the New Experiences Research group at NVIDIA and their collaborators, and is a compromise of many perspectives. This group includes Jonghyun Kim, Trey Greer, Rachel Albert, Joohwan Kim, Michael Stengel, Kaan Aksit, David Luebke, Alexander Majercik, Anjul Patney, Mark Kilgard, Morgan McGuire, Peter Shirley, Turner Whitted, Ward Lopes, Kishore Rathinavel, Praneeth Chakravarthula, Henry Fuchs and Fu-Chung Huang.

REFERENCES

Christine A Curcio and Kimberly A Allen. 1990. Topography of ganglion cells in human retina. *Journal of comparative Neurology* 300, 1 (1990), 5–25.
Brian Guenter, Mark Finch, Steven Drucker, Desney Tan, and John Snyder. 2012. Foveated 3D graphics. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 164.
Paul L Kaufman, Leonard A Levin, Francis Heed Adler, and Albert Alm. 2011. *Adler's Physiology of the Eye.* Elsevier Health Sciences.
Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. 2016. Towards foveated rendering for gaze-tracked virtual reality. *ACM Transactions on Graphics (TOG)* 35, 6 (2016).
Josef Spjut, Ben Boudaoud, Jonghyun Kim, Trey Greer, Rachel Albert, Michael Stengel, Kaan Aksit, and David Luebke. 2019. Toward Standardized Classification of Foveated Displays. In *Technical Report - to be published.*