Breaking the limits of display and fabrication using perception-aware optimizations

Piotr Didyk, Università della Svizzera italiana; Lugano, Switzerland

Abstract

Novel display devices and fabrication techniques enable highly tangible ways of creating, experiencing, and interacting with digital content. The capabilities offered by these new output devices, such as virtual and augmented reality head-mounted displays and new multi-material 3D printers, make them real game-changers in many fields. At the same time, the new possibilities offered by these devices impose many challenges for content creation techniques regarding quality and computational efficiency. This paper discusses the concept of perception-aware optimizations, which incorporate insights from human perception into computational methods to optimize content according to the capabilities of different output devices, e.g., displays, 3D printers, and requirements of the human sensory system. As demonstrated in this paper, the key advantage of such strategies is that tailoring computation to perceptually-relevant aspects of the content often reduces the computational cost related to the content creation or overcomes certain limitations of output devices. Besides discussing the general concept, the paper presents several specific applications where perception-aware optimization has been proven beneficial. The examples include methods for optimizing visual content for novel display devices that focus on perceived quality and new computational fabrication techniques for manufacturing objects that look and feel like real ones.

Introduction

In recent years, there has been a tremendous increase in the quality and the number of new output devices. Novel virtual and augmented reality headsets are being developed to provide new, interactive and immersive ways to explore the real and the virtual worlds. 3D printers empower millions of users to create, customize, and manufacture digital models. While many of these technologies have already proven their benefits and gained many users, others will follow this trend soon. Due to these developments, the borders between the virtual and the real worlds are diminishing, and new ways of interacting with digital data are created. The synergy between the digital and the real words will lead to a better understanding and analysis of the presented information and a more immersive and engaging experience. These new technologies have the potential to play a crucial role in applications such as entertainment, education, and communication. They also will improve user performance whenever understanding and analyzing visual and haptic information is crucial. Such technologies have only been experimental and accessible in very specific fields in the past. Due to technological developments and drastic price drops of hardware, these advances are becoming available to the mass market.

This rapid technology development significantly improves but also drastically changes the way we experience and interact with digital content. The final success of new technologies depends on the experience quality that a particular solution will provide, which brings us to several fundamental questions: *What* is the quality? How do different hardware designs affect the quality? How can we model and measure the quality? How can we use the models to improve the user experience? While these questions are critical for providing solutions that fully exploit the capabilities of the new technologies and provide an optimal experience to the users, answering them poses many challenges and requires significant research efforts. For example, novel virtual reality display devices offer an immersive experience by providing wide-field-of-view images at very high resolution. These capabilities impose significant constraints on the provided image quality as rendering such images becomes computationally expensive. However, at the same time, it has been demonstrated that human sensitivity to visual image artifacts and distortions is greatly reduced in the far periphery. Understanding exact quality requirements across the wide field of view will enable rendering quality to align better with human perception. This can not only bring significant computational savings but also improve image quality where necessary. Furthermore, often the abilities of some technologies outperform the capabilities of current methods and tools for creating content. For example, despite the high quality of multi-material 3D printers, reproducing hard copies of real objects is still a challenging task. One of the reasons is the lack of proper tools for creating content. These tools, however, cannot be developed without a proper understanding of how people interact with and judge the quality of 3D prints.

This paper describes the concept of perception-aware optimizations which combines insights from human perception, computation, and hardware design. The key idea behind these techniques is to take the properties of the human sensory into account while designing content creation and optimization techniques. Doing so allows the computational efforts to create content that is relevant to the perception while avoiding generating irrelevant content for the users' experience. Also, by exploiting the limitations of the human sensory system together with information about the limitations of the hardware, it is possible to trade certain quality aspects and achieve superior content reproduction. While the general concept of perception-aware optimizations is presented in the next section, the paper also presents several applications where such a strategy was successfully applied. These include optimizations of visual content for novel display devices and building new computational fabrication techniques for haptic properties reproduction.

Methodology

Developing perception-aware optimization usually requires addressing three stages: perceptual experiment, perceptual modeling, and devising optimization techniques that leverage perceptual models. Figure 1 presents a concept for building such optimization for graphical display content. Below, we briefly describe each of the stages, highlighting their challenges.



Figure 1. The main concept and building blocks of perceptual optimization on an example of optimizing visual content for novel display devices.

Perceptual experiments

Gaining insights about perception involves performing perceptual experiments which reveal essential for the application properties of the human sensory system. Many such experiments are conducted using simple stimuli that focus on a specific perceptual phenomenon, such as visibility and detectability. Here, examples may include measuring contrast sensitivity function (CSF) using sinusoidal patterns or Gabor patches. Focusing on a very narrow range of stimuli makes the experiments feasible and eases the data analysis. Unfortunately, it is often unclear how to extend such results to complex stimuli, such as images, videos, or complex 3D objects, which are the subject of perception-aware optimizations. Therefore, during designing perceptual experiments, care has to be taken that the obtained results can be later incorporated into perceptual models which can handle complex stimuli.

Perceptual modeling

While perceptual experiments are performed on a small number of simple stimuli, it is critical for perceptual models to extend the findings to more general and complex inputs. The tasks of the perceptual models can be different but usually involve estimating perceived quality or determining the detectability of specific types of signals by human senses, such as vision or touch. There are three main challenges when building such models. First, the extension to complex stimuli is often non-trivial. The models usually have to operate on a custom representation of stimuli. For example, for image content, often frequency decompositions are used. The second challenge is the accuracy of the models. It is clear that models serve only as an approximation of the processes performed by the human sensory system. Still, their prediction accuracy significantly impacts the final results of perceptual optimization. Therefore, their design has to be guided by the insights from the perception, and their accuracy carefully evaluated. The third challenge is computational efficiency. Perceptual models are often used in optimization loops or applications where real-time performance is critical, e.g., realtime rendering. Therefore, throughout the design process, it is critical to ensure that both the representations computation and the modeling can be performed efficiently.

Optimization

The final building block of perception-aware optimizations is the optimization itself. The goal here is to utilize the perceptual models to optimize content. Such a method can be classical optimization which defines a minimization problem solved using standards numerical solvers. Other types of optimizations may include direct methods that compute the optimal content. Regardless of the solution, the methods usually have at least two aspects in common. First, they utilize the perceptual model to guide the optimization. Here, the main goal is that the final stimuli presented to the user are not optimal according to an arbitrary numerical metric but a carefully designed perceptual metric that captures the content's perceptually-important features. Second, the optimization is guided by the capabilities of the output devices, e.g., display, 3D printer, and often computational constraints. As a result, the final output respects both human, hardware, and application limitations. Apart from that, similar to the perceptual models, the computational efficiency of the final optimizations is a significant concern. Keeping all the components of the perceptual optimization computationally efficient and accurate is one of the main challenges in the entire process.

Applications

This section provides several examples of how perceptionaware optimizations can be developed and applied to optimize visual content (perceptual display) and objects for digital fabrication (perceptual fabrication).

Perceptual display

Optimizations of content for display devices usually have to find a trade-off between different image qualities (e.g., spatial and temporal resolution, depth), limitations of the display, computational performance, and the requirements of the human visual system (HVS). Three examples here demonstrate applications of perception-aware optimizations to the problems of optimal spatial resolution and depth reproduction.

Apparent resolution enhancement Despite significant development in display technology, showing images with a spatial resolution close to the capabilities of the HVS is challenging. In our work, we demonstrated that it is possible to display images whose apparent resolution exceeds the physical resolution of a display device [1]. The key idea is to exploit the fact that the HVS integrates signal over time. More precisely, if the temporal changes on the screen exceed the so-called critical flicker frequency, the fluctuations are not perceived; instead, an averaged signal is perceived. This opens the doors for showing multiple sub-images that are integrated and perceived as one image. Our work demonstrated that this becomes particularly interesting for moving content. As it moves across the screen, so does the gaze of the human observer and the individual photoreceptors on the human retina that observe different display pixels. This allows for optimizing the values of the pixels to obtain apparent resolution enhancement. We propose a simple model that predicts the perceived image, given several sub-images and the motion. The model is then turned into an optimization that seeks the subimages that the HVS integrates into a high-resolution image. An example output of such optimization is shown in Figure 2. In that particular case, the method was used to optimize three subimages for 120 Hz display. The follow-up works extend this technique to videos [2] and improve the efficiency and applications to virtual reality headsets [3].



Figure 2. One of the results of the apparent resolution enhancement technique. On the left: three optimized sub-images; on the right: simulation of the image perceived by a viewer.

Foveated rendering While the human visual system can resolve very fine spatial details, this ability holds only for the central region of the vision, the fovea. In the periphery, the ability of the HVS to perceive high spatial frequency decreases. This phenomenon led to a family of so-called foveated rendering methods that improve rendering efficiency by reducing the quality of synthesized content for peripheral vision. An example of such techniques are methods that reduce spatial resolution, e.g., [4]. Most such techniques rely on modeling the reduction in perceived resolution using fall-off of receptor density on the human retina. This makes the required rendering resolution solely dependent on eccentricity. However, we have recently demonstrated that more accurate modeling of human perception can bring additional computational benefits [5]. Our method exploits the fact that the visibility of resolution loss also depends on the underlying content. It is much more likely that a viewer can detect the foveation for high-frequency and high-spatial-frequency image content than for low-contrast regions. The work demonstrates that this dependency can be efficiently modeled and used in realtime applications to optimize parameters of the foveated rendering. The example prediction of rendering resolution is presented in Figure 3.



Figure 3. Exemplar prediction of possible rendering resolution reduction. While the general trend is to lower the resolution as the distance to the gaze increases, the technique adapts to the underlying content. Lower resolution is needed for regions which do not contain high-contrast regions.

Depth reproduction New near-eye displays provide separate images to both eyes and, therefore, can reproduce one of the strongest cues for perceiving scene layout, binocular disparity. Unfortunately, most of the displays cannot reproduce accommodation cues. The viewer's eyes are constantly focusing at a fixed distance, the distance to the screen. The usual mismatch between the binocular disparity cue and the constant accommodation is a significant source of visual discomfort [7]. A solution

to the problem is to compress the depth information such that the reduced discrepancy between the binocular disparity limits viewing discomfort at the cost of depth reproduction. More advanced techniques use computational models of the HVS to access the magnitude of the perceived depth from binocular disparity. The estimation of perceived depth enables perception-aware optimizations of depth compression, which preserve visually important depth details [8, 10, 9]. It is possible to improve such techniques if information about gaze location is provided. The reduced sensitivity to distortions in the periphery also applies to depth perception. The sensitivity of the HVS to local depth variation decreases with eccentricity. Therefore, it is possible to apply more aggressive depth compression in the periphery while maintaining the original or less compressed depth in the fovea region [6]. Figure 4 presents an example of applying gaze-contingent depth manipulation to a stereoscopic screen.



Figure 4. The gaze-contingent technique limits viewing discomfort by compressing the depth of the scene in periphery while preserving it around the current gaze. The depth manipulations are performed by optimizing smooth depth mapping curve shown on the right.

Perceptual fabrication

Novel digital fabrication tools, such as 3D printers, enable the creation of highly-detailed objects from a wide range of different materials. Their capabilities enable fine control over the objects' mechanical, functional, appearance, and haptic properties. Unfortunately, due to their differences, e.g., printing resolution and used materials, it is often unclear how to best reproduce the designs' properties using a particular device. Perceptionaware optimizations for fabrication try to simultaneously account for the capabilities of the fabrication tools and the human sensory system. They enable finding optimal object designs, which both respect the limitations of the fabrication tools and the perceptual requirements. Below, we present two examples where such optimizations have proven to bring clear benefits.

Compliance properties Multi-material 3D printers enable fabricating materials with a wide range of compliance properties. While the physical stiffness is often characterized using a socalled force-displacement curve, it is unclear how humans interpret such information when investigating samples made of different materials. Therefore, in applications that involve touch interaction with printed samples, optimizing designs for digital fabrication by minimizing a simple error defined on the forcedisplacements may lead to suboptimal reproduction. Recently, we have demonstrated that it is possible to build a model that predicts perceived compliance for complex objects [11]. The model is based on experiments where participants compare samples made of different materials. The study results enabled building a so-called perceptual space and approximating it with a computational model which predicts perceived compliance based on the physical properties of the material and the geometry. The model can then be used to compare objects made of different materials (Figure 6), but more interestingly, to drive optimization of the designs. The experiments demonstrated that such optimizations could improve the reproduction of the compliance properties and



Figure 5. Overview of the procedure for developing styli that match real drawing tools. The process involved measuring and characterizing drawing tools, studying their perception, and designing a perception-aware space of the tools. The simulation technique generalized the space to new designs and enabled a perceptual-aware design process.

ease the design process by optimizing the material pallets used for rapid prototyping.



Figure 6. The perceptual model for compliance enables computing perceptual differences between two objects made of different materials. The process involves computing a map of perceived compliance for each object and subtracting the values corresponding to the same locations. The obtained values are scaled in just-noticeable-difference units (JNDs).

Tactile properties Perception-aware optimizations also found applications in the reproduction of more complex haptic properties. In this particular case, we considered the problem of reproducing tactile feedback of a drawing surface and a tool to mimic the feel of real drawing tools using digital styli [12, 13]. Despite the complexity of the phenomena governing the perception of drawing tools, we demonstrated that it is possible to investigate the perception of drawing tools, build the corresponding perceptual representation, and create a model which predicts the perceived differences between different designs. Using the model enabled building an optimization routine that finds the best digital counterparts of real drawing tools (Figure 5). Examples of the results are presented in Figure 7. One of the most exciting results demonstrated is that very different designs can often lead to very similar perceptions. This proves the purpose of perceptual-aware optimizations, which can find these non-trivial replicas.



Figure 7. Pairs of real drawing tools and surfaces together with our replicas. Despite significant differences in materials, shapes, and surface geometries, our tools provide very similar feedback to their real counterparts.

Conclusion

In this paper, we discussed the concept of perception-aware optimizations, which seek optimal reproduction of the different types of content in a perceptual sense, i.e., maximizing perceived quality while minimizing computational effort or mitigating limitations of the output devices. We also demonstrated several applications for which the approach provides benefits that could not be achieved if the processes governing perception were not considered. We argue that for many new output devices, such as head-mounted displays or 3D printers, such techniques will play a key role, and they will enable taking full advantage of these devices and increase their efficiency.

Acknowledgments

Projects described in this paper received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement $N \circ 804226$ PERDY).

References

- Piotr Didyk, et al., Apparent display resolution enhancement for moving images, ACM Transactions on Graphics (TOG) 29,4 (2010).
- [2] Krzysztof Templin, et al., Apparent resolution enhancement for animations, Proceedings of the 27th Spring Conference on Computer Graphics, (2011).
- [3] Haebom Lee and Piotr Didyk, Real-time apparent resolution enhancement for head-mounted displays, Proceedings of the ACM on Computer Graphics and Interactive Techniques, (2018).
- [4] Brian Guenter, et al., Foveated 3D graphics, ACM Transactions on Graphics (TOG) 31,6 (2012).
- [5] Okan Tarhan Tursun, et al., Luminance-contrast-aware foveated rendering, ACM Transactions on Graphics (TOG) 38,4 (2019).
- [6] Petr Kellnhofer, et al., GazeStereo3D: Seamless disparity manipulations, ACM Transactions on Graphics (TOG) 35,4 (2016).
- [7] Takashi Shibata, et al., The zone of comfort: Predicting visual discomfort with stereo displays, Journal of vision 11,8 (2011).
- [8] Piotr Didyk, et al., A perceptual model for disparity, ACM Transactions on Graphics (TOG) 30,4 (2011).
- [9] Petr Kellnhofer, et al., Motion parallax in stereo 3D: Model and applications, ACM Transactions on Graphics (TOG) 35,6 (2016).
- [10] Piotr Didyk, et al., A luminance-contrast-aware disparity model and applications, ACM Transactions on Graphics (TOG) 31,6 (2012).
- [11] Michal Piovarči, et al., An interaction-aware, perceptual model for non-linear elastic objects, ACM Transactions on Graphics (TOG) 35,4 (2016).
- [12] Michal Piovarči, et al., Perception-aware modeling and fabrication of digital drawing tools, ACM Transactions on Graphics (TOG) 37,4 (2018).

[13] Michal Piovarči, et al., Fabrication-in-the-loop co-optimization of surfaces and styli for drawing haptics, ACM Transactions on Graphics (TOG) 39,4 (2020).

Author Biography

Piotr Didyk is an Associate Professor at Università della Svizzera italiana (USI), Switzerland. He obtained his Ph.D. from Saarland University and MPI Informatics (2012). Before joining USI, he was an Independent Research Group Leader at Saarland University/MPI Informatics and a postdoctoral associate at MIT. In 2019, he was elected Junior Fellow of the European Association for Computer Graphics. His research combines perception, computation, and hardware design to create better display and fabrication techniques.