Eye tracking and dynamic foveated rendering—technologies that solve the VR challenge

Delivering the best possible user experience with available system resources

Abstract

VR users demand amazing content, low-priced devices, and fantastic experiences. To meet these demands and deliver the best possible experience with available system resources, device manufacturers and application developers leverage a range of smart and innovative technologies in their solutions.

But adding a performance-enhancing technology to an existing device architecture is rarely a simple matter of plug-and-play. Processing overhead can outweigh the benefit. Yet new technologies can provide a considerable advantage, especially for those who are first-to-market.

Dynamic foveated rendering is a unique performance-enhancing technology that can significantly reduce rendering loads in VR headsets. In practice, this technology relies on eye tracking to deliver the necessary real-time data about a user's eye movements.

This ebook presents eye tracking and dynamic foveated rendering, how they free up resources, and why you might consider including these technologies in your solution designs.

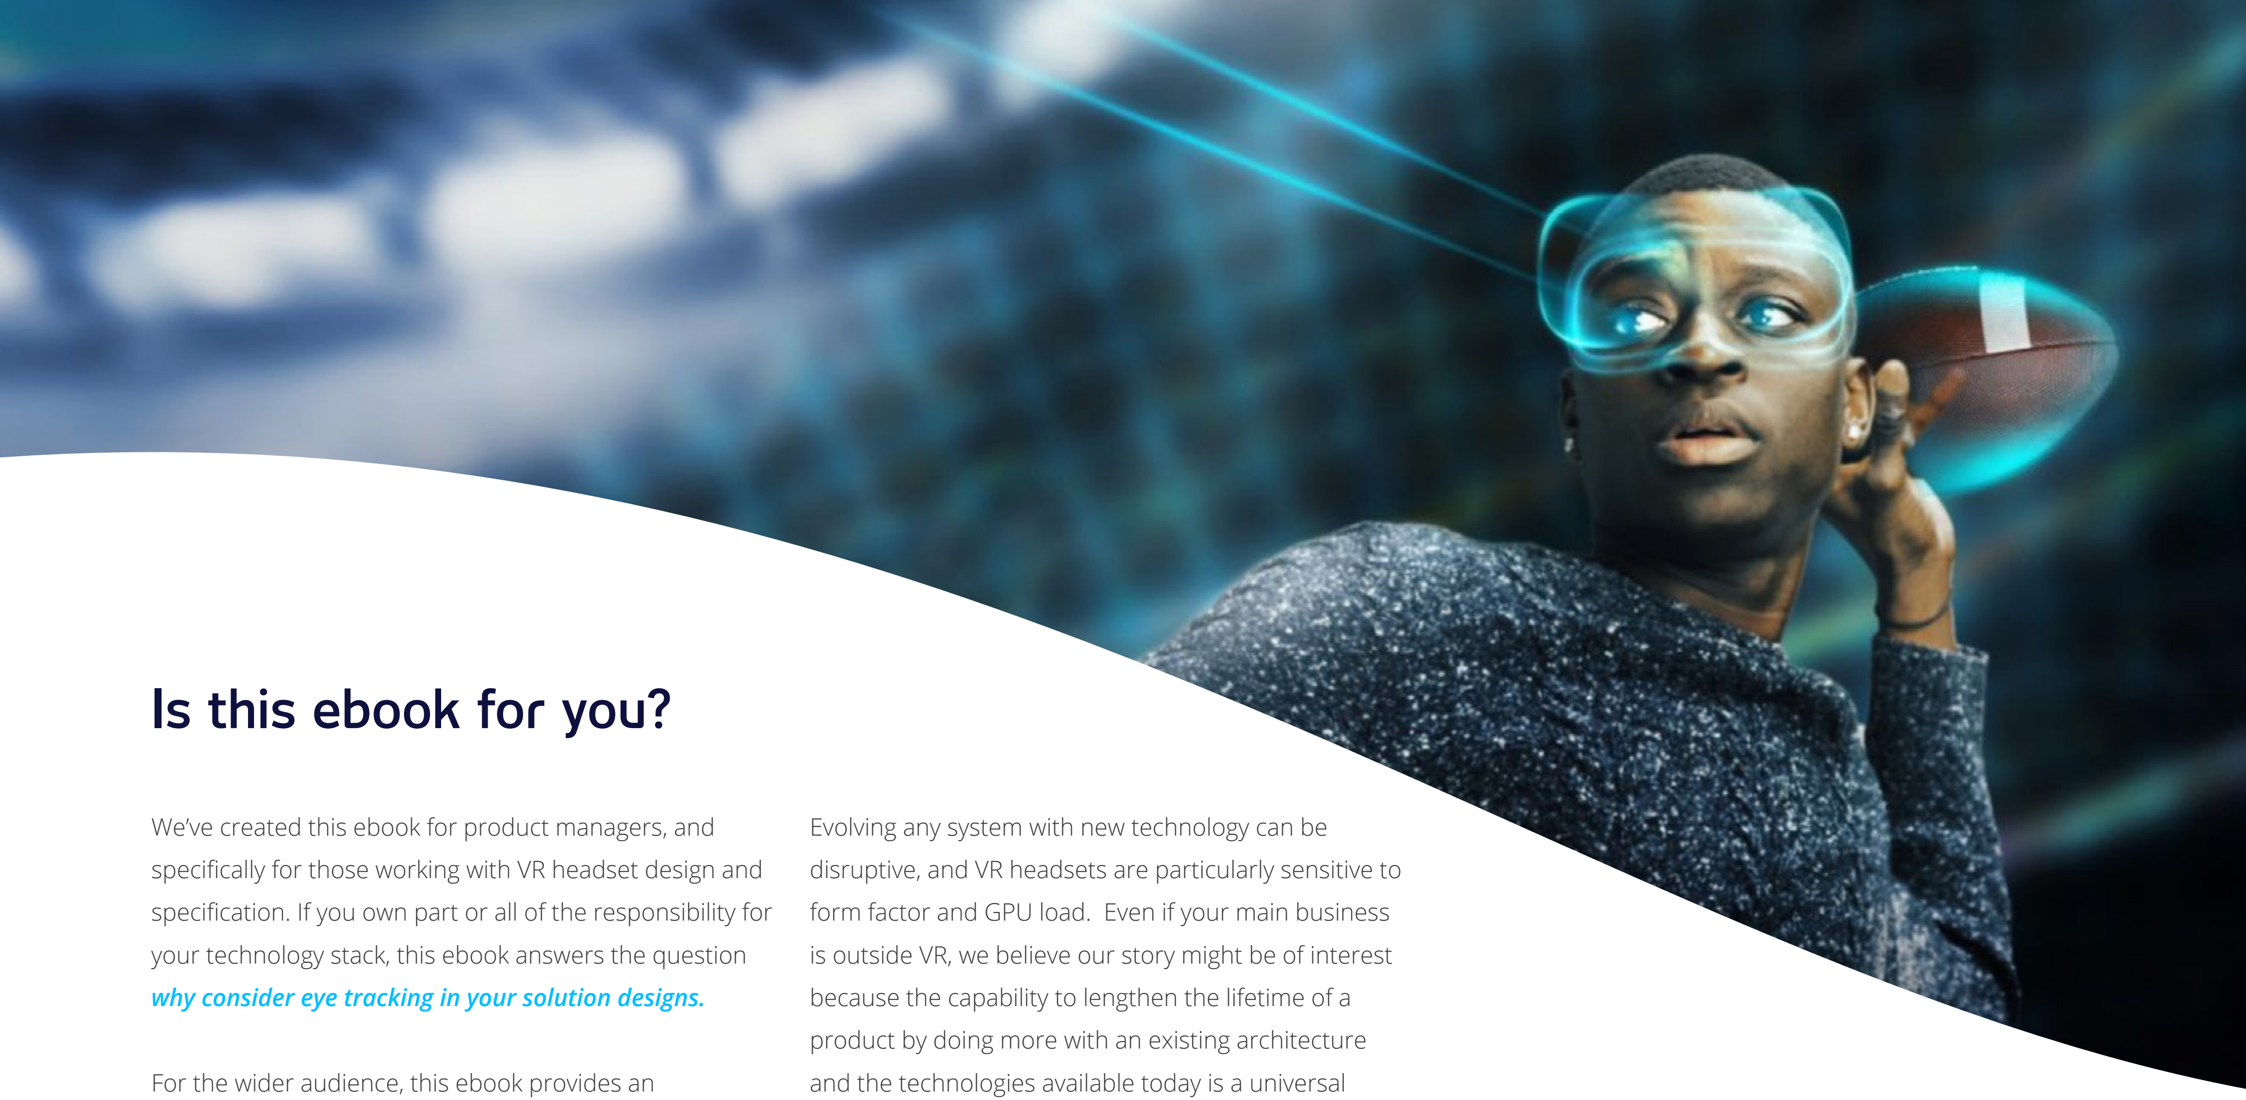

Is this ebook for you?

We've created this ebook for product managers, and specifically for those working with VR headset design and specification. If you own part or all of the responsibility for your technology stack, this ebook answers the question ***why consider eye tracking in your solution designs.***

For the wider audience, this ebook provides an overview of the VR ecosystem and some of its current challenges. It explains eye tracking and dynamic foveated rendering, and how these technologies optimize resources for better performance or for use in advanced processing techniques.

Evolving any system with new technology can be disruptive, and VR headsets are particularly sensitive to form factor and GPU load. Even if your main business is outside VR, we believe our story might be of interest because the capability to lengthen the lifetime of a product by doing more with an existing architecture and the technologies available today is a universal consideration in just about every system design.

The VR challenge

The aim of VR is ultimate immersion. To make a person forget where they are, to move in multiple other spaces where reality and simulation seamlessly combine. Achieving that requires VR applications with original features and good storytelling delivered on high-performance headsets that can ***balance load across subsystem components to achieve the best possible user experience.***

Significant development has taken place in VR over the past few years, not least in the emergence of enterprise-grade applications and standalone headsets. Specifically, development in display technologies is creating a demand for higher screen resolutions, pixel densities, and refresh rates with expanding field of view capabilities to improve the visual experience. Unfortunately, this comes at the cost of processing resources.

Frame rates of 90Hz are becoming a sort of gold standard for VR headsets. Streaming content at this level of quality helps to reduce motion sickness and headaches that can occur at lower frame rates. But frame rates are not just about what a headset can support. The ability to maintain a stable frame rate for a given application is fundamental to an immersive experience. Again, achieving this often comes at the expense of resources.

> The only thing that matters is where the user is looking

Solving the VR challenge

While haptic devices and controllers enhance VR, the experience is primarily a visual one. Consequently, one of the best ways to solve the VR challenge is to lower the impact of image processing.

Human vision focuses on a small area of interest at any one time, leaving the brain to fill in the surrounding details. Leveraging this fact makes it feasible to reduce the amount of image processing needed to deliver VR with the best experience possible.

Dynamic foveated rendering is a unique technology that leverages the physiology of human vision to free up resources and deliver stable frame rates.

With a dramatic drop in the amount of graphical data to process in the area of focus, GPU resources can be used to, for example, support higher screen resolutions or pixel densities.

But there is more to the story. Less graphical processing reduces the GPU computational time as well as the amount of video data transported from the GPU to the headset—supporting the requirements posed by enterprise-grade applications for low latency.

In practice, ***dynamic foveated rendering relies on eye tracking*** to deliver the necessary real-time data about a user's eye movements.

Eye tracking in a VR headset

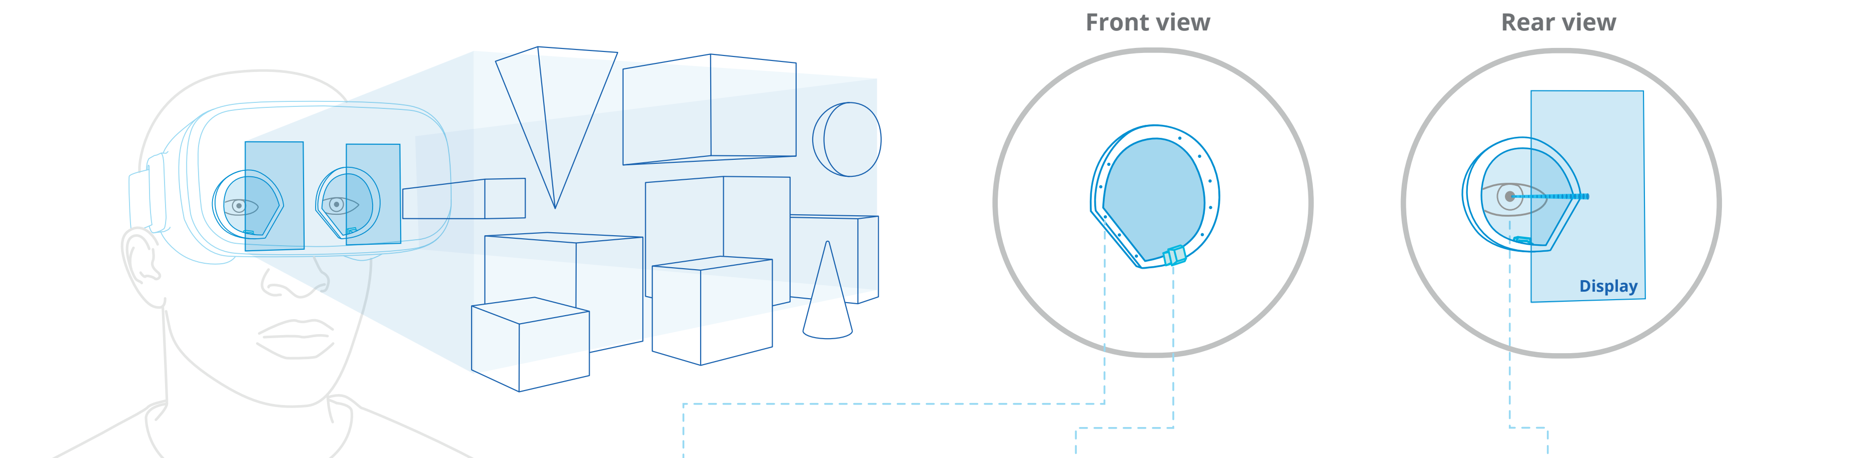

1 **An Eye Tracker** consists of cameras, illuminators, and algorithms.

2 **The illuminators** shine near-infrared light, creating reflections (or glints) on the eyes.

3 **The cameras** take high-resolution images of the eyes, including the glints.

4 **The image processing algorithms** identify the glints and pupils and calculate the gaze point.

Eye tracking components typically include a number of cameras, light sources (illuminators) and image processing algorithms. In a VR headset, these components, powerful enough to capture and process detailed eye movements, including millisecond variations in gaze direction, are typically integrated into a ring structure between the user and the display.

The illuminators shine near-infrared light (at frequencies that are practically invisible to humans) onto each eye, creating reflections—glints—on the user's corneas. The camera records each eye in real-time, capturing images of the pupils and glints.

Machine-learning algorithms interpret the captured images, using mathematical formulas to generate a stream of data points used by the high-level application. ***Eye tracking can be configured to deliver a range of parameters***, including the position of the pupil, gaze direction, pupil size, and eye openness.

The user's gaze coordinates tend to be the most useful data points for applications and performance enhancements. Dynamic foveated rendering relies on this information to free up headset resources.

What is foveated rendering

Foveated rendering is a resource optimization technique that leverages human vision to reduce the amount of GPU processing required to render a scene.

Our brains render what we see, by combining the point of our attention—illustrated by the apple—in high resolution, with the rest of what we see in medium and low resolutions.

Foveated rendering mimics the natural degradation of perception, fully rendering just the relatively small part of an image where the user's attention is focused.

The technique ***reduces the amount of graphical processing needed***, without a perceived reduction in quality in the periphery. If you think about it, fully rendering all parts of an image is a waste of resources.

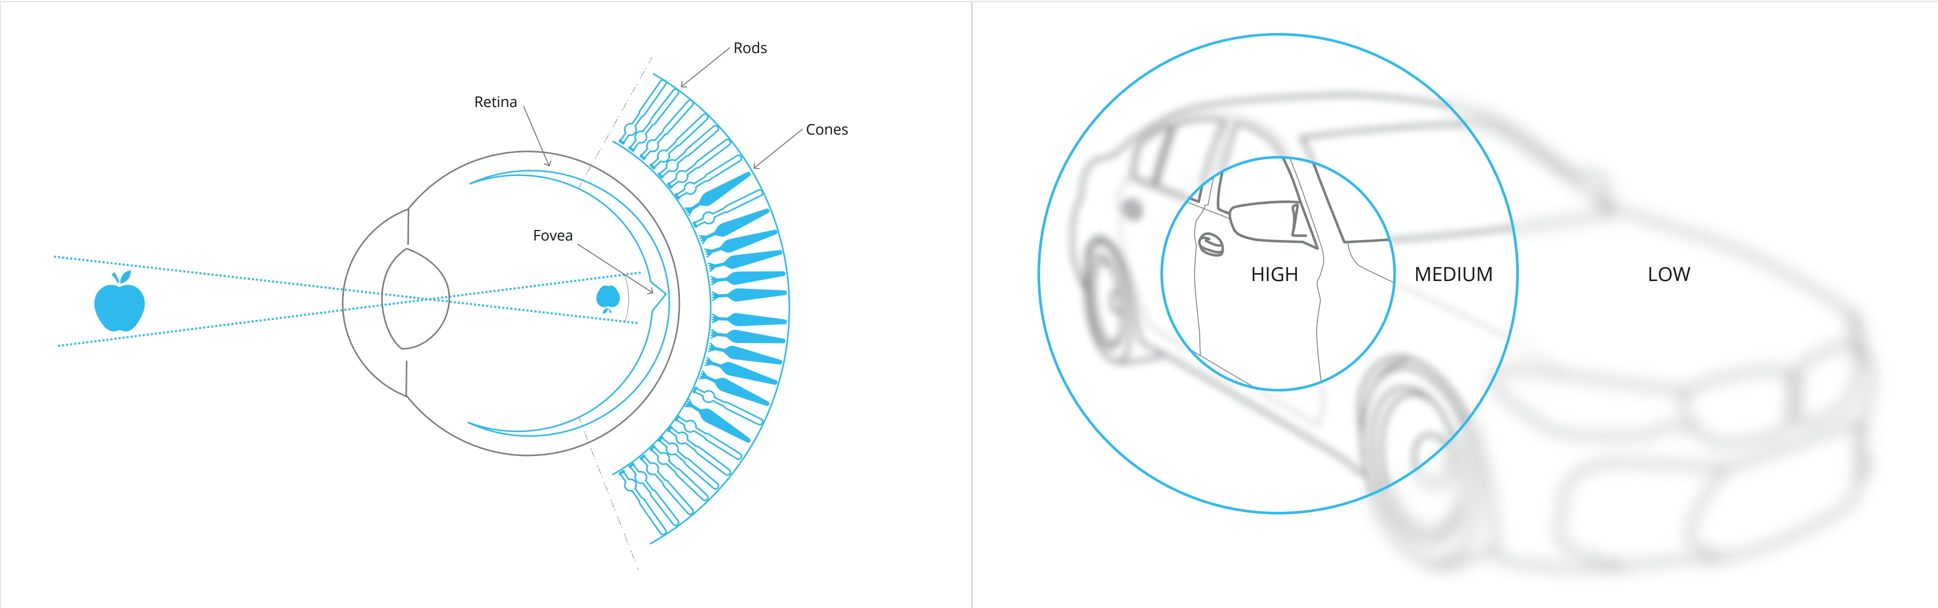

Foveated rendering mimics human vision—reducing the amount of GPU processing needed for rendering

Foveated rendering in a VR headset

Foveated rendering comes in two flavors: fixed (or static) and dynamic.

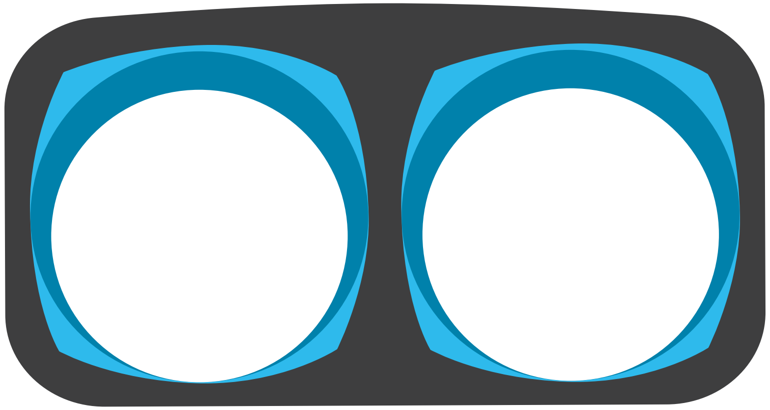

Fixed (or static) foveated rendering in a VR headset

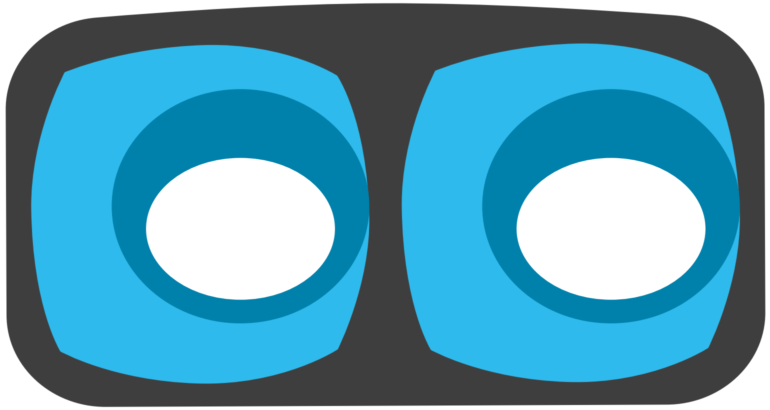

Dynamic foveated rendering in a VR headset

Fixed foveated rendering (FFR) reduces shading loads by portioning the screen into hard-coded zones. ***The assumed region of user attention***, indicated by the white area in the center of the screen, is fully rendered at 100%. The darker blue sections are rendered at medium resolution, and the small light blue areas with low resolution.

Naturally, applying this technique will result in some resource optimization, but it does not always yield an optimum user experience.

Dynamic foveated rendering (DFR) leverages the ***actual region of the user's attention*** to fully render just a small portion of the image (white circle), expanding outwards to medium (darker blue) and low resolution (light blue) with no degradation of quality or user experience.

Benchmark testing of dynamic foveated rendering shows a significant reduction in GPU shading loads (See Results section for details).

What to do with the extra resources

The GPU resources freed up through dynamic foveated rendering can serve other purposes. For example, to increase frame and refresh rates, reduce power consumption, or enable complex shading and forward-rendering—***delivering the best possible user experience with available system resources.***

Lower processing load makes it possible to reduce the clock speed of the GPU. By underclocking, the processor runs at lower temperatures, with a corresponding drop in power consumption, lessening the need for cooling, reducing ventilation-related noise, and prolonging battery life.

Limiting the region of full-resolution rendering reduces load on complex shaders, further reducing the time it takes to render a scene.

Freed-up resources can be used to deliver realistic shading and higher levels of scene complexity.

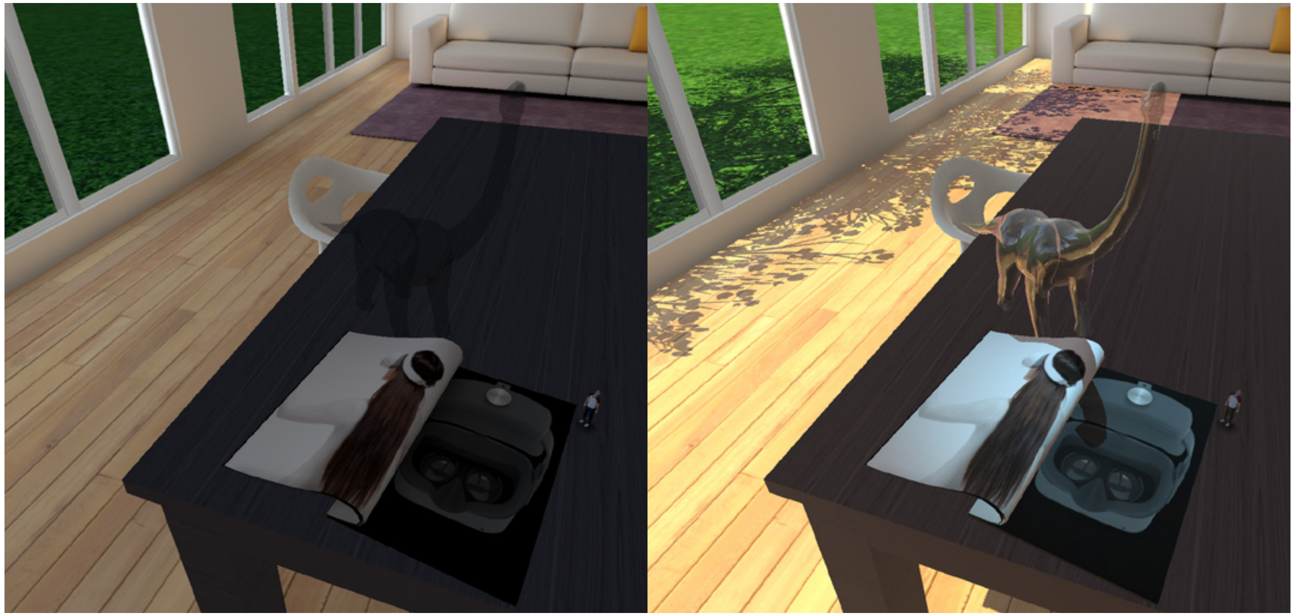

Without dynamic foveated rendering *DFR enabled*

Or even to improve forward rendering performance, which allows some freedom of choice in how scenes are rendered, supporting higher frame rates. For developers with access to features such as Qualcomm's Adreno foveation or NVIDIA's VRS, the cost of implementation and maintenance is low, and there are almost always guaranteed performance benefits.

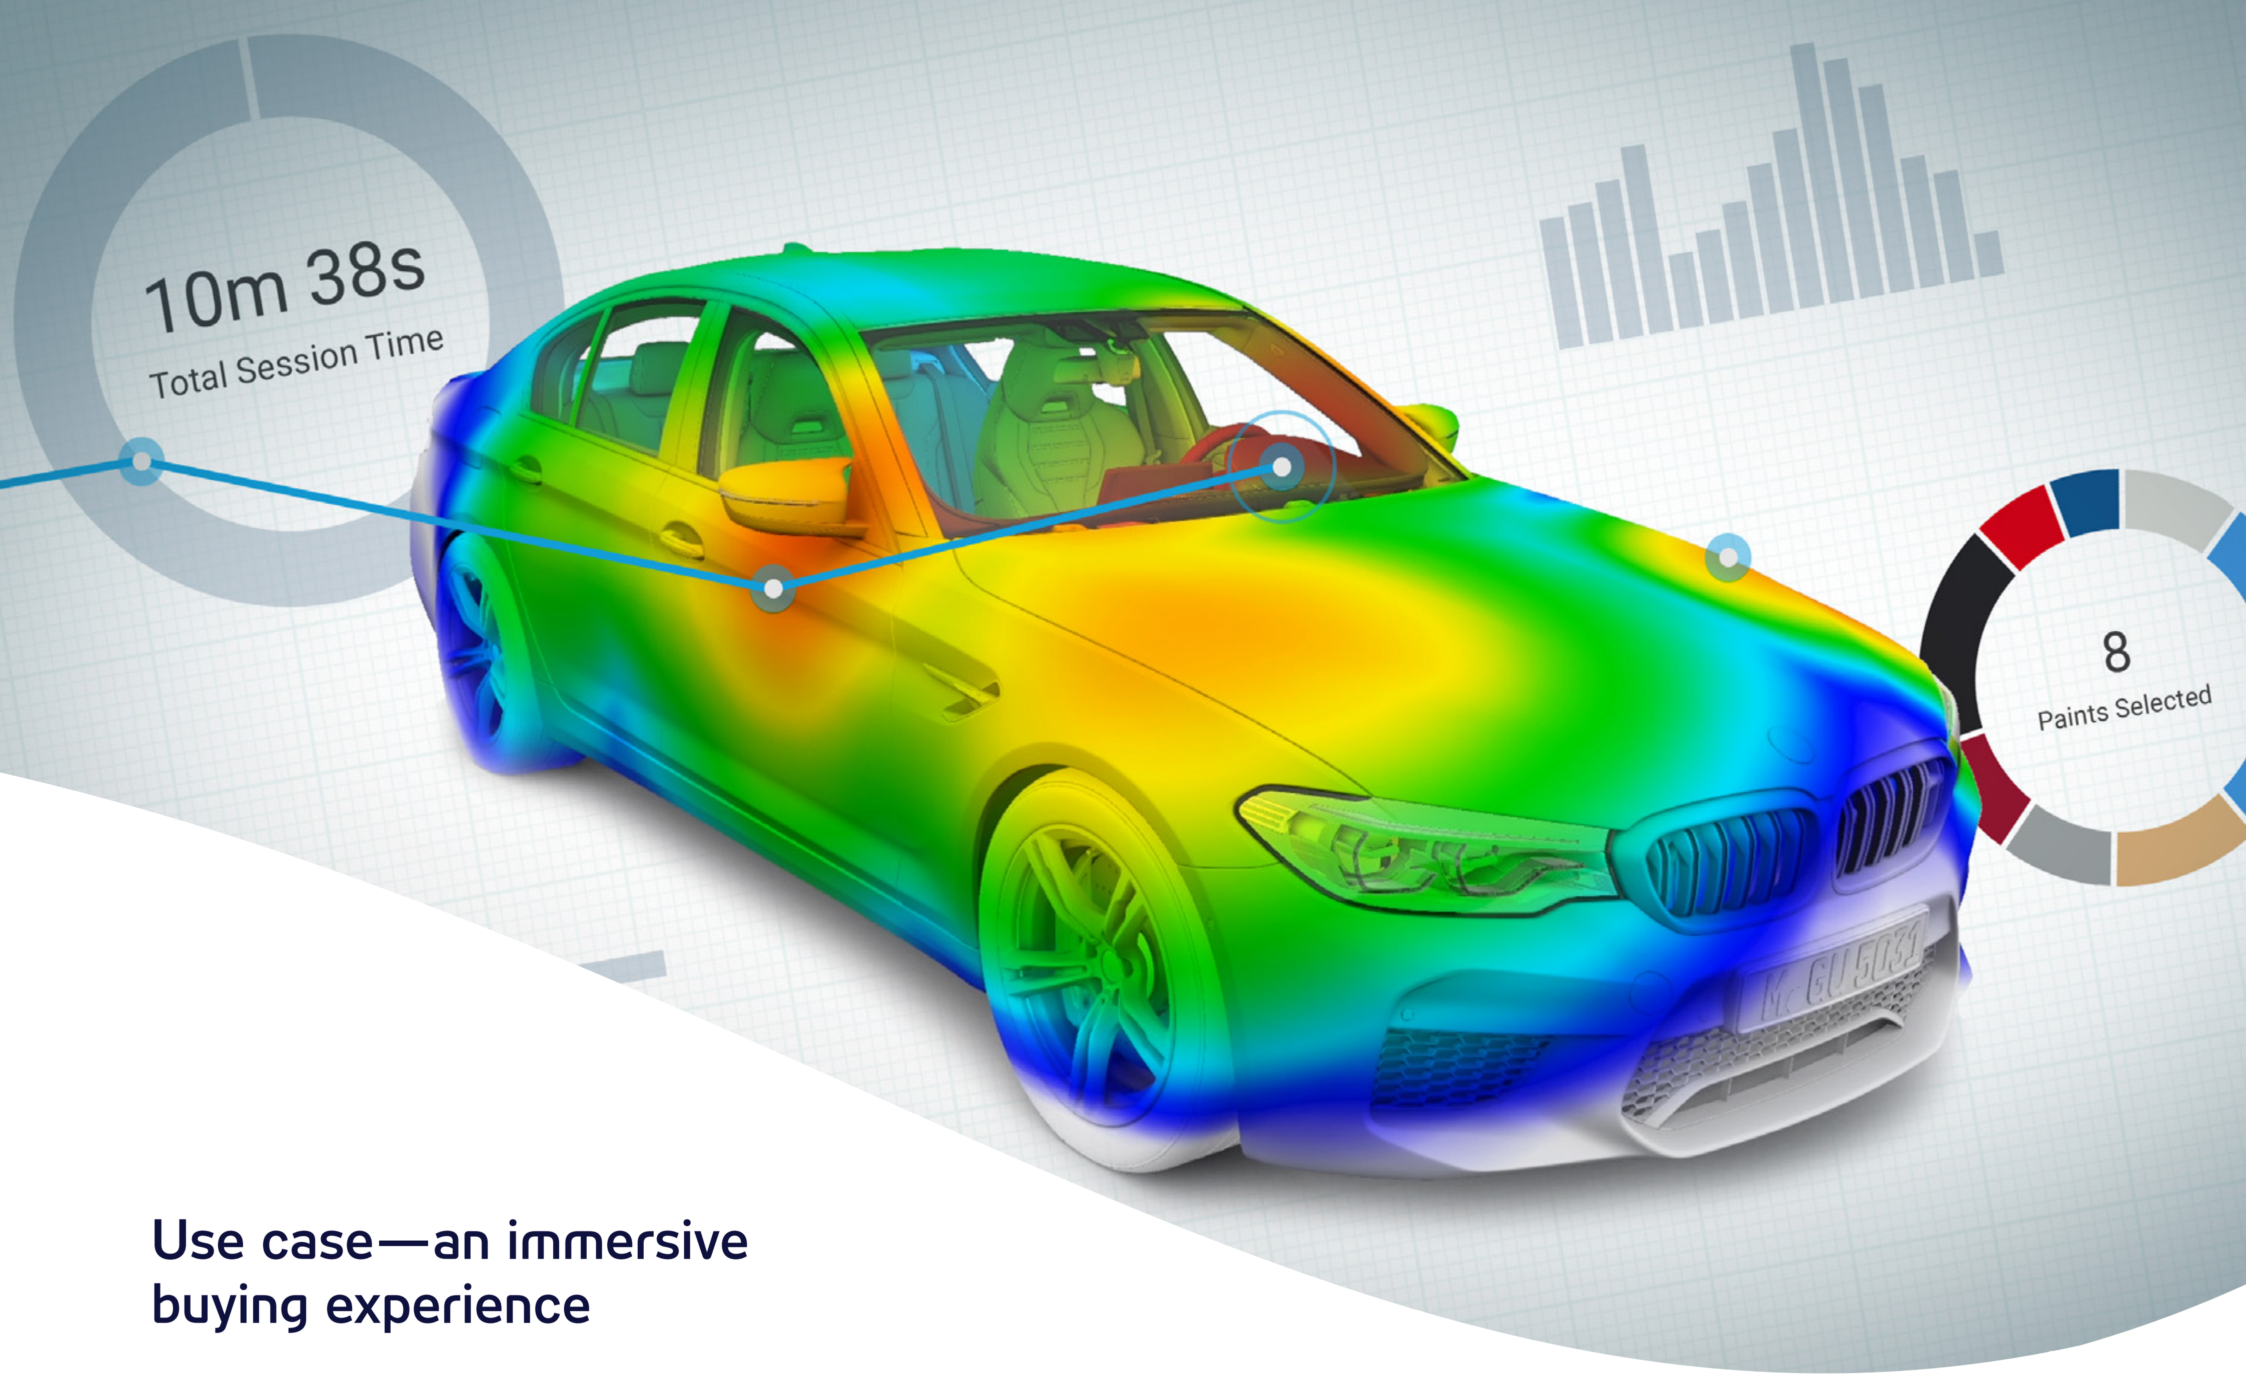

Use case—an immersive buying experience

Image courtesy of ZeroLight

Use case—an immersive buying experience

VR is an enhanced media platform that presents enterprises with a unique channel to woo consumers. The level of immersion VR offers allows potential buyers ***to experience a product in a deeply personal way***, unhindered by outside influence or expectations. To test VR on its customers, German automobile and engine manufacturer BMW turned to visualization specialists ZeroLight to create such an immersive VR experience for its high-performance M5 series.

Running on an HTC VIVE Pro Eye with eye tracking from Tobii and DFR enabled, the experience provides users with a means to explore and configure a simulated M5 sedan. Delivering the freedom to customize colors, finish, wheels, and interior details before going for a test drive on the Circuit de la Sarthe—renowned racetrack of the 24-hour Le Mans auto race.

With Tobii Spotlight Technology and NVIDIA's Variable Rate Shading (VRS) enabled, ZeroLight was able to increase the visual fidelity of the VR experience without needing additional graphical resources. What customers experience is sharp, realistic images—an individual interaction with BMW and the M5.

But there's more to eye tracking than the VR experiences it enables.

By capturing what the user is looking at in real-time, eye tracking can deliver insights into user behavior, analytics that BMW can leveraged in product development. For example, eye tracking can provide the data to generate heatmaps and graphs that show what aspects of the experience users spend the most time in and what grabs their attention—data-based business intelligence and unbiased user feedback.

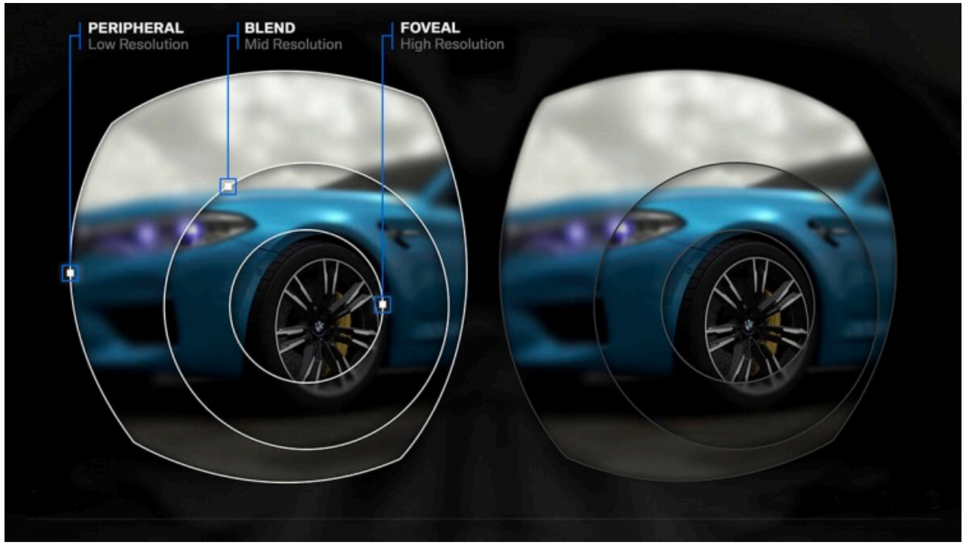

Image courtesy of ZeroLight

Benchmarking Tobii Spotlight Technology™

Tobii Spotlight is an advanced set of eye tracking technologies specialized for foveation in areas such as rendering, transport, and streaming.

To determine the performance gains delivered by our technology, we carried out a number of headset tests using the standalone Pico Neo 2 Eye and HTC's VIVE Pro Eye (tethered)—both of which are equipped with Tobii eye tracking and Tobii Spotlight technologies.

We tested the headsets under various conditions to gain a baseline with no foveation techniques applied and a best case scenario with dynamic foveated rendering delivered by Tobii Spotlight Technology.

These benchmarking tests were verified by global technology analyst Moor Insights and Strategy (MI&S).

For the HTC VIVE Pro Eye, we ran tests connected to an Alienware Aurora desktop configured with two different graphics cards. For each scenario, we ran a baseline performance test with NVIDIA VRS disabled. A second test to obtain performance for fixed foveation with NVIDIA VRS enabled. And a third test to obtain the performance for dynamic foveated rendering with Tobii Spotlight Technology and VRS enabled.

To establish repeatable and comparable test runs we:

- Tested with the Unreal Engine running the Showdown demo and Unity Engine running the Adam demo.
- Ran both demos on a preset path, to ensure that events, effects, and animations occurred at the same moment during each test run.
- Used the appropriate GPU profiler (Unreal or Unity) to capture frame rates and GPU shading loads.
- Placed the headsets on a stable surface with no user present.
- Repeated tests with the headset placed in different positions. The purpose being to simulate a range of head positions and assure that results are independent of how a user moves around.
- Simulated the user's gaze point.
- Used SteamVR supersampling to simulate higher resolution screens.
- Benchmarked with the Unity Adam demo to ensure results were engine independent.

For the Pico Neo 2 Eye, a similar test approach was adopted delivering a baseline test with no foveation techniques applied, a second test with Qualcomm's Adreno foveation, and a third test with both Tobii Spotlight Technology and Qualcomm's Adreno foveation enabled.

Benchmarking Tobii Spotlight Technology™

The results are the comparison between the baseline scenario with full shading load and the best case with maximum load reductions offered by dynamic foveated rendering.

Table 1 shows the full set of results of the benchmarking tests carried out on the HTC VIVE Pro Eye. The gain, measured in reduction in shading load, varies slightly depending on sampling rates, and the development engine used.

For example, running the Unreal Engine at 1x sampling rate, ***GPU shading load drops by 57% when dynamic foveated rendering is applied***. The average shading rate also drops, from 100% in the baseline scenario to 41.25% with FFR, and to 16%, with DFR enabled.

The complete analysis and verification are available in the Moor Insights and Strategy (MI&S) white paper:

Tobii Benchmark Testing Validates The Benefits of Dynamic Foveated Rendering.

Table 1: Results of benchmarking tests on HTC VIVE Pro Eye

Graphics card	Engine/demo	Sampling	Full shading load	DFR shading load	Shading load reduction
RTX 2070	Unreal/Showdown	1x	23.49%	10.06%	57.19%
	Unreal/Showdown	3x	59.68%	23.81%	60.10%
	Unity/Adam	1x	13.6%	5.46%	59.85%
RTX 2080 Ti	Unreal/Showdown	1x	5.34%	2.17%	59.37%
	Unity/Adam	1x	4.46%	1.79%	59.85%

Benchmarking Tobii Spotlight Technology™

The Pico Neo 2 Eye was tested with the Unity engine. Results reveal a max reduction in GPU shading load of 72.21% with a consistent frame rate of 75fps.

Shading loads with no rendering techniques and with DFR

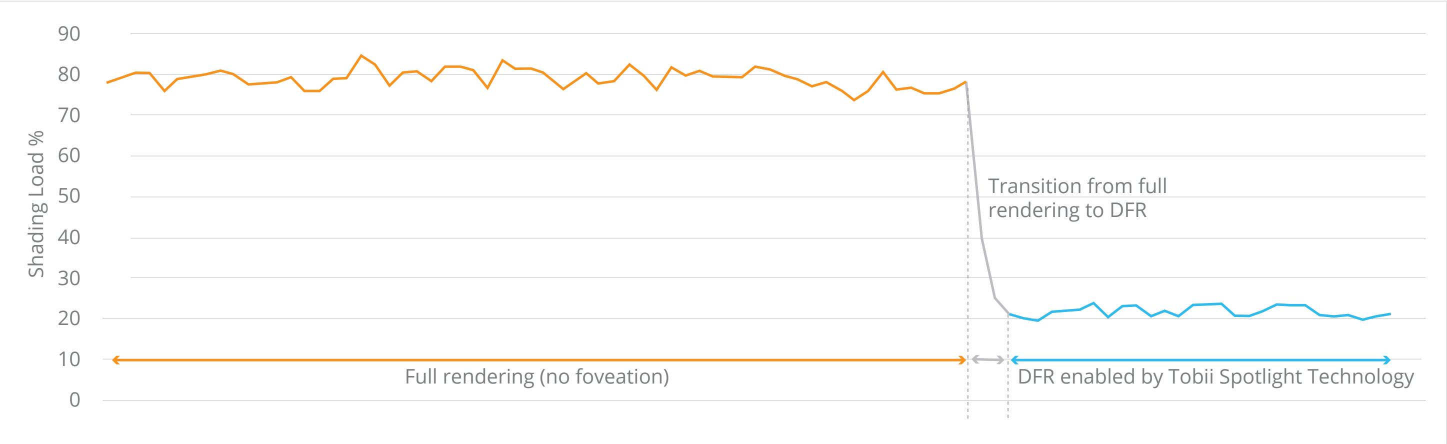

High and stable frame rates with DFR

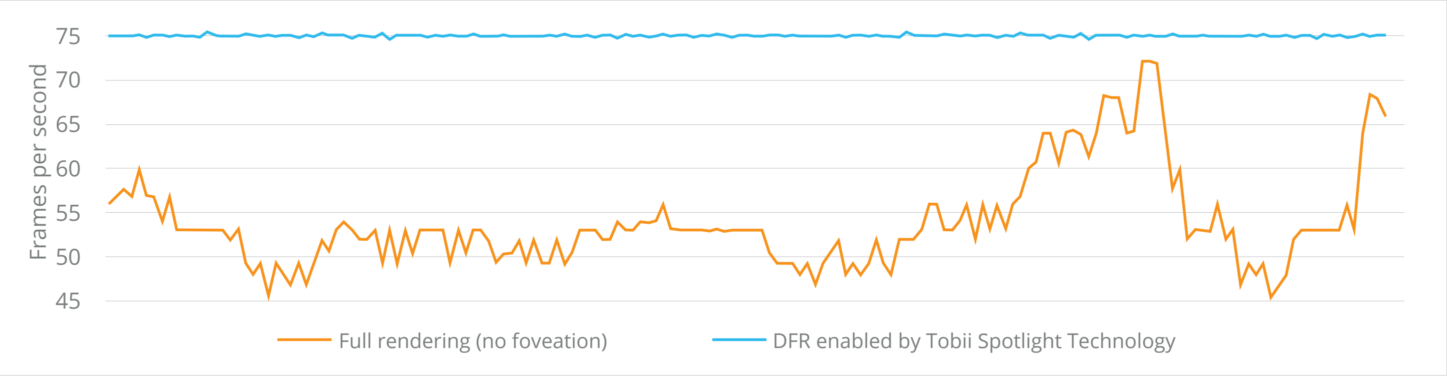

What else does eye tracking do for VR applications

Aside from being a prerequisite for dynamic foveated rendering, eye tracking delivers several additional benefits.

Owing to the technologies that have enabled avatars to evolve from gimmicky cartoons into realistic digital twins, VR applications are beginning to creep into the C-suite. With an eye tracking enabled headset, an avatar's eye movements will mimic the user's natural gaze, avoiding the uncanny stare and providing avatars with credibility.

Eye tracking raises the value of VR as it replaces traditional object selection — through the use of pointers or unnatural head movements—with intuitive eye movements.

Eye tracking calculates a person's interpupillary distance (IPD). Built-in IPD measurement allows users to adjust headset lenses to the optimum position for their eyes. Automatic adjustment could also be a headset feature.

Eye tracking can be used to help with biometric identification, ensuring that information created in a VR environment for your eyes only cannot be accessed by another user. Biometric identification supports instant profile loading, which is useful for headsets shared among a group of users.

Cost and safety are the main reasons why enterprises are looking to VR to train staff, especially people in high-risk professions such as surgeons, astronauts, and first-line responders. The addition of eye tracking to simulation-based VR training adds the capability of performance-assessment to the application. Training efficiency can be measured by comparing the eye movements of trainees with those of experts—ensuring that trainees are focused and observe specific events. With eye tracking, trainee performance assessments become data-driven and comparable.

For market research and user-centric design, eye tracking can raise the value of say a virtual store experience. By delivering precise data about where the user is looking, what catches their attention, and perhaps more significantly what people don't pay attention to, eye tracking provides insights that can be leveraged in product design, store layout, and premium product positioning.

Why you should consider adding eye tracking to your headset design

The VR challenge is to deliver VR applications with original features and good storytelling on high-performance headsets that can ***balance load across subsystem components to achieve the best possible user experience.***

Dynamic foveated rendering is a unique technology that leverages the physiology of human vision to limit full-resolution rendering to just a small portion of each scene. If you think about it, fully rendering all parts of an image is a waste, and with a potential drop in GPU shading load of nearly 60%, dynamic foveated rendering frees up a lot of resources. Resourced that can serve other purposes, such as higher frame and refresh rates, lower power consumption, or enable complex shading and improved forward-rendering—all of which contribute to an enriched user experience.

To enable such gains eye tracking is required in the headset to deliver the real-time data dynamic foveated rendering requires. But the benefits of eye tracking don't stop there. Eye tracking removes the uncanny stare, giving avatars natural eye movements and side glances. And by removing the need for unnatural head movements to, for example, select objects, eye tracking can improve the level of intuitiveness in VR applications.

Why choose Tobii

Tobii is the world leader in eye tracking with nearly 20 years of dedicated research and development in the field. Some of the top HMD manufacturers have adopted our technology because of its track record. Throughout the years, we have helped many customers to implement eye tracking in their product designs.

We work with manufacturers in many industries, including healthcare, construction, and gaming, helping them to implement eye tracking in their devices. By choosing Tobii, you gain from that experience, avoid some of the common integration pitfalls, such as how to mount components or software integrations.

Our solutions support the broadest range of users—delivering accurate data independent of the user's eye shape, ethnicity, cornea thickness, wearing of contact lenses or prescription glasses, wearing of makeup, or having dirt or smudges around the eyes.

We don't make headsets, so our success relies on our technology's ability to deliver innovative user interactivity, opportunities for application developers, and feedback analytics for the enterprise. Best illustrated by the BMW use case, Tobii eye tracking provides users with a simple input method to build their dream cars.

By enabling DFR, Tobii eye tracking frees resources to deliver high-fidelity visuals, ensuring the VR experience matches the luxuriousness of the BMW brand. And BMW receives unbiased data-based feedback from their potential buyers.

As a technology vendor, Tobii offers a full-scale solution. We deliver eye tracking as a partnership. We work closely with you to ensure that our components integrate with current designs, protecting investment, and lengthening the lifetime of products and applications by enabling more with existing architecture.

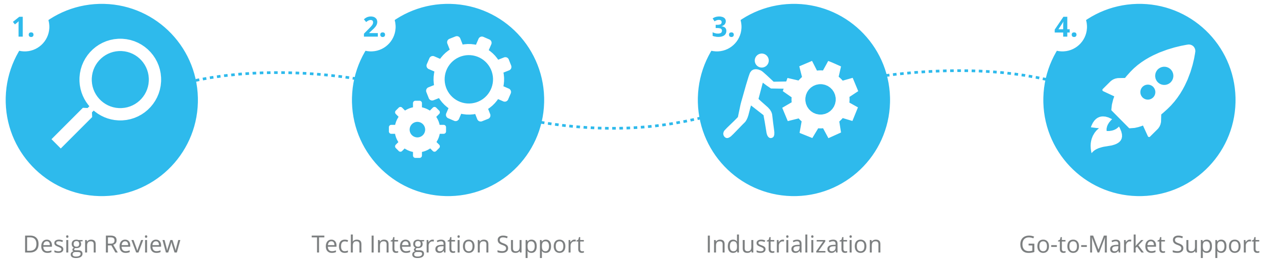

Tobii providing support at all stages

Tobii in brief

Tobii is the global leader in eye tracking. We envision a world where technology works in harmony with natural human behavior. The company employs over 1,000 people and operates through three business units: Tobii Dynavox, Tobii Pro, and Tobii Tech. Tobii Dynavox delivers specially-designed devices that use eye tracking and touch to help people with communication disabilities. Tobii Pro provides products and services that help businesses and science professionals gain insights into human behavior, helping them to make data-driven decisions. Tobii Tech develops eye tracking hardware and related software for the healthcare sector, PC manufacturers, virtual and augmented reality OEMs, as well as application developers in various fields, such as medical assessment and gaming.

Contact us

tobii

vr.tobii.com/foveated-rendering