Metameric Rendering

A technique to save bandwidth when transmitting images or video in head-mounted displays in areas not perceived by the human fovea

Category
Software

Learn more

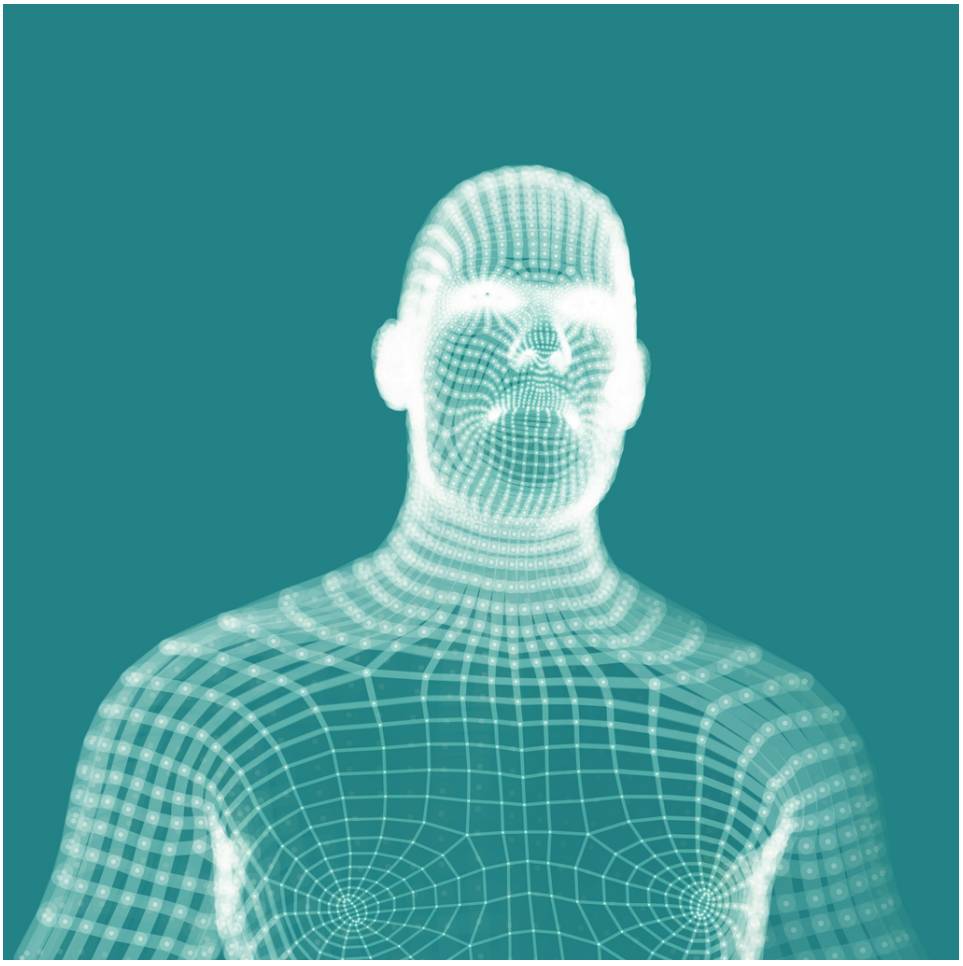

iStock image: StudioM1

Background

Graphics are usually rendered by a device incorporating a graphics processing unit (GPU) in an image processing pipeline, used in mobile devices, head-mounted displays (HMDs), personal computers, games consoles, Near-Eye Displays (NEDs), and more. In NEDs, a large proportion of the user's visual field must be rendered, which requires substantial computational power and bandwidth. Humans can only view fine spatial details in the fovea and not in the periphery -significant computational power and bandwidth can be saved by taking this into account in the rendering.

Today's solutions fall short in performance:

• Foveated rendering focuses computational effort to the foveal area and shows band-limited blurry images in the periphery. However, the blur can be unnatural and does not match well to what the user perceives.

• Ray tracing can cast more rays to the foveal area and update the view parameters during image generation. However, it is processing intensive and therefore slow in large/dynamic scenes.

• Neural networks can reconstruct images from samples of video frames that are denser in the fovea and sparser in the periphery. However, the loss is the same in the periphery and fovea, and does not consider different perceptual characteristics.

To optimise bandwidth of transmitted images, there is a need to more effectively and realistically render those images across a display whilst accounting for differences in a user's end perception of image in the fovea and periphery.

Technology Overview

UCL Computer Science has developed a novel and patented technique to reduce bandwidth when transmitting images or video in areas that are not perceived by the human fovea. The fovea perceives more detail than the periphery of the visual field. This technique can be used whenever an image is transmitted and it is known where the human observer is fixating their eyes in real-time. Whilst existing techniques for foveation simply blur the periphery, this technique uses image statistics to replace this blur with content that is perceived to be more realistic.

The technique is embodied in a set of algorithms and demonstration software to construct foveated images in real-time using the theory of metamers i.e:

1. Receive a first input image

2. Divide the image into multiple regions comprising of foveal region and peripheral region

3. Identify distribution of image statistics for peripheral region

4. For each peripheral region, identify a metamer with similar/identical distribution of statistics

5. Create an output image such that the peripheral region is perceived to be the same as the peripheral region in the input image.

The developed Unity implementation can use eye trackers to apply the method to any desired image content. The technique has been evaluated in perceptual experiments.

Benefits

The perception of peripheral areas is improved from standard foveated rendering techniques, from blurry to detailed and more realistic.

Applications

• Image compression in video streaming to head-mounted displays or other displays with eye tracking.

• Rendering new synthetic scenes based on feature statistics

Opportunity

Seeking partners for licensing opportunities or joint development work.

Seeking

Development partner, Commercial partner, Licensing

IP Status

Patented