

(51) International Patent Classification:
G06T 15/06 (2011.01)(21) International Application Number:
PCT/US2019/036220(22) International Filing Date:
10 June 2019 (10.06.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
16/017,747 25 June 2018 (25.06.2018) US

(71) Applicant: MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventor: CONNELL, Trebor Lee; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(74) Agent: MINHAS, Sandip S. et al.; MICROSOFT TECHNOLOGY LICENSING, LLC, One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,


AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

(54) Title: SIMULATING LENSES



A FIG. 6A

(57) Abstract: Examples are disclosed that relate to computing devices and methods for simulating light passing through one or more lenses. In one example, a method comprises obtaining a point spread function of the one or more lenses, obtaining a first input raster image comprising a plurality of pixels, and ray tracing the first input raster image using the point spread function to generate a first output image. Based on ray tracing the first input raster image, a look up table is generated by computing a contribution to a pixel in the first output image, wherein the contribution is from a pixel at each location of a subset of locations in the first input raster image. A second input raster image is obtained, and the look up table is used to generate a second output image from the second input raster image.

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

SIMULATING LENSES

BACKGROUND

[0001] Ray tracing can generate accurate images by calculating how light travels
5 through optical systems, such as lenses or mirrors. However, ray tracing can be
computationally expensive and may take a long time to simulate effects of such optical
systems on large numbers of images.

SUMMARY

[0002] This Summary is provided to introduce a selection of concepts in a
10 simplified form that are further described below in the Detailed Description. This
Summary is not intended to identify key features or essential features of the claimed
subject matter, nor is it intended to be used to limit the scope of the claimed subject
matter. Furthermore, the claimed subject matter is not limited to implementations that
solve any or all disadvantages noted in any part of this disclosure.

15 [0003] Examples are disclosed that relate to computing devices and methods for
simulating light passing through one or more lenses. In one example, a method comprises
at a precomputing stage, obtaining a point spread function of the one or more lenses and
obtaining a first input raster image comprising a plurality of pixels. The first input raster
image is ray traced using the point spread function to generate a first output image.

20 [0004] The method generates a look up table based on ray tracing the first input
raster image by computing a contribution to a pixel in the first output image, wherein the
contribution is from a pixel at each location of a subset of locations in the first input raster
image. At a runtime stage, a second input raster image comprising a plurality of pixels is
obtained, and using the look up table the method generates a second output image from the
25 second input raster image.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] FIG. 1 shows a block diagram illustrating an example system for simulating
light passing through one or more lenses according to examples of the present disclosure.

[0006] FIG. 2 is an illustrative example of a use case scenario in which a
30 computing device simulates images of a model environment according to examples of the
present disclosure.

[0007] FIG. 3 shows two examples of head-mounted display devices according to
examples of the present disclosure.

[0008] FIG. 4 is an illustrative example of a point spread function according to

examples of the present disclosure.

[0009] FIG. 5A illustrates a contribution from each pixel of a subset of pixels of an input raster image to a pixel in an output image according to examples of the present disclosure.

5 [0010] FIG. 5B is an illustrative example of combining contributions from two pixels of a subset of pixels according to examples of the present disclosure.

[0011] FIGS. 6A and 6B are a flow chart of an example method for simulating light passing through one or more lenses.

10 [0012] FIG. 7 shows a block diagram of a computing system according to examples of the present disclosure.

DETAILED DESCRIPTION

[0013] A wide variety of devices may utilize optical systems that process and/or generate light to produce images. Examples of such devices include cameras, microscopes, head-mounted displays, illumination systems, and the like. Simulating the
15 output of such optical systems using computer-generated images may be helpful in variety of situations.

[0014] For example, in some devices such as mixed reality head-mounted display (HMD) devices, cognitive algorithms may be utilized to perform image-based localization and/or other tasks. Such algorithms may be trained using sets of images representing an
20 environment in which the HMD device is located. In some examples, image-processing algorithms may be trained using actual images captured from a prototype of an image capture device or other optical hardware implemented along with the algorithms. However, it may be difficult to maintain sufficient production consistency of end user devices for the prototype's training to be valid for every device produced.

25 [0015] As such, it may be desirable to train cognitive algorithms on sets of simulated images that simulate a view of a scene from the image capture device. Generating images of real or simulated scenes also may be useful in a variety of other examples, including cinematography and gaming.

[0016] In some examples, effects of viewing an image via a lens, prism, mirror or
30 other optical hardware may be simulated using reference physical parameters of the optical hardware. For example, a simulated view through a lens may be generated by ray tracing how light will travel inside the lens based on given parameters. For example, the parameters may include a point spread function describing how light spreads when it passes through the lens.

[0017] Ray tracing using parameters such as a point spread function may generate accurate simulations of optical systems. However, ray tracing is computationally expensive and requires a long time to simulate the effects of an optical system on large numbers of images. For example, ray tracing a single image for some optical systems may take 10 minutes or longer. Further, in some examples an image-processing algorithm may be trained on sets of thousands to millions of images. As such, it may be inordinately time-consuming and expensive to ray trace an entire set of images. Further, ray tracing may process each image at a much slower rate than a frame rate at which images are obtained. As such, processed images may not be fed directly into applications requiring many frames to be processed at or near a frame capture rate.

[0018] Accordingly, examples are disclosed that relate to computing devices and methods for simulating light passing through one or more lenses. With reference now to FIG. 1, in one example a computing device 104, such as a network server, may comprise a processor 108 and a memory 112 holding instructions executable by the processor 108 to simulate light passing through one or more lenses as described herein. In different examples the computing device 104 may comprise a desktop, laptop or tablet computer, mobile computing device, mobile communication device (e.g., smart phone), and/or other computing device. Additional details regarding the components and computing aspects of the computing device 104 are described in more detail below with reference to FIG. 7.

[0019] As described in more detail below regarding FIG. 2, the computing device 104 may be communicatively coupled via network 116 with various user devices, such as a user computing device 120 or a head-mounted display (HMD) device 124. As described below, in some examples the computing device 104 may host a remote service that simulates light passing through one or more lenses as described herein.

[0020] For example, in FIG. 1, the user computing device 120 or the HMD device 124 may request that computing device 104 process a plurality of images to simulate the effect of viewing the images through one or more lenses or other optical system. In some examples, the computing device 104 may obtain a plurality of images from the user computing device 120. For example, the user computing device 120 may upload the images or provide a link to the images at a location on the network 116 or the Internet.

[0021] The computing device 104 may comprise a rasterizer 128, which may convert each image of the plurality of images into an input raster image 132. In other examples, the user computing device 120 may select a plurality of input raster images 132 already stored by the computing device 104.

[0022] In yet another example, discussed in more detail below, the user computing device 120 may request that computing device 104 generate a plurality of input raster images 132 from a model environment 136. The model environment 136 may comprise one or more predetermined three-dimensional worlds from which environment images may be generated.

[0023] In the example of FIG. 1, the computing device 104 may generate a plurality of output images 140 that may be used to train or validate performance of one or more cognitive algorithms 144. As illustrated in dashed blocks in FIG. 1, the cognitive algorithm(s) 144 may be implemented and/or stored on the user computing device 120, computing device 104, or other devices such as the HMD device 124.

[0024] In some examples, the HMD device 124 may use the cognitive algorithm(s) 144 to process image data from an image capture device 148. For example, the HMD device 124 may feed image data from the image capture device 148 to the cognitive algorithm(s) 144 to determine its location and orientation with respect to a real-world environment by analyzing the image data.

[0025] As described above, in some examples developers may desire a set of images to train or validate the cognitive algorithm(s) 144. In some examples, a developer may log into a service provided by the computing device 104 from the user computing device 120. In some examples, the developer may select a simulated model environment 136 from which the set of images may be generated. For example, as illustrated in FIG. 2, the simulated model environment 136 may comprise a simulated three-dimensional model of a room 200. The developer may then request images of the room 200 from one or more specified perspectives.

[0026] For example, the developer may define or select a path 204 from one or more pre-defined paths that simulate the positions and orientations of an HMD device 208 worn by a virtual user 212 walking around the room 200. The path 204 may describe the position and the orientation of the simulated HMD device 208 at a specified frame rate, such as 24 data frames per second, as the HMD device 208 progresses along the path 204.

[0027] With reference again to FIG. 1, the rasterizer 128 may generate one or more input raster images 132 that simulate a view of the room 200 from a perspective of the HMD device 208 at given positions and orientations along the path 204. In different examples, the rasterizer 128 may generate the input raster images 132 for a computer-generated film, video game, virtual or mixed reality application, etc.

[0028] For example, in FIG. 2, the room 200 may comprise a three-dimensional

mesh that models walls 216, floor 220, window 224, and other objects within the room 200 such as chair 228, couch 232, table 236, television 240 and shelf 244. In other examples, the room 200 may comprise a combination of three-dimensional shape and surface models, textures, depth values, light sources, and any other data that may be used to suitably generate images of a scene within the room 200.

[0029] Model environments 136, such as the room 200, also may comprise ground truth data 150 describing at least a portion of the model environment 136. For example, the room 200 may comprise depth data, orientation data, and/or identification data of objects and surfaces within the room 200. In one example, ground truth data 150 may comprise a label corresponding to a location within the room 200 that identifies the table 236. These ground-truth data 150 may be provided along with the output images 140 to describe the portion of the model environment 136 depicted in each output image 140. In one example, ground-truth data 150 may be provided in an image file's metadata description.

[0030] Ground-truth data 150 may be desirable for training the cognitive algorithm(s) 144 or validating an algorithm's performance by checking an output of the algorithm against the ground truth data. For example, an image-recognition cognitive algorithm may output "table" in response to processing an output image 140 of the room 200. A developer may check the ground truth data and determine that the cognitive algorithm correctly recognized the table 236.

[0031] In some examples involving an image capture device, the developer also may describe one or more optical systems of the image capture device so that images of the room 200 may be generated with proper distortions to accurately simulate images captured via the image capture device. For example, the developer may provide a point spread function 152 of one or more lenses, as described in more detail below regarding FIG. 4. In other examples, the computing device 104 may apply a pre-determined model of a lens, such as in cases when the developer may refrain from providing a point spread function.

[0032] Using the point spread function 152, the computing device 104 may utilize a ray tracer 156 to ray trace a first input raster image 132 and generate a first output image 140. Next, the computing device 104 may generate a look up table 172 based on ray tracing the first input raster image 132. For example, the first input raster image 132 may comprise a plurality of pixels 160. Each pixel of the plurality of pixels 160 may have a discrete location 164 within the image and a color value 168.

[0033] As described in more detail below, the computing device 104 may select a subset of the pixels 160 in the input raster image 132 to populate the look up table 172. For each pixel 160 in this subset of pixels, a contribution to a color value 180 or other value of each pixel in the first output image 140 may be computed. In this manner and for each pixel 160 in the subset of pixels in the input raster image 132, the look up table 172 may map the pixel's computed contribution to each pixel 176 in the first output image 140. Using these computed contributions in the look up table 172, the color value 180 for each pixel 176 in the output image 140 may be determined. In some examples, generating the look up table 172 as described above may be performed at a precomputing stage.

[0034] As described in more detail below, the computing device 104 may receive additional input raster images for processing, such as a second input raster image, third input raster image, etc. Accordingly, and in one potential advantage of the present disclosure, a second output image, third output image and/or additional output images may be generated, respectively, using the look up table 172. In this manner, the additional time requirements and computational expense of ray-tracing each additional image may be avoided.

[0035] With reference now to FIG. 3, and as described above, in some examples aspects of the present disclosure may be practiced by an HMD device that may take a variety of forms. For example, HMD device 124 may take the form of a virtual reality HMD device 304A that includes an opaque, non-see-through display 308. Another example of an HMD device is an augmented reality HMD device 304B that comprises an at least partially transparent display 312 that is configured to enable a wearer of the augmented reality HMD device to view physical, real-world objects in the physical environment through one or more partially transparent pixels displaying virtual object representations. It will be appreciated that the following descriptions of sensors and systems may apply to both the augmented reality HMD device 304B and the virtual reality HMD device 304A.

[0036] In the example of FIG. 3, each of the example HMD devices 304A and 304B has a construction that includes a frame 316 that wraps around the head of a user to position a display close to the user's eyes. The frame of virtual reality HMD device 304A may include a rigid portion and an elastic portion, whereas the frame 316 of augmented reality HMD device 304B may be substantially rigid around its circumference. The frame 316 may support additional components such as, for example, a processor 320 and an input device 324. The processor 320 may include logic and associated computer memory

328 configured to run a cognitive algorithm 144, to simulate light passing through one or more lenses, to interact with computing device 104 and/or user computing device 120 or other computing devices, to receive and process sensor data from input device 324, and/or to enact various other processes described herein. It will be appreciated that the processor 320 may include a logic processor and the two example HMD devices 304A and 304B may include volatile memory and non-volatile storage, as discussed in more detail below with respect to the example computing system 700 of FIG. 7.

[0037] With continued reference to FIG. 3, the input device 324 may include various sensors and related systems to provide information to the processor 320. Such sensors may include one or more image sensor(s). The one or more outward facing image sensor(s) may be configured to capture and/or measure physical environment attributes of the physical environment in which the augmented reality HMD device 304B is located. In one example, the one or more image sensor(s) may include a visible-light camera configured to collect a visible-light image of a physical space. The visible-light camera may comprise one or more optical systems, such as lenses, prisms or mirrors, configured to collect the visible light image.

[0038] In one example of the augmented reality HMD device 304B that includes the transparent display 312, a position and/or orientation of the augmented reality HMD device 304B relative to the physical environment may be assessed so that augmented-reality images may be accurately displayed in desired real-world locations with desired orientations. In both augmented reality HMD device 304B and virtual reality HMD device 304A, the image sensor(s) may be configured to provide image data to the cognitive algorithm 144, which may process the image data to generate a three-dimensional model of the physical environment and determine the position and/or orientation of the augmented reality HMD device 304B or virtual reality HMD device 304A.

[0039] With reference now to FIG. 4, and as mentioned above, it may be desirable to simulate distortions introduced into one or more images by an optical system, such as lenses of the image sensor(s) of HMD devices. As described above, distortions and aberrations introduced by such optical systems may be modeled using a point spread function 152 that describes how light spreads when it passes through a lens.

[0040] FIG. 4 illustrates one example of a point spread function 400 of a lens. In the example of FIG. 4, the point spread function 400 is modeled as a Gaussian function with a standard deviation of 0.5 that illustrates a relative intensity 404 of light focused at a pixel as it spreads out, or blurs, after passing through the lens. The x-axis of the point

spread function 400 is the relative position 408 to the pixel along an axis of an output image 140.

[0041] It will be appreciated that FIG. 4 illustrates the point spread function 400 as a two-dimensional function for simplicity, but that the point spread function 400 may also be applied to each axis of an image. For example, the point spread function 400 may be applied symmetrically to both an x- and y-axis of an image, or different point spread functions may be applied to each axis. In other examples, the point spread function 400 may be modeled by more complex functions, such as a three-dimensional Airy function, or may be modified to more accurately model optical aberrations of the lens. In yet other examples, different pixels of an output image 140 may have different point spread functions.

[0042] In some examples, a numerical input/output table may be used in place of, or in combination with, the point spread function 152 to describe how light spreads when passing through a lens. For example, an input/output table may be generated by shining a laser through one or more lenses and using an image sensor to detect the light after it passes through the lens. The image sensor may comprise a plurality of pixels, and may report a location of each pixel that absorbs the light, and the intensity of the light at said pixel(s). Like a point spread function, this information also may be used to simulate distortions introduced by the lens.

[0043] With reference again to FIG. 1 and as described above, using the point spread function 152 a first input raster image 132 may be ray traced to generate a first output image 140. In some examples, the ray tracer 156 may process the first input raster image 132 as a high-resolution pinhole image. For example, the pinhole image may have a resolution of 4000 X 3000 pixels, 12000 X 8000 pixels, or any other suitable resolution.

[0044] The ray-tracer 156 may project each pixel of the first input raster image 132 on a first side of a simulated lens onto a first output image 140 on an opposing side of the simulated lens by calling the point spread function 152 for every point. In another example, an input/output table may be called instead of the point spread function 152 to generate the first output image 140. It will also be appreciated that interpolation or any other suitable method may be used to project each pixel of the first input raster image 132 onto the output image 140. For example, mathematics used to generate the first output image 140 may be simplified, such as by using an analytical model for distortion in place of ray-tracing refraction of the simulated lens. In these examples, the first output image may be generated using a simplified line to plane calculation.

[0045] Next, the look up table 172 may be populated by mapping each pixel of the first input raster image 132 that contributes to each pixel of the first output image 140, and by computing how much each pixel of the first input raster image 132 contributes to each pixel of the first output image 140. For example, a contribution from each pixel of the first input raster image 132 to each pixel of the first output image 140 may comprise a weighted contribution value, such as 0.3 or 30%. It will also be appreciated that, in other examples, the look up table 172 may be populated or initialized using one or more images other than the first input raster image.

[0046] As previously mentioned, a look up table may specify one or more pixels at a subset of one or more locations in a first input raster image, with each such pixel of the subset contributing to each pixel of a first output image. For example, each output image of an image set may comprise 640 X 480 pixels. The subset of locations in the first input raster image may comprise locations of 512 different pixels in the first input raster image.

[0047] In one example, the look up table may comprise 512 entries for each pixel of the 640 X 480 pixels of the first output image. Each entry in the look up table may comprise a location of a pixel in the first input raster image and a contribution value from that pixel to a pixel of the first output image. In other examples, the look up table may utilize other quantities of entries for each pixel of an output image, such as 256, 1024, or other quantity.

[0048] If parameters such as the resolution and dimensions of each image and camera intrinsic parameters such as focal length and distortion remain unchanged, a pixel at a given location in an input raster image may contribute a same amount to a given pixel in an output image each time the input raster image is ray traced, regardless of the input raster image's content. For example, a pixel in an upper left corner of an input raster image may consistently project to a lower right corner in a corresponding output image each time the input raster image is ray traced. However, as the content of the input raster images changes, a color of a pixel at a given location in a first input raster image may not be the same as the color of the pixel at the same given location in a second input raster image.

[0049] Accordingly, in one potential advantage of the present disclosure, a color value of a given pixel in a second output image 140 may be computed using the contribution values that have been precomputed for the subset of locations in the first input raster image 132. For example, a color of a given pixel in a second input raster image may be determined at each location of the subset of locations. As described above,

the look up table 172 may specify the contribution from the pixel in the first input raster image at each location of the subset of locations. This contribution may be combined with the color of the pixel in the second input raster image. Accordingly and as described in more detail below, in some examples and at a runtime stage, a second input raster image comprising a plurality of pixels may be obtained. Using the look up table 172, a second output image may be generated from the second input raster image.

[0050] For example, FIG. 5A illustrates a subset 504 of pixels enlarged from a first input raster image 508 of the table 236 of FIG. 2. In the example of FIG. 5A, the subset 504 comprises pixels 512A, 512B, 512C and 512D. Based on ray tracing the first input raster image 508, it may be determined that pixels 512A, 512B, 512C and 512D each contribute to a color of a pixel 516 in a first output image 518.

[0051] It will be appreciated that pixels 512A, 512B, 512C, and 512D each correspond to a different location in the first input raster image 508. Pixel 516 corresponds to a location in the first output image 518. As such, pixels 512A, 512B, 512C, 512D and 516 as they are discussed herein may be present at the same location in each respective input or output image. Additionally, pixels 512A, 512B, 512C, 512D and 516 as they are discussed herein may have a different color among a first image, second image, third image, and/or additional images as each image's content changes.

[0052] In the example of FIG. 5A, pixels 512A, 512B, 512C and 512D may have different colors in the first input raster image 508 than at corresponding locations in a second input raster image. Accordingly, pixel 516 may have a different color in the first output image 518 as compared to a second output image generated from a second, different input raster image.

[0053] For example, the color of the pixel 516 in a second output image may be determined by mixing the colors of pixels 512A, 512B, 512C and 512D in the second input raster image according to a weighted contribution value from each pixel. The weighted contribution value may be computed by ray tracing the first input raster image 508 to generate the first output image 518, as discussed above, and saved in look up table 172.

[0054] In one example, the colors of pixels 512A, 512B, 512C and 512D in the first input raster image 508 may contribute 30%, 20%, 25%, and 25%, respectively, to the color of the pixel 516 in the first output image 518. Accordingly, the look up table 172 may specify that pixels 512A, 512B, 512C and 512D have weighted contributions of 0.30, 0.20, 0.25, and 0.25, respectively, to pixel 516. It will be appreciated that these weighted

contributions are provided by way of example. In other examples, weighted contribution values may not add to 100%, as relative illumination and quantum efficiency may prevent 100% of incoming light from being recorded in a final image.

[0055] In this example, a second input raster image may be obtained in which pixels 512A, 512B, 512C and 512D may have different color values than in the first input raster image 508. However, the pre-computed weighted contribution values from each pixel 512A, 512B, 512C and 512D of the subset 504 of pixels of the first input raster image 508 to pixel 516 in the first output image 518 are saved in look up table 172. Accordingly, these pre-computed contributions may be used to compute the color of pixel 516 in the second output image. For example, the color of the pixel 516 may be computed by mixing the colors of pixels 512A, 512B, 512C and 512D in the second input raster image in a ratio of 30%, 20%, 25%, and 25%, respectively, which is the same ratio in which these pixels contribute to pixel 516 in the first output image 518.

[0056] It will be appreciated that the subset 504 of pixels may comprise a different number of pixels in other examples, such as up to every pixel in the first input raster image 508. However, not every pixel of an input raster image may have a meaningful contribution to each pixel of an output image. Accordingly, a look up table 172 may be simplified by selectively saving contributions from a subset of pixels of an input raster image that contribute most to each pixel of an output image. In one example, the subset may comprise the 512 pixels of a first input raster image that have the highest-weighted contributions to each pixel in a first output image as compared to the other pixels of the first input raster image. In other examples, the subset may comprise a different number of pixels, such as 256, 1024, etc. In some examples, one of a plurality of subsets that have different number of pixels may be utilized for one or more pixels in the first output image.

[0057] Further, a graphics processing unit (GPU) may be configured to use the look up table 172 to process images faster than one or more central processing unit (CPU) cores. For example, a GPU may implement one or more texture filters that may accelerate image processing. In the example of FIG. 1, processor 108 of computing device 104 may comprise a GPU. In some examples, the GPU may comprise a bilinear filter that may be utilized to determine a color value of a given pixel in an output image. For example, the bilinear filter may be configured to look up and combine two color values of two pixels in an input raster image with two contributions to the given pixel from the two pixels in the input raster image. It will be appreciated that bilinear filtering and other methods discussed herein also may be implemented on a CPU or other suitable processor.

[0058] FIG. 5B shows an example of how a bilinear filter may be applied to combine contributions from two pixels. In the example of FIG. 5B, the bilinear filter may interpolate a color value of a virtual pixel 520 using the weighted contributions of pixels 512A and 512B. As described above, pixels 512A and 512B may have weighted
5 contributions of 0.3 and 0.2, respectively. Accordingly, the color value of the virtual pixel 520 may be determined by mixing the colors of pixels 512A and 512B in a ratio of 60% to 40%, respectively.

[0059] The bilinear filter also may interpolate a color value of virtual pixel 524 using the weighted contributions of pixels 512C and 512D. As described above, pixels
10 512C and 512D may have weighted contributions of 0.25 and 0.25, respectively. Accordingly, the color value of the virtual pixel 524 may be determined by mixing the colors of pixels 512C and 512D in a 1:1 ratio.

[0060] In this example, the virtual pixels 520 and 524 may be combined in a 1:1 ratio to yield the same color for pixel 516 as in the example of FIG. 5A. In other examples,
15 the virtual pixels 520 and 524 may be mixed with different weights and/or in different ratios. In some examples, using the bilinear filter to combine contributions from two pixels may approximately halve the computational expense of computing the color of pixel 516.

[0061] In some examples, output images may be generated in real-time or close to real-time. In one example, a look up table 172 may be populated as described above, such
20 as by ray tracing a first input raster image. Once the look up table 172 is populated, additional input raster images may be obtained at a frame rate. For example, the additional input raster images may stream directly from an image capture device, a database or an image generator. Advantageously, by utilizing the processes described herein, additional output images corresponding to each of the additional input raster images may be
25 generated at approximately the same frame rate by simply calling the look up table 172 as described above. In this manner, output images may be generated in real-time, which may further enable production of image sets at scale and may enable interactivity between consequent frames of an image set.

[0062] In some examples, bilinear filters may bias an output image by including
30 contributions from pixels outside the subset of highest-weighted pixels. Even a small bias that may be imperceptible to a human observer may negatively impact a cognitive algorithm. Accordingly, in some examples, bias may be reduced by instructing the GPU to sample pixels in a pattern, such as a square or a circle.

[0063] In some examples, the bilinear filter may combine the contribution of a

pixel with that of its highest weighted neighbor. In the example of FIG. 5B, pixel 512A has a weight of 0.3. Adjacent pixels 512B and 512C have weights of 0.2 and 0.25, respectively. Accordingly, the bilinear filter may combine pixel 512A with its highest weighted neighbor pixel 512C. In yet other examples, pixels may be combined about a default axis, and the default axis may be switched randomly or at a preset rate (e.g. to combine pixels 512A/512C and pixels 512B/512D instead of pixels 512A/512B and pixels 512C/512D). It will be appreciated that pixels may be processed in any other suitable manner to reduce bias.

[0064] With reference now to FIGS. 6A-6B, a flow chart is illustrated of an example method 600 for simulating light passing through one or more lenses. The following description of method 600 is provided with reference to the software and hardware components described herein and shown in FIGS. 1-5 and 7. It will be appreciated that method 600 also may be performed in other contexts using other suitable hardware and software components.

[0065] With reference to FIG. 6A, at 604, the method 600 may include, at a precomputing stage, obtaining a point spread function of the one or more lenses. At 608, the method 600 may include obtaining a first input raster image comprising a plurality of pixels. At 612, the method 600 may include ray tracing the first input raster image using the point spread function to generate a first output image.

[0066] At 616, the method 600 may include, based on ray tracing the first input raster image, generating a look up table by computing a contribution to a pixel in the first output image, wherein the contribution is from a pixel at each location of a subset of locations in the first input raster image. At 620, the method 600 may include, wherein the subset of locations comprises locations of at least 512 different pixels.

[0067] At 624, the method 600 may include, at a runtime stage, obtaining a second input raster image comprising a plurality of pixels. At 628, the method 600 may include using the look up table to generate a second output image from the second input raster image. At 632, the method 600 may include using at least the second output image to train a cognitive algorithm. At 636, the method 600 may include, wherein the first output image and the second output image comprise ground truth data describing at least a portion of a real world environment depicted in the first output image and the second output image. At 640, the method 600 may include, wherein the ground truth data comprise one or more of depth data, orientation data, and identification data.

[0068] At 644, the method 600 may include determining a color value of a given

pixel in the second output image by: (a) determining a color value of a pixel in the second input raster image at one of the locations of the subset of locations; (b) obtaining, from the look up table, the contribution to the given pixel in the second output image from the pixel at the one location of the subset of locations; (c) combining the color value of the pixel with the contribution to the given pixel; and repeating the processes (a), (b), and (c) above for each location of the subset of locations in the second input raster image. At 648, the method 600 may include, wherein determining the color of the pixel comprises using a bilinear filter to combine two color values of two pixels in the second input raster image with two contributions to the given pixel from the two pixels in the second input raster image.

[0069] With reference to FIG. 6B, at 652, the method 600 may include obtaining a third input raster image. At 656, the method 600 may include, wherein obtaining the second input raster image and the third input raster image comprises obtaining the second input raster image and the third input raster image at a frame rate. At 660, the method 600 may include using the look up table to generate a third output image from the third input raster image. At 664, the method 600 may include, wherein generating the second output image and the third output image comprises generating the second output image and the third output image at approximately the frame rate.

[0070] In some embodiments, the methods and processes described herein may be tied to a computing system of one or more computing devices. In particular, such methods and processes may be implemented as a computer-application program or service, an application-programming interface (API), a library, and/or other computer-program product.

[0071] FIG. 7 schematically shows a non-limiting embodiment of a computing system 700 that can enact one or more of the methods and processes described above. Computing system 700 is shown in simplified form. Computing system 700 may take the form of one or more gaming consoles, personal computers, server computers, tablet computers, home-entertainment computers, network computing devices, gaming devices, mobile computing devices, mobile communication devices (e.g., smart phones), and/or other computing devices, including wearable computing devices such as smart wristwatches and head mounted display devices. In the above examples, computing device 104, user computing device 120, HMD device 124, HMD device 208 and HMD devices 304A and 304B may comprise computing system 700 or one or more aspects of computing system 700.

[0072] Computing system 700 includes a logic processor 704, volatile memory 708, and a non-volatile storage device 712. Computing system 700 may optionally include a display subsystem 716, input subsystem 720, communication subsystem 724 and/or other components not shown in FIG. 7.

5 **[0073]** Logic processor 704 includes one or more physical devices configured to execute instructions. For example, the logic processor may be configured to execute instructions that are part of one or more applications, services, programs, routines, libraries, objects, components, data structures, or other logical constructs. Such instructions may be implemented to perform a task, implement a data type, transform the
10 state of one or more components, achieve a technical effect, or otherwise arrive at a desired result.

[0074] The logic processor 704 may include one or more physical processors (hardware) configured to execute software instructions. Additionally or alternatively, the logic processor may include one or more hardware logic circuits or firmware devices
15 configured to execute hardware-implemented logic or firmware instructions. Processors of the logic processor 704 may be single-core or multi-core, and the instructions executed thereon may be configured for sequential, parallel, and/or distributed processing. Individual components of the logic processor optionally may be distributed among two or more separate devices, which may be remotely located and/or configured for coordinated
20 processing. Aspects of the logic processor may be virtualized and executed by remotely accessible, networked computing devices configured in a cloud-computing configuration. In such a case, these virtualized aspects are run on different physical logic processors of various different machines, it will be understood.

[0075] Non-volatile storage device 712 includes one or more physical devices
25 configured to hold instructions executable by the logic processors to implement the methods and processes described herein. When such methods and processes are implemented, the state of non-volatile storage device 712 may be transformed—e.g., to hold different data.

[0076] Non-volatile storage device 712 may include physical devices that are
30 removable and/or built-in. Non-volatile storage device 712 may include optical memory (e.g., CD, DVD, HD-DVD, Blu-Ray Disc, etc.), semiconductor memory (e.g., ROM, EPROM, EEPROM, FLASH memory, etc.), and/or magnetic memory (e.g., hard-disk drive, floppy-disk drive, tape drive, MRAM, etc.), or other mass storage device technology. Non-volatile storage device 712 may include nonvolatile, dynamic, static,

read/write, read-only, sequential-access, location-addressable, file-addressable, and/or content-addressable devices. It will be appreciated that non-volatile storage device 712 is configured to hold instructions even when power is cut to the non-volatile storage device 712.

5 **[0077]** Volatile memory 708 may include physical devices that include random access memory. Volatile memory 708 is typically utilized by logic processor 704 to temporarily store information during processing of software instructions. It will be appreciated that volatile memory 708 typically does not continue to store instructions when power is cut to the volatile memory 708.

10 **[0078]** Aspects of logic processor 704, volatile memory 708, and non-volatile storage device 712 may be integrated together into one or more hardware-logic components. Such hardware-logic components may include field-programmable gate arrays (FPGAs), program- and application-specific integrated circuits (ASIC / ASICs), program- and application-specific standard products (PSSP / ASSPs), system-on-a-chip
15 (SOC), and complex programmable logic devices (CPLDs), for example.

[0079] The terms “program” and “application” may be used to describe an aspect of computing system 700 typically implemented in software by a processor to perform a particular function using portions of volatile memory, which function involves transformative processing that specially configures the processor to perform the function.
20 Thus, a program or application may be instantiated via logic processor 704 executing instructions held by non-volatile storage device 712, using portions of volatile memory 708. It will be understood that different programs and/or applications may be instantiated from the same application, service, code block, object, library, routine, API, function, etc. Likewise, the same program and/or application may be instantiated by different
25 applications, services, code blocks, objects, routines, APIs, functions, etc. The terms “program” and “application” may encompass individual or groups of executable files, data files, libraries, drivers, scripts, database records, etc.

[0080] It will be appreciated that a “service”, as used herein, is an application program executable across multiple user sessions. A service may be available to one or
30 more system components, programs, and/or other services. In some implementations, a service may run on one or more server-computing devices.

[0081] When included, display subsystem 716 may be used to present a visual representation of data held by non-volatile storage device 712. As the herein described methods and processes change the data held by the non-volatile storage device, and thus

transform the state of the non-volatile storage device, the state of display subsystem 716 may likewise be transformed to visually represent changes in the underlying data. Display subsystem 716 may include one or more display devices utilizing virtually any type of technology. Such display devices may be combined with logic processor 704, volatile
5 memory 708, and/or non-volatile storage device 712 in a shared enclosure, or such display devices may be peripheral display devices.

[0082] When included, input subsystem 720 may comprise or interface with one or more user-input devices such as a keyboard, mouse, touch screen, or game controller. In some embodiments, the input subsystem may comprise or interface with selected natural
10 user input (NUI) componentry. Such componentry may be integrated or peripheral, and the transduction and/or processing of input actions may be handled on- or off-board. Example NUI componentry may include a microphone for speech and/or voice recognition; an infrared, color, stereoscopic, and/or depth camera for machine vision and/or gesture recognition; a head tracker, eye tracker, accelerometer, and/or gyroscope
15 for motion detection and/or intent recognition; as well as electric-field sensing componentry for assessing brain activity; and/or any other suitable sensor.

[0083] When included, communication subsystem 724 may be configured to communicatively couple various computing devices described herein with each other, and with other devices. Communication subsystem 724 may include wired and/or wireless
20 communication devices compatible with one or more different communication protocols. As non-limiting examples, the communication subsystem may be configured for communication via a wireless telephone network, or a wired or wireless local- or wide-area network, such as a HDMI over Wi-Fi connection. In some embodiments, the communication subsystem may allow computing system 700 to send and/or receive
25 messages to and/or from other devices via a network such as the Internet.

[0084] The following paragraphs provide additional support for the claims of the subject application. One aspect provides a method for simulating light passing through one or more lenses, the method comprising: at a precomputing stage: obtaining a point spread function of the one or more lenses; obtaining a first input raster image comprising a
30 plurality of pixels; ray tracing the first input raster image using the point spread function to generate a first output image; and based on ray tracing the first input raster image, generating a look up table by computing a contribution to a pixel in the first output image, wherein the contribution is from a pixel at each location of a subset of locations in the first input raster image; at a runtime stage: obtaining a second input raster image comprising a

plurality of pixels; and using the look up table to generate a second output image from the second input raster image. The method may additionally or alternatively include using at least the second output image to train a cognitive algorithm. The method may additionally or alternatively include, wherein the first output image and the second output image
5 comprise ground truth data describing at least a portion of a real world environment depicted in the first output image and the second output image. The method may additionally or alternatively include, wherein the ground truth data comprise one or more of depth data, orientation data, and identification data. The method may additionally or alternatively include, obtaining a third input raster image; using the look up table to
10 generate a third output image from the third input raster image; wherein obtaining the second input raster image and the third input raster image comprises obtaining the second input raster image and the third input raster image at a frame rate; and wherein generating the second output image and the third output image comprises generating the second output image and the third output image at approximately the frame rate. The method may
15 additionally or alternatively include, determining a color value of a given pixel in the second output image by: (a) determining a color value of a pixel in the second input raster image at one of the locations of the subset of locations; (b) obtaining, from the look up table, the contribution to the given pixel in the second output image from the pixel at the one location of the subset of locations; (c) combining the color value of the pixel with the
20 contribution to the given pixel; and repeating the processes (a), (b), and (c) above for each location of the subset of locations in the second input raster image. The method may additionally or alternatively include, wherein determining the color value of the given pixel comprises using a bilinear filter to combine two color values of two pixels in the second input raster image with two contributions to the given pixel from the two pixels in
25 the second input raster image. The method may additionally or alternatively include, wherein the subset of locations comprises locations of at least 512 different pixels.

[0085] Another aspect provides a computing device for simulating light passing through one or more lenses, comprising: a processor; and a memory holding instructions executable by the processor to, obtain a point spread function of the one or more lenses;
30 obtain a first input raster image comprising a plurality of pixels; ray trace the first input raster image using the point spread function to generate a first output image; based on ray tracing the first input raster image, generate a look up table by computing a contribution to a pixel in the first output image, wherein the contribution is from a pixel at each location of a subset of locations in the first input raster image; obtain a second input raster

image comprising a plurality of pixels; and use the look up table to generate a second output image from the second input raster image. The computing device may additionally or alternatively include, wherein the instructions are further executable by the processor to use at least the second output image to train a cognitive algorithm. The computing device

5 may additionally or alternatively include, wherein the first output image and the second output image comprise ground truth data describing at least a portion of a real world environment depicted in the first output image and the second output image. The computing device may additionally or alternatively include, wherein the ground truth data comprises one or more of depth data, orientation data, and identification data. The

10 computing device may additionally or alternatively include, wherein the instructions are further executable by the processor to: obtain a third input raster image; use the look up table to generate a third output image from the third input raster image; wherein obtaining the second input raster image and the third input raster image comprises obtaining the second input raster image and the third input raster image at a frame rate of an image

15 capture device; and wherein generating the second output image and the third output image comprises generating the second output image and the third output image at approximately the frame rate. The computing device may additionally or alternatively include, wherein the instructions are further executable by the processor to determine a color value of a given pixel in the second output image by: (a) determining a color value

20 of a pixel in the second input raster image at one of the locations of the subset of locations; (b) obtaining, from the look up table, the contribution to the given pixel in the second output image from the pixel at the one location of the subset of locations; (c) combining the color value of the pixel with the contribution to the given pixel; and repeating the processes (a), (b) and (c) above for each location of the subset of locations in the second

25 input raster image. The computing device may additionally or alternatively include, wherein the instructions are further executable by the processor to use a bilinear filter to combine two color values of two pixels in the second input raster image with two contributions to the given pixel from the two pixels in the second input raster image. The computing device may additionally or alternatively include, wherein the subset of

30 locations comprises locations of at least 512 different pixels.

[0086] Another aspect provides a computing device for simulating light passing through one or more lenses, comprising: a processor; and a memory holding instructions executable by the processor to, obtain a point spread function of the one or more lenses; obtain a first input raster image; generate a pinhole image from the first input raster image,

the pinhole image comprising a plurality of pixels; ray trace the pinhole image using the point spread function to generate a first output image; based on ray tracing the pinhole image, generate a look up table by computing a contribution to a pixel in the first output image, wherein the contribution is from a pixel at each location of a subset of locations in the first input raster image; obtain additional input raster images at a frame rate; and use the look up table to generate additional output images from the additional input raster images at approximately the frame rate. The computing device may additionally or alternatively include, wherein the instructions are further executable by the processor to use at least the additional output images to train a cognitive algorithm. The computing device may additionally or alternatively include, wherein the instructions are further executable by the processor to determine a color value of a given pixel in the second output image by: (a) determining a color value of a pixel in the second input raster image at one of the locations of the subset of locations; (b) obtaining, from the look up table, the contribution to the given pixel in the second output image from the pixel at the one location of the subset of locations; (c) combining the color value of the pixel with the contribution to the given pixel; and repeating the processes (a), (b) and (c) above for each location of the subset of locations in the second input raster image. The computing device may additionally or alternatively include, wherein the instructions are further executable by the processor to use a bilinear filter to combine two color values of two pixels in the additional input images with two contributions to the given pixel from the two pixels in the second input raster image.

[0087] It will be understood that the configurations and/or approaches described herein are exemplary in nature, and that these specific embodiments or examples are not to be considered in a limiting sense, because numerous variations are possible. The specific routines or methods described herein may represent one or more of any number of processing strategies. As such, various acts illustrated and/or described may be performed in the sequence illustrated and/or described, in other sequences, in parallel, or omitted. Likewise, the order of the above-described processes may be changed.

[0088] The subject matter of the present disclosure includes all novel and non-obvious combinations and sub-combinations of the various processes, systems and configurations, and other features, functions, acts, and/or properties disclosed herein, as well as any and all equivalents thereof.

CLAIMS

1. A method for simulating light passing through one or more lenses, the method comprising:
 - at a precomputing stage:
 - obtaining a point spread function of the one or more lenses;
 - obtaining a first input raster image comprising a plurality of pixels;
 - ray tracing the first input raster image using the point spread function to generate a first output image; and
 - based on ray tracing the first input raster image, generating a look up table by computing a contribution to a pixel in the first output image, wherein the contribution is from a pixel at each location of a subset of locations in the first input raster image;
 - at a runtime stage:
 - obtaining a second input raster image comprising a plurality of pixels; and
 - using the look up table to generate a second output image from the second input raster image.
2. The method of claim 1, further comprising using at least the second output image to train a cognitive algorithm.
3. The method of claim 1, wherein the first output image and the second output image comprise ground truth data describing at least a portion of a real world environment depicted in the first output image and the second output image.
4. The method of claim 3, wherein the ground truth data comprise one or more of depth data, orientation data, and identification data.
5. The method of claim 1, further comprising:
 - obtaining a third input raster image;
 - using the look up table to generate a third output image from the third input raster image;wherein obtaining the second input raster image and the third input raster image comprises obtaining the second input raster image and the third input raster image at a frame rate; and
 - wherein generating the second output image and the third output image comprises generating the second output image and the third output image at approximately the frame rate.
6. The method of claim 1, further comprising determining a color value of a given

pixel in the second output image by:

(a) determining a color value of a pixel in the second input raster image at one of the locations of the subset of locations;

(b) obtaining, from the look up table, the contribution to the given pixel in the second output image from the pixel at the one location of the subset of locations;

(c) combining the color value of the pixel with the contribution to the given pixel; and

repeating the processes (a), (b), and (c) above for each location of the subset of locations in the second input raster image.

7. The method of claim 6, wherein determining the color value of the given pixel comprises using a bilinear filter to combine two color values of two pixels in the second input raster image with two contributions to the given pixel from the two pixels in the second input raster image.

8. The method of claim 1, wherein the subset of locations comprises locations of at least 512 different pixels.

9. A computing device for simulating light passing through one or more lenses, comprising:

a processor; and

a memory holding instructions executable by the processor to,

obtain a point spread function of the one or more lenses;

obtain a first input raster image comprising a plurality of pixels;

ray trace the first input raster image using the point spread function to generate a first output image;

based on ray tracing the first input raster image, generate a look up table by computing a contribution to a pixel in the first output image, wherein the contribution is from a pixel at each location of a subset of locations in the first input raster image;

obtain a second input raster image comprising a plurality of pixels; and

use the look up table to generate a second output image from the second input raster image.

10. The computing device of claim 9, wherein the instructions are further executable by the processor to use at least the second output image to train a cognitive algorithm.

11. The computing device of claim 9, wherein the first output image and the second output image comprise ground truth data describing at least a portion of a real world environment depicted in the first output image and the second output image.

12. The computing device of claim 9, wherein the instructions are further executable by the processor to:

obtain a third input raster image;

use the look up table to generate a third output image from the third input raster image;

wherein obtaining the second input raster image and the third input raster image comprises obtaining the second input raster image and the third input raster image at a frame rate of an image capture device; and

wherein generating the second output image and the third output image comprises generating the second output image and the third output image at approximately the frame rate.

13. The computing device of claim 9, wherein the instructions are further executable by the processor to determine a color value of a given pixel in the second output image by:

(a) determining a color value of a pixel in the second input raster image at one of the locations of the subset of locations;

(b) obtaining, from the look up table, the contribution to the given pixel in the second output image from the pixel at the one location of the subset of locations;

(c) combining the color value of the pixel with the contribution to the given pixel; and

repeating the processes (a), (b) and (c) above for each location of the subset of locations in the second input raster image.

14. The computing device of claim 13, wherein the instructions are further executable by the processor to use a bilinear filter to combine two color values of two pixels in the second input raster image with two contributions to the given pixel from the two pixels in the second input raster image.

15. The computing device of claim 9, wherein the subset of locations comprises locations of at least 512 different pixels.

1/8




FIG. 1




FIG. 2




FIG. 3




FIG. 4

5/8




FIG. 5A





FIG. 5B



7/8




FIG. 6B




FIG. 7

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2019/036220

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☒ Claims Nos.: 2-8, 11-15(completely); 1, 9, 10(partially)
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
see FURTHER INFORMATION sheet PCT/ISA/210
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/036220

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06T15/06
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EP0-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	THOMAS HACH ET AL: "Cinematic Bokeh rendering for real scenes", PROCEEDINGS OF THE 12TH EUROPEAN CONFERENCE ON VISUAL MEDIA PRODUCTION, CVMP '15, 1 January 2015 (2015-01-01), pages 1-10, XP055632848, New York, New York, USA DOI: 10.1145/2824840.2824842 ISBN: 978-1-4503-3560-7 Abstract, Figure 9, Sections 1,2.2,2.3, 3 and subsections, in particular 3.3, 5 -----	1,9,10
A	EP 2 687 892 A1 (SONY CORP [JP]) 22 January 2014 (2014-01-22) abstract; claims 1,2; figures 8,9 paragraphs [0001] - [0005] ----- -/--	1,9,10



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

18 October 2019

Date of mailing of the international search report

25/10/2019

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Almeida Garcia, B

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/036220

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 2 442 272 A1 (THOMSON LICENSING [FR]) 18 April 2012 (2012-04-18) abstract; claims 1,5-7 paragraphs [0001] - [0014], [0025] - [0027] -----	1,9,10
A	Brian A Barsky ET AL: "Introducing Vision-Realistic Rendering", Thirteenth Eurographics Workshop on Rendering, 1 January 2002 (2002-01-01), XP055633877, Retrieved from the Internet: URL: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.117.6493&rep=rep1&type=pdf Abstract, section 2.2 -----	1,9,10

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2019/036220

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
EP 2687892	A1	22-01-2014	CA 2820005 A1	19-01-2014
			CN 103578100 A	12-02-2014
			EP 2687892 A1	22-01-2014
			JP 2014021487 A	03-02-2014
			US 2014022347 A1	23-01-2014

EP 2442272	A1	18-04-2012	EP 2442272 A1	18-04-2012
			WO 2012049022 A1	19-04-2012

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

Continuation of Box II.2

Claims Nos.: 2-8, 11-15(completely); 1, 9, 10(partially)

Claim 1 is directed to a method for simulating light passing through one or more lenses including some steps, inter alia:

obtaining a point spread function of the one or more lenses;

obtaining a first input raster image comprising a plurality of pixels;

ray tracing the first input raster image using the point spread function to generate a first output image ; and

based on ray tracing the first input raster image, generating a look up table by computing a contribution to a pixel in the first output image, wherein the contribution is from a pixel at each location of a subset of locations in the first input raster image;
at a runtime stage:

Firstly, it is unclear what is meant with "ray tracing" in the sense of the present claims with regard to the disclosure on paragraph [0044] which provides other means "in place of ray-tracing refraction of the simulated lens" and provides other means (e.g. interpolation) that can hardly be considered to amount to ray tracing.

Secondly, and most fundamentally, it is unclear how the simulation of light passing through one or more lenses can be reasonably performed by the steps of the claim: whereas it is acknowledged that ray tracing a three-dimensional scene considering a particular point spread function will yield a simulation of light passing through one or more lenses, the claimed ray tracing of a raster image (see also figures 5A, 5B and 6) will amount at most to the equivalent of a commonplace blurring since its output is not depth dependent as that of a lens would be. Therefore, it is unclear how the claimed subject-matter convincingly solves the problem of simulation of light passing through one or more lenses and it seems that no technical effect arises from the subject-matter as now claimed. Paragraph [0001] hints at a problem of reducing computational requirements for the simulation of lenses, however it is disputed that said simulation of lenses is taking place to in a meaningful way.

Since it is unclear how the problem can be seen to be solved by the claims and furthermore by the description as a whole, no meaningful complete search could be carried out.

What was searched was what could be seen as solving a technical problem, namely the following features:

A method for simulating light passing through one or more lenses, the method comprising:
at a precomputing stage:
obtaining a

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

point spread function of the one or more lenses;
 obtaining a first input
 raster image comprising a plurality of pixels;
 ray tracing the first
 input raster image using the point spread function to generate a first
 output image;

The remaining features were not searched, namely

based on ray tracing the first input raster image, generating a look up
 table by computing a contribution to a pixel in the first output image,
 wherein the contribution is from a pixel at each location of a subset of
 locations in the first input raster image;

at a runtime stage:

obtaining a second input raster image comprising a plurality of pixels;
 and

using the look up table to generate a second output image from the
 second input raster image.

The cited documents are all examples where
 a simulation of light passing through one or more lenses is taking place.
 All of the documents act in a depth-dependent manner. The claimed
 features generate a second output image from a second input image and a
 look-up table but it is unclear how the problem to be solved by the
 application can be reasonably solved merely by using this raster image.
 According to the description on paragraph [0032],

" [0032]Using the
 point spread function 152, the computing device 104 may utilize a ray
 tracer 156 to ray trace a first input raster image 132 and generate a
 first output image 140. Next, the computing device 104 may generate a
 look up table 172 based on ray tracing the first input raster image 132.
 For example, the first input raster image 132 may comprise a plurality of
 pixels 160. Each pixel of the plurality of pixels 160 may have a
 discrete location 164 within the image and a color value 168 "

The
 description is silent as to how depth data is acquired and considered to
 solve the problem of simulating light passing through one or more lenses.
 The last sentence of said paragraph even leads away from including depth
 data.

Therefore, it is considered that it is unclear how the problem can
 be solved by the claimed features and by the disclosure in the
 description. Also no likely fallback position is presently envisioned.
 Therefore, for the unsearched features, the provisions of Article
 17(2)(a)(ii) PCT apply and no meaningful search could be carried out.
 The subject-matter of claims 2-8,11-15 refers to and thus implies the
 presence of at least some of the said unsearched features and was
 therefore accordingly also not searched.

The applicant's attention is drawn to the fact that claims relating to
 inventions in respect of which no international search report has been
 established need not be the subject of an international preliminary
 examination (Rule 66.1(e) PCT). The applicant is advised that the EPO
 policy when acting as an International Preliminary Examining Authority is
 normally not to carry out a preliminary examination on matter which has
 not been searched. This is the case irrespective of whether or not the

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

claims are amended following receipt of the search report or during any Chapter II procedure. If the application proceeds into the regional phase before the EPO, the applicant is reminded that a search may be carried out during examination before the EPO (see EPO Guidelines C-IV, 7.2), should the problems which led to the Article 17(2) declaration be overcome.