



(12) **DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION**

(13) **A1**

(86) Date de dépôt PCT/PCT Filing Date: 2017/12/01
(87) Date publication PCT/PCT Publication Date: 2018/06/14
(85) Entrée phase nationale/National Entry: 2019/05/30
(86) N° demande PCT/PCT Application No.: US 2017/064311
(87) N° publication PCT/PCT Publication No.: 2018/106542
(30) Priorités/Priorities: 2016/12/05 (US62/430,308);
2016/12/05 (US62/430,279)

(51) Cl.Int./Int.Cl. *G06T 19/00* (2011.01)
(71) Demandeur/Applicant:
MAGIC LEAP, INC., US
(72) Inventeurs/Inventors:
KAEHLER, ADRIAN, US;
CROSTON, JOHN ADAM, US...
(74) Agent: RICHES, MCKENZIE & HERBERT LLP

(54) Titre : COMMANDES D'ENTREE D'UTILISATEUR VIRTUEL DANS UN ENVIRONNEMENT DE REALITE MIXTE
(54) Title: VIRTUAL USER INPUT CONTROLS IN A MIXED REALITY ENVIRONMENT




FIG. 13

(57) **Abrégé/Abstract:**

A wearable display system can automatically recognize a physical remote or a device that the remote serves using computer vision techniques. The wearable system can generate a virtual remote with a virtual control panel viewable and interactable by user of the

(57) **Abrégé(suite)/Abstract(continued):**

wearable system. The virtual remote can emulate the functionality of the physical remote. The user can select a virtual remote for interaction, for example, by looking or pointing at the parent device or its remote control, or by selecting from a menu of known devices. The virtual remote may include a virtual button, which is associated with a volume in the physical space. The wearable system can detect that a virtual button is actuated by determining whether a portion of the users body (e.g., the user's finger) has penetrated, the volume associated with the virtual button.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau(43) International Publication Date
14 June 2018 (14.06.2018)(10) International Publication Number
WO 2018/106542 A1(51) International Patent Classification:
G06T 19/00 (2011.01)(21) International Application Number:
PCT/US2017/064311(22) International Filing Date:
01 December 2017 (01.12.2017)

(25) Filing Language: English

(26) Publication Language: English


(30) Priority Data:
62/430,279 05 December 2016 (05.12.2016) US
62/430,308 05 December 2016 (05.12.2016) US(71) Applicant: **MAGIC LEAP, INC.** [US/US]; 7500 W. Sunrise Blvd., Plantation, FL 33322 (US).(72) Inventors: **KAEHLER, Adrian**; 1940 N. Western Avenue, Los Angeles, CA 90027 (US). **CROSTON, John**,

Adam; 1800 Eastlake Ave. East #310, Seattle, WA 98102 (US).

(74) Agent: **ALTMAN, Daniel, E.**; Knobbe, Martens, Olson & Bear LLP, 2040 Main Street 14th Floor, Irvine, CA 92614 (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.

(54) Title: VIRTUAL USER INPUT CONTROLS IN A MIXED REALITY ENVIRONMENT

**FIG. 13**

(57) Abstract: A wearable display system can automatically recognize a physical remote or a device that the remote serves using computer vision techniques. The wearable system can generate a virtual remote with a virtual control panel viewable and interactable by user of the wearable system. The virtual remote can emulate the functionality of the physical remote. The user can select a virtual remote for interaction, for example, by looking or pointing at the parent device or its remote control, or by selecting from a menu of known devices. The virtual remote may include a virtual button, which is associated with a volume in the physical space. The wearable system can detect that a virtual button is actuated by determining whether a portion of the users body (e.g., the user's finger) has penetrated, the volume associated with the virtual button.

[Continued on next page]

WO 2018/106542 A1

(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).

Published:
— *with international search report (Art. 21(3))*

VIRTUAL USER INPUT CONTROLS IN A MIXED REALITY ENVIRONMENT

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of priority under 35 U.S.C. § 119(e) to U.S. Provisional Application No. 62/430,308, filed on December 5, 2016, entitled "UNIVERSAL VISUAL REMOTE," and U.S. Provisional Application No. 62/430,279, filed on December 5, 2016, entitled "MECHANISM FOR A VIRTUAL BUTTON", the disclosures of which are hereby incorporated by reference herein in their entireties.

FIELD

[0002] The present disclosure relates to virtual reality and augmented reality imaging and visualization systems and more particularly to rendering virtual user input controls and detecting interactions with the virtual user input controls.

BACKGROUND

[0003] Modern computing and display technologies have facilitated the development of systems for so called "virtual reality", "augmented reality", or "mixed reality" experiences, wherein digitally reproduced images or portions thereof are presented to a user in a manner wherein they seem to be, or may be perceived as, real. A virtual reality, or "VR", scenario typically involves presentation of digital or virtual image information without transparency to other actual real-world visual input; an augmented reality, or "AR", scenario typically involves presentation of digital or virtual image information as an augmentation to visualization of the actual world around the user; a mixed reality, or "MR", related to merging real and virtual worlds to produce new environments where physical and virtual objects co-exist and interact in real time. As it turns out, the human visual perception system is very complex, and producing a VR, AR, or MR technology that facilitates a comfortable, natural-feeling, rich presentation of virtual image elements amongst other virtual or real-world imagery elements is challenging. Systems and methods disclosed herein address various challenges related to VR, AR and MR technology.

SUMMARY

[0004] Various examples of rendering virtual user input control and detecting an actuation of the virtual input control are described.

[0005] As one example, a wearable display system can automatically recognize a physical remote or a device that the remote serves using computer vision techniques. The wearable system can generate a virtual remote with a virtual control panel viewable and interactable by user of the wearable system. The virtual remote can emulate the functionality of the physical remote. The user can select a virtual remote for interaction, for example, by looking or pointing at the parent device or its remote control, or by selecting from a menu of known devices. The virtual remote may include a virtual button, which is associated with a volume in the physical space. The wearable system can detect that a virtual button is actuated by determining whether a portion of the user's body (e.g., the user's finger) has penetrated the volume associated with the virtual button.

[0006] Details of one or more implementations of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages will become apparent from the description, the drawings, and the claims. Neither this summary nor the following detailed description purports to define or limit the scope of the inventive subject matter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 depicts an illustration of a mixed reality scenario with certain virtual reality objects, and certain physical objects viewed by a person.

[0008] FIG. 2 schematically illustrates an example of a wearable system.

[0009] FIG. 3 schematically illustrates aspects of an approach for simulating three-dimensional imagery using multiple depth planes.

[0010] FIG. 4 schematically illustrates an example of a waveguide stack for outputting image information to a user.

[0011] FIG. 5 shows example exit beams that may be outputted by a waveguide.

[0012] FIG. 6 is a schematic diagram showing an optical system including a waveguide apparatus, an optical coupler subsystem to optically couple light to or from the

waveguide apparatus, and a control subsystem, used in the generation of a multi-focal volumetric display, image, or light field.

[0013] FIG. 7 is a block diagram of an example of a wearable system.

[0014] FIG. 8 is a process flow diagram of an example of a method of rendering virtual content in relation to recognized objects.

[0015] FIG. 9 is a block diagram of another example of a wearable system.

[0016] FIG. 10 is a process flow diagram of an example of a method for determining user input to a wearable system.

[0017] FIG. 11 is a process flow diagram of an example of a method for interacting with a virtual user interface.

[0018] FIG. 12 illustrates an example of a physical environment as perceived by a user of an augmented reality device.

[0019] FIG. 13 illustrates an example of initiating a selection event associated with a virtual remote using a hand gesture.

[0020] FIG. 14 illustrates an example of actuating a virtual remote during a selection event.

[0021] FIG. 15 illustrates an example of automatically hiding a virtual remote or terminating the selection event if a threshold condition is met.

[0022] FIG. 16 illustrates an example process of rendering a virtual remote.

[0023] FIG. 17 illustrates an example of detecting an act of pressing a virtual button by a user.

[0024] FIG. 18 illustrates an example where a user actuates a virtual button by sliding the virtual button up and down.

[0025] FIG. 19 illustrates an example where a user actuates a virtual touch surface.

[0026] FIG. 20 illustrates an example process for initiating an interaction event with a virtual element in a mixed reality user interface.

[0027] Throughout the drawings, reference numbers may be re-used to indicate correspondence between referenced elements. The drawings are provided to illustrate example embodiments described herein and are not intended to limit the scope of the disclosure.

DETAILED DESCRIPTION

Overview

[0028] Many devices (also referred to as parent devices) such as televisions, audio systems, home theaters, home security systems, door bells or door locks, air conditioners or heaters, thermostats, lighting systems (e.g., smart light bulbs), garage door openers, or other household appliances or smart devices can be operated by a remote control. The remote control may be a physical remote comprising a small hand held device. The physical remote can comprise buttons (or a touch surface) that allow a user of a parent device to perform basic functions of the device from a distance. The physical remote can communicate with the parent device using various wireless channels such as an infrared (IR) signal or a radio frequency (RF) signal such as Bluetooth, Wi-Fi, etc. The communication between the physical remote and the parent device may be governed by one or more industry standards or protocols such as those created by the Infrared Data Association (IrDA) or by the Institute of Electrical and Electronics Engineers (IEEE).

[0029] Over time, the number of household appliances and their corresponding remotes has grown in a consumer's home. Increasingly many devices are connected to, and controllable from, a network (e.g., devices connected to the Internet of Things) via a remote control. This situation led to the introduction of universal remote control devices that could be programmed to emulate the signals generated by some or even all of the remotes that had been accumulated for different devices. The universal remote control devices may be part of a smartphone or a tablet. For example, the smartphone can use software to emulate a physical remote and generate an IR signal, in conjunction with a control interface appropriate to the emulated physical remote.

[0030] In an augmented reality (AR) or mixed reality (MR) environment, however, emulating a physical remote on a physical device (such as a smartphone or a tablet) may be impracticable because a head-mounted device (HMD) may have a small field of view and the user may need to hold the universal remote control device at an uncomfortable location to interact with the parent device. In addition, the user in the AR/MR environment can also interact with objects using body poses and gestures and therefore, may not need a separate physical device as the universal remote control.

[0031] The systems and methods described herein are directed at least in part to these challenges. For example, a wearable system can automatically recognize the physical remote or the parent device the remote serves using computer vision techniques. In addition, the wearable system can generate a virtual remote with a control panel viewable and interactable by the user of the wearable system. The user can identify a parent device or physical remote using body pose or eye gaze (e.g., by looking at or pointing at the parent device or the physical remote). Upon detecting the body pose or eye gaze, the wearable system can bring up an associated virtual remote. The user can also select a virtual remote from a menu of known devices.

[0032] The virtual remote may include one or more virtual elements on its control panel. The virtual elements may include virtual buttons, virtual sliders, virtual toggles, or a virtual touch surface. Unlike a physical remote or user input device, the virtual remote does not have a user control element (e.g., a button or a touch surface) such that the actuation of the user control element can directly trigger an electrical signal to be sent from the physical remote or the user input device to the associated computing device. Advantageously, in some embodiments, to detect that a user has actuated a virtual element, the wearable system can detect whether a portion of the user's body has intersected with an active surface of the virtual element. The wearable system can calculate the active surface's movement based the user's movement. For example, the speed and displacement of the active surface's movement may be a portion of the movement of the user's finger, which is being used to virtually touch a button. When the active surface reaches a certain threshold position, the wearable system can initiate an interaction event where the wearable system generates a signal indicating that the user has activated the virtual element. For example, the wearable system can determine that the user has pressed a virtual button when the displacement of the active surface has passed a threshold condition. When the user retracts his finger, the wearable system may move the active surface back toward its original position (before the user interaction). When the active surface moves sufficiently far back (such as passing another threshold position, the wearable system may generate another interaction event associated with releasing the virtual button.

Examples of 3D Display of a Wearable System

[0033] A wearable system (also referred to herein as an augmented reality (AR) system) can be configured to present 2D or 3D virtual images to a user. The images may be still images, frames of a video, or a video, in combination or the like. The wearable system can include a wearable device that can present a VR, AR, or MR environment, alone or in combination, for user interaction. The wearable device can be a head-mounted device (HMD) which is used interchangeably as an AR device (ARD). Further, for the purpose of the present disclosure, the term "AR" is used interchangeably with the term "MR".

[0034] FIG. 1 depicts an illustration of a mixed reality scenario with certain virtual reality objects, and certain physical objects viewed by a person. In FIG. 1, an MR scene 100 is depicted wherein a user of an MR technology sees a real-world park-like setting 110 featuring people, trees, buildings in the background, and a concrete platform 120. In addition to these items, the user of the MR technology also perceives that he "sees" a robot statue 130 standing upon the real-world platform 120, and a cartoon-like avatar character 140 flying by which seems to be a personification of a bumble bee, even though these elements do not exist in the real world.

[0035] In order for the 3D display to produce a true sensation of depth, and more specifically, a simulated sensation of surface depth, it may be desirable for each point in the display's visual field to generate an accommodative response corresponding to its virtual depth. If the accommodative response to a display point does not correspond to the virtual depth of that point, as determined by the binocular depth cues of convergence and stereopsis, the human eye may experience an accommodation conflict, resulting in unstable imaging, harmful eye strain, headaches, and, in the absence of accommodation information, almost a complete lack of surface depth.

[0036] VR, AR, and MR experiences can be provided by display systems having displays in which images corresponding to a plurality of depth planes are provided to a viewer. The images may be different for each depth plane (e.g., provide slightly different presentations of a scene or object) and may be separately focused by the viewer's eyes, thereby helping to provide the user with depth cues based on the accommodation of the eye required to bring into focus different image features for the scene located on different depth

plane or based on observing different image features on different depth planes being out of focus. As discussed elsewhere herein, such depth cues provide credible perceptions of depth.

[0037] FIG. 2 illustrates an example of wearable system 200. The wearable system 200 includes a display 220, and various mechanical and electronic modules and systems to support the functioning of display 220. The display 220 may be coupled to a frame 230, which is wearable by a user, wearer, or viewer 210. The display 220 can be positioned in front of the eyes of the user 210. The display 220 can present AR/VR/MR content to a user. The display 220 can comprise a head mounted display that is worn on the head of the user. In some embodiments, a speaker 240 is coupled to the frame 230 and positioned adjacent the ear canal of the user (in some embodiments, another speaker, not shown, is positioned adjacent the other ear canal of the user to provide for stereo/shapeable sound control). The display 220 can include an audio sensor (e.g., a microphone) for detecting an audio stream from the environment on which to perform voice recognition.

[0038] The wearable system 200 can include an outward-facing imaging system 464 (shown in FIG. 4) which observes the world in the environment around the user. The wearable system 200 can also include an inward-facing imaging system 462 (shown in FIG. 4) which can track the eye movements of the user. The inward-facing imaging system may track either one eye's movements or both eyes' movements. The inward-facing imaging system 462 may be attached to the frame 230 and may be in electrical communication with the processing modules 260 or 270, which may process image information acquired by the inward-facing imaging system to determine, e.g., the pupil diameters or orientations of the eyes, eye movements or eye pose of the user 210.

[0039] As an example, the wearable system 200 can use the outward-facing imaging system 464 or the inward-facing imaging system 462 to acquire images of a pose of the user. The images may be still images, frames of a video, or a video, in combination or the like.

[0040] The wearable system 200 can also include a sensor 232 configured to communicate with a physical device (such as a television, an air conditioner, or other household appliances). For example, the wearable system 200 may serve as a virtual remote for controlling the physical device. The sensor 232 can be an infra-red light source controllable by the local processing and data module 270. The sensor 232 can also emit other

optical signals or electromagnetic signals for communications with the physical device. The communications between the sensor 232 and the physical device may conform to one or more industry standards. For example, where the sensor 232 is configured to emit infrared signals, the communication between the sensor 232 and the physical device may be in accordance with a protocol designed by Infrared Data Association (IrDA). The sensor 232 can be located at the temple (such as the earpieces) of glasses, the nose pad of the wearable system 200, or other locations on a head-mounted display or remote to the head-mounted display (such as e.g., on a backpack). Although the example in FIG. 2 only illustrates one sensor 232, in some embodiments, the display system may be associated with multiple sensors 232.

[0041] The display 220 can be operatively coupled 250, such as by a wired lead or wireless connectivity, to a local data processing module 260 which may be mounted in a variety of configurations, such as fixedly attached to the frame 230, fixedly attached to a helmet or hat worn by the user, embedded in headphones, or otherwise removably attached to the user 210 (e.g., in a backpack-style configuration, in a belt-coupling style configuration).

[0042] The local processing and data module 260 may comprise a hardware processor, as well as digital memory, such as non-volatile memory (e.g., flash memory), both of which may be utilized to assist in the processing, caching, and storage of data. The data may include data a) captured from sensors (which may be, e.g., operatively coupled to the frame 230 or otherwise attached to the user 210), such as image capture devices (e.g., cameras in the inward-facing imaging system or the outward-facing imaging system), audio sensors (e.g., microphones), inertial measurement units (IMUs), accelerometers, compasses, global positioning system (GPS) units, radio devices, or gyroscopes; or b) acquired or processed using remote processing module 270 or remote data repository 280, possibly for passage to the display 220 after such processing or retrieval. The local processing and data module 260 may be operatively coupled by communication links 262 or 264, such as via wired or wireless communication links, to the remote processing module 270 or remote data repository 280 such that these remote modules are available as resources to the local processing and data module 260. In addition, remote processing module 270 and remote data repository 280 may be operatively coupled to each other.

[0043] In some embodiments, the remote processing module 270 may comprise one or more processors configured to analyze and process data or image information. In some embodiments, the remote data repository 280 may comprise a digital data storage facility, which may be available through the internet or other networking configuration in a "cloud" resource configuration. In some embodiments, all data is stored and all computations are performed in the local processing and data module, allowing fully autonomous use from a remote module.

[0044] The human visual system is complicated and providing a realistic perception of depth is challenging. Without being limited by theory, it is believed that viewers of an object may perceive the object as being three-dimensional due to a combination of vergence and accommodation. Vergence movements (i.e., rolling movements of the pupils toward or away from each other to converge the lines of sight of the eyes to fixate upon an object) of the two eyes relative to each other are closely associated with focusing (or "accommodation") of the lenses of the eyes. Under normal conditions, changing the focus of the lenses of the eyes, or accommodating the eyes, to change focus from one object to another object at a different distance will automatically cause a matching change in vergence to the same distance, under a relationship known as the "accommodation-vergence reflex." Likewise, a change in vergence will trigger a matching change in accommodation, under normal conditions. Display systems that provide a better match between accommodation and vergence may form more realistic and comfortable simulations of three-dimensional imagery.

[0045] FIG. 3 illustrates aspects of an approach for simulating a three-dimensional imagery using multiple depth planes. With reference to FIG. 3, objects at various distances from eyes 302 and 304 on the z-axis are accommodated by the eyes 302 and 304 so that those objects are in focus. The eyes 302 and 304 assume particular accommodated states to bring into focus objects at different distances along the z-axis. Consequently, a particular accommodated state may be said to be associated with a particular one of depth planes 306, which has an associated focal distance, such that objects or parts of objects in a particular depth plane are in focus when the eye is in the accommodated state for that depth plane. In some embodiments, three-dimensional imagery may be simulated by providing different presentations of an image for each of the eyes 302 and 304, and also by

providing different presentations of the image corresponding to each of the depth planes. While shown as being separate for clarity of illustration, it will be appreciated that the fields of view of the eyes 302 and 304 may overlap, for example, as distance along the z-axis increases. In addition, while shown as flat for the ease of illustration, it will be appreciated that the contours of a depth plane may be curved in physical space, such that all features in a depth plane are in focus with the eye in a particular accommodated state. Without being limited by theory, it is believed that the human eye typically can interpret a finite number of depth planes to provide depth perception. Consequently, a highly believable simulation of perceived depth may be achieved by providing, to the eye, different presentations of an image corresponding to each of these limited number of depth planes.

Waveguide Stack Assembly

[0046] FIG. 4 illustrates an example of a waveguide stack for outputting image information to a user. A wearable system 400 includes a stack of waveguides, or stacked waveguide assembly 480 that may be utilized to provide three-dimensional perception to the eye/brain using a plurality of waveguides 432b, 434b, 436b, 438b, 4400b. In some embodiments, the wearable system 400 may correspond to wearable system 200 of FIG. 2, with FIG. 4 schematically showing some parts of that wearable system 200 in greater detail. For example, in some embodiments, the waveguide assembly 480 may be integrated into the display 220 of FIG. 2.

[0047] With continued reference to FIG. 4, the waveguide assembly 480 may also include a plurality of features 458, 456, 454, 452 between the waveguides. In some embodiments, the features 458, 456, 454, 452 may be lenses. In other embodiments, the features 458, 456, 454, 452 may not be lenses. Rather, they may simply be spacers (e.g., cladding layers or structures for forming air gaps).

[0048] The waveguides 432b, 434b, 436b, 438b, 440b or the plurality of lenses 458, 456, 454, 452 may be configured to send image information to the eye with various levels of wavefront curvature or light ray divergence. Each waveguide level may be associated with a particular depth plane and may be configured to output image information corresponding to that depth plane. Image injection devices 420, 422, 424, 426, 428 may be utilized to inject image information into the waveguides 440b, 438b, 436b, 434b, 432b, each

of which may be configured to distribute incoming light across each respective waveguide, for output toward the eye 410. Light exits an output surface of the image injection devices 420, 422, 424, 426, 428 and is injected into a corresponding input edge of the waveguides 440b, 438b, 436b, 434b, 432b. In some embodiments, a single beam of light (e.g., a collimated beam) may be injected into each waveguide to output an entire field of cloned collimated beams that are directed toward the eye 410 at particular angles (and amounts of divergence) corresponding to the depth plane associated with a particular waveguide.

[0049] In some embodiments, the image injection devices 420, 422, 424, 426, 428 are discrete displays that each produce image information for injection into a corresponding waveguide 440b, 438b, 436b, 434b, 432b, respectively. In some other embodiments, the image injection devices 420, 422, 424, 426, 428 are the output ends of a single multiplexed display which may, e.g., pipe image information via one or more optical conduits (such as fiber optic cables) to each of the image injection devices 420, 422, 424, 426, 428.

[0050] A controller 460 controls the operation of the stacked waveguide assembly 480 and the image injection devices 420, 422, 424, 426, 428. The controller 460 includes programming (e.g., instructions in a non-transitory computer-readable medium) that regulates the timing and provision of image information to the waveguides 440b, 438b, 436b, 434b, 432b. In some embodiments, the controller 460 may be a single integral device, or a distributed system connected by wired or wireless communication channels. The controller 460 may be part of the processing modules 260 or 270 (illustrated in FIG. 2) in some embodiments.

[0051] The waveguides 440b, 438b, 436b, 434b, 432b may be configured to propagate light within each respective waveguide by total internal reflection (TIR). The waveguides 440b, 438b, 436b, 434b, 432b may each be planar or have another shape (e.g., curved), with major top and bottom surfaces and edges extending between those major top and bottom surfaces. In the illustrated configuration, the waveguides 440b, 438b, 436b, 434b, 432b may each include light extracting optical elements 440a, 438a, 436a, 434a, 432a that are configured to extract light out of a waveguide by redirecting the light, propagating within each respective waveguide, out of the waveguide to output image information to the eye 410. Extracted light may also be referred to as outcoupled light, and light extracting optical elements may also be referred to as outcoupling optical elements. An extracted beam of light

is outputted by the waveguide at locations at which the light propagating in the waveguide strikes a light redirecting element. The light extracting optical elements (440a, 438a, 436a, 434a, 432a) may, for example, be reflective or diffractive optical features. While illustrated disposed at the bottom major surfaces of the waveguides 440b, 438b, 436b, 434b, 432b for ease of description and drawing clarity, in some embodiments, the light extracting optical elements 440a, 438a, 436a, 434a, 432a may be disposed at the top or bottom major surfaces, or may be disposed directly in the volume of the waveguides 440b, 438b, 436b, 434b, 432b. In some embodiments, the light extracting optical elements 440a, 438a, 436a, 434a, 432a may be formed in a layer of material that is attached to a transparent substrate to form the waveguides 440b, 438b, 436b, 434b, 432b. In some other embodiments, the waveguides 440b, 438b, 436b, 434b, 432b may be a monolithic piece of material and the light extracting optical elements 440a, 438a, 436a, 434a, 432a may be formed on a surface or in the interior of that piece of material.

[0052] With continued reference to FIG. 4, as discussed herein, each waveguide 440b, 438b, 436b, 434b, 432b is configured to output light to form an image corresponding to a particular depth plane. For example, the waveguide 432b nearest the eye may be configured to deliver collimated light, as injected into such waveguide 432b, to the eye 410. The collimated light may be representative of the optical infinity focal plane. The next waveguide up 434b may be configured to send out collimated light which passes through the first lens 452 (e.g., a negative lens) before it can reach the eye 410. First lens 452 may be configured to create a slight convex wavefront curvature so that the eye/brain interprets light coming from that next waveguide up 434b as coming from a first focal plane closer inward toward the eye 410 from optical infinity. Similarly, the third up waveguide 436b passes its output light through both the first lens 452 and second lens 454 before reaching the eye 410. The combined optical power of the first and second lenses 452 and 454 may be configured to create another incremental amount of wavefront curvature so that the eye/brain interprets light coming from the third waveguide 436b as coming from a second focal plane that is even closer inward toward the person from optical infinity than was light from the next waveguide up 434b.

[0053] The other waveguide layers (e.g., waveguides 438b, 440b) and lenses (e.g., lenses 456, 458) are similarly configured, with the highest waveguide 440b in the stack

sending its output through all of the lenses between it and the eye for an aggregate focal power representative of the closest focal plane to the person. To compensate for the stack of lenses 458, 456, 454, 452 when viewing/interpreting light coming from the world 470 on the other side of the stacked waveguide assembly 480, a compensating lens layer 430 may be disposed at the top of the stack to compensate for the aggregate power of the lens stack 458, 456, 454, 452 below. Such a configuration provides as many perceived focal planes as there are available waveguide/lens pairings. Both the light extracting optical elements of the waveguides and the focusing aspects of the lenses may be static (e.g., not dynamic or electro-active). In some alternative embodiments, either or both may be dynamic using electro-active features.

[0054] With continued reference to FIG. 4, the light extracting optical elements 440a, 438a, 436a, 434a, 432a may be configured to both redirect light out of their respective waveguides and to output this light with the appropriate amount of divergence or collimation for a particular depth plane associated with the waveguide. As a result, waveguides having different associated depth planes may have different configurations of light extracting optical elements, which output light with a different amount of divergence depending on the associated depth plane. In some embodiments, as discussed herein, the light extracting optical elements 440a, 438a, 436a, 434a, 432a may be volumetric or surface features, which may be configured to output light at specific angles. For example, the light extracting optical elements 440a, 438a, 436a, 434a, 432a may be volume holograms, surface holograms, and/or diffraction gratings. Light extracting optical elements, such as diffraction gratings, are described in U.S. Patent Publication No. 2015/0178939, published June 25, 2015, which is incorporated by reference herein in its entirety.

[0055] In some embodiments, the light extracting optical elements 440a, 438a, 436a, 434a, 432a are diffractive features that form a diffraction pattern, or “diffractive optical element” (also referred to herein as a “DOE”). Preferably, the DOE has a relatively low diffraction efficiency so that only a portion of the light of the beam is deflected away toward the eye 410 with each intersection of the DOE, while the rest continues to move through a waveguide via total internal reflection. The light carrying the image information can thus be divided into a number of related exit beams that exit the waveguide at a multiplicity of

locations and the result is a fairly uniform pattern of exit emission toward the eye 304 for this particular collimated beam bouncing around within a waveguide.

[0056] In some embodiments, one or more DOEs may be switchable between “on” state in which they actively diffract, and “off” state in which they do not significantly diffract. For instance, a switchable DOE may comprise a layer of polymer dispersed liquid crystal, in which microdroplets comprise a diffraction pattern in a host medium, and the refractive index of the microdroplets can be switched to substantially match the refractive index of the host material (in which case the pattern does not appreciably diffract incident light) or the microdroplet can be switched to an index that does not match that of the host medium (in which case the pattern actively diffracts incident light).

[0057] In some embodiments, the number and distribution of depth planes or depth of field may be varied dynamically based on the pupil sizes or orientations of the eyes of the viewer. Depth of field may change inversely with a viewer’s pupil size. As a result, as the sizes of the pupils of the viewer’s eyes decrease, the depth of field increases such that one plane that is not discernible because the location of that plane is beyond the depth of focus of the eye may become discernible and appear more in focus with reduction of pupil size and commensurate with the increase in depth of field. Likewise, the number of spaced apart depth planes used to present different images to the viewer may be decreased with the decreased pupil size. For example, a viewer may not be able to clearly perceive the details of both a first depth plane and a second depth plane at one pupil size without adjusting the accommodation of the eye away from one depth plane and to the other depth plane. These two depth planes may, however, be sufficiently in focus at the same time to the user at another pupil size without changing accommodation.

[0058] In some embodiments, the display system may vary the number of waveguides receiving image information based upon determinations of pupil size or orientation, or upon receiving electrical signals indicative of particular pupil size or orientation. For example, if the user’s eyes are unable to distinguish between two depth planes associated with two waveguides, then the controller 460 (which may be an embodiment of the local processing and data module 260) can be configured or programmed to cease providing image information to one of these waveguides. Advantageously, this may reduce the processing burden on the system, thereby increasing the responsiveness of the

system. In embodiments in which the DOEs for a waveguide are switchable between the on and off states, the DOEs may be switched to the off state when the waveguide does receive image information.

[0059] In some embodiments, it may be desirable to have an exit beam meet the condition of having a diameter that is less than the diameter of the eye of a viewer. However, meeting this condition may be challenging in view of the variability in size of the viewer's pupils. In some embodiments, this condition is met over a wide range of pupil sizes by varying the size of the exit beam in response to determinations of the size of the viewer's pupil. For example, as the pupil size decreases, the size of the exit beam may also decrease. In some embodiments, the exit beam size may be varied using a variable aperture.

[0060] The wearable system 400 can include an outward-facing imaging system 464 (e.g., a digital camera) that images a portion of the world 470. This portion of the world 470 may be referred to as the field of view (FOV) of a world camera and the imaging system 464 is sometimes referred to as an FOV camera. The entire region available for viewing or imaging by a viewer may be referred to as the field of regard (FOR). The FOR may include 4π steradians of solid angle surrounding the wearable system 400 because the wearer can move his body, head, or eyes to perceive substantially any direction in space. In other contexts, the wearer's movements may be more constricted, and accordingly the wearer's FOR may subtend a smaller solid angle. Images obtained from the outward-facing imaging system 464 can be used to track gestures made by the user (e.g., hand or finger gestures), detect objects in the world 470 in front of the user, and so forth.

[0061] The wearable system 400 can also include an inward-facing imaging system 466 (e.g., a digital camera), which observes the movements of the user, such as the eye movements and the facial movements. The inward-facing imaging system 466 may be used to capture images of the eye 410 to determine the size and/or orientation of the pupil of the eye 304. The inward-facing imaging system 466 can be used to obtain images for use in determining the direction the user is looking (e.g., eye pose) or for biometric identification of the user (e.g., via iris identification). In some embodiments, at least one camera may be utilized for each eye, to separately determine the pupil size or eye pose of each eye independently, thereby allowing the presentation of image information to each eye to be dynamically tailored to that eye. In some other embodiments, the pupil diameter or

orientation of only a single eye 410 (e.g., using only a single camera per pair of eyes) is determined and assumed to be similar for both eyes of the user. The images obtained by the inward-facing imaging system 466 may be analyzed to determine the user's eye pose or mood, which can be used by the wearable system 400 to decide which audio or visual content should be presented to the user. The wearable system 400 may also determine head pose (e.g., head position or head orientation) using sensors such as IMUs, accelerometers, gyroscopes, etc.

[0062] The wearable system 400 can include a user input device 466 by which the user can input commands to the controller 460 to interact with the wearable system 400. For example, the user input device 466 can include a trackpad, a touchscreen, a joystick, a multiple degree-of-freedom (DOF) controller, a capacitive sensing device, a game controller, a keyboard, a mouse, a directional pad (D-pad), a wand, a haptic device, a totem (e.g., functioning as a virtual user input device), and so forth. A multi-DOF controller can sense user input in some or all possible translations (e.g., left/right, forward/backward, or up/down) or rotations (e.g., yaw, pitch, or roll) of the controller. A multi-DOF controller which supports the translation movements may be referred to as a 3DOF while a multi-DOF controller which supports the translations and rotations may be referred to as 6DOF. In some cases, the user may use a finger (e.g., a thumb) to press or swipe on a touch-sensitive input device to provide input to the wearable system 400 (e.g., to provide user input to a user interface provided by the wearable system 400). The user input device 466 may be held by the user's hand during the use of the wearable system 400. The user input device 466 can be in wired or wireless communication with the wearable system 400.

[0063] FIG. 5 shows an example of exit beams outputted by a waveguide. One waveguide is illustrated, but it will be appreciated that other waveguides in the waveguide assembly 480 may function similarly, where the waveguide assembly 480 includes multiple waveguides. Light 520 is injected into the waveguide 432b at the input edge 432c of the waveguide 432b and propagates within the waveguide 432b by TIR. At points where the light 520 impinges on the DOE 432a, a portion of the light exits the waveguide as exit beams 510. The exit beams 510 are illustrated as substantially parallel but they may also be redirected to propagate to the eye 410 at an angle (e.g., forming divergent exit beams), depending on the depth plane associated with the waveguide 432b. It will be appreciated that

substantially parallel exit beams may be indicative of a waveguide with light extracting optical elements that outcouple light to form images that appear to be set on a depth plane at a large distance (e.g., optical infinity) from the eye 410. Other waveguides or other sets of light extracting optical elements may output an exit beam pattern that is more divergent, which would require the eye 410 to accommodate to a closer distance to bring it into focus on the retina and would be interpreted by the brain as light from a distance closer to the eye 410 than optical infinity.

[0064] FIG. 6 is a schematic diagram showing an optical system including a waveguide apparatus, an optical coupler subsystem to optically couple light to or from the waveguide apparatus, and a control subsystem, used in the generation of a multi-focal volumetric display, image, or light field. The optical system can include a waveguide apparatus, an optical coupler subsystem to optically couple light to or from the waveguide apparatus, and a control subsystem. The optical system can be used to generate a multi-focal volumetric, image, or light field. The optical system can include one or more primary planar waveguides 632a (only one is shown in FIG. 6) and one or more DOEs 632b associated with each of at least some of the primary waveguides 632a. The planar waveguides 632b can be similar to the waveguides 432b, 434b, 436b, 438b, 440b discussed with reference to FIG. 4. The optical system may employ a distribution waveguide apparatus to relay light along a first axis (vertical or Y-axis in view of FIG. 6), and expand the light's effective exit pupil along the first axis (e.g., Y-axis). The distribution waveguide apparatus may, for example, include a distribution planar waveguide 622b and at least one DOE 622a (illustrated by double dash-dot line) associated with the distribution planar waveguide 622b. The distribution planar waveguide 622b may be similar or identical in at least some respects to the primary planar waveguide 632b, having a different orientation therefrom. Likewise, at least one DOE 622a may be similar to or identical in at least some respects to the DOE 632a. For example, the distribution planar waveguide 622b or DOE 622a may be comprised of the same materials as the primary planar waveguide 632b or DOE 632a, respectively. Embodiments of the optical display system 600 shown in FIG. 6 can be integrated into the wearable system 200 shown in FIG. 2.

[0065] The relayed and exit-pupil expanded light may be optically coupled from the distribution waveguide apparatus into the one or more primary planar waveguides 632b.

The primary planar waveguide 632b can relay light along a second axis, preferably orthogonal to first axis (e.g., horizontal or X-axis in view of FIG. 6). Notably, the second axis can be a non-orthogonal axis to the first axis. The primary planar waveguide 632b expands the light's effective exit pupil along that second axis (e.g., X-axis). For example, the distribution planar waveguide 622b can relay and expand light along the vertical or Y-axis, and pass that light to the primary planar waveguide 632b which can relay and expand light along the horizontal or X-axis.

[0066] The optical system may include one or more sources of colored light (e.g., red, green, and blue laser light) 610 which may be optically coupled into a proximal end of a single mode optical fiber 640. A distal end of the optical fiber 640 may be threaded or received through a hollow tube 642 of piezoelectric material. The distal end protrudes from the tube 642 as fixed-free flexible cantilever 644. The piezoelectric tube 642 can be associated with four quadrant electrodes (not illustrated). The electrodes may, for example, be plated on the outside, outer surface or outer periphery or diameter of the tube 642. A core electrode (not illustrated) may also be located in a core, center, inner periphery or inner diameter of the tube 642.

[0067] Drive electronics 650, for example electrically coupled via wires 660, drive opposing pairs of electrodes to bend the piezoelectric tube 642 in two axes independently. The protruding distal tip of the optical fiber 644 has mechanical modes of resonance. The frequencies of resonance can depend upon a diameter, length, and material properties of the optical fiber 644. By vibrating the piezoelectric tube 642 near a first mode of mechanical resonance of the fiber cantilever 644, the fiber cantilever 644 can be caused to vibrate, and can sweep through large deflections.

[0068] By stimulating resonant vibration in two axes, the tip of the fiber cantilever 644 is scanned biaxially in an area filling two-dimensional (2D) scan. By modulating an intensity of light source(s) 610 in synchrony with the scan of the fiber cantilever 644, light emerging from the fiber cantilever 644 can form an image. Descriptions of such a set up are provided in U.S. Patent Publication No. 2014/0003762, which is incorporated by reference herein in its entirety.

[0069] A component of an optical coupler subsystem can collimate the light emerging from the scanning fiber cantilever 644. The collimated light can be reflected by

mirrored surface 648 into the narrow distribution planar waveguide 622b which contains the at least one diffractive optical element (DOE) 622a. The collimated light can propagate vertically (relative to the view of FIG. 6) along the distribution planar waveguide 622b by TIR, and in doing so repeatedly intersects with the DOE 622a. The DOE 622a preferably has a low diffraction efficiency. This can cause a fraction (e.g., 10%) of the light to be diffracted toward an edge of the larger primary planar waveguide 632b at each point of intersection with the DOE 622a, and a fraction of the light to continue on its original trajectory down the length of the distribution planar waveguide 622b via TIR.

[0070] At each point of intersection with the DOE 622a, additional light can be diffracted toward the entrance of the primary waveguide 632b. By dividing the incoming light into multiple outcoupled sets, the exit pupil of the light can be expanded vertically by the DOE 622a in the distribution planar waveguide 622b. This vertically expanded light coupled out of distribution planar waveguide 622b can enter the edge of the primary planar waveguide 632b.

[0071] Light entering primary waveguide 632b can propagate horizontally (relative to the view of FIG. 6) along the primary waveguide 632b via TIR. As the light intersects with DOE 632a at multiple points as it propagates horizontally along at least a portion of the length of the primary waveguide 632b via TIR. The DOE 632a may advantageously be designed or configured to have a phase profile that is a summation of a linear diffraction pattern and a radially symmetric diffractive pattern, to produce both deflection and focusing of the light. The DOE 632a may advantageously have a low diffraction efficiency (e.g., 10%), so that only a portion of the light of the beam is deflected toward the eye of the view with each intersection of the DOE 632a while the rest of the light continues to propagate through the primary waveguide 632b via TIR.

[0072] At each point of intersection between the propagating light and the DOE 632a, a fraction of the light is diffracted toward the adjacent face of the primary waveguide 632b allowing the light to escape the TIR, and emerge from the face of the primary waveguide 632b. In some embodiments, the radially symmetric diffraction pattern of the DOE 632a additionally imparts a focus level to the diffracted light, both shaping the light wavefront (e.g., imparting a curvature) of the individual beam as well as steering the beam at an angle that matches the designed focus level.

[0073] Accordingly, these different pathways can cause the light to be coupled out of the primary planar waveguide 632b by a multiplicity of DOEs 632a at different angles, focus levels, or yielding different fill patterns at the exit pupil. Different fill patterns at the exit pupil can be beneficially used to create a light field display with multiple depth planes. Each layer in the waveguide assembly or a set of layers (e.g., 3 layers) in the stack may be employed to generate a respective color (e.g., red, blue, green). Thus, for example, a first set of three adjacent layers may be employed to respectively produce red, blue and green light at a first focal depth. A second set of three adjacent layers may be employed to respectively produce red, blue and green light at a second focal depth. Multiple sets may be employed to generate a full 3D or 4D color image light field with various focal depths.

Other Components of the Wearable System

[0074] In many implementations, the wearable system may include other components in addition or in alternative to the components of the wearable system described above. The wearable system may, for example, include one or more haptic devices or components. The haptic devices or components may be operable to provide a tactile sensation to a user. For example, the haptic devices or components may provide a tactile sensation of pressure or texture when touching virtual content (e.g., virtual objects, virtual tools, other virtual constructs). The tactile sensation may replicate a feel of a physical object which a virtual object represents, or may replicate a feel of an imagined object or character (e.g., a dragon) which the virtual content represents. In some implementations, haptic devices or components may be worn by the user (e.g., a user wearable glove). In some implementations, haptic devices or components may be held by the user.

[0075] The wearable system may, for example, include one or more physical objects which are manipulable by the user to allow input or interaction with the wearable system. These physical objects may be referred to herein as totems. Some totems may take the form of inanimate objects, such as for example, a piece of metal or plastic, a wall, a surface of table. In certain implementations, the totems may not actually have any physical input structures (e.g., keys, triggers, joystick, trackball, rocker switch). Instead, the totem may simply provide a physical surface, and the wearable system may render a user interface so as to appear to a user to be on one or more surfaces of the totem. For example, the

wearable system may render an image of a computer keyboard and trackpad to appear to reside on one or more surfaces of a totem. For example, the wearable system may render a virtual computer keyboard and virtual trackpad to appear on a surface of a thin rectangular plate of aluminum which serves as a totem. The rectangular plate does not itself have any physical keys or trackpad or sensors. However, the wearable system may detect user manipulation or interaction or touches with the rectangular plate as selections or inputs made via the virtual keyboard or virtual trackpad. The user input device 466 (shown in FIG. 4) may be an embodiment of a totem, which may include a trackpad, a touchpad, a trigger, a joystick, a trackball, a rocker or virtual switch, a mouse, a keyboard, a multi-degree-of-freedom controller, or another physical input device. A user may use the totem, alone or in combination with poses, to interact with the wearable system or other users.

[0076] Examples of haptic devices and totems usable with the wearable devices, HMD, and display systems of the present disclosure are described in U.S. Patent Publication No. 2015/0016777, which is incorporated by reference herein in its entirety.

Example Wearable Systems, Environments, and Interfaces

[0077] A wearable system may employ various mapping related techniques in order to achieve high depth of field in the rendered light fields. In mapping out the virtual world, it is advantageous to know all the features and points in the real world to accurately portray virtual objects in relation to the real world. To this end, FOV images captured from users of the wearable system can be added to a world model by including new pictures that convey information about various points and features of the real world. For example, the wearable system can collect a set of map points (such as 2D points or 3D points) and find new map points to render a more accurate version of the world model. The world model of a first user can be communicated (e.g., over a network such as a cloud network) to a second user so that the second user can experience the world surrounding the first user.

[0078] FIG. 7 is a block diagram of an example of an MR environment 700. The MR environment 700 may be configured to receive input (e.g., visual input 702 from the user's wearable system, stationary input 704 such as room cameras, sensory input 706 from various sensors, gestures, totems, eye tracking, user input from the user input device 466 etc.) from one or more user wearable systems (e.g., wearable system 200 or display system

220) or stationary room systems (e.g., room cameras, etc.). The wearable systems can use various sensors (e.g., accelerometers, gyroscopes, temperature sensors, movement sensors, depth sensors, GPS sensors, inward-facing imaging system, outward-facing imaging system, etc.) to determine the location and various other attributes of the environment of the user. This information may further be supplemented with information from stationary cameras in the room that may provide images or various cues from a different point of view. The image data acquired by the cameras (such as the room cameras and/or the cameras of the outward-facing imaging system) may be reduced to a set of mapping points.

[0079] One or more object recognizers 708 can crawl through the received data (e.g., the collection of points) and recognize or map points, tag images, attach semantic information to objects with the help of a map database 710. The map database 710 may comprise various points collected over time and their corresponding objects. The various devices and the map database can be connected to each other through a network (e.g., LAN, WAN, etc.) to access the cloud.

[0080] Based on this information and collection of points in the map database, the object recognizers 708a to 708n may recognize objects in an environment. For example, the object recognizers can recognize faces, persons, windows, walls, user input devices, televisions, documents (e.g., travel tickets, driver's license, passport as described in the security examples herein), other objects in the user's environment, etc. One or more object recognizers may be specialized for object with certain characteristics. For example, the object recognizer 708a may be used to recognizer faces, while another object recognizer may be used recognize documents.

[0081] The object recognitions may be performed using a variety of computer vision techniques. For example, the wearable system can analyze the images acquired by the outward-facing imaging system 464 (shown in FIG. 4) to perform scene reconstruction, event detection, video tracking, object recognition (e.g., persons or documents), object pose estimation, facial recognition (e.g., from a person in the environment or an image on a document), learning, indexing, motion estimation, or image analysis (e.g., identifying indicia within documents such as photos, signatures, identification information, travel information, etc.), and so forth. One or more computer vision algorithms may be used to perform these tasks. Non-limiting examples of computer vision algorithms include: Scale-invariant feature

transform (SIFT), speeded up robust features (SURF), oriented FAST and rotated BRIEF (ORB), binary robust invariant scalable keypoints (BRISK), fast retina keypoint (FREAK), Viola-Jones algorithm, Eigenfaces approach, Lucas-Kanade algorithm, Horn-Schunk algorithm, Mean-shift algorithm, visual simultaneous location and mapping (vSLAM) techniques, a sequential Bayesian estimator (e.g., Kalman filter, extended Kalman filter, etc.), bundle adjustment, Adaptive thresholding (and other thresholding techniques), Iterative Closest Point (ICP), Semi Global Matching (SGM), Semi Global Block Matching (SGBM), Feature Point Histograms, various machine learning algorithms (such as e.g., support vector machine, k-nearest neighbors algorithm, Naive Bayes, neural network (including convolutional or deep neural networks), or other supervised/unsupervised models, etc.), and so forth.

[0082] The object recognitions can additionally or alternatively be performed by a variety of machine learning algorithms. Once trained, the machine learning algorithm can be stored by the HMD. Some examples of machine learning algorithms can include supervised or non-supervised machine learning algorithms, including regression algorithms (such as, for example, Ordinary Least Squares Regression), instance-based algorithms (such as, for example, Learning Vector Quantization), decision tree algorithms (such as, for example, classification and regression trees), Bayesian algorithms (such as, for example, Naive Bayes), clustering algorithms (such as, for example, k-means clustering), association rule learning algorithms (such as, for example, a-priori algorithms), artificial neural network algorithms (such as, for example, Perceptron), deep learning algorithms (such as, for example, Deep Boltzmann Machine, or deep neural network), dimensionality reduction algorithms (such as, for example, Principal Component Analysis), ensemble algorithms (such as, for example, Stacked Generalization), and/or other machine learning algorithms. In some embodiments, individual models can be customized for individual data sets. For example, the wearable device can generate or store a base model. The base model may be used as a starting point to generate additional models specific to a data type (e.g., a particular user in the telepresence session), a data set (e.g., a set of additional images obtained of the user in the telepresence session), conditional situations, or other variations. In some embodiments, the wearable HMD can be configured to utilize a plurality of techniques to generate models

for analysis of the aggregated data. Other techniques may include using pre-defined thresholds or data values.

[0083] Based on this information and collection of points in the map database, the object recognizers 708a to 708n may recognize objects and supplement objects with semantic information to give life to the objects. For example, if the object recognizer recognizes a set of points to be a door, the system may attach some semantic information (e.g., the door has a hinge and has a 90 degree movement about the hinge). If the object recognizer recognizes a set of points to be a mirror, the system may attach semantic information that the mirror has a reflective surface that can reflect images of objects in the room. The semantic information can include affordances of the objects as described herein. For example, the semantic information may include a normal of the object. The system can assign a vector whose direction indicates the normal of the object. Over time the map database grows as the system (which may reside locally or may be accessible through a wireless network) accumulates more data from the world. Once the objects are recognized, the information may be transmitted to one or more wearable systems. For example, the MR environment 700 may include information about a scene happening in California. The environment 700 may be transmitted to one or more users in New York. Based on data received from an FOV camera and other inputs, the object recognizers and other software components can map the points collected from the various images, recognize objects etc., such that the scene may be accurately "passed over" to a second user, who may be in a different part of the world. The environment 700 may also use a topological map for localization purposes.

[0084] FIG. 8 is a process flow diagram of an example of a method 800 of rendering virtual content in relation to recognized objects. The method 800 describes how a virtual scene may be presented to a user of the wearable system. The user may be geographically remote from the scene. For example, the user may be in New York, but may want to view a scene that is presently going on in California, or may want to go on a walk with a friend who resides in California.

[0085] At block 810, the wearable system may receive input from the user and other users regarding the environment of the user. This may be achieved through various input devices, and knowledge already possessed in the map database. The user's FOV camera, sensors, GPS, eye tracking, etc., convey information to the system at block 810. The

system may determine sparse points based on this information at block 820. The sparse points may be used in determining pose data (e.g., head pose, eye pose, body pose, or hand gestures) that can be used in displaying and understanding the orientation and position of various objects in the user's surroundings. The object recognizers 708a-708n may crawl through these collected points and recognize one or more objects using a map database at block 830. This information may then be conveyed to the user's individual wearable system at block 840, and the desired virtual scene may be accordingly displayed to the user at block 850. For example, the desired virtual scene (e.g., user in CA) may be displayed at the appropriate orientation, position, etc., in relation to the various objects and other surroundings of the user in New York.

[0086] FIG. 9 is a block diagram of another example of a wearable system. In this example, the wearable system 900 comprises a map, which may include map data for the world. The map may partly reside locally on the wearable system, and may partly reside at networked storage locations accessible by wired or wireless network (e.g., in a cloud system). A pose process 910 may be executed on the wearable computing architecture (e.g., processing module 260 or controller 460) and utilize data from the map to determine position and orientation of the wearable computing hardware or user. Pose data may be computed from data collected on the fly as the user is experiencing the system and operating in the world. The data may comprise images, data from sensors (such as inertial measurement units, which generally comprise accelerometer and gyroscope components) and surface information pertinent to objects in the real or virtual environment.

[0087] A sparse point representation may be the output of a simultaneous localization and mapping (e.g., SLAM or vSLAM, referring to a configuration wherein the input is images/visual only) process. The system can be configured to not only find out where in the world the various components are, but what the world is made of. Pose may be a building block that achieves many goals, including populating the map and using the data from the map.

[0088] In one embodiment, a sparse point position may not be completely adequate on its own, and further information may be needed to produce a multifocal AR, VR, or MR experience. Dense representations, generally referring to depth map information, may be utilized to fill this gap at least in part. Such information may be computed from a process

referred to as Stereo 940, wherein depth information is determined using a technique such as triangulation or time-of-flight sensing. Image information and active patterns (such as infrared patterns created using active projectors) may serve as input to the Stereo process 940. A significant amount of depth map information may be fused together, and some of this may be summarized with a surface representation. For example, mathematically definable surfaces may be efficient (e.g., relative to a large point cloud) and digestible inputs to other processing devices like game engines. Thus, the output of the stereo process (e.g., a depth map) 940 may be combined in the fusion process 930. Pose 950 may be an input to this fusion process 930 as well, and the output of fusion 930 becomes an input to populating the map process 920. Sub-surfaces may connect with each other, such as in topographical mapping, to form larger surfaces, and the map becomes a large hybrid of points and surfaces.

[0089] To resolve various aspects in a mixed reality process 960, various inputs may be utilized. For example, in the embodiment depicted in FIG. 9, Game parameters may be inputs to determine that the user of the system is playing a monster battling game with one or more monsters at various locations, monsters dying or running away under various conditions (such as if the user shoots the monster), walls or other objects at various locations, and the like. The world map may include information regarding where such objects are relative to each other, to be another valuable input to mixed reality. Pose relative to the world becomes an input as well and plays a key role to almost any interactive system.

[0090] Controls or inputs from the user are another input to the wearable system 900. As described herein, user inputs can include visual input, gestures, totems, audio input, sensory input, etc. In order to move around or play a game, for example, the user may need to instruct the wearable system 900 regarding what he or she wants to do. Beyond just moving oneself in space, there are various forms of user controls that may be utilized. In one embodiment, a totem (e.g. a user input device), or an object such as a toy gun may be held by the user and tracked by the system. The system preferably will be configured to know that the user is holding the item and understand what kind of interaction the user is having with the item (e.g., if the totem or object is a gun, the system may be configured to understand location and orientation, as well as whether the user is clicking a trigger or other sensed button or element which may be equipped with a sensor, such as an IMU, which may assist

in determining what is going on, even when such activity is not within the field of view of any of the cameras.)

[0091] Hand gesture tracking or recognition may also provide input information. The wearable system 900 may be configured to track and interpret hand gestures for button presses, for gesturing left or right, stop, grab, hold, etc. For example, in one configuration, the user may want to flip through emails or a calendar in a non-gaming environment, or do a "fist bump" with another person or player. The wearable system 900 may be configured to leverage a minimum amount of hand gesture, which may or may not be dynamic. For example, the gestures may be simple static gestures like open hand for stop, thumbs up for ok, thumbs down for not ok; or a hand flip right, or left, or up/down for directional commands.

[0092] Eye tracking is another input (e.g., tracking where the user is looking to control the display technology to render at a specific depth or range). In one embodiment, vergence of the eyes may be determined using triangulation, and then using a vergence/accommodation model developed for that particular person, accommodation may be determined. Eye tracking can be performed by the eye camera(s) to determine eye gaze (e.g., direction or orientation of one or both eyes). Other techniques can be used for eye tracking such as, e.g., measurement of electrical potentials by electrodes placed near the eye(s) (e.g., electrooculography).

[0093] Voice recognition can be another input, which can be used alone or in combination with other inputs (e.g., totem tracking, eye tracking, gesture tracking, etc.). The system 900 can include an audio sensor (e.g., a microphone) that receives an audio stream from the environment. The received audio stream can be processed (e.g., by processing modules 260, 270 or central server 1650) to recognize a user's voice (from other voices or background audio), to extract commands, parameters, etc. from the audio stream. For example, the system 900 may identify from an audio stream that the phrase "show me your identification" was said, identify that this phrase was said by the wearer of the system 900 (e.g., a security inspector rather than another person in the inspector's environment), and extract from the phrase and the context of the situation (e.g., a security checkpoint) that there is an executable command to be performed (e.g., computer vision analysis of something in the wearer's FOV) and an object for which the command is to be performed on ("your

identification"). The system 900 can incorporate speaker recognition technology to determine who is speaking (e.g., whether the speech is from the wearer of the HMD or another person or voice (e.g., a recorded voice transmitted by a loudspeaker in the environment)) as well as speech recognition technology to determine what is being said. Voice recognition techniques can include frequency estimation, hidden Markov models, Gaussian mixture models, pattern matching algorithms, neural networks, matrix representation, Vector Quantization, speaker diarisation, decision trees, and dynamic time warping (DTW) technique. Voice recognition techniques can also include anti-speaker techniques, such as cohort models, and world models. Spectral features may be used in representing speaker characteristics.

[0094] With regard to the camera systems, the example wearable system 900 shown in FIG. 9 can include three pairs of cameras: a relative wide FOV or passive SLAM pair of cameras arranged to the sides of the user's face, a different pair of cameras oriented in front of the user to handle the stereo imaging process 940 and also to capture hand gestures and totem/object tracking in front of the user's face. The FOV cameras and the pair of cameras for the stereo process 940 may be a part of the outward-facing imaging system 464 (shown in FIG. 4). The wearable system 900 can include eye tracking cameras (which may be a part of an inward-facing imaging system 462 shown in FIG. 4) oriented toward the eyes of the user in order to triangulate eye vectors and other information. The wearable system 900 may also comprise one or more textured light projectors (such as infrared (IR) projectors) to inject texture into a scene.

[0095] FIG. 10 is a process flow diagram of an example of a method 1000 for determining user input to a wearable system. In this example, the user may interact with a totem. The user may have multiple totems. For example, the user may have designated one totem for a social media application, another totem for playing games, etc. At block 1010, the wearable system may detect a motion of a totem. The movement of the totem may be recognized through the outward-facing imaging system or may be detected through sensors (e.g., haptic glove, image sensors, hand tracking devices, eye-tracking cameras, head pose sensors, etc.).

[0096] Based at least partly on the detected gesture, eye pose, head pose, or input through the totem, the wearable system detects a position, orientation, or movement of the

totem (or the user's eyes or head or gestures) with respect to a reference frame, at block 1020. The reference frame may be a set of map points based on which the wearable system translates the movement of the totem (or the user) to an action or command. At block 1030, the user's interaction with the totem is mapped. Based on the mapping of the user interaction with respect to the reference frame 1020, the system determines the user input at block 1040.

[0097] For example, the user may move a totem or physical object back and forth to signify turning a virtual page and moving on to a next page or moving from one user interface (UI) display screen to another UI screen. As another example, the user may move their head or eyes to look at different real or virtual objects in the user's FOR. If the user's gaze at a particular real or virtual object is longer than a threshold time, the real or virtual object may be selected as the user input. In some implementations, the vergence of the user's eyes can be tracked and an accommodation/vergence model can be used to determine the accommodation state of the user's eyes, which provides information on a depth plane on which the user is focusing. In some implementations, the wearable system can use ray casting techniques to determine which real or virtual objects are along the direction of the user's head pose or eye pose. In various implementations, the ray casting techniques can include casting thin, pencil rays with substantially little transverse width or casting rays with substantial transverse width (e.g., cones or frustums).

[0098] The user interface may be projected by the display system as described herein (such as the display 220 in FIG. 2). It may also be displayed using a variety of other techniques such as one or more projectors. The projectors may project images onto a physical object such as a canvas or a globe. Interactions with user interface may be tracked using one or more cameras external to the system or part of the system (such as, e.g., using the inward-facing imaging system 462 or the outward-facing imaging system 464).

[0099] FIG. 11 is a process flow diagram of an example of a method 1100 for interacting with a virtual user interface. The method 1100 may be performed by the wearable system described herein. Embodiments of the method 1100 can be used by the wearable system to detect persons or documents in the FOV of the wearable system.

[0100] At block 1110, the wearable system may identify a particular UI. The type of UI may be predetermined by the user. The wearable system may identify that a particular UI needs to be populated based on a user input (e.g., gesture, visual data, audio data, sensory

data, direct command, etc.). The UI can be specific to a security scenario where the wearer of the system is observing users who present documents to the wearer (e.g., at a travel checkpoint). At block 1120, the wearable system may generate data for the virtual UI. For example, data associated with the confines, general structure, shape of the UI etc., may be generated. In addition, the wearable system may determine map coordinates of the user's physical location so that the wearable system can display the UI in relation to the user's physical location. For example, if the UI is body centric, the wearable system may determine the coordinates of the user's physical stance, head pose, or eye pose such that a ring UI can be displayed around the user or a planar UI can be displayed on a wall or in front of the user. In the security context described herein, the UI may be displayed as if the UI were surrounding the traveler who is presenting documents to the wearer of the system, so that the wearer can readily view the UI while looking at the traveler and the traveler's documents. If the UI is hand centric, the map coordinates of the user's hands may be determined. These map points may be derived through data received through the FOV cameras, sensory input, or any other type of collected data.

[0101] At block 1130, the wearable system may send the data to the display from the cloud or the data may be sent from a local database to the display components. At block 1140, the UI is displayed to the user based on the sent data. For example, a light field display can project the virtual UI into one or both of the user's eyes. Once the virtual UI has been created, the wearable system may simply wait for a command from the user to generate more virtual content on the virtual UI at block 1150. For example, the UI may be a body centric ring around the user's body or the body of a person in the user's environment (e.g., a traveler). The wearable system may then wait for the command (a gesture, a head or eye movement, voice command, input from a user input device, etc.), and if it is recognized (block 1160), virtual content associated with the command may be displayed to the user (block 1170).

[0102] Additional examples of wearable systems, UIs, and user experiences (UX) are described in U.S. Patent Publication No. 2015/0016777, which is incorporated by reference herein in its entirety.

Examples of Rendering a Virtual Remote Control

[0103] Besides interacting with the user's virtual environment, a user can also control a physical object in the user's environment using a wearable device whose display 230 may be configured to render a virtual remote. As further described herein, the wearable system can identify a target device such as a physical remote or a parent device which is controlled by the physical remote. The wearable system can generate a virtual remote based on the characteristics of the target device or the physical remote associated with the target device, present the virtual remote to the user in an AR/MR environment, and allow the user to control the parent device with the virtual remote. Although many of the examples described herein are in the context of using a virtual remote control to control a physical device (e.g., a physical television, a DVD or CD player or an audio system, a thermostat, a coffee machine, refrigerator, or other appliance, an electric lamp or light bulb, a door entry system or home security system, etc.), this is for illustration and is not a limitation. Embodiments of the virtual remote control can be used, additionally or alternatively, to control virtual devices. For example, a user of the wearable system can use a virtual remote control to control a virtual television that is rendered by the wearable device and displayed to the user.

Examples of Initiating a Selection Event on a Virtual Remote

[0104] FIG. 12 illustrates an example of a physical environment as perceived by a user of a wearable system. The example environment 1200 includes a living room of a user's home. The environment 1200 has physical objects such as a television (TV) 1210, a physical remote control 1220 (sometimes simply referred to as a remote), a TV stand 1230, and a window 1240. While the user is wearing a wearable device, the user can perceive the physical objects and interact with the physical objects. For example, the user may watch the TV 1210 while wearing the wearable device. The user can control the TV using the physical remote 1220. For example, the user can control the physical remote 1220 to turn the TV 1210 on/off or change the channel or volume of the TV 1210. The user can also interact with the TV 1210 using a virtual remote. The virtual remote may be generated based on the functions of the physical remote 1220. For example, the virtual remote may emulate some or all of the

functions of the physical remote 1220 (and may provide additional or alternative functionality as well).

[0105] The user can initiate a selection event on a virtual remote associated with a target device. The selection event may cause a user interface action associated with the virtual remote to be performed. For example, upon initiation of the selection event, the wearable device may render the virtual remote in the user's FOV. The target device may be a physical remote which the wearable system can emulate to generate a virtual remote, which emulates functions of the physical remote. The target device may also be a parent device such as a TV, a coffeemaker, a thermostat, or other appliances or electronic devices, etc., which may be controlled by a physical remote or a virtual remote. The user can initiate the selection event by actuating a user input device such as, e.g., clicking on a mouse, tapping on a touch pad, swiping on a touch screen, hovering over or touching a capacitive button, pressing a key on a keyboard or a game controller (e.g., a 5-way d-pad), pointing a joystick, wand or totem toward the object, pressing a button on a remote control, other interactions with a user input device, etc. The user may also initiate the selection event using head, eye, or body poses, such as e.g., by gazing or pointing at a target object for a period of time.

[0106] To initiate a selection event on the virtual remote, the user can indicate a selection of the target device associated with the virtual remote. For example, the user can indicate a physical remote to select the virtual remote associated with the physical remote. As shown in FIG. 12, if the user wants to interact with a virtual remote which is based on the functions of the physical remote 1220, the user may indicate the physical remote 1220 by hand gestures such as touching, pointing with finger, visually enclosing the objects by, for example, pinching, or using other hand gestures. As an example, the user may point in the direction of the physical remote 1220 for an extended period of time. As another example, the user may select a virtual remote associated with the physical remote 1220 by making a hand gesture for grabbing the physical remote 1220. The user may also indicate the physical remote 1220 using a user input device (such as, e.g., the user input device 466 shown in FIG. 4). For example, the user may point at the physical remote using a stylus.

[0107] The user can also select a virtual remote by selecting the parent device that the virtual remote controls. The user can use the hand gestures and actuate the user input device for such selection. FIG. 13 illustrates an example of initiating a selection event

associated with a virtual remote using a hand gesture. The user can perceive, via the wearable system, the environment 1300 as shown in FIG. 13. The environment 1300 can include physical objects such as the TV 1210, the physical remote 1220, the TV stand 1230, and the window 1240. As illustrated, the user's left arm 1250a is making a pinch gesture indicating the TV 1210. The wearable system may recognize this pinch gesture as a command to render and present the virtual remote 1260 associated with the TV 1210 to the user. As another example, if the user wants to select a virtual remote for controlling the TV 1210, the user can use a body pose (such as grabbing the TV 1210 or pointing at the TV 1210) to indicate a selection of the TV 1210.

[0108] In addition to or in alternative to hand gestures and actuations of user input device, a user can also select a virtual remote using eye gaze. The wearable system may perform a cone cast based on the direction of the user's eye gaze (as tracked by the inward-facing imaging system 462). The wearable system can identify an object which intersects the user's direction of gaze and render a remote associated with such object. For example, the wearable system can identify that the TV 1210 based on the user's direction of gaze, and thus present a virtual remote associated with the TV 1210. The virtual remote may be generated based on the physical remote 1220, where the virtual remote may have similar functions and control panels as the physical remote 1220.

Examples of Recognizing a Target Device

[0109] The wearable system can recognize a target device, such as the physical remote 1220 or the parent device (such as the TV 1210) using a variety of techniques, alone or in combination. As an example, the wearable system may query the target device directly. This is possible where the device is capable of identifying itself, e.g., by RF wireless protocols such as Wi-Fi or Bluetooth, and where protocols already exist to support such a query. For example, the wearable system can pair with the TV 1210 or the physical remote 1220 using Bluetooth technology.

[0110] The wearable system may also identify the target device visually based on the characteristics of the target device. For example, the wearable system may use the outward-facing imaging system to acquire an image of the target device and identify the type target device using computer vision algorithms such as feature keypoint methods (e.g. SIFT)

combined with a Bag of Words type search, or through the application of an artificial neural network (e.g. "Alex-net") capable of categorizing devices. The wearable system may identify the target device in a hierarchical fashion. For example, a target device may first be identified as a television using a first algorithm, and then as a specific product, e.g., a ManuCo 55" 530U Series Premium 4K UHD Slim Direct-Lit LED TV, using a second algorithm.

[0111] The user can also specify the target device using other poses or user input device. For example, the user may identify a TV with a waving gesture while identifying another type of device (e.g., a thermostat) by pointing his finger at the thermostat. In some embodiments, the user can identify a type of the target device (such as TV or thermostat) using a pose. The wearable system may perform other techniques to further identify a subtype of the type indicated by the user's pose. For example, the user may use a waving gesture to identify that the target device is a TV. The wearable system may run a computer vision algorithm to further identify that the TV is a ManuCo 55" 530U Series Premium 4K UHD Slim Direct-Lit LED TV.

[0112] The wearable system may identify the target device by visual inspection of the existing target device for one or more characteristics. For example, the target device may include an optically-readable label such as a bar code or a QR code. The wearable system may include an optical sensor (e.g., an outward facing camera) for scanning or imaging the label and extracting information from the label. The extracted information may include device identifiers such as manufacturer, type of the device, etc. The label may be stuck to or imprinted onto the target device. For example, a TV may have a sticker or placard explicitly indicating the make, model, serial number, etc. of the TV. As another example, the TV's brand may be imprinted on a surface of the TV panel.

[0113] In some embodiments, when the target device is the physical remote and the wearable system has recognized the physical remote, the wearable system may be configured to automatically identify the associated parent device. For example, once the wearable system has recognized the physical remote 1220 for the TV 1210, the wearable system can retrieve or search for information associated with the parent device such as the appearance of the parent device, the communication protocols associated with the parent device, or the brand/model of the parent device. The wearable system can search and

recognize the parent device in the user's environment based on the images acquired by the outward-facing imaging system. The recognition may be achieved using the computer vision algorithms described herein. The wearable system can also send a query signal to the environment and identify the parent device based on the received response to the query signal. These techniques for identifying the parent device based on the physical remote can also be applied when the target device is the parent device. For example, the wearable system can use similar techniques to identify a physical remote in the user's environment based on a recognized parent device.

[0114] The parent device may be associated with multiple physical remotes in some situations. For example, a gaming station may be associated with multiple game controllers. The wearable system may prompt a notification to the user indicating that there is more than one virtual remote associated with the parent device (where each virtual remote may correspond to a physical remote). The wearable system may provide the user the option to select one virtual remote among the multiple virtual remotes with which the user wants to use.

[0115] In another example, the wearable system may be uncertain whether the parent device or the physical remote selected by the user is of a certain type. For example, some TVs may be identical in appearance but compatible with different types of physical remotes. Similarly, certain physical remotes may appear to have the same button or shape, but serve different parent devices. The wearable system may present a list of possible options of the virtual remote. The user can then select the correct virtual remote among the list of possible options.

[0116] Additionally or alternatively, the wearable system can be configured to allow a user to select a virtual remote associated with a target device using a virtual user interface of the wearable system. For example, the wearable system may present a list of available virtual remotes in a menu. The user can select the virtual remote with which to interact with from the menu. The wearable system can also present a menu including a list of parent devices or physical remotes associated with the parent devices for which the wearable system can provide the virtual remote. For example, the parent devices or physical remotes on the list may be previously identified by the wearable system using one or more object

recognizers 2208a – 2208n shown in FIG. 7. The user can accordingly select the virtual remote by selecting the parent device or the physical remote from the list.

[0117] The virtual remotes themselves may be stored in a data store such as the remote data repository 280 shown in FIG. 2. The wearable system can recognize the target device using the techniques described herein and fetch the specifications for the virtual remote based on the target device. For example, the wearable system may submit a description of the recognized target device to the remote processing module 270 to retrieve the specifications for the associated virtual remote. In some embodiments, the processing module may redirect the request for the specifications to another computing device (such as a server for maintaining the specification for a specific type of device) based on the characteristics of the target device such as brand, type, year of manufacture, etc.

[0118] The specifications may include instructions on how to render the virtual remote (e.g. specifying the layout, buttons and other controls, etc.), the communications channel needed to emulate the remote (e.g. from the IRDA specification), the actual codes to be emulated on that channel (e.g. the exact IR pulse sequence associated with the selection of “Channel 2”, etc.), and so on. The specification may also include a mapping of a control element to a corresponding function performed on the parent device. For example, an actuation of an upward arrow button (on the physical and the virtual remotes) may be mapped to a function of increasing volume by the parent device.

Examples of Rendering a Virtual Remote

[0119] The virtual remote may mirror some or all of the functions or at least a portion of the layout of the physical remote. This may make it easier for a user, who is familiar with the functionality or layout of the physical remote, to operate the virtual remote. Upon the selection of the virtual remote (e.g., by selecting the parent device or the physical remote associated with the virtual remote), the wearable device can render the virtual remote so that it appears visually near the user. For example, as shown in FIG. 14, the wearable device can render a virtual control panel 1264 of the virtual remote 1260 within the user’s arm’s reach (such as, e.g., between 10cm and 1m from the user) so that the user can conveniently interact with the virtual remote. In some embodiments, as the user moves around in his environment, the virtual remote may accordingly move with the user. For

example, the wearable device may present the virtual remote at a certain distance from the user regardless of the user's current position in his environment.

[0120] Additionally or alternatively, the wearable device can render the virtual remote near the parent device or the physical remote. For example, the wearable device can render a virtual remote for controlling the TV where the virtual control panel of the virtual remote is located next to the TV (e.g., as shown in FIGS. 13 and 14). In some embodiments, the rendering location of the virtual remote may be based on the type of the parent device. For example, if the parent device is a TV, the wearable device may render the virtual remote near the TV because the user's interaction with the TV is likely to happen when the user is in the vicinity of the TV. As another example, if the parent device is a thermostat, the wearable device may render a virtual remote near the user (which may be far away from the thermostat). A user may control a thermostat upstairs in a house using the virtual remote even though the user is downstairs.

[0121] The wearable device can render the virtual remote as superimposed onto the user's physical environment. For example, the wearable device may render the virtual remote 1260 as if it is in front of a wall. The virtual remote may be rendered as non-transparent such that the user can perceive the virtual remote occluding a portion of the user's physical environment, so that the virtual remote appears as if in front of the portion of the environment. For example, as shown in FIG. 14, the user may perceive that the virtual remote occludes a portion of the window 1240. In some implementations, the virtual remote may be rendered as at least partially transparent so that the user may see through the virtual remote. For example, as shown in FIG. 15, the user may see the window sill as well as the wall even though the user may perceive the virtual remote 1260 as being in front of the window and the wall. Portions of the virtual remote (e.g., user actuable elements such as the buttons) may be rendered less transparently than other portions (e.g., a body or frame) so that the virtual remote occludes less of the background environment.

[0122] The user can also move the rendering location, size, or orientation of the virtual remote. For example, the user can move the virtual remote closer (or farther) to him, upward/downward, left/right, and so on. The user can also fix the rendering location of the virtual remote to be at a certain distance from him or be at a certain location in his environment.

[0123] In some embodiments, the wearable device may present a focus indicator associated with the TV 1210 or the virtual remote 1260. The focus indicator may assist the user in seeing or finding the virtual remote. The focus indicator can comprise a halo, a color, a perceived size or depth change (e.g., causing the target object to appear closer and/or larger when selected), a change in the shape of the cursor graphic (e.g., the cursor is changed from a circle to an arrow, where the cursor may indicate the position of the user's hand or finger and the change in the cursor graphic may indicate a user's actuation of a control element or indicate that a user is about to actuate the control element on the virtual remote), or other audible, tactile, or visual effects which draw the user's attention. For example, in FIG. 13, the wearable device can present a focus indicator 1262c around the TV 1210 and a focus indicator 1262a around the virtual remote 1260. The wearable device can also present a focus indicator 1262b connecting the virtual remote 1260 with the TV 1210 to show that the virtual remote 1260 is associated with the TV 1210. In some embodiments, some or all of the focus indicators can be displayed, or not displayed, depending on contextual information, timing, etc. For example, all three focus indicators 1262a, 1262b, and 1262c may be displayed when the virtual remote 1260 is first selected by the user, which may assist the user in recognizing that the virtual remote 1260 is associated with the TV 1210. After a time period, some or all of the focus indicators may disappear or fade away, as they may be no longer needed to show the user that the remote 1260 is associated with the TV 1210.

Examples of Interacting with a Virtual Remote

[0124] During a selection event, the user can actuate a virtual remote to interact with a target device. FIG. 14 illustrates an example of actuating a virtual remote during a selection event. In FIG. 14, the user can perceive the environment 1400 using his wearable device. The environment 1400 can include physical objects such as the TV 1210, the physical remote 1220 for controlling the TV 1210, the TV stand 1230, and the window 1240. The environment 1400 can also include a virtual remote 1260. The virtual remote 1260 can emulate the functions of the physical remote 1220 to control the TV 1210. For example, the layout and functions of virtual buttons of the virtual remote 1260 may be substantially the same as the physical keys on the physical remote 1210.

[0125] The virtual remote 1260 may include virtual elements such as virtual keyboards, virtual buttons, virtual switches or toggles or sliders, or virtual touch surfaces. These virtual elements may be part of the control panel 1264 of the virtual remote 1260. To actuate the virtual remote, the user may initiate an interaction event on a virtual element. For example, the user can interact with the virtual element by touching, pressing, releasing, sliding up/down or left/right, moving along a trajectory, or other types of movements in the 3D space.

[0126] Upon actuation of the virtual remote 1260, the wearable system may communicate with the TV 1210 as if it were the physical remote 1220. As an example, in FIG. 14, the user can use his right arm 1250b to actuate the virtual remote 1260. The wearable system can use the outward-facing imaging system to image the user's poses. As further described with reference to FIGS. 16 – 18, based on the user's poses, the wearable system can compute which virtual element the user has actuated. In the example depicted in FIG. 14, the wearable system can determine that the user's right index finger is actuating the button 1266. In some embodiments, the wearable system may present a focus indicator for the button 1266 indicating that this actuation by the user.

[0127] Once the wearable system detects that the user has actuated a virtual element of the virtual remote 1260, the wearable system can accordingly send a signal via the sensor 232 (shown in FIG. 2) to the corresponding parent device to instruct the parent device to perform an action based on the virtual element actuated by the user. For example, the user can touch the virtual button 1266 on the virtual remote. Because this button is associated with increasing the volume of the TV 1210, the wearable system can accordingly generate a signal (such as an IR signal generated by an IR emitter on a wearable device) and communicate the signal to the TV 1210 (which may have an IR detector), thereby causing the TV to increase its volume. The signal generated by the wearable system can be the same signal that would be generated by the corresponding physical remote control.

[0128] In cases where the signal is a line-of sight signal (such as an IR signal that must be directed at the IR detector of the TV), the emitter on the wearable system may need to be directed toward the device (just as a physical remote control must be pointed at its associated device). Advantageously, the wearable system may be configured to determine whether the requested command (e.g., to increase the volume of the TV or change a channel)

has occurred (e.g., by using a microphone on the wearable system to determine an increase in sound intensity or an outward-facing camera to determine the display of the TV has changed, respectively). If the effect of the command has not been produced by the parent device, the wearable system may instruct the user to change the user's pose so that the wearable system emitter is directed toward the parent device. For example, the wearable system may generate a visual graphic (or audible instruction) to suggest the user to point the user's head toward the parent device. In some implementations, and for some types of parent devices, the communication between the wearable system and the parent device may not be line-of-sight sensitive (e.g., when wireless RF signals or ultrasonic acoustic signals are used), and the foregoing functionality may be optional.

[0129] Although the examples in FIG. 14 illustrate actuating the virtual remote using hand gestures, the user can also actuate the virtual remote using other poses such as head poses, eye poses (including, e.g., eye gaze directions), or other body poses. In addition to or in alternative to various poses, the user may interact with the virtual remote using a user input device. For example, the user may actuate the virtual remote using a stylus or a wand or a totem.

[0130] In some embodiments, the wearable system may communicate with the physical remote to control the functions of the parent device. For example, upon detection of actuating the virtual remote, the wearable system may communicate the information associated with the actuation (such as which key is pressed, etc.) to the physical remote. The physical remote can then communicate the information associated with the actuation to the parent device.

[0131] When the user is done with the virtual remote, the user may use a hand gesture or actuate the virtual remote to terminate the selection event. As an example, while the user is watching a TV program, the user may decide that he does not need the virtual remote any more. As a result, the user may wave his hand to indicate that he is done with the virtual remote. The user may also press the dismiss button on the virtual remote. The wearable system may, in response, cease displaying the virtual remote or display the virtual remote so that it is substantially less visually perceptible (e.g., with increased transparency), which may assist the user in later selecting the virtual remote.

[0132] In certain implementations, rather than terminating the selection event, the wearable system may temporarily hide the virtual remote from the user's FOV or move the virtual remote outside of the user's FOV or to an edge of the user's FOV. For example, the user may start watching a movie on the TV 1210 and decide that he may not need the virtual remote for a while. The user can use a swipe gesture to indicate that he does not need the virtual remote right now. The wearable system can accordingly hide the virtual remote from the user's FOV or move the virtual remote outside of the user's FOV or minimize the virtual remote.

[0133] The wearable system can also automatically hide a virtual remote or terminate the selection event if a threshold condition is met. FIG. 15 illustrates an example of such features. In FIG. 15, the user can perceive the environment 1500 which includes physical objects such as the TV 1210, the physical remote 1220, the TV stand 1230, and the window 1240. The user can also perceive the virtual remote 1260. As described herein, the wearable system can emulate the signal of the physical remote 1220 and present a virtual user remote 1260 which includes the functions of the physical remote 1220. In the example environment 1500, however, the user is currently watching a program on the TV 1210 through the wearable system. The wearable system can detect that the virtual remote is inactive for a threshold period of time (such as 10 seconds, 1 minute, 2 minutes, etc.), because the user has not actuated the virtual remote during the threshold period of time. The wearable system can gradually fade out the virtual remote such as by increasing the transparency of the virtual remote. For example, the virtual remote 1260 may change from non-transparent to transparent as part of the fading out process. The wearable system can also fade out the virtual remote by decreasing the visibility of the virtual remote. For example, the wearable system can gradually reduce the size of the virtual remote or change the color of the virtual remote from a dark color to a light color (or change the color to be a color in the surrounding environment). In some embodiments, the wearable system may remove the focus indicator associated with the virtual remote when the virtual remote is hidden. In other embodiments, the wearable system may still display the focus indicator (such as, e.g., the halo around the virtual remote), even though the control panel of the virtual remote is hidden.

[0134] In addition to or as an alternative to a period of inactivity, the threshold condition for hiding the virtual remote or terminating the selection event may be based on the contextual information associated with the target device. For example, when the wearable system detects that the user is watching a movie, the wearable system may automatically hide the virtual remote (or terminate the selection event) after the user selects the movie. As another example, the wearable system may automatically hide the virtual remote or terminate the selection event when based on the type of the parent device. For example, the wearable system may be configured to automatically hide the virtual remote or terminate the selection event after a period of inactivity when the parent device associated with the virtual remote is a TV or a thermostat. However, when the parent device is a game controller, the wearable system may be configured to leave the virtual remote there unless specifically indicated by the user.

[0135] Although the examples described herein are with reference to one virtual remote, in some embodiments, the wearable system may render multiple virtual remotes associated with one or more target devices in the user's FOV. The wearable system may allow the user to interact and switch among these virtual remotes using poses or a user input device.

Example Processes of Rendering a Virtual Remote

[0136] FIG. 16 illustrates an example process of rendering a virtual remote. The process 1600 may be performed by the wearable system 200 described herein. In some embodiments, the process 1600 may be part of a selection event.

[0137] At block 1610, the wearable system can receive an indication of a target device. The indication may include a change in the user's pose or an actuation of a user input device. For example, the user can make certain hand gestures that indicate a selection of a target device. The target device may be a physical remote or a parent device of the physical remote. For example, the user can point at a TV or the physical remote of the TV indicating that the user wants to interact with the TV via a virtual remote.

[0138] At block 1620, the wearable system can recognize the target device based at least in part on characteristics of the target device. The characteristics may include appearances of the target device, such as the shape, size, color, layout of control panel, and

so on. For example, the wearable system can acquire an image of the targeted device and apply a computer vision algorithm which can identify appearances of the target device. The wearable system can communicate with a data store and use the appearance of the target device to identify the brand and model of the target device. Characteristics may also include other visual cues such as a label. For example, the wearable system can identify the label associate with the target device, which may include information such as the brand, model, manufacturer, date of manufacture, compatibility, etc., of the target device. The wearable system can accordingly image and parse the label to recognize the target device. In certain implementations, the wearable system may also send a query signal to the target device to recognize the type associated with the target device based on the response to the query signal. For example, the wearable system can send a query signal wirelessly to the target device requesting a device identifier from the target device. The target device may respond with its device identifier. The wearable system can thereby associate the device identifier with the target device. The wearable system can use the identifier to access the characteristics of the target device such as the specifications, appearances of the target device, etc. The wearable system and the target device can use wireless advertising techniques in which, for example, one of these devices communicates advertisements, and the other device scans for such advertisements and initiates wireless connections to advertising devices.

[0139] At block 1630, the wearable system can identify a virtual remote associated with the target device. For example, based on the recognized target device, the wearable system may access a data store storing which virtual remote is associated with the target device. The wearable system can accordingly fetch the virtual remote (together with the specification of the virtual remote) as shown in block 1640. In some situations, where the target device is a parent device (such as, e.g., a TV), the wearable system may identify a physical remote associated with the parent device after the wearable system recognizes the target device. The wearable system can generate a virtual remote emulating the functions of the physical remote. Where the target device is a physical remote, the wearable system can identify the associated parent device. The wearable system can establish a connection to the parent device and communicate with the parent device as if it is the physical remote.

[0140] At block 1640, the wearable system can access a specification associated with the virtual remote. The specification may include protocols for communicating with the parent device (such as, e.g., the channel and frequency used for communication). The specification may also include layout of virtual elements on the control panel and the functions of each virtual element of virtual remote. The specification may be based on the physical remote used for controlling the parent device such that the virtual remote may have the same specification as the physical remote and the virtual remote may substitute the physical remote for communications with the parent device. In some implementations, the wearable system can perform the block 1640 right after the target device is recognized. For example, the wearable system can perform the block 1640 before the block 1630. The wearable system can automatically fetch the specification from a data store after the wearable system has identified the model of a physical remote or a parent device. In certain implementations, the wearable system can access a virtual remote based on the specification.

[0141] At block 1650, the wearable system can render the virtual remote in an AR/MR environment based at least in part on the specification. The wearable system can render the virtual remote and the virtual elements of the virtual remote based on the retrieved layout. In some embodiments, the wearable system may provide one or more focus indicators associated with the virtual remote, the virtual elements of the virtual remote, or the parent device. For example, the focus indicator may include a visual connection (such as a line) connecting the parent device and the virtual remote. This focus indicator may show that the virtual remote controls the parent device.

[0142] As further described herein, the user can actuate the virtual remote (and the virtual elements of the virtual remote) using poses or a user input device. The wearable system can generate a signal causing one or more user interface actions to be performed based on the user's action.

[0143] Although the examples herein are described with reference to rendering a virtual remote based on a physical remote or a parent device, in various embodiments, the virtual remote can also be rendered based on a user input device 466 (such as, e.g., a physical keyboard).

Example Interactions with a Virtual Element on a Virtual User Interface

[0144] The wearable system can be configured to detect actuation of a virtual element based on the user's hand or finger movement or the movement of a stylus or wand. The virtual element may be part of a virtual remote described herein, although in some embodiments, the virtual element may be a standalone user interface element rendered in the AR/VR/MR environment. The virtual element may be a virtual button, toggle, slider, or a virtual touch surface, or elements with other types of user interface appearances.

[0145] FIG. 17 illustrates an example of detecting an act of pressing a virtual button by a user. The virtual button 1700 in FIG. 17 includes a proximal surface 1712 and a distal surface 1714, where the proximal surface 1712 is the side of the virtual button that is close to the user while the distal surface 1714 is the side that is away from the user.

[0146] The virtual button 1700 can include an active volume 1710. The active volume 1710 may represent a volume of space in the user's environment in which the button can be actuated. When the wearable system detects that the user's finger has entered into this volume of space, the wearable system may initiate an interaction event on the virtual button 1700 associated with this volume of space. Because the location of the virtual remote (or the virtual button) may change based on the user's position or interaction (such as when the user moves around in his room or when the user moves the virtual remote to a different location), the active volume 1710 associated with the virtual button 1700 may not necessarily be fixed to a certain location in the user's environment.

[0147] In FIG. 17, the active volume 1710 associated with the virtual button 1700 may be the volume of space in-between the proximal surface 1712 and the distal surface 1714 of the virtual button 1700. Although in this example, the volume of space associated with the active volume 1710 is the same as the volume of space associated with the virtual button 1700, in some embodiments, the volume of space associated with the active volume 1710 may be a portion of the active volume associated with the virtual button 1700.

[0148] The active volume 1710 can include a trigger surface 1750 and a release surface 1760. The trigger surface 1750 or the release surface 1760 may cause the wearable system to initiate an interaction event on the virtual element.

[0149] The active volume 1710 may also include an active surface 1730. The active surface 1730 may be in a parallel relationship with the proximal surface 1712 or the

distal surface 1714. The position of the active surface 1730 may change based on the movement of the user's finger. For example, the initial position of the active surface 1730 may be at the proximal surface 1712. However, when the user's finger 1720 enters into the active volume 1710, the wearable system may calculate a penetration distance 1740 into the active volume 1710 based on the movement of the user's finger 1720. This penetration distance 1740 may be used to calculate the displacement of the active surface 1730. For example, the active surface 1730 may move away from the proximal surface 1712 by the amount of the penetration distance 1740. In some embodiments, at least the active surface 1730 is displayed to the user so that the user can visually perceive the virtual button 1700 being depressed or released. Further, although the finger 1720 is illustrated in FIG. 17 as actuating the button 1700, this is for illustration and in other embodiments, an end of a stylus or wand can be used to actuate the button 1700. Further, the user's finger 1720 can be a thumb, forefinger, or other finger.

[0150] The wearable system may also calculate a velocity associated with the active surface 1730. For example, the user's finger 1720 may enter the active volume 1710 with a velocity $\vec{V}$. The velocity $\vec{V}$ may include two components: $\vec{V}_n$ and $\vec{V}_t$, where $\vec{V}_n$ is normal to the active surface 1730 and $\vec{V}_t$ is parallel to the active surface 1730. The velocity $\vec{V}_n$ may be imparted to the active surface 1730. As a result, active surface 1730 can co-move with the finger 1720 when the finger 1720 is moving toward the back of the active surface. The wearable system can utilize an outward-facing imaging system 464 to image the movement of the user's finger 1720 and determine a position or the velocity $\vec{V}$ using computer vision techniques. In implementations in which the velocity is measured, the penetration distance 1740 can be determined via numerical integration (e.g., Simpson's rule, Newton-Cotes algorithms, or Gauss quadrature formulas) as $\int \vec{V}_n dt$.

[0151] As the finger 1720 continues moving toward the distal surface 1714, the active surface 1730 may reach the trigger surface 1750. In some embodiments, the trigger surface 1750 may be located at the distal surface 1714 of the active volume 1710. In other embodiments, the trigger surface 1750 may be in-between the proximal surface 1712 and the

distal surface 1744. Upon determining that the active surface 1730 has reached the position of the trigger surface 1750, the wearable system may initiate an interaction event associated with the trigger surface 1750. For example, the wearable system may generate a signal indicating that the user has pressed the button 1710 as part of the interaction event. In some embodiments, when an interaction event is triggered, the wearable system may prevent another interaction event from being generated, at least temporarily. For example, a virtual button may be associated with multiple interaction events such as pressing, releasing, or sliding. When a pressing interaction is triggered, the wearable system may prevent the sliding or the releasing interaction event to be triggered. The distance between the proximal surface 1712 and the trigger surface 1750 may be set to reflect the sensitivity of the button 1700. For example, a very sensitive button 1700 may have a relatively small distance between the proximal surface 1712 and the trigger surface 1750 (so that the button is easy to trigger with a small movement of the fingertip). A less sensitive button 1700 may have a relatively large distance between the proximal surface 1712 and the trigger surface 1750 (so that greater fingertip movement is needed to actuate the button). The distance between the proximal surface 1712 and the trigger surface 1750 may be in a range from about 1 mm to 5 cm or some other distance, in various embodiments.

[0152] In some implementations, the active surface 1730 may not be able to move past the trigger surface 1750. For example, the active surface 1730 may stop moving past the trigger surface 1750 even if the user's finger 1720 continues moving in the direction toward or past the distal surface 1714.

[0153] The user can terminate the interaction event, for example, by retracting his finger. As an example, the user may move his finger back toward the proximal surface 1712. The active surface 1730 may accordingly move towards the proximal surface 1712 from the trigger surface 1750. In some embodiments, when the active surface 1730 reaches the release surface 1760, the wearable system may terminate the interaction event associated with pressing the virtual button 1700 and thereby allow another interaction event to be initiated (e.g., by the user pressing the button toward the distal surface 1714 again). Additionally or alternatively, when the active surface 1730 reaches the release surface 1760, the wearable system may initiate another interaction event associated with releasing the button. For example, the virtual button 1700 may be associated with an arrow key such that when the

user presses the virtual button 1700, the user's avatar is moving in a certain direction in the virtual environment rendered by the wearable system. When the user releases the virtual button 1700, the wearable system may generate a signal indicating that the user's avatar will stop moving in that direction.

[0154] The active surface 1730 may recede toward the proximal surface 1712 at a pre-determined return velocity. The return velocity may be a constant or a function of the position of the active surface. For example, the return velocity may gradually increase until it reaches the proximal surface 1712 (where the return velocity reaches 0 because the active surface 1730 may be configured not to pass the proximal surface 1712 of the virtual button). As another example, the return velocity may change when it reaches a certain threshold position. For example, the return velocity may be slow while the active surface is in-between the trigger surface 1750 and the release surface 1760 but may increase when the active surface moves past the release surface 1760 toward the proximal surface 1712. The return velocity may be greater than, less than, or equal to the velocity of the receding finger.

[0155] In some embodiments, active surface 1730 may co-move with the finger 1720 when active surface 1730 is receding. For example, the active surface 1730 may recede with a velocity $\overline{V}_n$ while the finger recedes with a velocity $\overline{V}$.

[0156] In addition to or in alternative to pressing a virtual button, a user can also interact with a virtual button by sliding the virtual button. FIG. 18 illustrates an example where a user actuates a virtual button by sliding the virtual button up and down. The virtual button 1800 may include an active surface 1830. The active surface 1830 may be an embodiment of the active surface 1730 described in FIG. 17. The active surface 1830 may be initially located at the position 1832.

[0157] Although in this example, the active surface 1730 moves in a direction parallel to the distal surface 1714 and the proximal surface 1712, the movement of the active surface 1730 may move in any direction in the 3D space (such as a direction that is not parallel to the distal surface 1714 or the proximal surface 1712).

[0158] In FIG. 18, the wearable system may determine that the finger 1720 is initially at the position 1722 which corresponds to the position 1832 of the active surface 1830. The wearable system may further detect that a portion of the finger 1720 intersects

with the active surface 1830. The wearable system can detect that the finger 1720 moves from the position 1722 to the position 1724 with a velocity $\vec{V}$. The velocity $\vec{V}$ is composed of velocities $\vec{V}_n$ and $\vec{V}_t$, where $\vec{V}_n$ is normal to the active surface 1830 and $\vec{V}_t$ is parallel to the active surface 1830. Accordingly, the wearable system can move the active surface 1830 with a velocity $\vec{V}_t$.

[0159] In this example, the displacement of the active surface 1830 may also be calculated based on the movement of the finger 1720. For example, the wearable system can calculate that the finger 1720 has moved for a distance 1836 in the direction that is parallel to the active surface 1830. The wearable system may then determine that the active surface may be displaced for the distance 1836. Accordingly, the wearable system may move the active surface upward from the position 1832 to the position 1834. In some embodiments, the wearable system may set one or more threshold positions where an interaction event may be triggered. For example, when the active surface reaches the position 1834, the wearable system may generate a signal to move a virtual object in the user's FOV upward. Similarly as described with reference to FIG. 17, the wearable system can utilize an outward-facing imaging system 464 to image the movement of the user's finger 1720 and determine a position or the velocity $\vec{V}$ using computer vision techniques. In implementations in which the velocity is measured, the displacement distance 1836 can be determined via numerical integration (e.g., Simpson's rule, Newton-Cotes algorithms, or Gauss quadrature formulas) as $\int \vec{V}_t dt$.

[0160] When the user's finger 1720 no longer touches the active surface 1830 or when the user's finger 1720 slides down, the wearable system may recede the active surface to its initial position (such as from the position 1834 to the position 1832). In some embodiments, the velocity of receding may be a pre-determined velocity which may not correspond to the finger's movement. In other embodiments, the active surface may co-move with the finger 1720.

[0161] FIG. 19 illustrates an example where a user actuates a virtual touch surface. The virtual touch surface 1900 may be part of a virtual remote described herein or be

a standalone virtual element. The user can actuates the virtual touch surface 1900 using the finger 1720. The finger 1720 may move along a trajectory 1772. The wearable system can accordingly initiate an interaction event based on the trajectory 1772. For example, the wearable system can calculate the tangent 1774 of the trajectory 1772. The wearable system can initiate an interaction event if the movement (velocity or displacement) along the tangent 1774 exceeds a threshold condition. During the interaction event, the wearable system can provide an instruction for moving a virtual object along the trajectory 1772.

[0162] In some embodiments, the wearable system may present a focus indicator associated with a virtual element. For example, the wearable system may present the color or depth change based on the movement of the active plane, where the intensity of the color or the perceived depth increases as the active surface moves towards the distal surface of the virtual button. As another example, the wearable system may provide a sound or a glow around a virtual button when the virtual button is pressed. The wearable system may also illuminate the proximal surface of the virtual button indicating that the virtual button is pressed. The focus indicator may also be provided when the user releases the button. For example, as the button moves closer to the proximal surface, the illumination on the button may decrease. As another example, the size (or the font) associated with the virtual button may decrease as it recedes back to its initial position.

[0163] The appearances of the virtual elements (e.g., the virtual buttons 1700, 1800, and the virtual touch surface 1900) illustrated in FIGS. 16 – 18 are merely examples and are not intended to be limiting. For example, the active surfaces as well as the virtual elements may be in any 2D or 3D shapes such as a rectangle, triangle, circle, oval, sphere, cuboid, pyramid, irregular, etc.

[0164] Furthermore, although the examples are described with reference to actuating the virtual elements with a finger of the user, in some implementations, the user can also actuate the virtual elements using other hand poses, body poses, or more than one figure. Additionally or alternatively, the user can actuate the virtual elements using a user input device, such as a stylus or a wand.

Example Processes for Interacting with a Virtual Button

[0165] FIG. 20 illustrates an example process for initiating an interaction event with a virtual element in a virtual user interface. The process 2000 in FIG. 20 may be performed by the wearable system 200 described herein.

[0166] At block 2010, the wearable system can identify a virtual element in a virtual user interface. The virtual user interface may be part of an AR/VR/MR environment. The VR or the MR environment may be rendered by the wearable system alone or in combination with another computing device. The virtual element may be a virtual button or a virtual touch surface. The virtual element may be part of the virtual remote described herein or be a standalone virtual user interface element. The virtual element may include an active surface whose movements may be associated with the user's movement. In some embodiments, the active surface may be part of an active volume which can cause an interaction event to be triggered upon detecting that a portion of the user's body (such as a finger) has entered into the active volume.

[0167] At block 2020, the wearable system can identify a first position of the active surface. The first position may be the initial position of the active surface before the user interaction. For example, the first position may be located at the proximal surface of a virtual button (or of an active volume of the virtual button). However, in some implementations, when the user has already pressed the button, the first position may be at the trigger surface or at the distal surface of the virtual button (or the active volume of the virtual button).

[0168] At block 2030, the wearable system can receive an indication of a user interaction with the virtual element. The indication may be that a portion of the user's body intersects with the virtual element (or the active surface). For example, a user may press a virtual button or touch the active surface.

[0169] At block 2040, the wearable system can calculate a movement of the active surface based at least partly on the indication 2040. For example, the wearable system can calculate a velocity vector of the user's finger when the user presses a virtual button. Based on the velocity vector, the wearable system can calculate a velocity in a direction normal to the active surface and use this velocity as the velocity for the active surface. As another example, when the user is sliding a virtual surface, the wearable system can calculate

a velocity in the direction that is parallel to the active surface for determining the velocity of the active surface. The wearable system can also calculate a displacement of the active surface based on the amount of distance that the user's finger has moved.

[0170] At block 2050, the wearable system can calculate a second position of the active surface based at least partly on the movement 2050. For example, the wearable system can calculate the second position based on the displacement of the active surface. In some embodiments, once the active surface reaches a certain threshold position, the active surface will not continue moving. For example, when the active surface reaches the distal end or the proximal end of the active volume, the active surface may stop moving so that it will move pass the distal end or the proximal end of the active volume.

[0171] At block 2060, the wearable system can initiate an interaction event based at least partly on the first position and the second position of the active surface. The second position may be associated with triggering a certain user interface action. For example, when the active surface reaches a triggering surface, the wearable system may determine that the user has pressed the virtual button. As another example, when the user's finger recedes, the active surface may move back toward the proximal surface of the active volume. Once the active surface reaches a release surface, the wearable system may generate a signal indicating that the user has released the virtual button.

Additional Aspects

[0172] In a 1st aspect, a system for providing a visual remote by an augmented reality device (ARD), the system comprising: an augmented reality (AR) display for presenting a virtual image to a user of an ARD; an outward-facing imaging system configured to image an environment of the user; a sensor configured to communicate with a parent device; and a hardware processor programmed to: identify, based on data from at least one of the outward-facing imaging system or the sensor, a physical remote associated with the parent device in the environment of the user; access specifications of the physical remote; emulate the physical remote based at least partly on the accessed specification to generate a virtual remote; receive a selection of the virtual remote; render, by the display, the virtual remote in response to the selection; receive an indication of an interaction with the parent

device via the virtual remote; and transmit an instruction comprising the indication of the interaction to the parent device.

[0173] In a 2nd aspect, the system of aspect 1, wherein the AR display comprises a light field display.

[0174] In a 3rd aspect, the system of any one of aspects 1 – 2, wherein the sensor comprises an infrared light source.

[0175] In a 4th aspect, the system of any one of aspects 1 – 3, wherein the parent device comprises at least one of: a television, a thermostat, an audio system, a home theater system, a home security system, a door bell, a door lock, an air conditioner, a heater, a lighting system, or a garage door opener.

[0176] In a 5th aspect, the system of any one of aspects 1 – 4, wherein to identify, based on the data from at least one of the outward-facing imaging system or the sensor, the physical remote, the hardware processor is programmed to: receive an image of the environment from the outward-facing imaging system; and identify at least one of the parent device or the physical remote in the image using a computer vision algorithm.

[0177] In a 6th aspect, the system of aspect 5, wherein the image comprises at least one of: a still image, a frame of a video, or a video.

[0178] In a 7th aspect, the system of any one of aspects 5 – 6, to identify at least one of the parent device or the physical remote, the hardware processor is programmed to: recognize existence of at least one of the parent device or the physical remote; and identify a type of at least one of the parent device or the physical remote.

[0179] In an 8th aspect, the system of any one of aspects 5 – 7, wherein the computer vision algorithm comprises at least one of: a feature keypoint algorithm, a bag of words type search, or a neural network algorithm.

[0180] In a 9th aspect, the system of any one of aspects 1 – 8, wherein to identify, based on the data from at least one of the outward-facing imaging system or the sensor, the physical remote comprises: receive, by the sensor, a signal indicating a presence of the physical remote or the parent device in the environment of the user.

[0181] In a 10th aspect, the system of any one of aspects 1 – 9, wherein the specifications of the physical remote comprise at least one of: a plurality of buttons, a layout of buttons, a communication channel, or a signal associated with the communication channel.

[0182] In an 11th aspect, the system of any one of aspects 1 – 10, wherein to emulate the physical remote, the hardware processor is programmed to associate functionality of a button on the physical remote with a button on the virtual remote.

[0183] In a 12th aspect, the system of any one of aspects 1 – 11, wherein to receive a selection of the virtual remote, the hardware processor is configured to: detect a pose of the user, wherein the pose indicates the selection of the virtual remote; or receive an indication of the selection from a user input device.

[0184] In a 13th aspect, the system of any one of aspects 1 – 12, wherein the virtual remote is rendered based at least partly on a layout of buttons of the physical remote.

[0185] In a 14th aspect, the system of any one of aspects 1 – 13, wherein to receive an indication of an interaction, the hardware processor is programmed to: detect a pose of the user, wherein the pose indicates the selection of the virtual remote; or receive an indication of the selection from a user input device.

[0186] In a 15th aspect, the system of aspect 14, wherein the pose comprises at least one of a head pose, a hand gesture, a body pose, or a foot pose.

[0187] In a 16th aspect, the system of any one of aspects 1 – 15, wherein the instruction is transmitted in accordance with an infrared data association communication standard or a radio frequency communication standard.

[0188] In a 17th aspect, a wearable display device comprising: a display for presenting a virtual image to a user of the wearable display device; an outward-facing imaging system configured to image an environment of the user; a sensor configured to communicate with an electronic device; and a hardware processor in communication with the display, the outward-facing imaging system, and the sensor, the hardware processor programmed to: access remote control specifications of the electronic device; render a virtual remote having a user interface permitting user interaction with the electronic device, the user interface providing functionality based at least partly on the remote control specifications; analyze an image obtained by the outward-facing camera to determine a gesture by the user; analyze the gesture to determine an interaction by the user with the user interface of the virtual remote; and transmit, via the sensor to the electronic device, an instruction associated with the interaction, the instruction based at least partly on the remote control specifications of the electronic device.

[0189] In an 18th aspect, the wearable display device of aspect 17, wherein the display comprises a light field display.

[0190] In a 19th aspect, the wearable display device of aspect 17 or 18, wherein the sensor comprises an infrared emitter or a radio frequency emitter.

[0191] In a 20th aspect, the wearable display device of any one of aspects 17 – 19, wherein: the hardware processor is programmed to analyze an image obtained by the outward-facing imaging system to determine a specific model of the electronic device.

[0192] In a 21st aspect, the wearable display device of any one of aspects 17 – 20, wherein the hardware processor is programmed to render a focus indicator associated with the virtual remote.

[0193] In a 22nd aspect, a method for providing a visual remote by an augmented reality device (ARD), the method comprising: under control of an ARD comprising a hardware processor, a display for presenting a virtual image to a user of the ARD, an outward-facing imaging system configured to image an environment of the user, and a sensor configured to communicate with a parent device: identifying, based on data from at least one of the outward-facing imaging system or the sensor, a physical remote associated with the parent device in the environment of the user; accessing specifications of the physical remote; emulating the physical remote based at least partly on the accessed specification to generate a virtual remote; receiving a selection of the virtual remote; rendering, by the display, the virtual remote in response to the selection; receiving an indication of an interaction with the parent device via the virtual remote; and transmitting an instruction comprising the indication of the interaction to the parent device.

[0194] In a 23rd aspect, the method of aspect 22, wherein identifying, based on the data from at least one of the outward-facing imaging system or the sensor, the physical remote comprises: receiving an image of the environment from the outward-facing imaging system; and identifying at least one of the parent device or the physical remote in the image using a computer vision algorithm.

[0195] In a 24th aspect, the method of aspect 23, wherein identifying at least one of the parent device or the physical remote comprises: recognizing existence of at least one of the parent device or the physical remote; and identifying a type of at least one of the parent device or the physical remote.

[0196] In a 25th aspect, the method of aspect 23 or 24, wherein the computer vision algorithm comprises at least one of: a feature keypoint algorithm, a bag of words type search, or a neural network algorithm.

[0197] In a 26th aspect, the method of any one of aspects 22 – 25, wherein identifying, based on the data from at least one of the outward-facing imaging system or the sensor, the physical remote comprises: receiving, by the sensor, a signal indicating a presence of the physical remote or the parent device in the environment of the user.

[0198] In a 27th aspect, the method of any one of aspects 22 – 26, wherein the specifications of the physical remote comprise at least one of: a plurality of buttons, a layout of buttons, a communication channel, or a signal associated with the communication channel.

[0199] In a 28th aspect, the method of any one of aspects 22 – 27, wherein emulating the physical remote comprises associating functionality of a button on the physical remote with a button on the virtual remote.

[0200] In a 29th aspect, the method of any one of aspects 22 – 28, wherein receiving a selection of the virtual remote comprises at least one of: detecting a pose of the user, wherein the pose indicates the selection of the virtual remote; or receiving an indication of the selection from a user input device.

[0201] In a 30th aspect, the method of any one of aspects 22 – 29, wherein the virtual remote is rendered based at least partly on a layout of buttons of the physical remote.

[0202] In a 31st aspect, the method of any one of aspects 22 – 30, wherein receiving an indication of an interaction comprises: detecting a pose of the user, wherein the pose indicates the selection of the virtual remote; or receiving an indication of the selection from a user input device.

[0203] In a 32nd aspect, the method of any one of aspects 22 – 31, wherein the instruction is transmitted in accordance with an infrared data association communication standard or a radio frequency communication standard.

[0204] In a 33rd aspect, a system for detecting actuation of a virtual button, the system comprising: a wearable display system configured to display a virtual image to a user; an outward-facing imaging system configured to obtain images of an environment of the user; and a hardware processor programmed to: identify a virtual button of a virtual user interface, the virtual button comprising an active surface and a trigger surface; identify a first

position of the active surface; determine a user interaction indicating movement of the virtual button; calculate a movement of the active surface based at least partly on the user interaction; calculate a second position of the active surface based at least partly on the movement; and initiate an interaction event based at least partly on the first position and the second position of the active surface, wherein the interaction event is associated with actuating the virtual button when the movement of the active surface is toward the trigger surface and the second position of the active surface intersects with at least a portion of the trigger surface, and wherein the interaction event is associated with releasing the virtual button when the movement of the active surface is away from the trigger surface and the first position of the active surface intersects with at least a portion of the trigger surface.

[0205] In a 34th aspect, the system of aspect 33, wherein the active surface is in a parallel relationship with the trigger surface.

[0206] In a 35th aspect, the system of aspect 33 or 34, wherein the virtual button further comprises a release surface and wherein the interaction event is associated with releasing the virtual button when the second position of the active surface intersects with at least a portion of the release surface.

[0207] In a 36th aspect, the system of any one of aspects 33 – 35, wherein the user interaction indicating movement of the virtual button comprises an intersection between the active plane and a physical object in the environment of the user, the physical object comprising at least one of: a portion of a user or a portion of a user input device.

[0208] In a 37th aspect, the system of aspect 36, wherein to calculate the movement of the active surface, the hardware processor is configured to: calculate a velocity vector of the physical object; and calculate a velocity of the active surface using at least one of: a value the velocity vector of the physical object normal to the active surface, a value of the velocity vector of the physical object parallel to the active surface, or a value of the velocity vector tangent to the active surface.

[0209] In a 38th aspect, the system of any one of aspects 36 – 37, wherein to calculate the second position of the active surface, the hardware processor is configured to: determine the physical object has receded from the trigger surface at a recession speed; and move the active surface at speed less than or equal to the recession speed.

[0210] In a 39th aspect, the system of any one of aspects 33 – 38, wherein calculating the movement of the active surface comprises calculating a displacement of the active surface.

[0211] In a 40th aspect, the system of any one of aspects 33 – 39, wherein the hardware processor is further configured to provide a visualization of the virtual button based at least partly on the movement of the active surface.

[0212] In a 41st aspect, the system of aspect 40, wherein the visualization comprises a focus indicator of the virtual button, the focus indicator being provided when the movement of the virtual button passes a threshold condition.

[0213] In a 42nd aspect, the system of aspect 41, wherein the threshold condition comprises at least one of: a threshold position, a threshold duration, or a threshold velocity.

[0214] In a 43rd aspect, the system of any one of aspects 40 – 42, wherein the visualization comprises a focus indicator of the virtual button, wherein the focus indicator changes in accordance with the movement of the active plane.

[0215] In a 44th aspect, the system of any one of aspects 33 – 43, wherein the interaction event comprises at least one of: triggering a user interface function on the virtual user interface or disabling a user interface function.

[0216] In a 45th aspect, the system of any one of aspects 33 – 44, wherein to determine the user interaction indicating movement of the virtual button, the hardware processor is configured to analyzing images obtained by the outward-facing imaging system.

[0217] In a 46th aspect, the system of any one of aspects 33 – 45, wherein the hardware processor is further configured to generate a sound associated with initiating the interaction event.

[0218] In a 47th aspect, a method for detecting actuation of a virtual button, the method comprising: under control of a wearable display system configured to display a virtual image to a user, a hardware processor, and an outward-facing imaging system configured to obtain images of an environment of the user: identifying a virtual button of a virtual user interface, the virtual button comprising an active surface and a trigger surface; identifying a first position of the active surface; determining a user interaction indicating movement of the virtual button; calculating a movement of the active surface based at least partly on the user interaction; calculating a second position of the active surface based at

least partly on the movement; and initiating an interaction event based at least partly on the first position and the second position of the active surface, wherein the interaction event is associated with actuating the virtual button when the movement of the active surface is toward the trigger surface and the second position of the active surface intersects with at least a portion of the trigger surface, and wherein the interaction event is associated with releasing the virtual button when the movement of the active surface is away from the trigger surface and the first position of the active surface intersects with at least a portion of the trigger surface.

[0219] In a 48th aspect, the method of aspect 47, wherein the active surface is in a parallel relationship with the trigger surface.

[0220] In a 49th aspect, the method of any one of aspects 47 – 48, wherein the virtual button further comprises a release surface and wherein the interaction event is associated with releasing the virtual button when the second position of the active surface intersects with at least a portion of the release surface.

[0221] In a 50th aspect, the method of any one of aspects 47 – 49, wherein the user interaction indicating movement of the virtual button comprises an intersection between the active plane and a physical object in the environment of the user, the physical object comprising at least one of: a portion of a user or a portion of a user input device.

[0222] In a 51st aspect, the method of aspect 50, wherein calculating the movement of the active surface comprises: calculating a velocity vector of the physical object; and calculating a velocity of the active surface using at least one of: a value the velocity vector of the physical object normal to the active surface, a value of the velocity vector of the physical object parallel to the active surface, or a value of the velocity vector tangent to the active surface.

[0223] In a 52nd aspect, the method of aspect 50 or 51, wherein calculating the second position of the active surface comprises: determining the physical object has receded from the trigger surface at a recession speed; and moving the active surface at speed less than or equal to the recession speed.

[0224] In a 53rd aspect, the method of any one of aspects 47 – 52, further comprising providing a visualization of the virtual button based at least partly on the movement of the active surface.

[0225] In a 54th aspect, the method of aspect 53, wherein the visualization comprises a focus indicator of the virtual button, the focus indicator being provided when the movement of the virtual button passes a threshold condition.

[0226] In a 55th aspect, the method of aspect 54, wherein the threshold condition comprises at least one of: a threshold position, a threshold duration, or a threshold velocity.

[0227] In a 56th aspect, the method of aspect 54 or 55, wherein the visualization comprises a focus indicator of the virtual button, wherein the focus indicator changes in accordance with the movement of the active plane.

[0228] In a 57th aspect, the method of any one of aspects 47 – 56, wherein the interaction event comprises at least one of: triggering a user interface function on the virtual user interface or disabling a user interface function.

[0229] In a 58th aspect, a wearable system for providing a virtual remote control in a mixed reality environment, the wearable system comprising: a mixed reality display for presenting a virtual image to a user; an outward-facing imaging system configured to image an environment of a user; a sensor configured to communicate with a parent device; and a hardware processor programmed to: access images acquired by the outward-facing imaging system; analyze the images to identify a physical remote associated with the parent device; access specifications of the physical remote; generate a virtual remote based at least partly on the specifications of the physical remote; cause the display to render the virtual remote; receive an actuation of the virtual remote for an interaction with the parent device; and generate and transmit an instruction to the parent device causing the parent device to perform a function as if the physical remote is actuated.

[0230] In a 59th aspect, the wearable system of aspect 58, wherein to identify the physical remote, the hardware processor is programmed to: analyze the images to recognize the physical remote using one or more computer vision algorithms comprising at least one of: a feature keypoint algorithm, a bag of words type search, or a neural network algorithm.

[0231] In a 60th aspect, the wearable system of aspect 59, wherein the hardware processor is further programmed to identify the presence of the parent device based at least partly on the images acquired by the outward-facing imaging system or a signal received by the sensor which indicates the presence of the parent device.

[0232] In a 61st aspect, the wearable system of any one of aspects 58 – 60, wherein the specifications include at least one of a mapping of controls elements of the physical remote to corresponding functions of the parent device, or communication protocols between the physical remote and the parent device.

[0233] In a 62nd aspect, the wearable system of aspect 61, wherein to generate the virtual remote, the hardware processor is programmed to: determine a layout of virtual elements for the virtual remote based at least partly on the specification; and associate the virtual elements of the virtual remote with the control elements of the physical remote such that an actuation of a virtual element on the virtual remote would cause the parent device react as if the control element corresponding to the virtual element is actuated on the physical remote.

[0234] In a 63rd aspect, the wearable system of any one of aspects 58 – 62, wherein the hardware processor is further programmed to: determine a plurality of candidate virtual remotes associated with the parent device; and select the virtual remote among the plurality of candidate virtual remotes based at least partly on the user's pose or an indication from a user input device.

[0235] In a 64th aspect, the wearable system of any one of aspects 58 – 63, wherein the sensor comprises an infrared light source, and wherein the instruction is generated and transmitted to the parent device in accordance with an infrared data association communication standard or a radio frequency communication standard.

[0236] In a 65th aspect, the wearable system of any one of aspects 58 – 64, wherein the virtual remote comprises a virtual button, the virtual button is associated with a volume of space and has an active surface configured to track the user's interaction with the virtual button and a trigger surface for triggering a user interface action.

[0237] In a 66th aspect, the wearable system of aspect 65, wherein the active surface is in a parallel relationship with the trigger surface.

[0238] In a 67th aspect, the wearable system of aspect 65 or 66, wherein to receive an actuation of the virtual remote, the hardware processor is configured to: identify a first position of the active surface; calculate a second position and a movement of the active surface based at least partly on the user's movement; and actuate the virtual button when the movement of the active surface is toward the trigger surface and the second position of the

active surface intersects with at least a portion of the trigger surface, or release the virtual button when the movement of the active surface is away from the trigger surface and the first position of the active surface intersects with at least a portion of the trigger surface.

[0239] In a 68th aspect, the wearable system of aspect 67, wherein the virtual button further comprises a release surface and wherein the hardware processor is programmed to release the virtual button when the second position of the active surface intersects with at least a portion of the release surface.

[0240] In a 69th aspect, the wearable system of aspect 67 or 68, wherein to calculate the movement of the active surface, the hardware processor is programmed to: calculate a velocity vector of the user's movement; and calculate a velocity of the active surface based on at least one of: a value the velocity vector normal to the active surface, a value of the velocity vector parallel to the active surface, or a value of the velocity vector tangent to the active surface.

[0241] In a 70th aspect, the wearable system of aspect 65, wherein the hardware processor is programmed to provide a visualization of the virtual button for presentation to the user by the mixed reality display, the visualization based at least partly on a movement of the active surface.

[0242] In a 71st aspect, a method for providing a virtual remote control in a mixed reality environment, the method comprising: accessing images of a user's environment acquired by a wearable device; analyzing the images to identify a target device in the user's environment; accessing specifications associated with the target device, wherein the specifications comprises at least a mapping between control elements of a physical remote to functions of a parent device; generating a virtual remote based at least partly on the specifications; causing the wearable device to render the virtual remote in a mixed reality environment; detecting an actuation of the virtual remote for an interaction with the parent device; and generating and transmitting an instruction to the parent device causing the parent device to perform a function as if the physical remote is actuated.

[0243] In a 72nd aspect, the method of aspect 71, wherein the target device comprises the physical remote or the parent device.

[0244] In a 73rd aspect, the method of aspect 71 or 72, wherein identifying the target device comprises: analyzing the images to recognize the target device using one or

more computer vision algorithms comprising at least one of: a feature keypoint algorithm, a bag of words type search, or a neural network algorithm.

[0245] In a 74th aspect, the method of any one of aspects 71 – 73, wherein generating the virtual remote comprises: determining a layout of virtual elements for the virtual remote based at least partly on the specification; and associating the virtual elements of the virtual remote with the control elements of the physical remote such that an actuation of a virtual element on the virtual remote would cause the parent device react as if the control element corresponding to the virtual element is actuated on the physical remote.

[0246] In a 75th aspect, the method of any one of aspects 71 – 74, wherein the virtual remote comprises a virtual button, the virtual button is associated with a volume of space and has an active surface configured to track the user's interaction with the virtual button and a trigger surface for triggering a user interface action.

[0247] In a 76th aspect, the method of any one of aspects 71 – 75, wherein detecting an actuation of the virtual remote comprises: identifying a first position of the active surface; calculating a second position and a movement of the active surface based at least partly on the user's movement; and actuating the virtual button when the movement of the active surface is toward the trigger surface and the second position of the active surface intersects with at least a portion of the trigger surface, or release the virtual button when the movement of the active surface is away from the trigger surface and the first position of the active surface intersects with at least a portion of the trigger surface.

[0248] In a 77th aspect, the method of aspect 76, wherein the virtual button further comprises a release surface and the method further comprises releasing the virtual button when the second position of the active surface intersects with at least a portion of the release surface.

Other Considerations

[0249] Each of the processes, methods, and algorithms described herein and/or depicted in the attached figures may be embodied in, and fully or partially automated by, code modules executed by one or more physical computing systems, hardware computer processors, application-specific circuitry, and/or electronic hardware configured to execute specific and particular computer instructions. For example, computing systems can include

general purpose computers (e.g., servers) programmed with specific computer instructions or special purpose computers, special purpose circuitry, and so forth. A code module may be compiled and linked into an executable program, installed in a dynamic link library, or may be written in an interpreted programming language. In some implementations, particular operations and methods may be performed by circuitry that is specific to a given function.

[0250] Further, certain implementations of the functionality of the present disclosure are sufficiently mathematically, computationally, or technically complex that application-specific hardware or one or more physical computing devices (utilizing appropriate specialized executable instructions) may be necessary to perform the functionality, for example, due to the volume or complexity of the calculations involved or to provide results substantially in real-time. For example, a video may include many frames, with each frame having millions of pixels, and specifically programmed computer hardware is necessary to process the video data to provide a desired image processing task or application in a commercially reasonable amount of time.

[0251] Code modules or any type of data may be stored on any type of non-transitory computer-readable medium, such as physical computer storage including hard drives, solid state memory, random access memory (RAM), read only memory (ROM), optical disc, volatile or non-volatile storage, combinations of the same and/or the like. The methods and modules (or data) may also be transmitted as generated data signals (e.g., as part of a carrier wave or other analog or digital propagated signal) on a variety of computer-readable transmission mediums, including wireless-based and wired/cable-based mediums, and may take a variety of forms (e.g., as part of a single or multiplexed analog signal, or as multiple discrete digital packets or frames). The results of the disclosed processes or process steps may be stored, persistently or otherwise, in any type of non-transitory, tangible computer storage or may be communicated via a computer-readable transmission medium.

[0252] Any processes, blocks, states, steps, or functionalities in flow diagrams described herein and/or depicted in the attached figures should be understood as potentially representing code modules, segments, or portions of code which include one or more executable instructions for implementing specific functions (e.g., logical or arithmetical) or steps in the process. The various processes, blocks, states, steps, or functionalities can be combined, rearranged, added to, deleted from, modified, or otherwise changed from the

illustrative examples provided herein. In some embodiments, additional or different computing systems or code modules may perform some or all of the functionalities described herein. The methods and processes described herein are also not limited to any particular sequence, and the blocks, steps, or states relating thereto can be performed in other sequences that are appropriate, for example, in serial, in parallel, or in some other manner. Tasks or events may be added to or removed from the disclosed example embodiments. Moreover, the separation of various system components in the implementations described herein is for illustrative purposes and should not be understood as requiring such separation in all implementations. It should be understood that the described program components, methods, and systems can generally be integrated together in a single computer product or packaged into multiple computer products. Many implementation variations are possible.

[0253] The processes, methods, and systems may be implemented in a network (or distributed) computing environment. Network environments include enterprise-wide computer networks, intranets, local area networks (LAN), wide area networks (WAN), personal area networks (PAN), cloud computing networks, crowd-sourced computing networks, the Internet, and the World Wide Web. The network may be a wired or a wireless network or any other type of communication network.

[0254] The systems and methods of the disclosure each have several innovative aspects, no single one of which is solely responsible or required for the desirable attributes disclosed herein. The various features and processes described above may be used independently of one another, or may be combined in various ways. All possible combinations and subcombinations are intended to fall within the scope of this disclosure. Various modifications to the implementations described in this disclosure may be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other implementations without departing from the spirit or scope of this disclosure. Thus, the claims are not intended to be limited to the implementations shown herein, but are to be accorded the widest scope consistent with this disclosure, the principles and the novel features disclosed herein.

[0255] Certain features that are described in this specification in the context of separate implementations also can be implemented in combination in a single implementation. Conversely, various features that are described in the context of a single

implementation also can be implemented in multiple implementations separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination. No single feature or group of features is necessary or indispensable to each and every embodiment.

[0256] Conditional language used herein, such as, among others, “can,” “could,” “might,” “may,” “e.g.,” and the like, unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements and/or steps. Thus, such conditional language is not generally intended to imply that features, elements and/or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without author input or prompting, whether these features, elements and/or steps are included or are to be performed in any particular embodiment. The terms “comprising,” “including,” “having,” and the like are synonymous and are used inclusively, in an open-ended fashion, and do not exclude additional elements, features, acts, operations, and so forth. Also, the term “or” is used in its inclusive sense (and not in its exclusive sense) so that when used, for example, to connect a list of elements, the term “or” means one, some, or all of the elements in the list. In addition, the articles “a,” “an,” and “the” as used in this application and the appended claims are to be construed to mean “one or more” or “at least one” unless specified otherwise.

[0257] As used herein, a phrase referring to “at least one of” a list of items refers to any combination of those items, including single members. As an example, “at least one of: A, B, or C” is intended to cover: A, B, C, A and B, A and C, B and C, and A, B, and C. Conjunctive language such as the phrase “at least one of X, Y and Z,” unless specifically stated otherwise, is otherwise understood with the context as used in general to convey that an item, term, etc. may be at least one of X, Y or Z. Thus, such conjunctive language is not generally intended to imply that certain embodiments require at least one of X, at least one of Y and at least one of Z to each be present.

[0258] Similarly, while operations may be depicted in the drawings in a particular order, it is to be recognized that such operations need not be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. Further, the drawings may schematically depict one more example processes in the form of a flowchart. However, other operations that are not depicted can be incorporated in the example methods and processes that are schematically illustrated. For example, one or more additional operations can be performed before, after, simultaneously, or between any of the illustrated operations. Additionally, the operations may be rearranged or reordered in other implementations. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system components in the implementations described above should not be understood as requiring such separation in all implementations, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products. Additionally, other implementations are within the scope of the following claims. In some cases, the actions recited in the claims can be performed in a different order and still achieve desirable results.

WHAT IS CLAIMED IS:

1. A wearable system for providing a virtual remote control in a mixed reality environment, the wearable system comprising:

a mixed reality display for presenting a virtual image to a user;

an outward-facing imaging system configured to image an environment of a user;

a sensor configured to communicate with a parent device; and

a hardware processor programmed to:

access images acquired by the outward-facing imaging system;

analyze the images to identify a physical remote associated with the parent device;

access specifications of the physical remote;

generate a virtual remote based at least partly on the specifications of the physical remote;

cause the display to render the virtual remote;

receive an actuation of the virtual remote for an interaction with the parent device; and

generate and transmit an instruction to the parent device causing the parent device to perform a function as if the physical remote is actuated.

2. The wearable system of claim 1, wherein to identify the physical remote, the hardware processor is programmed to:

analyze the images to recognize the physical remote using one or more computer vision algorithms comprising at least one of: a feature keypoint algorithm, a bag of words type search, or a neural network algorithm.

3. The wearable system of claim 2, wherein the hardware processor is further programmed to identify the presence of the parent device based at least partly on the images acquired by the outward-facing imaging system or a signal received by the sensor which indicates the presence of the parent device.

4. The wearable system of claim 1, wherein the specifications include at least one of a mapping of controls elements of the physical remote to corresponding functions of the

parent device, or communication protocols between the physical remote and the parent device.

5. The wearable system of claim 4, wherein to generate the virtual remote, the hardware processor is programmed to:

determine a layout of virtual elements for the virtual remote based at least partly on the specification; and

associate the virtual elements of the virtual remote with the control elements of the physical remote such that an actuation of a virtual element on the virtual remote would cause the parent device react as if the control element corresponding to the virtual element is actuated on the physical remote.

6. The wearable system of claim 1, wherein the hardware processor is further programmed to:

determine a plurality of candidate virtual remotes associated with the parent device; and

select the virtual remote among the plurality of candidate virtual remotes based at least partly on the user's pose or an indication from a user input device.

7. The wearable system of claim 1, wherein the sensor comprises an infrared light source, and wherein the instruction is generated and transmitted to the parent device in accordance with an infrared data association communication standard or a radio frequency communication standard.

8. The wearable system of claim 1, wherein the virtual remote comprises a virtual button, the virtual button is associated with a volume of space and has an active surface configured to track the user's interaction with the virtual button and a trigger surface for triggering a user interface action.

9. The wearable system of claim 8, wherein the active surface is in a parallel relationship with the trigger surface.

10. The wearable system of claim 8, wherein to receive an actuation of the virtual remote, the hardware processor is configured to:

identify a first position of the active surface;

calculate a second position and a movement of the active surface based at least partly on the user's movement; and

actuate the virtual button when the movement of the active surface is toward the trigger surface and the second position of the active surface intersects with at least a portion of the trigger surface, or release the virtual button when the movement of the active surface is away from the trigger surface and the first position of the active surface intersects with at least a portion of the trigger surface.

11. The wearable system of claim 10, wherein the virtual button further comprises a release surface and wherein the hardware processor is programmed to release the virtual button when the second position of the active surface intersects with at least a portion of the release surface.

12. The wearable system of claim 10, wherein to calculate the movement of the active surface, the hardware processor is programmed to:

calculate a velocity vector of the user's movement; and

calculate a velocity of the active surface based on at least one of: a value the velocity vector normal to the active surface, a value of the velocity vector parallel to the active surface, or a value of the velocity vector tangent to the active surface.

13. The wearable system of claim 8, wherein the hardware processor is programmed to provide a visualization of the virtual button for presentation to the user by the mixed reality display, the visualization based at least partly on a movement of the active surface.

14. A method for providing a virtual remote control in a mixed reality environment, the method comprising:

accessing images of a user's environment acquired by a wearable device;

analyzing the images to identify a target device in the user's environment;

accessing specifications associated with the target device, wherein the specifications comprises at least a mapping between control elements of a physical remote to functions of a parent device;

generating a virtual remote based at least partly on the specifications;

causing the wearable device to render the virtual remote in a mixed reality environment;

detecting an actuation of the virtual remote for an interaction with the parent device; and

generating and transmitting an instruction to the parent device causing the parent device to perform a function as if the physical remote is actuated.

15. The method of claim 14, wherein the target device comprises the physical remote or the parent device.

16. The method of claim 14, wherein identifying the target device comprises: analyzing the images to recognize the target device using one or more computer vision algorithms comprising at least one of: a feature keypoint algorithm, a bag of words type search, or a neural network algorithm.

17. The method of claim 14, wherein generating the virtual remote comprises:

determining a layout of virtual elements for the virtual remote based at least partly on the specification; and

associating the virtual elements of the virtual remote with the control elements of the physical remote such that an actuation of a virtual element on the virtual remote would cause the parent device react as if the control element corresponding to the virtual element is actuated on the physical remote.

18. The method of claim 14, wherein the virtual remote comprises a virtual button, the virtual button is associated with a volume of space and has an active surface configured to track the user's interaction with the virtual button and a trigger surface for triggering a user interface action.


19. The method of claim 14, wherein detecting an actuation of the virtual remote comprises:


identifying a first position of the active surface;


calculating a second position and a movement of the active surface based at least partly on the user's movement; and


actuating the virtual button when the movement of the active surface is toward the trigger surface and the second position of the active surface intersects with at least a portion of the trigger surface, or release the virtual button when the movement of the active surface is away from the trigger surface and the first position of the active surface intersects with at least a portion of the trigger surface.


20. The method of claim 19, wherein the virtual button further comprises a release surface and the method further comprises releasing the virtual button when the second position of the active surface intersects with at least a portion of the release surface.


**FIG. 1**


**FIG. 2**


**FIG. 3**





**FIG. 5**


**FIG. 6**

**FIG. 7**

**FIG. 8**

**FIG. 9**

**FIG. 10**

**FIG. 11**




FIG. 12




FIG. 13




FIG. 14





FIG. 15

**FIG. 16**




FIG. 17






FIG. 18

**FIG. 19**

**FIG. 20**




FIG. 13