



- (51) **International Patent Classification:**
G06F 3/048 (2013.01)
- (21) **International Application Number:**
PCT/CN2019/092147
- (22) **International Filing Date:**
20 June 2019 (20.06.2019)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
16/156,979 10 October 2018 (10.10.2018) US
- (71) **Applicant: MIDEA GROUP CO., LTD.** [CN/CN];
B26-28F, Midea Headquarter Building, No.6 Midea Avenue, Beijijiao, Shunde, Foshan, Guangdong 528311 (CN).
- (72) **Inventors: TIAN, Yuan;** 250 W Tasman Dr. San Jose, California 95134 (US). **MANI, Suresh;** 250 W Tasman Dr. San Jose, California 95134 (US).
- (74) **Agent: CHINA PAT INTELLECTUAL PROPERTY OFFICE;** 2nd Floor, Zhongguancun Intellectual Property Building, Block B, No.21 Haidian South Road, Haidian District, Beijing 100080 (CN).
- (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,

(54) **Title:** METHOD AND SYSTEM FOR PROVIDING REMOTE ROBOTIC CONTROL




Figure 1

(57) **Abstract:** A virtual pointer object (118) is displayed within a virtualized environment corresponding to a physical environment currently surrounding a robot (102). The virtualized environment is generated and updated in accordance with streaming environment data received from sensors (106) collocated with the robot (102). First user input is detected via a haptic-enabled input device (114) that causes the virtual pointer object (118) to move along a movement path in the virtualized environment, where the movement path is constrained by simulated surfaces in the virtualized environment. Haptic feedback is generated via the haptic-enabled input device (114) in accordance with simulated material and/or structural characteristics of the movement path. The virtualized environment is modified at the locations of marking inputs along the movement path and affects path planning for the robot (102) within the physical environment.

SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

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

METHOD AND SYSTEM FOR PROVIDING REMOTE ROBOTIC CONTROL

CROSS REFERENCE TO RELATED APPLICATION

[0001] The present application claims the priority to the US Patent Application No. 16/156,979, filed on October 10, 2018, the disclosure of which is hereby incorporated by reference in its entirety.

TECHNICAL FIELD

[0002] This relates generally to robotic control, including but not limited to providing mixed-initiative remote robotic control.

BACKGROUND

[0003] Today, robots are being used to perform or assist in many tasks that previously were performed by human beings. In some instances, robots are better suited for certain tasks than humans due to their precision in movement and control and their highly customizable designs and operation modes. Some suitable tasks for remotely controlled robots include tasks in hazardous environments, such as natural disaster areas, toxic environment, virus or radiation contaminated environment, etc. In some examples, tasks that are highly repetitive and tedious, and/or physically strenuous tasks, such as performing household chores, organizing warehouse inventories, routine care and assistance to the disabled and/or elderly, etc., can benefit greatly from robotic assistance.

[0004] Although teleoperation and complete automation are two modes of robotic control that are widely used today, the kinds of tasks that can be performed using these two modes of control are limited. For example, with teleoperation, a human operator precisely controls every move of the robot by performing a corresponding control movement in real-time while watching a streaming video of the scene in which the robot is operating. The teleoperation mode of control requires a highly skilled human operator to be fully engaged (e.g., physically and mentally) in real-time throughout performance of a task, which severely limits its usefulness and applications (e.g., due to the burdens and requirements placed on the human operators). The full automation mode of control requires artificial intelligence to be built into the robot control programs, which cannot adapt to highly varied operation scenarios and limits its usefulness and applications.

[0005] In a mixed-initiative robotic control mode, a human operator provides high-level task-oriented instructions, and relies on the robot's pre-programming to determine the exact movement and operations to be performed to achieve the high-level instructions. Mixed-initiative robotic control alleviates some burdens and requirements placed on the human operators, and takes advantage of the robot's capabilities in terms of precision, stamina, robustness, and relative immunity to hazards.

[0006] It is challenging to provide an intuitive and efficient human-machine interface that facilitates remote mixed-initiative robotic control. Good techniques for supporting indirect and/or intermittent intervention in a robot's operation (e.g., navigation and/or object manipulation) in real-time over a network are in great need.

SUMMARY

[0007] Accordingly, there is a great need for mixed-initiative robotic control techniques where a human operator controlling the robot visualizes a 3D scene as perceived by the robot in real-time and guides the robot's actions on a high-level (as opposed to controlling the robot via teleoperation) using indirect and/or intermittent instructions and inputs. When providing remote mixed-initiative control to robots, a mixed reality environment is utilized for the user to perceive the 3D scene and provide guidance inputs. The mixed reality environment is a combination of a virtualized representation of the physical environment as perceived by the robot as well as virtual markers and virtual objects that are placed into the virtualized representation of the physical environment. This technique provides the user more freedom and information during the interaction with the physical environment via the robot, such as switching viewpoints in real-time, inspecting selected portions of the environment, receiving augmentation of visual information (e.g., texture information, physical property information, motion and appearance history, etc.).

[0008] In this disclosure, embodiments of a haptic-enabled mixed reality system that provides haptic interfaces to interact with a virtualized representation of a physical environment and to provide remote high-level guidance to mobile robots in performing high-level tasks are described. The system includes a local site with a mobile robot equipped with Red Green Blue Depth (RGBD) sensors, and a remote site with a user operating a haptic-enabled input/output (I/O) device. A three-dimensional virtualized real-world scene is generated from streaming RGBD data using real-time dense mapping. The user can use the haptic-enabled input device to "touch" and experience physical characteristics of surfaces and objects within the scene, to mark the scene to indicate virtual routes, passages, obstacles, and/or targets for the robot in the scene, and to add virtual objects to the scene that alter the scene's composition from the perspective of the robot and thereby indirectly influence the exact operation that is performed by the robot in accordance with preprogrammed instructions. In some embodiments, the system performs physics simulation for a virtual object's interactions with other objects in the scene, to help the user experience the scene and more accurately determine where to place the virtual object in the scene. In some embodiments, real-time user assisted object segmentation is performed. The user uses the haptic-enabled input device to mark a location in the scene, and an object is segmented from the scene by a region growing method from the marked location. In some embodiments, a new processing pipeline for haptic rendering is used to handle proxy update. The proposed proxy update method with force shading that is reflective of friction and haptic textures is more stable and more realistically handles intersecting boundaries of different planes in the scene. In some embodiments, prediction of actual robot position is made in light of expected network latency, such that real-time operation and control of the robot over a latency network is more smooth (e.g., with less jerky, back-forth movement of the robot).

[0009] In one aspect, a method of providing mixed-initiative robotic control is performed at a computing device having one or more processors and memory, wherein the computing device is communicably coupled to a robot and is configured to generate a planned path for the robot in accordance with a first set of preprogrammed path-planning instructions, and the robot is configured to navigate within a physical environment in accordance with the planned path received from the computing device and locally-stored path-execution instructions. The method includes: displaying a control user interface via a display generation component coupled to the computing device, including displaying a virtual pointer object within a virtualized environment corresponding to a first physical environment currently surrounding the robot, wherein the virtualized environment is generated and updated in accordance with streaming environment data received from a first set of sensors collocated with the robot; while displaying the virtual pointer object in the virtualized

environment, detecting first user input via a haptic-enabled input device, including detecting a first movement input and one or more marking inputs at one or more locations along a movement path of the virtual pointer object in the virtualized environment; and in response to detecting the first user input: moving the virtual pointer object along a first movement path in the virtualized environment in accordance with the first movement input, wherein the first movement path is constrained by one or more simulated surfaces in the virtualized environment; generating haptic feedback via the haptic-enabled input device in accordance with simulated material and/or structural characteristics of the first movement path in the virtualized environment; and modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment, wherein the modified virtualized environment replaces the virtualized environment as environment basis for path planning for the robot within the first physical environment that is performed by the computing device in accordance with the first set of preprogrammed path-planning instructions.

[0010] In one aspect, a method of performing interactive object segmentation from streaming surfaces is performed at a computing device having one or more processors and memory, wherein the computing device is communicably coupled, via a network, to a set of sensors and a robot that are collocated in a first physical environment. The method includes: receiving an environment data stream from the set of sensors, the environment data stream including correlated image and depth data capturing the first physical environment; displaying, via a display generation component, a virtualized representation of the first physical environment, in a user interface region, wherein the virtualized representation of the first physical environment is generated and updated in accordance with the environment data stream; while displaying the virtualized representation of the first physical environment, receiving a marking input from a haptic-enabled input device, wherein a position of the marking input in the virtualized representation of the first physical environment is determined in accordance with the marking input and is constrained by the a first virtualized surface in the virtualized representation of the first physical environment; and in response to receiving the marking input, performing object segmentation from the position of the marking input on the correlated image and depth data, including, for a respective time-step in updating the virtualized representation of the first physical environment in accordance with the environment data stream: clustering pixels from a first position in the image data that corresponds to the position of the marking input to form a first cluster; segmenting out a first labeled object based on the first cluster and the depth data corresponding to the first cluster; fusing the first labeled object and the depth data into a first Truncated Signed Distance Function (TSDF) representation of the first physical environment corresponding to the respective time-step; and obtaining a first segmented point cloud from the first TSDF representation fused with the first labeled object and the depth data corresponding to the first cluster.

[0011] In an aspect, a method of providing mixed-initiative robotic control is performed at a computing device having one or more processors and memory, wherein the computing device is communicably coupled to a robot and is configured to generate a planned path for the robot in accordance with a first set of preprogrammed path-planning instructions, and the robot is configured to navigate within a physical environment in accordance with the planned path received from the computing device and locally-stored path-execution instructions. The method includes: displaying a control user interface via a display generation component coupled to the computing device, including displaying a virtualized environment corresponding to a first physical environment currently surrounding the robot, wherein the virtualized environment is generated and updated in accordance with streaming environment data received from a first set of sensors collocated with the robot; while displaying the virtualized environment, detecting a first user input inserting a first

virtual object at a first location in the virtualized environment; in response to detecting the first user input, modifying the first virtualized environment in accordance with the insertion of the first virtual object at the first location, wherein the first virtual object at the first location causes the robot to execute a first navigation path in the physical environment that is generated in accordance with the first set of pre-programmed path-planning instructions; while displaying the first virtual object at the first location in the virtualized environment and while the robot is executing the first navigation path in the physical environment, detecting a second user input, including detecting a first movement input directed to the first virtual object via a haptic-enabled input device; and in response to detecting the second user input: moving the first virtual object along a first movement path to a second location in the virtualized environment in accordance with the first movement input, wherein the first movement path is constrained by one or more simulated surfaces in the virtualized environment, wherein the first virtual object at the second location causes the robot to execute a modified navigation path in the physical environment that is generated in accordance with the first set of pre-programmed path-planning instructions.

[0012] In accordance with some implementations, a computing system includes one or more processors, memory, and one or more programs; the one or more programs are stored in the memory and configured to be executed by the one or more processors and the one or more programs include instructions for performing the operations of any of the methods described above. In accordance with some implementations, a non-transitory computer readable storage medium has stored therein instructions which when executed by a computing system with one or more processors, cause the computing system to perform the operations of any of the methods described above. In accordance with some implementations, a computing system includes means for performing the operations of any of the methods described above.

[0013] Additional advantages of the disclosed systems and methods are described throughout this disclosure, and/or are apparent to a person skilled in the art in light of the disclosure provided herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] For a better understanding of the various described implementations, reference should be made to the Description of Implementations below, in conjunction with the following drawings in which like reference numerals refer to corresponding parts throughout the figures.

[0015] Figure 1 is an exemplary mixed-initiative robotic control environment in accordance with some implementations.

[0016] Figure 2 is a block diagram illustrating an exemplary architecture that includes the mixed-initiative robotic control environment in accordance with some implementations.

[0017] Figure 3 illustrates an exemplary processing pipeline for real-time visual and haptic rendering on the remote site of the mixed-initiative robotic control environment, in accordance with some implementations.

[0018] Figure 4 illustrates a proxy update method using force shading, in accordance with some embodiments.

[0019] Figure 5 illustrates proxy update for friction, in accordance with some embodiments.

[0020] Figure 6 illustrates an exemplary processing pipeline for interactive region growing 2D segmentation method, in accordance with some embodiments.

- [0021] Figures 7A-7C illustrate various examples of haptic markings and virtual objects to influence route planning and operation of the robot, in accordance with some embodiments.
- [0022] Figure 8 is a flowchart of a method of providing mixed-initiative robotic control, in accordance with some embodiments.
- [0023] Figure 9 is a flowchart of a method of performing interactive object segmentation from streaming surfaces, in accordance with some embodiments.
- [0024] Figure 10 is a flowchart of a method of providing mixed-initiative robotic control, in accordance with some embodiments.
- [0025] Figure 11 is a block diagram illustrating an exemplary robot in accordance with some embodiments.
- [0026] Figure 12 is a block diagram illustrating an exemplary server in accordance with some implementations.
- [0027] Figure 13 is a block diagram illustrating an exemplary user-side device in accordance with some implementations.
- [0028] Like reference numerals refer to corresponding parts throughout the several views of the drawings.

DESCRIPTION OF IMPLEMENTATIONS

- [0029] This disclosure provides system and method for providing mixed-initiative robotic control in accordance with various embodiments.
- [0030] Mixed-initiative robotic control has advantages over teleoperation and fully autonomous robotic control using preprogrammed instructions due to its flexibility and real-time adaptability to new situations with intermittent aid of a remotely located human operator. On the one hand, the robot is programmed to perform various low level or intermediate level tasks in accordance with environment input collected via on-board sensors and/or through a network, on the other hand, a human user can intervene indirectly at suitable times by modifying the perceived environment input, in accordance with actual real-time visual and haptic rendering of the scene based on streaming RGBD data received from the sensors collocated with robot. The haptic feedback provided to the user via a haptic-enabled input device allows the user to explore, touch, and experience the scene via a haptic cursor or to interact with the scene through simulated physical interactions between a user controlled virtual object and the scene. The flexibility and adaptability of haptic-enabled mixed-initiative remote robotic control allows the system to be used in many situations that are too tedious, strenuous, and/or dangerous for human user and/or situations that are too complex and rapidly changing for fully autonomous robots. In addition, the level of user intervention is adjustable based on an actual situation in real-time, further improving the performance and adaptability of the system.
- [0031] Below are some example scenarios in which the haptic-enabled mixed-initiative remote robotic control system and method can be utilized with advantage over conventional systems.
- [0032] In one example, in taking care of patients in field hospitals in remote epidemic disease areas, robots can assist in caring for patients that are highly contagious. Simple tasks such as moving from patient to patient to check their vitals, delivering medication and water to patients at scheduled times, removing soiled linens, cleaning up trash on the floors or bedside tables, etc. Due to the varied conditions in the field hospital rooms, artificial intelligence of autonomous robots

based on machine learning and preprogramming may not work very well in such situations (e.g., for lack of training samples similar to the field hospital rooms, and lack of sufficient amount of time for training and perfecting the robot's decision-making programs). At the same time, regular hospital staff and doctors are busy with their regular duties, and may not have the physical stamina and/or undivided attention required to remotely control robots using the teleoperation control mode. In such cases, mixed-initiative remote robotic control with haptic feedback is particularly useful. It relies on the robot's existing programming for low level tasks that are well trained and programmed (e.g., navigating to an identified target along an obstacle-free path, planning and executing a path to avoid obstacles that have been identified in its field of view, and grabbing and putting down an identified target object) and it can also benefit greatly from a human user's help in identifying which objects in the scene is a target, an obstacle, a passage, and/or the best route to approach a target and/or avoid possible obstacles.

[0033] In another example, for domestic robots that help with assisting elderly people and managing household chores, the same customized robots can be shipped and delivered to different households, and the preprogrammed operation and navigation instructions may work well for standardized tasks and room layouts, such as sweeping the floor, wiping countertops, taking out trash from the trash can, etc. However, when the rooms are littered after a party, or when the household members randomly leave various types items around the house, or when there are children or pets moving about in and out of the rooms, the limited image processing and decision making capabilities afforded by the standard preprogrammed instructions will not be adequate in properly characterizing and treating these highly varied environment inputs. Furthermore, the same un-customized robots shipped to different houses (e.g., houses with different appearances and layouts, different degrees of organization and cleanliness requirements, and different sensitivity to accidental errors and collisions between the robot and the environment), haptic-enabled mixed-initiative remote robotic control allows a human user to tailor the amount and types of control and intervention provided to the un-customized robot, such that the robot performs customized tasks in its actual customized environment. In some embodiments, the human guidance history is collected overtime, such that the un-customized robot may utilize the history to self-training through machine learning and other artificial intelligence training methods to become more and more autonomous in each household.

[0034] In another example, in an industrial manufacturing setting, manual labor is largely replaced by computer-controlled machines and autonomous robots. However, certain tasks remain manually driven, due to the infrequency of the tasks and/or high sensitivities to machine errors and accidents. For infrequent tasks, it is not financially feasible to design and program the robot to specialized in performing those tasks and it is difficult to know when those specially designed robots should be deployed due to the unpredictability of when the needs for each task would arise. Furthermore, with expensive equipment, dangerous chemicals, and/or highly sensitive machinery or measurements in the operating environment, fully autonomous robots are too unreliable to be deployed and utilized without human supervision. In such cases, haptic-enabled mixed-initiative remote robotic control utilizes a human user to provide high-level guidance in real-time and to structure and guide the low-level navigation and operation capabilities of the robot, such that the robot operates with sufficient guidelines and ample margin for machine errors when performing these highly infrequent and/or varied tasks in highly sensitive environments.

[0035] There are innumerable example scenarios where haptic-enabled mixed-initiative remote robotic control would outperform conventional teleoperation, or fully autonomous robots, e.g., in terms of design, manufacturing and maintenance cost, ease of programming, flexibility and

adaptability, and ease of use. Haptic-feedback provided to the human operator through the haptic-enabled input device when the human operator interacts with the virtualized environment directly (e.g. using a virtual pointer) or through a virtual object allows the user to more accurately assess the physical scene of the robot, and more accurately provides guidance to the robot by changing, marking, or otherwise augmenting its environment inputs.

[0036] Individual features or combinations of features of the haptic-enabled mixed initiative remote robotic control system as described herein in various embodiments will further the advantages set forth above and/or provide additional advantages which will be elaborated on in more detail or will be apparent to a person skilled in the art in light the present disclosure.

[0037] Reference will now be made in detail to implementations, examples of which are illustrated in the accompanying drawings. In the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of the various described implementations. However, it will be apparent to one of ordinary skill in the art that the various described implementations may be practiced without these specific details. In other instances, well-known methods, procedures, components, circuits, and networks have not been described in detail so as not to unnecessarily obscure aspects of the implementations.

[0038] It will also be understood that, although the terms first, second, etc. are, in some instances, used herein to describe various elements, these elements should not be limited by these terms. These terms are only used to distinguish one element from another. For example, a first user interface could be termed a second user interface, and, similarly, a second user interface could be termed a first user interface, without departing from the scope of the various described implementations. The first user interface and the second user interface are both user interfaces, but they are not the same user interface.

[0039] The terminology used in the description of the various described implementations herein is for the purpose of describing particular implementations only and is not intended to be limiting. As used in the description of the various described implementations and the appended claims, the singular forms “a”, “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will also be understood that the term “and/or” as used herein refers to and encompasses any and all possible combinations of one or more of the associated listed items. It will be further understood that the terms “includes,” “including,” “comprises,” and/or “comprising,” when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

[0040] As used herein, the term “if” is, optionally, construed to mean “when” or “upon” or “in response to determining” or “in response to detecting” or “in accordance with a determination that,” depending on the context. Similarly, the phrase “if it is determined” or “if [a stated condition or event] is detected” is, optionally, construed to mean “upon determining” or “in response to determining” or “upon detecting [the stated condition or event]” or “in response to detecting [the stated condition or event]” or “in accordance with a determination that [a stated condition or event] is detected,” depending on the context.

[0041] It is to be appreciated that “remote control” may refer to control over a wired or wireless network, such as the Internet, or an intranet within an organization or household network, the exact distance or location from which the inputs from the human user is provided is not specifically limited, and the scope of the present teachings is not so limited.

[0042] It is also to be appreciated that while the terms user, human operator, administrator, technician, person, and the like may be used to refer to the person or persons acting in the context of some particularly situations described herein, these references do not limit the scope of the present teachings with respect to the person or persons who are performing such actions.

[0043] Figure 1 is an exemplary mixed-initiative robotic control environment 100 in accordance with some implementations.

[0044] In the environment 100, a robot 102 is located at a first location (e.g., local site 104). The robot 102 or the local site is equipped with a set of sensors 106 for capturing visual and depth data of the local site (e.g., Light Detection and Ranging (LIDAR), or RGBD cameras for capturing RGB colored image data and corresponding depth data of the scene). The visualized area is the region within the field of view of the set of sensors. In some embodiments, the sensors 106 are fixed in location relative to the entire scene. In some embodiments, the sensors 106 are fixed to the robot, and the field of view of the sensors will change when the robot moves within the physical environment of the local site. The robot has onboard processing capabilities to perform low level functions such as moving, grabbing, translating and rotating a grabbed object, and optionally obstacle avoidance during movement. In some embodiments, the robot has onboard processing capabilities to perform intermedia level tasks, such as folding, sweeping, wiping, pressing, opening/closing doors, shaking, stirring, tapping, loading/unloading, etc. sorting, etc. The types of tasks that the robot 102 is designed and preprogrammed to perform independently (without human intervention) are tailored to the particular application scenarios. For example, a domestic robot will be programmed with capabilities to sweep and mop a target area when the target area is clearly identified to the robot, and to pick up trash and send them to trash can when the trash items are clearly identified to the robot and the route to the trash is clear and unblocked by other unidentified objects. The robot 102 and the set of sensors 106 are further equipped with network communication capabilities with sufficient bandwidth to transmit the RBGD data collected at the local site, and to receive control instructions from a remotely located human operator 108 and/or a server 110.

[0045] At the local site, various physical objects, physical boundaries, physical surfaces, physical paths, and physical passage ways exist at different locations. The different physical objects, physical surfaces, boundaries, paths, passage ways, may also change in appearance, shape, location, and state. For example, a curtain may be open or closed, or sway in a gust of wind; a door may be opened or closed, a liquid spill may gradually spread on the floor, a patient may change his/her posture, a pet may move around, etc. The physical characteristics of the objects, boundaries, surfaces, paths, and passages include the shape, size, surface texture, weight, friction, flexibility, elasticity, plasticity, response to touch or impact (e.g., breaking, moving, deforming with or without the ability for restoration, or being immobile), etc. The objects may be identified as targets (e.g., a recipient or receptacle of an article that is being delivered, an item that is to be picked up and/or moved, an area that is to be cleaned, etc.), obstacles (e.g., furniture, pet, walls, etc.), things to steer around (e.g., puddles, toys on the floor, wires, sensitive equipment, expensive home decors, etc.), or things that can be run over, pushed away, open or through (e.g., curtains, doors, balls, carpet, etc.) during normal navigation movement of the robot, and/or things that can only withstand certain actions but other actions (e.g., carpeted area can be swept and vacuumed, but not mopped), etc.

[0046] At the remote site 105, a display generation component 112, such as a display, a projector, a heads up display or the like, is used to display a virtualized version of the physical environment captured by the set of sensors 106. The image and depth data that are streamed from the set of sensors 106 are used to generate the virtualized version of the physical environment in the form of a point cloud or other three-dimensional representation of the physical environment (e.g.,

models including representations of virtual surfaces and objects), such that the virtualized version of the physical environment that are visually presented at the remote site 105 via the display generation component 112 corresponds closely to the state of the physical environment at the local site at substantially the same time.

[0047] In addition to the display generation component 112, one or more input devices (e.g., a touch-sensitive surface, such as a touch-sensitive remote control, or a touch-screen display that also serves as the display generation component, a mouse, a joystick, a wand controller, and/or cameras tracking the position of one or more features of the user such as the user's hands) is utilized by the human operator 108 to provide inputs and instructions that will be utilized in controlling the robot 102. The one or more input devices include a haptic-enabled input device 114 (e.g., a three-dimensional haptic-enabled pointing device, a haptic-enabled glove, etc.) that generates force, motion, and/or texture feedback to the hand(s) of the human operator in accordance with simulated physical characteristics and physical interactions that occurs at a location in the virtualized environment that corresponds to the current movement and position inputs provided via the input device. For example, when the movement and position inputs provided via the input device 114 corresponds to movement along a simulated surface in the virtualized environment corresponding to the floor at the local site 104, the haptic feedback generated on the haptic-enabled input device 114 will elicit haptic sensations in the user's hands that correspond to the friction, texture, and hardness of the physical floor. When the movement and position inputs provided via the input device 114 corresponds to movement to the edge of the floor and up an abutting wall, the force feedback provided on the input device will inform the user of the physical structural relationship between the floor and the wall. When the movement and position inputs provided via the input device 114 corresponds to movement on a gravel path, the haptic feedback generated on the input device 114 will allow the user to get a sense of whether the robot will be able to navigate on the gravel path or topple over due to its unevenness. The real-time haptic feedback in conjunction with the visual rendering of the virtualized environment of the local site allows the human operator to accurately experience and assess the scene, and to provide more accurate and prudent guidance to the robot's navigation and operation. The virtual representation of the physical world combined with the haptic feedback provided during interaction with the scene give rise to a sense of "mixed-reality" because the scene is neither purely "virtual" in the conventional "virtual reality" sense (e.g., as it closely reflects the physical characteristics of the physical environment (e.g., in terms of geometry and material properties)) nor "reality" in the conventional "augmented reality" sense (e.g., as it is not a visual image of the physical environment, but rather a computed representation that combines RGB image data and corresponding depth data).

[0048] In some embodiments, the input devices at the remote site optionally supports teleoperation as well, and the human operation can temporarily switch to the teleoperation mode at any time for particular difficult or complex tasks. At other times, the human operator allows the robot to operate fully autonomously by not providing any guidance inputs or alteration to the virtualized environment inputs processed by the robot's program logic.

[0049] In some embodiments, the human operator 108 uses the haptic-enabled input device 114 to interact with virtualized surfaces and/or objects in the virtualized environment, and receive corresponding haptic feedback. Based on the haptic feedback, the user can select a location in the virtualized environment as the starting location for object segmentation. The segmentation with the aid of the user input is more accurate and less error prone. In some embodiments, some visual features in the virtualized environment are artifacts of lighting and shadows, and/or decorative patterns on surfaces. These are easily distracting to a purely computer-based object segmentation

method with no real-time human assistance, leading to inaccurate results. With the aid of real-time human guidance, the three-dimensional object segmentation is less likely to be side-tracked by these visual distractions, as object segmentation is anchored around the point of interest identified by the human operator and assisted with corresponding depth data.

[0050] In some embodiments, the human operator 108 uses the haptic-enabled input device to insert one or more virtual objects (e.g., virtual object 116) into the virtualized environment. The human operation may interact with the virtualized environment through manipulation of the virtual objects in the virtualized environment, such as moving a virtual object along a surface and/or pushing a virtual object against another surface or object in the virtualized environment. A virtual object may be given specific shape, size, and surface characteristics that result in different simulated haptic feedback to the human operator on the haptic-enabled input device 114, when the virtual object is moved along or pushed against surfaces or objects (virtualized physical objects and virtual objects) in the virtualized environment. The haptic feedback is generated based on physics simulation in accordance with the physical characteristics of the virtualized environment at locations of the interactions.

[0051] In some embodiments, the virtual objects can take on different types of roles that change the virtualized environment in different ways to influence the navigation and operation decisions of the robot in different ways. In some embodiments, a first type of virtual objects can serve as a protective shield over a virtualized physical object that is either stationary or mobile in the scene. For example, a protective dome is placed at the location of a small child or pet and moves with the child or pet in the virtualized environment. In some embodiments, a second type of virtual objects can serve as a passage indicator that overrides the presence of a physical object that is normally treated as an obstacle or barrier by the robot. For example, the robot will treat the drapery blocking a door way as an obstacle, and placing a virtual object that is a passage at the location of the drapery allows the robot to ignore the presence of the drapery and push through the drapery in navigation. In some embodiments, a third type of virtual objects can serve as a measuring aid for measuring the virtualized physical objects in the virtualized environment and/or provide more accurate route guidance. For example, the user can place down a first anchor point of a virtual object at a first corner where the floor meets two adjacent walls (e.g., the user will know where the corner is by touching the scene via the haptic enabled input device 114 and receiving the haptic feedback), and then extend the virtual object along the edge between the floor and one of the walls until reaching the other corner. The size of the virtual object will inform the user of the dimension of the room along the wall. In another example, the user needs to determine whether there is enough room between two physical objects in the scene for the robot to pass through with enough safety clearance. As it is difficult to accurately make this determination visually, the user can place a virtual object that has dimensions of the robot plus the added clearance, and the user can try dragging the virtual object through the gap between the two physical objects in the virtualized environment, the physics simulation of this interaction will produce haptic feedback to the user via the haptic-enabled input device that the virtual object is or is not blocked by the two objects in the scene. The user can determine an alternative route or plan other tasks (e.g., request the robot to moving one of the objects first) based on the haptic feedback. In another example, if a one-meter gap is needed between the hot oven and the robot, the user can place a $1 \times 1 \text{m}^2$ virtual cube in the scene (on the floor near the oven) and then push the virtual cube until it is pushed against the virtualized oven (e.g., the user will feel the resistance of the virtual cube when the virtual cube has come into contact with the virtualized oven in the scene), and virtual cube will serve as the virtual wall between the oven and the robot when the robot plans its route). In some embodiments, a fourth type of virtual objects may be a series of waypoints that the robot should reach before reaching a selected target

(e.g., a designated virtual object or virtualized physical object in the scene). As the user moves a pointer object 118 along one or more virtualized surfaces in the scene and experiencing simulated haptic feedback via the haptic-enabled input device 114, the user can choose to mark one or more points along the movement path, as guidance for the robot 102 to follow. The route guidance provided with the haptic-feedback generated in accordance with actual physical characteristics (e.g., structure, size, texture, etc.) of the environment is more accurate and less difficult for the robot to execute in most cases. For example, the user will follow the path of least resistance when dragging his/her finger or hand toward a target object in the scene, and the path will be followed easily by the robot. In other types of interfaces where such haptic feedback is not available, the user may have to draw lines based on visual information alone, and does not take into account the surface texture, slope, etc. of the path.

[0052] In addition to the equipment collocated with the robot 102 and the equipment collocated with the human operator 108, the haptic-enabled mixed initiative robotic control system includes a central computing device 120 (e.g., the server 110) that handles the extensive computation tasks related to visual and haptic data processing and rendering, and generating intermediate instructions and workflows that bridge the gap between the high-level instructions and environment inputs received from the human operator 108 and the low-level instructions executable by the robot 102. The central computing device 120 is connected to the local site equipment and/or the remote site equipment via one or more networks. In some embodiments, the central computing device 120 is collocated with the robot 102 (e.g., in an industrial manufacturing application setting, the human operator is remotely located from the robot and the factory server). In some embodiments, the central computing device 120 is collocated with the human operator 108 (e.g., in a field hospital application setting, the robot is remotely located from the doctor and the hospital server). In some embodiments, the central computing device 120 is not collocated with either the robot 102 or the human operator 108 (e.g., in a domestic assistance setting, the server is remotely located from the robot and the human operator).

[0053] In some embodiments, the central computing device 120 handles the computation related to real-time, simultaneous localization and mapping (SLAM) using real-time dense surface mapping and tracking techniques, such as KinectFusion. In some embodiments, other real-time three-dimensional modeling methods are used to generate a virtualized three-dimensional or pseudo-three-dimensional representation of the physical environment based on the streaming image and depth data collected from the local site 104. The central computing device 120 generates a point cloud during every time step of the streaming RGBD data using KinectFusion for visual rendering. In some embodiments, real-time, user-directed object segmentation is incorporated into the KinectFusion data processing pipeline, as needed. The central computing device 120 also performs haptic render and physics simulation for interactions between the user 108 and the virtualized environment, e.g., via a virtual pointer or via manipulation of a virtual object inside the virtualized environment. In some embodiments, the virtualized environment is represented by a three-dimensional point cloud. In some embodiments, the virtualized environment is represented by a three-dimensional mesh that includes simulated surfaces that correspond to physical surfaces detected in the physical environment.

[0054] In some embodiments, the central computing device 120 includes a path planner that generates a path based on the user's marking or virtual objects inserted by the user. The path planner takes into account both the characteristics of the virtualized physical environment and the user's modification and/or augmentation to the virtualized physical environment, and generates the path in accordance with preprogrammed path planning instructions. In some embodiments, the path planner

is implemented on the robot. In some embodiments, the path planner generates predicted positions and commands for the robot to follow.

[0055] Figure 2 is a block diagram illustrating an exemplary architecture 200 that includes the mixed-initiative robotic control environment in accordance with some implementations.

[0056] As described herein, reconstruction of a representation of the three-dimensional physical space in which the robot is operating needs to be performed in real-time with fair degree of accuracy in order for the user to provide appropriate and timely intervention. In order to facilitate real-time three-dimensional exploration, segmentation, marking, and interactions with the environment, haptic-rendering and physics simulation need to be performed in real-time in conjunction with the visualization of the environment. Since the visual and depth data that needs to be transmitted and processed is voluminous in nature, network latency needs to be managed in order to provide smooth, real-time operation and control based on such data communicated over networks with various latencies.

[0057] To address the above technical requirements, a real-time efficient and robust mixed reality platform for mixed-initiative control is proposed to enable haptic interactions with streaming data. A Truncated Signed Distance Function (TSDF)-based haptic rendering method with streaming surfaces is proposed to ensure the smooth and robust haptic interaction with virtualized static scenes. An interactive object segmentation method is proposed to segment objects quickly and accurately. Different types of haptic interactions are introduced in the mixed reality platform, and a robot state prediction method is proposed to compensate network delays.

[0058] In some embodiments, Fusion-based methods (e.g., KinectFusion) are used to perform localization and mapping of streaming visual and depth data. According to KinectFusion, the streaming RGBD data from the three-dimensional camera are fused and saved as a Truncated Signed Distance Function (TSDF). KinectFusion can provide the full-scene dense geometry to enable mixed reality. The present disclosure describes visual and haptic rendering with streaming surfaces generated by KinectFusion using an improved haptic rendering pipeline including collision detection, proxy update and force computation. This improved method is computationally efficient and integrates well with the KinectFusion framework. The improved method also works well for intersecting boundary of two or more planes, ameliorating the instability of previously disclosed methods for such calculations. This is especially significant when working with real-world scenes which include complex geometry with many intersecting boundaries. When performing object segmentation, the improved method allows user interaction with the reconstructed object surface in the scene, and performs segmentation in real-time, rather than through offline semantic classifications. In the present disclosure, haptic interaction is used to select a target object, an interactive region growing method is used for object segmentation using both a color image and a depth image, and a resulting labeled image is fused into the TSDF data structure.

[0059] As shown in the Figure 2, an exemplary mixed reality system architecture 200 comprises three layers: a robot layer 202, a user layer 204, and an execution layer 206.

[0060] In some embodiments, the robot layer 202 corresponds to a mobile robot (e.g., robot 102) and a set of RGBD 3D sensors (e.g., 3D camera 106) placed on the top of the robot. This layer is responsible for collecting the color and depth images in real-time, and sending the data to the execution layer 206 for scene reconstruction, and visual and haptic rendering. A low-level task executor in this layer executes the control commands that are sent by the controller in the execution layer 206 to perform navigation and object manipulation tasks.

[0061] In some embodiments, the execution layer 206 receives the RGBD images (e.g., streamed from the RGBD sensors on top of the robot 102 over a network) and performs simultaneous localization and mapping (SLAM) using a KinectFusion module. The KinectFusion module generates a point cloud every time step for visual rendering on a display monitor. The execution layer 206 combines an object segmentation module with the KinectFusion module to segment and mark an object in the scene as required by the user. The execution layer 206 also includes a separate thread for haptic rendering that is performed in parallel with the visual rendering (e.g., generation of the point cloud). A physics simulation module handles the situation that the haptic interaction interface is used to interact with a virtual object in the scene and/or to interact with the scene through manipulation of a virtual object in the scene. The haptic rendering and physics simulation produces haptic feedback data that controls the haptic-enabled input device (e.g., input device 114) held by the user (e.g., user 108). The haptic feedback includes position, motion, force (e.g., reactive force such as resistance to pressure and impact, and frictions), texture, vibrations, etc., output to the user's hand via the haptic-enabled input device (e.g., input device 114). In some embodiments, the execution layer 206 includes a path planner that generates a path based on the user's marking inputs and/or the virtual objects placed within the scene, in addition to the characteristics of the virtualized environment deduced from the image and depth data (e.g., as represented by the TSDF or point cloud). In some embodiments, the execution layer 206 includes a controller module that is used to generate a predicted position and commands for the robot to follow. In some embodiments, the path planner and the controller are optionally implemented in the robot layer 202 and receives the user input data and the environment data from the execution layer 206. In some embodiments, the execution layer 206 is implemented on the central computing system 120 (e.g., a server 110).

[0062] In some embodiments, the user layer 204 provides all the interaction interfaces and outputs to the user (e.g., user 108). The user can either use a teleoperation interface to directly operate the robot (e.g., directly providing low level commands and instructions), or use the haptic interfaces to interact with the virtualized three-dimensional environment. In some embodiments, a haptic guided object segmentation interface is provided to the user and is only used for segmentation. The haptic interaction interface enables the user to use haptic input to experience and realize placement and movement of virtual objects in the scene, and assigning various types of roles to the virtual objects (e.g., target, passage, obstacle, shield, guide, measure, etc.). Haptic marking is used to either define a path on a surface in the scene or mark one or more waypoints (e.g., virtual object and/or virtualized physical object) for the robot to follow, approach, and/or track. In some embodiments, the marking and movement of the virtual object is processed in real-time in the virtualized scene and followed by the robot in real-time in the physical scene. The virtual obstacle interface enables the user to add virtual objects of any form of geometries into the scene, then the path planner will search a new path in accordance with the positions and the assigned roles of the virtual objects, in conjunction with the characteristics of the scene. In addition to haptic feedback, the visual rendering also reflects the interactions with the scene by the user, either directly through a virtual pointer or indirectly through interaction with a virtual object in the scene.

[0063] Figure 3 illustrates an exemplary processing pipeline 300 for real-time visual and haptic rendering at the remote site of the mixed-initiative robotic control environment, in accordance with some implementations.

[0064] As shown in Figure 3, RGBD data 302 received from the sensors (e.g., camera 106) collocated with the robot (e.g., robot 102) is streamed over a network and provided to a graphical processing unit 302 (GPU) on a central computing system 120 (e.g., server 110). Since

KinectFusion is applied for dense mapping and localization, dense geometry is generated as the streaming surfaces of the 3D virtualized environment. In the GPU 302, TSDF update is performed, followed by collision detection, and followed by proxy update with force shading and ray casting. The resulting data includes a point cloud 306 and proxy values 308 for all haptic interaction points (HIPs). The point cloud 306 and the proxy values 308 are utilized in visual rendering on a display (e.g., display 112), and the proxy values 308 are utilized in force rendering on the haptic-enabled input device (e.g., input device 114).

[0065] In the present disclosure, a novel proxy update method with force shading is proposed, which is more efficient and guarantees the stable rendering at intersecting boundaries of different planes. Furthermore, the method allows addition of surface properties such as friction and haptic textures in haptic rendering.

[0066] Proxy update is a key-part of constraint-based haptic rendering, since the proxy is not only used to compute the force at an interaction point, but also rendered visually to the viewers. If the proxy update is not stable and smooth, the force rendering and visual rendering will not be smooth. Previously, a proxy update method that has been published uses gradient-based method to find the nearest surface point. As shown in Figure 4, Figure 4(1) on the left shows a scenario that the haptic interaction occurs on a surface with a sharp change in direction (e.g., interaction goes around a corner of a curved surface or the boundary of two intersecting planes). In this scenario, the haptic interaction point (HIP) is moved by the user from h_{i-1} to h_i , and the proxy position is changed from p_{i-1} to p_i , in accordance with previously disclosed haptic rendering methods. Since the proxy is always the nearest surface point according to the HIP in the previously disclosed method, the proxy undergoes a sudden change in position. In terms of user experience, it would feel as though the user's finger has suddenly "jumps" to the other side of the surface, and computed force is changed drastically to an almost reversed direction. This feels distracting, unrealistic, and confusing to the user.

[0067] Force shading is an algorithm for smooth haptic rendering. In this disclosure, a novel TSDF-based proxy update method with force shading is proposed. Different from previous force shading methods, the presently disclosed method focus on the TSDF data structure, which can be used in all of the fusion-based 3D reconstructions. Two scenarios are handled in the improved proxy updating method:

[0068] a. If the HIP is the first contact with the surface in an interaction, the proxy is to find the nearest surface point. Instead of the gradient-based iterative method proposed in previously disclosed methods, the task of finding the nearest surface point is integrated into the ray casting step in KinectFusion. The reason for this change is that the deformable property for the surface is not considered, and therefore the ray casting is performed after the haptic rendering. Per-pixel ray marches in TSDF to generate the point cloud for the whole surface. During this procedure, the distances between the HIP and every point on the surface are computed and saved. The nearest surface point finding problem now becomes a parallel problem that finds the minimum in the distance array. This problem can be solved through parallel reduction as disclosed in "SC07, High Performance Computing with CUDA (2007) by Mark Harris." In some embodiments, nearest surface point finding in ray casting require the following: Given a starting point h : 1. Parallelized thread: each pixel's corresponding ray; 2. Marches from minimum depth, stop when zero crossing to get surface points; 3. Compute the distance $d = |s-h|$; 4. Parallel reduction to get minimum distance, and return the corresponding surface point as the nearest point.

[0069] b. After the HIP penetrates into the surface, the subsequent proxy position needs to be updated since the HIP will penetrate further into the volume. As shown in Figure 4, the nearest surface point is not appropriate for this scenario, a more correct way is to constrain the succeeding proxy. In the presently disclosed method, the previous time step normal n_{i-1} is used to define a tangent plane (as shown in Figure 4(2)), the normal of the proxy will be computed every time step. Tracking this normal is like tracking a tangent gliding plane over the surface physically. As shown in Figure 4(2), the tangent plane T_{i-1} is "dragged" by the new proxy position h_i while attached on the surface. So, the tangent plane can be treated as a constraint plane for the proxy. First, a perpendicular is dropped from h_i to this constraint plane to get a goal position g_i , which is the first approximation of the proxy. Then, the nearest surface finding in the ray casting step (see description in section a above) will be used to find the new proxy p_i . The core of this method is to use the tangent plane to constraint the new proxy in a physically plausible way, then to refine it as the nearest surface point. In some embodiments, the whole procedure for subsequent proxy update with force shading includes: 1. Based on the normal n_{i-1} , get the tangent plane T_{i-1} ; 2. Drop a perpendicular from the current HIP h_i to T_{i-1} to get g_i ; and 3. Use the nearest surface point finding in ray casting algorithm described above, with g_i being used as the first point, to compute the final proxy p_i .

[0070] In some embodiments, the improved method described herein simulate surface properties to generate simulated friction forces, and different haptic sensations of textured surfaces. In some embodiments, the friction force can be simulated by a simple change using a known friction cone. The angle α defines a cone starting from the current HIP h_i , as shown in Figure 5. The friction cone forms an interaction circle with the tangent plane from the previous proxy point p_{i-1} . In some embodiments, $\alpha = \arctan(\mu)$, where μ is a user defined friction coefficient. If the previous time step proxy p_{i-1} is inside the interaction circle (e.g., as shown in Figure 5(1)), then the new proxy will be directly set to the same value as before: $p_i = p_{i-1}$. If the previous time step proxy p_{i-1} is outside of the interaction circle (e.g., as shown in Figure 5(2)), then the goal position (approximated proxy) $g_i = c_i$, where c_i is the point closest to p_{i-1} on the interaction circle. These two scenarios correspond to simulated static friction and simulated dynamic friction. In some embodiments, the haptic texture is implemented by using a bump texture method which can generate the constraint for each point on the surface to change the normal.

[0071] In some embodiments, an interaction region growing 2D segmentation method is used to perform object segmentation with haptic-enabled user guidance. To segment a three-dimensional object from the reconstructed three-dimensional scene, a haptic-enabled segmentation interface is provided. Such an interface enables more flexible haptic interaction, e.g. haptic texture, material properties for different objects, and also provides the object position and orientation for robot grasping tasks. Many previous researches combine object detection and semantic classification into dense mapping. The presently disclosed system aims to build haptic-enabled interfaces for the mixed-initiative control, therefore the high-level semantic segmentation is not required. The presently disclosed interactive 3D object segmentation method is not only efficient, but also compatible with the popular high-level object semantic algorithms as the input.

[0072] In some cases, the straightforward way to perform object segmentation is to segment the 3D object from the 3D point cloud. It is also possible to use the KD-tree to speed up the neighbor search for points. However, these conventional methods take extra processing time. An alternative way is to perform the segmentation based on TSDF data, and save the segmentation information into the TSDF.

[0073] In the presently proposed KinectFusion pipeline (as shown in Figure 3), the depth image is fused for surface reconstruction at each time step. Thus, a two-phase algorithm as follows may be proposed, as shown in Figure 6. In the first phase, 2D segmentation is performed from both the depth image and color image. After the 2D segmentation, a label image L_i is generated. In the second phase, the segmentation is fused into the TSDF together with the depth image (e.g., in the TSDF update step in Figure 3). In this way, the segmentation is seamlessly integrated into the KinectFusion processing and reduces the overall computation time cost. Moreover, the segmentation information will be fused by weight, which generates a robust segmentation result.

[0074] In some embodiments, during the first phase of the improved method, the user uses the haptic avatar (e.g., a virtual pointer) to “touch” and mark an object of interest in the 3D scene. Then the 3D mark (e.g., with its 3D coordinates) is transformed into coordinates in the current color image. In the next time step, starting from the marked point in the image, the pixels are clustered through a region growing method until there are no pixels to be added to the region of the marked point. The region is treated as a cluster, then the distances between a neighbor pixel and the cluster center is computed as the combination of two Euclidean distances as shown Equation 1 below:

$$[0075] \quad d(x_i, S) = \|I(x_i) - I(S)\|_2 + \beta \|P(x_i) - P(S)\|_2 \quad (1)$$

[0076] where x_i is the neighbor pixel’s position, and S is the center of the region. I is the CIELAB color space value of the pixel in the color image, which is widely considered as perceptually uniform for small color distances. P is the 3D coordinates which are computed from the depth image. The values for cluster center: $I(S)$ and $P(S)$ are computed as the averages of the values of all pixels in this cluster. $\beta = m/g$ is a parameter that controls the compactness of a region. m is the variable to control the compactness, g is the grid interval. The region growing with RGBD data provides better boundary of objects that only with RGB data. The greater the value of m , the more spatial proximity is emphasized and the more compact the cluster. This value can be in the range [1, 20], e.g., a value of 10 is commonly used. The distance threshold can be chosen by the user based on characteristics of the scene and/or the object of interest in the scene at the time of segmentation.

[0077] Figures 7A-7C illustrate various examples of using haptic markings and virtual objects to influence route planning, in accordance with some embodiments.

[0078] In most previous works, haptic force feedback is used to generate path guidance forces, collision preventing forces, or environmental force to improve the performance of high-level tasks operations. However, the presently disclosed system uses a haptic device in a different way. A haptic device is used as the 3D avatar to remotely touch, explore and interact with the virtualized real-world environment. The haptic interaction provides more flexible operations similar to using “virtual hands”. The haptic interfaces can intervene the robot control procedure, and add a new path or change destinations. These interfaces will not influence the velocity of the robot, but only the paths and target points.

[0079] In some embodiments, a haptic marking interface (e.g., user interface 700) is provided for providing path guidance to the robot, since haptic rendering for surfaces in the scene is performed in real-time and is computationally efficient. Through the haptic marking interface, a user can use an HIP to touch a surface (e.g., a floor) in the scene to mark a path for the robot to follow. Then, the control manager in the execution layer takes this marked path as input to invoke the path planning. In some embodiments, the marking is saved as an ordering point set, and saved separately in the remote server. In some embodiments, the guidance is provided as an option to the path planner, for the path planner to consider in light of other options available in the scene. In some embodiments, when multiple options are available to the path planner, the user’s marking is given

more weight or dominant weight relative to other options available to the path planner. In this manner, the preprogrammed path planning instructions can be used normally, with or without user intervention, and the robot control can slide between fully autonomous to teleoperation, or anywhere in between, at any time, thereby improving the flexibility and effectiveness of the robotic control.

[0080] In some embodiments, the presently disclosed system provides a haptic marking interface (e.g., user interface 700) that further allows the user to add one or more virtual objects to the 3D virtualized scene to indirectly influence path and/or operations of the robot, in accordance with some embodiments. It is much easier to use the haptic-enabled device to place a virtual object at a desired location in the scene because of the haptic feedback that conveys the absolute physical characteristics of the scene at different locations (e.g., dimensions, texture, slope, etc.), and the characteristics relative to a virtual object or the robot. This enables the user to locate a suitable 3D position for the virtual object quickly and accurately, or to guide the user's the placement of the markings that are used to provide path guidance to the robot.

[0081] As shown in Figure 7A, the example haptic marking interface 700 includes a visual rendering region 702 that displays the virtual representation of the physical environment in front of the robot. The virtualized physical environment represented in the visual rendering region 702 is optionally a point cloud that is dynamically updated in real-time with the streaming image and depth data received from the sensors collocated with the robot. In some embodiments, the points displayed in the visual rendering region 702 have been segmented (e.g., using the user-assisted region growth method disclosed herein), and surfaces and boundaries associated with physical objects in the point cloud have been labeled and visually represented as such. For example, the floor is segmented from a user-specified point 805, after the user has explored the floor region in the image using the haptic enabled input device.

[0082] As shown in Figure 7A, a cubic box has been segmented out of the point cloud, and an outline 704 of the box is shown around the box. The user has marked the box as a target for the robot's next task, by selecting the outline 704 and selecting an object type of "target" for the outline. For example, by invoking a contextual menu while the box is selected, and selecting the "target" object type from the contextual menu. When moving the virtual pointer 705 on the surface of the box using a haptic-enabled input device, the haptic feedback provided on the haptic-enabled input devices convey the geometry of the box (e.g., turning around the intersection of abutting planes of the box).

[0083] As shown in Figure 7A, the user has also placed a series of waypoints 708 on the surface of the floor shown in the visual rendering region 702. For example, the user can select the object type "waypoint" from the object type selection region 706 (e.g., a drop-down menu of various virtual object types, including waypoint, target, obstacle, shield, passage, measure, etc.). Then, the user can move the virtual pointer along the floor surface and selectively putting down a series of waypoints along a movement path of the virtual pointer. The haptic feedback provided on the haptic-enabled input devices convey the friction and surface geometry of the floor along the movement path of the virtual pointer, and also constrain the movement path of pointer (e.g., the pointer cannot go through the floor, and the smoothness and speed of the pointer movement is constrained by the roughness of the floor).

[0084] In an example, the proposed system provides a haptic-enabled interface (700) from which the user can select or specify one or more virtual objects (e.g., virtual objects with preset dimensions or shapes, such as spheres, domes, cubes, cones, lines, walls, tables, chairs, etc., of different sizes) that user can drag into the 3D rendering of the virtualized scene (e.g., visual

rendering region 702). In some embodiments, a virtual object with the shape and size and/or other physical characteristics of the robot is also available in the interface for the user to drag into the 3D scene. In some embodiments, the user can select the virtual object and resize it within the scene by setting an anchor point on one side of the virtual object and drag another anchor point on another side of the virtual object within the 3D scene to resize it.

[0085] As shown in Figure 7A, the user can select an object type of a virtual object from the object type selection region 706. The user can also select an object shape from the object shape selection region 710. The object shape selection region displays a listing of various shapes, and an option for the user to draw a customized shape. In some embodiments, the object shape selection region includes affordances (e.g., check boxes) for specifying whether the specified shape is a solid shape or a hollow shell.

[0086] As shown in Figure 7A, a user can also specify the surface material and surface texture of the virtual object in surface material selection region 712 and surface texture selection region 714. The surface material and surface texture information is used when generating the force feedback provided to the haptic enabled input device, when the user interacts with the virtual object or when the user interacts with the environment through manipulation of the virtual object. In some embodiments, the user can also specify a weight (e.g., a mass) for the virtual object that affects the haptic rendering involving the virtual object. In some embodiments, a virtual object may be a composite object with different portions corresponding to different shapes, materials, textures, and/or weights, respectively.

[0087] In some embodiments, as shown in Figure 7A, the user interface 700 also includes a staging region 716. The user can drag a default shape into the staging region 716 and adjust the dimensions and orientation of the shape. The user can also compose a composite object by dragging multiple object shapes into the staging region and adjust the relative size, location, and orientation of the shapes. Once the user is satisfied with the size, shape, and orientation of the object in the staging region 716, the user can drag and drop it into the visual rendering region 702 (e.g., onto the floor in the scene).

[0088] In some embodiments, after the user has selected a virtual object to place into the 3D scene, the user may explore the 3D scene via the haptic-enabled input device before dropping the virtual object into the scene at a specific location in the scene. For example, after the user has selected to drop a dome into the scene, an outline of the dome is shown in the 3D scene around the user's 3D avatar (e.g., virtual pointer) in the scene. The user can drag the 3D avatar around various surfaces in the scene, to experience the slope, texture, dimensions of the scene via the haptic-enabled input device. When a user has identified a suitable location, the user can provide an input to drop the virtual dome at that location. As shown in Figure 7B, a virtual dome 718 has been dropped into the virtualized environment shown in visual rendering region 702 of the user interface 700.

[0089] In some embodiments, if the user has moved into a region that cannot accommodate the virtual dome (e.g., the space is too narrow for the dimension of the virtual dome), or if the location is already occupied by a physical object that is too large to be placed under the virtual dome, a warning feedback (e.g., visual alert) is provided to the user visually on the visual rendering of the 3D scene, and/or through haptics on haptic-enabled input device. In some embodiments, different types of virtual objects are associated with different types of restrictions on what types of space or physical characteristics of space are compatible with the placement of the virtual object. For example, a solid object can only be placed within an empty space that is larger than the solid object's dimensions; while a shell object can be placed within an empty space that is larger than the

shell object's dimensions, or a space that is occupied by a physical object that is smaller than the inner dimensions of the shell object. Also, the orientation of the virtual object is adjustable by the user before it is placed into the scene. For example, when a user tries to drag a virtual couch through a doorway in the 3D scene, the virtual couch will be blocked until the user rotates the virtual couch to its side. This is useful when the user wants to determine the robot can actually make a real couch of the same dimensions go through the doorway if provided with correct instructions (e.g., rotation first before pushing).

[0090] In some embodiments, physics simulation may be performed in addition to haptic rendering at the location of the virtual object. The user may drag a virtual robot up an incline in the 3D scene and experience the slope of the surface, and the surface friction, and optionally, some representation of the weight of the robot (e.g., visual indication of the robot's power usage for driving up the incline). When the virtual robot is set down on the surface, the physics simulation may show the robot staying in place, or slide down the incline, or topple over depending on how steep the incline is.

[0091] In some embodiments, when the virtual object is placed on a simulated physical surface in the 3D scene that corresponds to a physical surface in the real-world, the ground plane is located and saved at the first several time steps. The virtual object can be assigned one or more roles, such as obstacle, shield, wall, target, passage, safety cushion, etc. The path planner will regenerate the new path to accommodate the changes to the virtualized scene with the addition of the virtual object. For example, an insertion of an obstacle type virtual object may cause the path planner to bend an original path for the robot to go around the obstacle. An insertion of a shield type virtual object that is placed over a physical object moves with the physical object in the 3D scene, and the path planner ensures that the robot is always outside of the shield. An insertion of a wall means that the path planner will allow the robot to be on either side of the wall, but cannot move across the wall. An insertion of a target type virtual object causes the path planner to modify a path to approach the target virtual object. In some embodiments, when multiple virtual objects are designated as target objects, the user optionally assign an order of the target objects via the haptic marking interface such that the path planner approaches the target objects in the specified order. In some embodiments, when the order is not explicitly specified otherwise, the path planner plans the path in accordance with predefined criteria set in the preprogrammed path planning instructions (e.g., based on the order that the virtual objects are inserted into the scene, based on the shortest path to reach all target objects, based on increasing/decreasing sizes of the virtual objects, etc.). In some embodiments, a target type virtual object is placed on a physical object or a portion of a physical object in the 3D scene, such that when the robot is instructed to manipulate the virtual object (e.g., grab and lift, or grab and rotate), the physical object or the portion of the physical object at the location of the virtual object is manipulated accordingly by the robot. In a more specific example, a target virtual object is placed on an edge of a revolving door, and when the robot approaches the edge of the revolving door and pushes on the target virtual object in accordance with instructions received from the path planner and controller, the revolving door will swing open. If no such virtual object is placed on the edge of the revolving door or when the virtual object is placed in the middle of the revolving door, the robot will not approach the revolving door (e.g., for seeing it as an obstacle) or will crash into the revolving door (e.g., for failing to opening it by pushing in the middle of the revolving door as opposed to the edge of the revolving door). In the above example, the preprogrammed algorithms and logic for generating the path plan and object manipulation instructions remain unchanged, and the resulting instructions sent to the robot are different because of the difference in environmental information received by the robot (e.g., the difference introduced by the virtual object(s)' presence/location(s)). In some embodiments, an insertion of a passage type of virtual object at a

location of a virtualized physical object in the 3D scene causes the path planner to disregard the physical object as an obstacle in path planning. For example, the path planner will avoid going through a doorway covered by drapery in accordance with its preprogrammed obstacle avoidance instructions; however, when a passage type virtual object is placed on the drapery via the haptic marking interface by a user, the path planner will no longer treat the drapery as an obstacle when planning a path toward a target object on the other side of the door way. If the drapery happens to partially block the robot, the robot will push through the drapery to go through the doorway (e.g., the closed or half open drapery will be pushed open to let through the robot) when executing the path. In some embodiments, an insertion of a passage type virtual object that is placed over a physical object moves with the physical object in the 3D scene, such that the path planner does not treat the physical object as an obstacle irrespective of the change in location of the physical object. For example, the path planner will avoid a small rubber ball lying on the floor in accordance with its preprogrammed obstacle avoidance instructions; however, when a passage type virtual object is placed and attached on the virtualized rubber ball via the haptic marking interface by a user, the path planner will no longer treat the rubber ball as an obstacle when planning a path toward a target object in the scene. If the rubber ball happens to be in the path, the robot will approach and push against the ball (which will roll away or be pushed along) when executing the path.

[0092] In some embodiments, the haptic-marking interface allows the user to manipulate a virtual object already placed in the virtualized 3D scene, and provide haptic feedback to the user via the haptic-enabled input device that mimics reactions from the 3D environment through the virtual object.

[0093] In some embodiments, after a virtual object is placed on or against a surface in the virtualized physical environment, the user can interact with the virtual object, and experience different haptic feedback depending on the manner by which the user is interacting with the virtual object, and the spatial relationship between the virtual object and the virtualized physical environment, as well as physical characteristics of the virtual object and the virtualized physical environment. For example, a virtual object can be placed on a floor surface in the virtualized physical environment, and pushed or dragged on the floor surface; and the haptic-enabled input device provides haptic feedback that simulates the static and dynamic frictions between the virtual object and the floor surface. In another example, a virtual object can be pushed against a surface (e.g., sideways against a wall or another object, or downward against the floor, etc.), and the haptic-enabled input device provides haptic feedback that simulates the reactive force from the surface against the virtual object.

[0094] In Figure 7B, a few examples are provided to illustrate the influence to path planning by the placement of various types of virtual objects (e.g., dome 718) into a virtualized physical environment, in accordance with some embodiments. Figure 7B(a) illustrates a path 720-1 (e.g., a route of least resistance and distance) of a robot 102 would take to reach a target object 704. The path 720-1 is generated by the preprogrammed path-planning instructions of the path planner implemented at the robot or the server in accordance with the environment information provided to the path planner. Figure 7B(b) illustrates a path 720-2 of the robot 102, when a series of waypoints 708 are laid down on the floor of the virtualized environment. The path 720-2 first goes through the series of waypoints and then turns toward the target object 704. The environment information provided to the path planner is changed from the scenario in Figure 7B(a), and resulting in the change in the planned path. Figure 7B(c) illustrates that a dome shaped virtual object 718-1 is placed in the scene, and the dome is of the “obstacle” object type. The path 720-3 is generated to avoid the obstacle 718-1 when moving toward the target 704. Even though the virtual dome does not

exist in the physical environment, the environment information provided to the path planner is modified to include the presence of the dome, and as a result, the path planner generates a path that goes around the dome. Figure 7B(d) illustrates that a dome-shaped shield 718-2 is placed over a pet that moves about the room. The dome-shaped shield 718-2 follows the pet's movement in real-time, and the path planner adjusts the path in accordance with the location of the dome shield 718-2. The environment information is continuously updated with the location of the pet (e.g., geometric center of the pet), but the calculation is greatly simplified in real-time, because even though the pet's shape continues to change (e.g., moving, running, rolling, etc.), the details of the pet's shape is replaced with the dome shield's shape, and the path planner will plan a path to steer clear of the dome, as opposed to the pet. Figure 7B(e) illustrates that, there is a crack 722 on the floor which appears to be an obstacle to the path planner because the continuity of the floor is disrupted at the crack. However, when a passage dome object 718-3 is placed on a portion of the crack 722, the environment information provided to the path planner is modified, and the gap in the floor surface is bridged, and the path 720-5 is goes through the dome passable object over the crack, and goes directly toward the target 704. Figure 7B(f) illustrates that, there are some obstacles 724 on the way to the target. The user uses the dome shaped measure object 718-4 to push through the gaps between obstacles. The size of the measure object 718-4 is selected to be the larger of the robot and the target, if the task is to move the target to a different location. In one location, the measure object 718-4 is stuck (e.g., the user will not be able to push through), and in another location, the measure object can pass through without problem. Thus, the user can, after trying out the two locations by pushing the dome measure object 718-4 against the two openings, either reset the object type at the narrow opening to obstacle, or reset the object type at the wide opening to passage, to cause the robot to move to the target through the wide opening. The path 720-6 is generated based on the modified environment information provided to the path planner.

[0095] Figure 7C illustrates movement of a virtual object (e.g., a cube 726) in the virtualized physical environment, in accordance with some embodiments. As shown in Figure 7C, the virtual object is pushed along the floor and ultimately against the box 704 in the virtualized physical environment. As the user pushes the virtual object, the movement of the virtual object is constrained by the geometry of the floor surface in the virtualized environment, and when the virtual object is pushed against the physical box in the scene, the movement of the virtual object is stopped. The haptic enabled input device provides force feedback that corresponds to the friction between the virtual object and the floor surface when the virtual object is pushed along the floor surface, and the reactive force from the physical box 704 when the virtual object is pushed against the physical box.

[0096] In some embodiments, the presently proposed mixed initiative remote robotic control system includes mechanisms that take into consideration of network latency when performing visual and haptic rendering, path planning, and/or physics simulations. In some embodiments, the latency-based adjustment is automatically triggered by the detection of latency exceeding a threshold amount (e.g., the network latency is monitored continuously) or when the user provides a request for adjusting for latency after seeing the effect of network latency (e.g., unexpected and/or unpredictable robot movement, or dithering movement of the robot) or an indicator value of the network latency (e.g., 100ms, or 200ms) through other latency monitoring systems.

[0097] The reasons that latency is an important consideration in path planning and operation control of robot for mixed-initiative remote robotic control include the following: (1) the amount of image and depth data that need to be transmitted over the network is very large; and (2) the user intervention is not continuous, but intermittent and unscheduled, and the robot has much time

executing autonomous motion. To accommodate these characteristics, the control architecture of the system will be distributed in both the robot layer and the execution layer.

[0098] In some embodiments, the execution layer (e.g., execution layer 206) includes a task planner, a path supervisor, a path planner, and a primitive supervisor. The robot layer (e.g., robot layer 204) includes a trajectory planner, a controller and the robot. A user uses haptic interfaces to invoke high-level tasks, including haptic marking a position or insertion of a virtual object in the 3D scene, and haptic interaction with a virtual object. These operations are passed to the task planner. The task planner is a high-level manager to communicate with the plan supervisor. It can parse the task into micro-action plans, and receive the re-planning request. The plan supervisor can request and receive the path between two points from the path planner. In the presently disclosed framework, the path generation is based on a Rapidly-exploring Random Tree algorithm, in some embodiments. In a mixed-initiative control scenario, the human-in-the-loop interactions happen in the primitive supervisor module. The low level primitive supervisor receives the path information such as waypoints and micro actions from the plan supervisor. It will receive the planned path, and also the haptic marking path, and generate a goal position for the robot motion. In the robot layer, the trajectory planner monitors and controls the trajectory towards the goal position. The haptic marking path provides a marking point x^m , and the planned path provides a path point x^p . The goal position x^g is chosen from these two points by choosing the maximal distance between the point and current robot position.

[0099] Network delays may influence the mapping and localization from KinectFusion. To compensate the delay, a method to generate a predicted goal position is proposed. Assuming current velocity of robot is $v_i = (a_i, b_i, c_i)$ at i_{th} time step, a way to predict the next velocity is to compute the velocity and acceleration with the last several frames. Most Kalman filters are based on an empirical model of this linear form. A general linear model used to predict next velocity v_{i+1} is Equation 2:

$$\begin{aligned} a_{i+1} &= \alpha_0 a_i + \alpha_1 a_{i-1} + \dots + \alpha_m a_{i-m} \\ b_{i+1} &= \beta_0 b_i + \beta_1 b_{i-1} + \dots + \beta_m b_{i-m} \\ c_{i+1} &= \gamma_0 c_i + \gamma_1 c_{i-1} + \dots + \gamma_m c_{i-m} \end{aligned} \quad (2)$$

[00100] For a given time series of points in a path, the matrix V is defined as Equation 3:

$$V = \begin{bmatrix} a_0 & \dots & a_m & b_0 & \dots & b_m & c_0 & \dots & c_m \\ a_1 & \dots & a_{m+1} & b_1 & \dots & b_{m+1} & c_1 & \dots & c_{m+1} \\ \vdots & & & & & & & & \\ a_i & \dots & a_{i+m} & b_i & \dots & b_{i+m} & c_i & \dots & c_{i+m} \end{bmatrix} \quad (3)$$

[00101] Let v be the predicted position $(v_{m+1}, v_{m+2}, \dots, v_{i+m+1}, \dots)^T$. The problem now is to solve and obtain three parameter vectors α , β , and γ . The general solution of these linear problems are shown as follows in Equation 4:

$$\begin{aligned} \alpha &= (V^T V)^{-1} V^T a \\ \beta &= (V^T V)^{-1} V^T b \\ \gamma &= (V^T V)^{-1} V^T c \end{aligned} \quad (4)$$

[00102] Every time step, this linear prediction model generates new parameters, and then predicts the next goal position $x_{i+1}^g = x_i^g + v_{i+1} t$, where t is the round time delay. This goal position will be sent to the trajectory planner for the low-level autonomous control of the robot.

[00103] Figure 8 is a flowchart of a method 800 of providing mixed-initiative robotic control, in accordance with some embodiments. The method 800 of providing mixed-initiative robotic control is performed (802) at a computing device having one or more processors and memory. The computing device communicably coupled to a robot and is configured to generate a planned path for the robot in accordance with a first set of preprogrammed path-planning instructions. The robot is configured to navigate within a physical environment in accordance with the planned path received from the computing device and locally-stored path-execution instructions. In some embodiments, the computing device is a server that is collocated with the robot (e.g., using the same local area network or in the same room). In some embodiments, the computing device is a server that is remotely located from the robot, and communicates with the robot over a wide area network, and may experience network delays when communicating with the robot over the network. In some embodiments, the computing device is a computer that is collocated with a human operator that facilitates remote interaction between the user and the robot.

[00104] In accordance with the method 800, the computing device displays (804) a control user interface (e.g., a control interface 700 that includes a graphical rendering 702 of the virtualized scene (e.g., displaying the point cloud or surfaces identified based on the image and depth data collected from the sensors at the location of the robot, e.g., with a frame rate corresponding to substantially real-time update) and various user interface controls (e.g., controls in 706, 710, 712, 714, and 716 in user interface 700 in Figure 7A) for marking the virtualized scene and inserting virtual objects into the virtualized scene) via a display generation component (e.g., a display, a heads-up display, a projector, a touch-screen display, etc.) coupled to the computing device. The control user interface displays a virtual pointer object (e.g., virtual pointer 705) (e.g., a virtual hand, a cursor, an arrow-shaped position indicator, etc. that moves in the virtualized scene in accordance with a movement input of the user (e.g., provided by a mouse, joystick, touch-pad, haptic-enabled pointing device, etc.)) within a virtualized environment corresponding to a first physical environment currently surrounding the robot, wherein the virtualized environment is generated and updated in accordance with streaming environment data (e.g., streaming RGB image data and corresponding depth data, with a video frame rate (e.g., 10, 12, 24 frames per second)) received from a first set of sensors (e.g., RGBD camera 106) collocated with the robot (e.g., robot 102) (e.g., the first set of sensors may be affixed to the robot (e.g., on top of the robot, or at the “eye”-level of the robot, or on the ceiling of the room in which the robot is located)).

[00105] While displaying the virtual pointer object in the virtualized environment, the computing device detects (806) first user input via a haptic-enabled input device. The first user input includes a first movement input and one or more marking inputs at one or more locations along a movement path of the virtual pointer object in the virtualized environment. For example, a user may operate the haptic-enabled input device (e.g., a haptic-enabled 3D mouse 114) by moving a positioning portion of the input device in 3D, and cause corresponding movement of the virtual pointer object (e.g., pointer 705) in the virtualized environment. Based on the location of the virtual pointer object in the virtualized environment (e.g., as shown in the visual rendering region 702 of the control user interface 700) and based the haptic feedback provided to the user via the haptic-enabled input device (e.g., the haptic feedback causes the user to experience force and/or vibrations simulating reactive force (e.g., from impact, resistance to push/pull/press, etc.), dynamic and/or static friction between user’s finger and the surfaces in the physical environment, textures of surfaces, slope of the surfaces, other physical geometrical characteristics of the surfaces, etc.), the user can determine when to mark the location as a waypoint or destination point or other types of marking points (e.g., handle for object manipulation, passage for going through a physical obstacle, etc.) and provides the required marking input (e.g., clicking a button, double tapping on a touch

surface, activating a marking control in the control user interface, etc.) via the haptic-enabled input device (e.g., haptic-enabled input device 114).

[00106] In response to detecting the first user input (810): the computing device moves (812) the virtual pointer object along a first movement path in the virtualized environment in accordance with the first movement input, where the first movement path is constrained by one or more simulated surfaces in the virtualized environment. For example, the virtual pointer object can only move on the surface of the floor, and cannot penetrate the floor when moving along the floor represented in the virtualized environment.

[00107] In response to detecting the first user input (810): the computing device generates (814) haptic feedback via the haptic-enabled input device in accordance with simulated material and/or structural characteristics of the first movement path in the virtualized environment. The computing device modifies (816) at least a portion of the virtualized environment (e.g., marking the point cloud, saved as an ordered point set) that corresponds to the locations of the one or more marking inputs in the virtualized environment. The modified virtualized environment replaces the virtualized environment as environment basis for path planning for the robot within the first physical environment that is performed by the computing device in accordance with the first set of preprogrammed path-planning instructions.

[00108] In some embodiments, to generate haptic feedback via the haptic-enabled input device in accordance with simulated material and/or structural characteristics of the first movement path in the virtualized environment, the computing device performs the following operations. The computing device, in accordance with a movement of the first pointer object from a first location to a second location in the virtualized environment, updates a proxy position for generating a force feedback during the movement of the first pointer object. Updating the proxy position includes: in accordance with a determination that the first location corresponds to an initial contact with the virtualized environment along the first movement path, updating the proxy position to locate a nearest surface point from the first location; and in accordance with a determination that the first location is not the initial contact with the virtualized environment along the first movement path, updating the proxy position to locate a nearest surface point from an intersection point between a tangential constraint plane corresponding the first location and a normal line dropped to the tangential constraint plane from the first location. This is illustrated in Figure 4 and accompanying descriptions. In some embodiments, updating the proxy position to locate a nearest surface point from the first location utilizes intermediate results computed for ray-casting for visual rendering of the one or more surfaces, and is completed in parallel to the ray-casting for visual rendering of the one or more surfaces.

[00109] In some embodiments, modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes marking a path that is to be followed by the robot (e.g., by marking 308), and the planned path generated in accordance with the first set of preprogrammed path-planning instructions is the marked path (e.g., as illustrated in Figure 7B(b)).

[00110] In some embodiments, modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes marking a boundary that prevents crossing by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions does not cross the boundary.

[00111] In some embodiments, modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes placing one or more virtual obstacles (e.g., obstacle 718-1) that are avoided by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions precludes contact with the one or more virtual obstacles (e.g., as illustrated in Figure 7B(c)).

[00112] In some embodiments, modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes placing one or more virtual passages (e.g., 718-3) at locations corresponding to physical objects in the first physical environment, wherein a path generated in accordance with the first set of preprogrammed path-planning instructions when the one or more virtual passages are not present precludes contact with the physical objects by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions when the one or more virtual passages are present does not preclude contact with the physical objects by the robot during navigation (e.g., as illustrated in Figure 7B(e)).

[00113] In some embodiments, modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes segmenting the virtualized environment from the locations of the one or more marking inputs in the virtualized environment, and marking one or more physical objects corresponding to result of the segmenting as targets for the navigation in accordance with the first set of preprogrammed path-planning instructions (e.g., as illustrated in Figure 7A, box 704 is marked as a target). In some embodiments, the computing device provides a high-level instruction regarding manipulating the one or more physical objects in accordance with a second set of preprogrammed object-manipulation instructions, wherein the robot manipulates the one or more physical objects in accordance with the high-level instructions and locally-stored task-execution instructions. For example, the high-level instruction is “rotate the box”, and the robot rotates the box in accordance with locally-stored task-execution instructions that includes instructions for approaching the box, identifying surfaces of the box, identifying firm locations on the surfaces of the box for pushing the box, and pushing the identified locations to rotate the box to a desired amount.

[00114] It should be understood that the particular order in which the operations in method 800 have been described is merely an example and is not intended to indicate that the described order is the only order in which the operations could be performed. One of ordinary skill in the art would recognize various ways to reorder the operations described herein. Additionally, it should be noted that details of other processes described herein with respect to other methods described herein (e.g., methods 900 and 1000) are also applicable in an analogous manner to method 800 described above with respect to Figure 8. For brevity, these details are not repeated here.

[00115] The operations in the information processing methods described above are, optionally implemented by running one or more functional modules in information processing apparatus such as general purpose processors or application specific chips.

[00116] Figure 9 is a flowchart of a method 900 of performing interactive object segmentation from streaming surfaces, in accordance with some embodiments. The method 900 of performing interactive object segmentation from streaming surfaces is performed (902) at a computing device having one or more processors and memory. The computing device is communicably coupled, via a network, to a set of sensors and a robot collocated in a first physical environment. The computing device communicably coupled to a robot and is configured to generate

a planned path for the robot in accordance with a first set of preprogrammed path-planning instructions. The robot is configured to navigate within a physical environment in accordance with the planned path received from the computing device and locally-stored path-execution instructions. In some embodiments, the computing device is a server that is collocated with the robot (e.g., using the same local area network or in the same room). In some embodiments, the computing device is a server that is remotely located from the robot, and communicates with the robot over a wide area network, and may experience network delays when communicating with the robot over the network. In some embodiments, the computing device is a computer that is collocated with a human operator that facilitates remote interaction between the user and the robot.

[00117] In method 900, the computing device receives (904) an environment data stream from the set of sensors, the environment data stream including correlated image and depth data capturing the first physical environment. The computing device displays (906), via a display generation component (e.g., display 112), a virtualized representation of the first physical environment, in a user interface region (e.g., user interface region 702 in Figures 7A-7C), wherein the virtualized representation of the first physical environment is generated and updated in accordance with the environment data stream. While displaying the virtualized representation of the first physical environment, the computing device receives (908) a marking input from a haptic-enabled input device. A position of the marking input in the virtualized representation of the first physical environment is determined in accordance with the marking input and is constrained by the a first virtualized surface in the virtualized representation of the first physical environment. In response to receiving the marking input (910), the computing device performs (912) object segmentation from the position of the marking input on the correlated image and depth data, including, for a respective time-step in updating the virtualized representation of the first physical environment in accordance with the environment data stream: clustering (912) pixels from a first position in the image data that corresponds to the position of the marking input to form a first cluster; segmenting out a first labeled object based on the first cluster and the depth data corresponding to the first cluster; fusing the first labeled object and the depth data into a first Truncated Signed Distance Function (TSDF) representation of the first physical environment corresponding to the respective time-step; and obtaining a first segmented point cloud from the first TSDF representation fused with the first labeled object and the depth data corresponding to the first cluster. This is illustrated in Figure 6 and accompanying descriptions.

[00118] In some embodiments, providing the first segmented point cloud including the first labeled object as input for path planning and/or task execution to be performed by a robot with respect to a first physical object corresponding to the first labeled object, in accordance with preprogrammed instructions.

[00119] In some embodiments, the computing device generates haptic feedback via the haptic-enabled input device in response to the marking input, including: in accordance with a determination that the location of the marking input corresponds to an initial contact with the virtualized environment, updating a proxy position to locate a nearest surface point from the location of the marking input; and in accordance with a determination that the location of the marking input is a successive location of contact with the virtualized environment along a movement path, updating the proxy position to locate a nearest surface point from an intersection point between a tangential constraint plane corresponding the first location and a normal line dropped to the tangential constraint plane from the first location. This is illustrated in Figure 4 and accompanying descriptions. In some embodiments, updating the proxy position to locate a nearest surface point

from the first location is performed using the first TSDF representation of the first physical environment corresponding to the respective time-step.

[00120] In some embodiments, the computing device marks the first labeled object as an obstacle in accordance with the marking input, wherein the robot avoids the first labeled object during navigation on a path generated in accordance with the first set of preprogrammed instructions (e.g., as illustrated in Figure 7B(c)).

[00121] In some embodiments, the computing device marks the first labeled object as a target in accordance with the marking input, wherein the robot approaches the first labeled object during navigation on a path generated in accordance with a first set of preprogrammed path-planning instructions (e.g., as illustrated in Figures 7B(a)-7B(f)).

[00122] In some embodiments, the computing device marks the first labeled object as a passage in accordance with the marking input, wherein the robot avoids the first labeled object during navigation on a path generated in accordance with a first set of preprogrammed path-planning instructions when the first labeled object is not marked as a passage, and the robot will not avoid the first labeled object during navigation on a path generated in accordance with the first set of preprogrammed path-planning instructions when the first labeled object is marked as a passage (e.g., as illustrated in Figure 7B(e)).

[00123] In some embodiments, the computing device marks the first labeled object as a subject of manipulation in accordance with the marking input. The computing device provides a high-level instruction regarding manipulating the first labeled object to the robot, wherein the robot manipulates a physical object corresponding to the first labeled object the high-level instructions generated in accordance with a second set of pre-programmed object-manipulation instructions.

[00124] It should be understood that the particular order in which the operations in method 900 have been described is merely an example and is not intended to indicate that the described order is the only order in which the operations could be performed. One of ordinary skill in the art would recognize various ways to reorder the operations described herein. Additionally, it should be noted that details of other processes described herein with respect to other methods described herein (e.g., methods 800 and 1000) are also applicable in an analogous manner to method 900 described above with respect to Figure 8. For brevity, these details are not repeated here.

[00125] The operations in the information processing methods described above are, optionally implemented by running one or more functional modules in information processing apparatus such as general purpose processors or application specific chips.

[00126] Figure 10 is a flowchart of a method of providing mixed-initiative robotic control, in accordance with some embodiments. The method 1000 of providing mixed-initiative robotic control is performed (1002) at a computing device having one or more processors and memory. The computing device (e.g., server 110) communicably coupled to a robot and is configured to generate a planned path for the robot in accordance with a first set of preprogrammed path-planning instructions. The robot (e.g., robot 102) is configured to navigate within a physical environment in accordance with the planned path received from the computing device and locally-stored path-execution instructions.

[00127] In method 1000, the computing device displays (1004) a control user interface (e.g., control user interface 700) via a display generation component coupled to the computing device, including displaying a virtualized environment (e.g., in region 702) corresponding to a first physical environment currently surrounding the robot (e.g., robot 102), wherein the virtualized environment

is generated and updated in accordance with streaming environment data received from a first set of sensors (e.g., 3D camera 106) collocated with the robot. While displaying the virtualized environment, the computing device detects (1006) a first user input inserting a first virtual object at a first location in the virtualized environment (e.g., via the control user interface 700). In response to detecting the first user input, the computing device modifies (1008) the first virtualized environment in accordance with the insertion of the first virtual object at the first location, wherein the first virtual object at the first location causes the robot to execute a first navigation path in the physical environment that is generated in accordance with the first set of pre-programmed path-planning instructions. While displaying the first virtual object at the first location in the virtualized environment and while the robot is executing the first navigation path in the physical environment, the computing device detects (1010) a second user input, including detecting a first movement input directed to the first virtual object via a haptic-enabled input device. In response to detecting the second user input, the computing device moves (1012) the first virtual object along a first movement path to a second location in the virtualized environment in accordance with the first movement input, wherein the first movement path is constrained by one or more simulated surfaces in the virtualized environment, wherein the first virtual object at the second location causes the robot to execute a modified navigation path in the physical environment that is generated in accordance with the first set of pre-programmed path-planning instructions. This is illustrated in Figures 7B-7C and accompanying descriptions.

[00128] In some embodiments, during movement of the first virtual object along the first movement path to the second location in the virtualized environment, the computing device provides haptic feedback via the haptic-enabled input device (e.g., input device 114) in accordance with simulated friction forces between the first virtual object and a simulated surface at a current location of the first virtual object.

[00129] In some embodiments, the first virtual object has a predefined size and shape. During movement of the first virtual object along the first movement path to the second location in the virtualized environment, the computing device detects that the first virtual object has come in contact with a first simulated surface in a current direction of movement of the first virtual object. In response to detecting that the first virtual object has come in contact with the first simulated surface in the current direction of movement of the first virtual object, the computing device provides haptic feedback via the haptic-enabled input device in accordance with simulated reactive forces from the first simulated surface to the first virtual object. This is illustrated in Figure 4, for example.

[00130] In some embodiments, the first virtual object has a predefined size and shape. During movement of the first virtual object along the first movement path to the second location in the virtualized environment, the computing device detects that the first virtual object has come in contact with a second virtual object that was previously placed in the simulated environment in the current direction of movement of the first virtual object. In response to detecting that the first virtual object has come in contact with the second virtual object in the current direction of movement of the first virtual object, the computing device provides haptic feedback via the haptic-enabled input device in accordance with simulated reactive forces from the second virtual object to the first virtual object. This is illustrated in Figure 5, for example.

[00131] In some embodiments, the first virtual object is marked as an obstacle and the robot avoids the first virtual object during navigation on a path generated in accordance with the first set of pre-programmed instructions (e.g., as illustrated in Figure 7B(c)).

[00132] In some embodiments, the first virtual object is marked as a target and the robot approaches the first labeled object during navigation on a path generated in accordance with the first set of pre-programmed instructions (e.g., as illustrated in Figures 7B(a)-(f)).

[00133] In some embodiments, the first virtual object is marked as a passage, the presence of the first virtual object overrides presence of a physical obstacle during the robot's navigation on a path generated in accordance with the first set of preprogrammed path-planning instructions (e.g., as illustrated in Figure 7B(e)).

[00134] In some embodiments, to detect the first user input inserting the first virtual object at the first location in the virtualized environment, the computing device performs the following operations: detecting a first input directed to a first position in the virtualized environment that corresponds to a first simulated surface; in response to detecting the first input directed to the first position, providing first haptic feedback via the haptic-enabled device indicating that the first position is on the first simulated surface; detecting a second input marking the first position in the virtualized environment as part of a first boundary of the first virtual object; detecting a third input directed to a second position in the virtualized environment that corresponds to a second simulated surface; in response to detecting the second input directed to the second position, providing second haptic feedback via the haptic-enabled device indicating that the second position is on the second simulated surface; detecting a fourth input marking the second position in the virtualized environment as part of a second boundary of the first virtual object; and generating the first virtual object in accordance with the first position and the second position.

[00135] In some embodiments, the computing device marks the first labeled object as a subject of manipulation in accordance with the marking input. The computing device provides a high-level instruction regarding manipulating the first labeled object to the robot, wherein the robot manipulates a physical object corresponding to the first labeled object the high-level instructions generated in accordance with a second set of pre-programmed object-manipulation instructions.

[00136] It should be understood that the particular order in which the operations in method 1000 have been described is merely an example and is not intended to indicate that the described order is the only order in which the operations could be performed. One of ordinary skill in the art would recognize various ways to reorder the operations described herein. Additionally, it should be noted that details of other processes described herein with respect to other methods described herein (e.g., methods 800 and 900) are also applicable in an analogous manner to method 1000 described above with respect to Figure 8. For brevity, these details are not repeated here.

[00137] The operations in the information processing methods described above are, optionally implemented by running one or more functional modules in information processing apparatus such as general purpose processors or application specific chips.

[00138] Figure 11 is a block diagram illustrating an exemplary robot 102 in accordance with some embodiments.

[00139] The robot 102 includes one or more processing units (CPUs) 1102, one or more network interfaces 1104 (e.g., including the I/O interface to server 110), memory 1106, and one or more communication buses 1108 for interconnecting these components (sometimes called a chipset). The memory 1106 includes high-speed random access memory, such as DRAM, SRAM, DDR RAM, or other random access solid state memory devices; and, optionally, includes non-volatile memory, such as one or more magnetic disk storage devices, one or more optical disk storage devices, one or more flash memory devices, or one or more other non-volatile solid state storage devices. The memory 1106, optionally, includes one or more storage devices remotely located from

the one or more processing units 1102. The memory 1106, or alternatively the non-volatile memory within the memory 1106, includes a non-transitory computer readable storage medium. In some implementations, the memory 1106, or the non-transitory computer readable storage medium of the memory 1106, stores the following programs, modules, and data structures, or a subset or superset thereof:

- Operating system 1116 including procedures for handling various basic system services and for performing hardware dependent tasks;
- Network communication module 1118 for connecting the robot 102 to other computing devices;
- Presentation module 1120 for enabling presentation of information at the robot 102 via the one or more output devices 1112 (e.g., displays, speakers, etc.) associated with the user interface 1110;
- Input processing module 1122 for detecting one or more user inputs or interactions from one of the one or more input devices 1114 and interpreting the detected input or interaction;
- Robot Control Unit 1124 for controlling functions of the robot 102, including camera control unit 1113 for controlling RGBD camera 106 of the robot 102, robot function control unit 1117 for controlling the robot operation unit 1126 of the robot, path execution unit 1119 for executing a path in accordance with a high level instructions of a path planner, and other units for implementing the functions of the robot 102 as described herein.

[00140] Each of the above identified elements may be stored in one or more of the previously mentioned memory devices, and corresponds to a set of instructions for performing a function described above. The above identified modules or programs (i.e., sets of instructions) need not be implemented as separate software programs, procedures, modules or data structures, and thus various subsets of these modules may be combined or otherwise re-arranged in various implementations. In some implementations, memory 1106, optionally, stores a subset of the modules and data structures identified above. Furthermore, the memory 1106, optionally, stores additional modules and data structures not described above.

[00141] In some implementations, at least some of the functions of the robot 102 are performed by the server 110, and the corresponding sub-modules of these functions may be located within the server 110 rather than the robot 102. The robot 102 shown in Figure 11 is merely illustrative, and different configurations of the modules for implementing the functions described herein are possible in various implementations.

[00142] Each of the above identified elements may be stored in one or more of the previously mentioned memory devices, and corresponds to a set of instructions for performing a function described above. The above identified modules or programs (i.e., sets of instructions) need not be implemented as separate software programs, procedures, or modules, and thus various subsets of these modules may be combined or otherwise re-arranged in various implementations. In some implementations, the memory 1106, optionally, stores a subset of the modules and data structures identified above. Furthermore, the memory 1106, optionally, stores additional modules and data structures not described above.

[00143] Figure 12 is a block diagram illustrating an exemplary server 110 in accordance with some implementations. The server system 110, typically, includes one or more processing units (CPUs) 1202, one or more network interfaces 1204 (e.g., including the I/O interface to one or more robots 102 and the I/O interface to one or more user-side device 1300), memory 1206, and one or

more communication buses 1208 for interconnecting these components (sometimes called a chipset). The memory 1206 includes high-speed random access memory, such as DRAM, SRAM, DDR RAM, or other random access solid state memory devices; and, optionally, includes non-volatile memory, such as one or more magnetic disk storage devices, one or more optical disk storage devices, one or more flash memory devices, or one or more other non-volatile solid state storage devices. The memory 1206, optionally, includes one or more storage devices remotely located from the one or more processing units 1202. The memory 1206, or alternatively the non-volatile memory within the memory 1206, includes a non-transitory computer readable storage medium. In some implementations, the memory 1206, or the non-transitory computer readable storage medium of the memory 1206, stores the following programs, modules, and data structures, or a subset or superset thereof:

- Operating system 1216 including procedures for handling various basic system services and for performing hardware dependent tasks;
- Network communication module 1218 for connecting the server 110 to other computing devices (e.g., the user-side device 1300 and the robots 102 (including e.g., camera(s) 106));
- Presentation module 1220 for enabling presentation of information at the server 110 via the one or more output devices 1212 (e.g., displays, speakers, etc.) associated with the user interface 1210;
- Input processing module 1222 for detecting one or more user inputs or interactions from one of the one or more input devices 1214 and interpreting the detected input or interaction;
- Server-side modules 1226 for controlling functions of the server 110, including object segmentation module 1228 for performing object segmentation in the virtualized environment, KinectFusion module 1230 for generating and updating the virtualized environment based on the image and depth data stream received from the robot, Haptic rendering module 1232 for generating haptic feedback based on the user's input provided via the haptic enabled input device and based on the location of the input in the virtualized environment, physics simulation module for generating reaction and friction force rendering as well as object interaction models in the virtualized environment, path planning module 1236 for generating a planned path based on the virtualized environment and haptics markings and virtual objects present in the virtualized environment, controller module 1238 for controlling path execution by the robot, network latency compensation module 1240 for adjusting path planning based on network delays, and other modules 1242 for implementing other functions of the server 110 as described herein.

[00144] Each of the above identified elements may be stored in one or more of the previously mentioned memory devices, and corresponds to a set of instructions for performing a function described above. The above identified modules or programs (i.e., sets of instructions) need not be implemented as separate software programs, procedures, modules or data structures, and thus various subsets of these modules may be combined or otherwise re-arranged in various implementations. In some implementations, memory 1206, optionally, stores a subset of the modules and data structures identified above. Furthermore, the memory 1206, optionally, stores additional modules and data structures not described above.

[00145] In some implementations, at least some of the functions of the server 110 are performed by the robot 110 or the user-side device 1300, and the corresponding sub-modules of these functions may be located within the server robot or user-side device rather than the server 110.

The server 110 shown in Figure 12 is merely illustrative, and different configurations of the modules for implementing the functions described herein are possible in various implementations.

[00146] Each of the above identified elements may be stored in one or more of the previously mentioned memory devices, and corresponds to a set of instructions for performing a function described above. The above identified modules or programs (i.e., sets of instructions) need not be implemented as separate software programs, procedures, or modules, and thus various subsets of these modules may be combined or otherwise re-arranged in various implementations. In some implementations, the memory 1206, optionally, stores a subset of the modules and data structures identified above. Furthermore, the memory 1206, optionally, stores additional modules and data structures not described above.

[00147] Figure 13 is a block diagram illustrating an exemplary user-side device 1300 in accordance with some implementations. The user-side device 1300, typically, includes one or more processing units (CPUs) 1302, one or more network interfaces 1304 (e.g., including the I/O interface to server 110), memory 1306, and one or more communication buses 1308 for interconnecting these components (sometimes called a chipset). The memory 1306 includes high-speed random access memory, such as DRAM, SRAM, DDR RAM, or other random access solid state memory devices; and, optionally, includes non-volatile memory, such as one or more magnetic disk storage devices, one or more optical disk storage devices, one or more flash memory devices, or one or more other non-volatile solid state storage devices. The memory 1306, optionally, includes one or more storage devices remotely located from the one or more processing units 1302. The memory 1306, or alternatively the non-volatile memory within the memory 1306, includes a non-transitory computer readable storage medium. In some implementations, the memory 1306, or the non-transitory computer readable storage medium of the memory 1306, stores the following programs, modules, and data structures, or a subset or superset thereof:

- Operating system 1316 including procedures for handling various basic system services and for performing hardware dependent tasks;
- Network communication module 1318 for connecting the server 110 to other computing devices;
- Presentation module 1320 for enabling presentation of information at the user-side device 1300 via the one or more output devices 1312 (e.g., displays 112, speakers, haptic-enabled input device 114, etc.) associated with the user interface 1310;
- Input processing module 1322 for detecting one or more user inputs or interactions from one of the one or more input devices 1314 (e.g., haptic-enabled input device 114) and interpreting the detected input or interaction;
- Applications 1324 for implementing various user-level functions, such as word processing, drawing, etc.
- User-side modules 1326 for controlling functions of the user-side devices, including user layer module 1328 and other modules 1330 for implementing other functions of the user-side device as described herein.

[00148] Each of the above identified elements may be stored in one or more of the previously mentioned memory devices, and corresponds to a set of instructions for performing a function described above. The above identified modules or programs (i.e., sets of instructions) need not be implemented as separate software programs, procedures, modules or data structures, and thus various subsets of these modules may be combined or otherwise re-arranged in various

implementations. In some implementations, memory 1306, optionally, stores a subset of the modules and data structures identified above. Furthermore, the memory 1306, optionally, stores additional modules and data structures not described above.

[00149] In some implementations, at least some of the functions of the server 110 are performed by the user-side device 1300, and the corresponding sub-modules of these functions may be located within the server rather than the user-side device 1300. The user-side device 1300 shown in Figure 13 is merely illustrative, and different configurations of the modules for implementing the functions described herein are possible in various implementations.

[00150] Each of the above identified elements may be stored in one or more of the previously mentioned memory devices, and corresponds to a set of instructions for performing a function described above. The above identified modules or programs (i.e., sets of instructions) need not be implemented as separate software programs, procedures, or modules, and thus various subsets of these modules may be combined or otherwise re-arranged in various implementations. In some implementations, the memory 1306, optionally, stores a subset of the modules and data structures identified above. Furthermore, the memory 1306, optionally, stores additional modules and data structures not described above.

CLAIMS

1. A method of providing mixed-initiative robotic control, including:
 - at a computing device having one or more processors and memory, wherein the computing device is communicably coupled to a robot and is configured to generate a planned path for the robot in accordance with a first set of preprogrammed path-planning instructions, and the robot is configured to navigate within a physical environment in accordance with the planned path received from the computing device and locally-stored path-execution instructions;
 - displaying a control user interface via a display generation component coupled to the computing device, including displaying a virtual pointer object within a virtualized environment corresponding to a first physical environment currently surrounding the robot, wherein the virtualized environment is generated and updated in accordance with streaming environment data received from a first set of sensors collocated with the robot;
 - while displaying the virtual pointer object in the virtualized environment, detecting first user input via a haptic-enabled input device, including detecting a first movement input and one or more marking inputs at one or more locations along a movement path of the virtual pointer object in the virtualized environment;
 - in response to detecting the first user input:
 - moving the virtual pointer object along a first movement path in the virtualized environment in accordance with the first movement input, wherein the first movement path is constrained by one or more simulated surfaces in the virtualized environment;
 - generating haptic feedback via the haptic-enabled input device in accordance with simulated material and/or structural characteristics of the first movement path in the virtualized environment; and
 - modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment, wherein the modified virtualized environment replaces the virtualized environment as environment basis for path planning for the robot within the first physical environment that is performed by the computing device in accordance with the first set of preprogrammed path-planning instructions.
2. The method of claim 1, wherein generating haptic feedback via the haptic-enabled input device in accordance with simulated material and/or structural characteristics of the first movement path in the virtualized environment includes:
 - in accordance with a movement of the first pointer object from a first location to a second location in the virtualized environment, updating a proxy position for generating a force feedback during the movement of the first pointer object, including:
 - in accordance with a determination that the first location corresponds to an initial contact with the virtualized environment along the first movement path, updating the proxy position to locate a nearest surface point from the first location; and
 - in accordance with a determination that the first location is not the initial contact with the virtualized environment along the first movement path, updating the proxy position to locate a nearest surface point from an intersection point between a tangential constraint plane corresponding the first location and a normal line dropped to the tangential constraint plane from the first location.
3. The method of claim 1, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes marking a path that is to be followed by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions is the marked path.

4. The method of claim 1, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes marking a boundary that prevents crossing by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions does not cross the boundary.

5. The method of claim 1, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes placing one or more virtual obstacles that are avoided by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions precludes contact with the one or more virtual obstacles by the robot.

6. The method of claim 1, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes placing one or more virtual passages at locations corresponding to physical objects in the first physical environment, wherein a path generated in accordance with the first set of preprogrammed path-planning instructions when the one or more virtual passages are not present precludes contact with the physical objects by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions when the one or more virtual passages are present does not preclude contact with the physical objects by the robot during navigation.

7. The method of claim 1, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes segmenting the virtualized environment from the locations of the one or more marking inputs in the virtualized environment, and marking one or more physical objects corresponding to result of the segmenting as targets for the navigation in accordance with the first set of preprogrammed path-planning instructions.

8. A computing device for providing mixed-initiative robotic control, comprising:
one or more processors; and
memory storing instructions, wherein:
the computing device is communicably coupled to a robot and is configured to generate a planned path for the robot in accordance with a first set of preprogrammed path-planning instructions,
the robot configured to navigate within a physical environment in accordance with the planned path received from the computing device and locally-stored path-execution instructions, and
the instructions, when executed by the one or more processors, cause the processors to perform operations comprising:
displaying a control user interface via a display generation component coupled to the computing device, including displaying a virtual pointer object within a virtualized environment corresponding to a first physical environment currently surrounding the robot, wherein the virtualized environment is generated and updated in accordance with streaming environment data received from a first set of sensors collocated with the robot;
while displaying the virtual pointer object in the virtualized environment, detecting first user input via a haptic-enabled input device, including detecting a first movement

input and one or more marking inputs at one or more locations along a movement path of the virtual pointer object in the virtualized environment; and
in response to detecting the first user input:
moving the virtual pointer object along a first movement path in the virtualized environment in accordance with the first movement input, wherein the first movement path is constrained by one or more simulated surfaces in the virtualized environment;
generating haptic feedback via the haptic-enabled input device in accordance with simulated material and/or structural characteristics of the first movement path in the virtualized environment; and
modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment, wherein the modified virtualized environment replaces the virtualized environment as environment basis for path planning for the robot within the first physical environment that is performed by the computing device in accordance with the first set of preprogrammed path-planning instructions.

9. The computing device of claim 8, wherein generating haptic feedback via the haptic-enabled input device in accordance with simulated material and/or structural characteristics of the first movement path in the virtualized environment includes:

in accordance with a movement of the first pointer object from a first location to a second location in the virtualized environment, updating a proxy position for generating a force feedback during the movement of the first pointer object, including:

in accordance with a determination that the first location corresponds to an initial contact with the virtualized environment along the first movement path, updating the proxy position to locate a nearest surface point from the first location; and

in accordance with a determination that the first location is not the initial contact with the virtualized environment along the first movement path, updating the proxy position to locate a nearest surface point from an intersection point between a tangential constraint plane corresponding the first location and a normal line dropped to the tangential constraint plane from the first location.

10. The computing device of claim 8, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes marking a path that is to be followed by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions is the marked path.

11. The computing device of claim 8, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes marking a boundary that prevents crossing by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions does not cross the boundary.

12. The computing device of claim 8, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes placing one or more virtual obstacles that are avoided by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions precludes contact with the one or more virtual obstacles by the robot.

13. The computing device of claim 8, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes placing one or more virtual passages at locations corresponding to physical objects in the first physical environment, wherein a path generated in accordance with the first set of preprogrammed path-planning instructions when the one or more virtual passages are not present precludes contact with the physical objects by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions when the one or more virtual passages are present does not preclude contact with the physical objects by the robot during navigation.

14. The computing device of claim 8, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes segmenting the virtualized environment from the locations of the one or more marking inputs in the virtualized environment, and marking one or more physical objects corresponding to result of the segmenting as targets for the navigation in accordance with the first set of preprogrammed path-planning instructions.

15. A computer-readable storage medium for providing mixed-initiative robotic control, wherein the computer-readable storage medium stores instructions, the instructions, when executed by one or more processors of a computing device, cause the computing device to perform operations, wherein:

the computing device is communicably coupled to a robot and is configured to generate a planned path for the robot in accordance with a first set of preprogrammed path-planning instructions,

the robot configured to navigate within a physical environment in accordance with the planned path received from the computing device and locally-stored path-execution instructions, and the operations include:

displaying a control user interface via a display generation component coupled to the computing device, including displaying a virtual pointer object within a virtualized environment corresponding to a first physical environment currently surrounding the robot, wherein the virtualized environment is generated and updated in accordance with streaming environment data received from a first set of sensors collocated with the robot;

while displaying the virtual pointer object in the virtualized environment, detecting first user input via a haptic-enabled input device, including detecting a first movement input and one or more marking inputs at one or more locations along a movement path of the virtual pointer object in the virtualized environment; and

in response to detecting the first user input:

moving the virtual pointer object along a first movement path in the virtualized environment in accordance with the first movement input, wherein the first movement path is constrained by one or more simulated surfaces in the virtualized environment;

generating haptic feedback via the haptic-enabled input device in accordance with simulated material and/or structural characteristics of the first movement path in the virtualized environment; and

modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment, wherein the modified virtualized environment replaces the virtualized environment as environment basis for path planning for the robot within the first physical environment that is performed by the computing device in accordance with the first set of preprogrammed path-planning instructions.

16. The computer-readable storage medium of claim 15, wherein generating haptic feedback via the haptic-enabled input device in accordance with simulated material and/or structural characteristics of the first movement path in the virtualized environment includes:

in accordance with a movement of the first pointer object from a first location to a second location in the virtualized environment, updating a proxy position for generating a force feedback during the movement of the first pointer object, including:

in accordance with a determination that the first location corresponds to an initial contact with the virtualized environment along the first movement path, updating the proxy position to locate a nearest surface point from the first location; and

in accordance with a determination that the first location is not the initial contact with the virtualized environment along the first movement path, updating the proxy position to locate a nearest surface point from an intersection point between a tangential constraint plane corresponding to the first location and a normal line dropped to the tangential constraint plane from the first location.

17. The computer-readable storage medium of claim 15, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes marking a path that is to be followed by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions is the marked path.

18. The computer-readable storage medium of claim 15, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes marking a boundary that prevents crossing by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions does not cross the boundary.

19. The computer-readable storage medium of claim 15, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes placing one or more virtual obstacles that are avoided by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions precludes contact with the one or more virtual obstacles by the robot.

20. The computer-readable storage medium of claim 15, wherein modifying at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment includes placing one or more virtual passages at locations corresponding to physical objects in the first physical environment, wherein a path generated in accordance with the first set of preprogrammed path-planning instructions when the one or more virtual passages are not present precludes contact with the physical objects by the robot, and the planned path generated in accordance with the first set of preprogrammed path-planning instructions when the one or more virtual passages are present does not preclude contact with the physical objects by the robot during navigation.

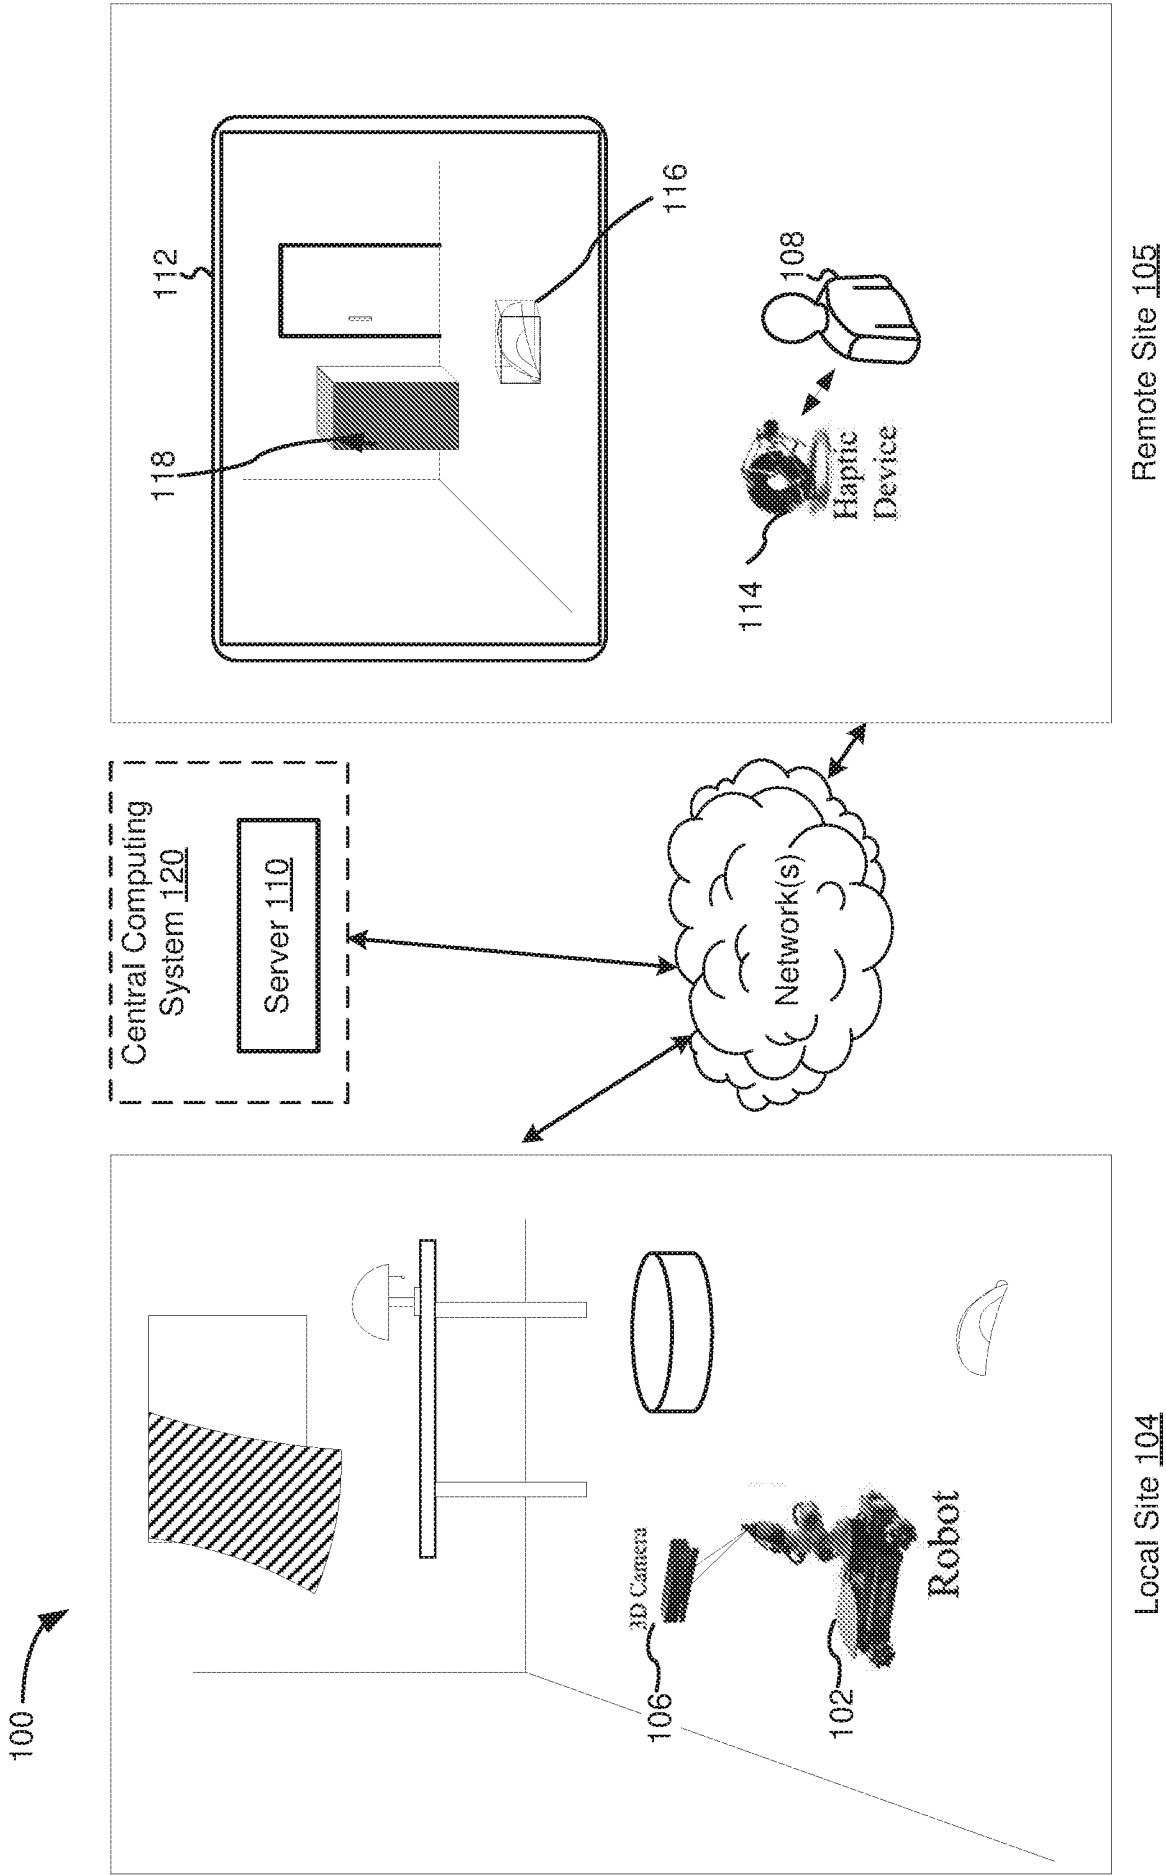


Figure 1

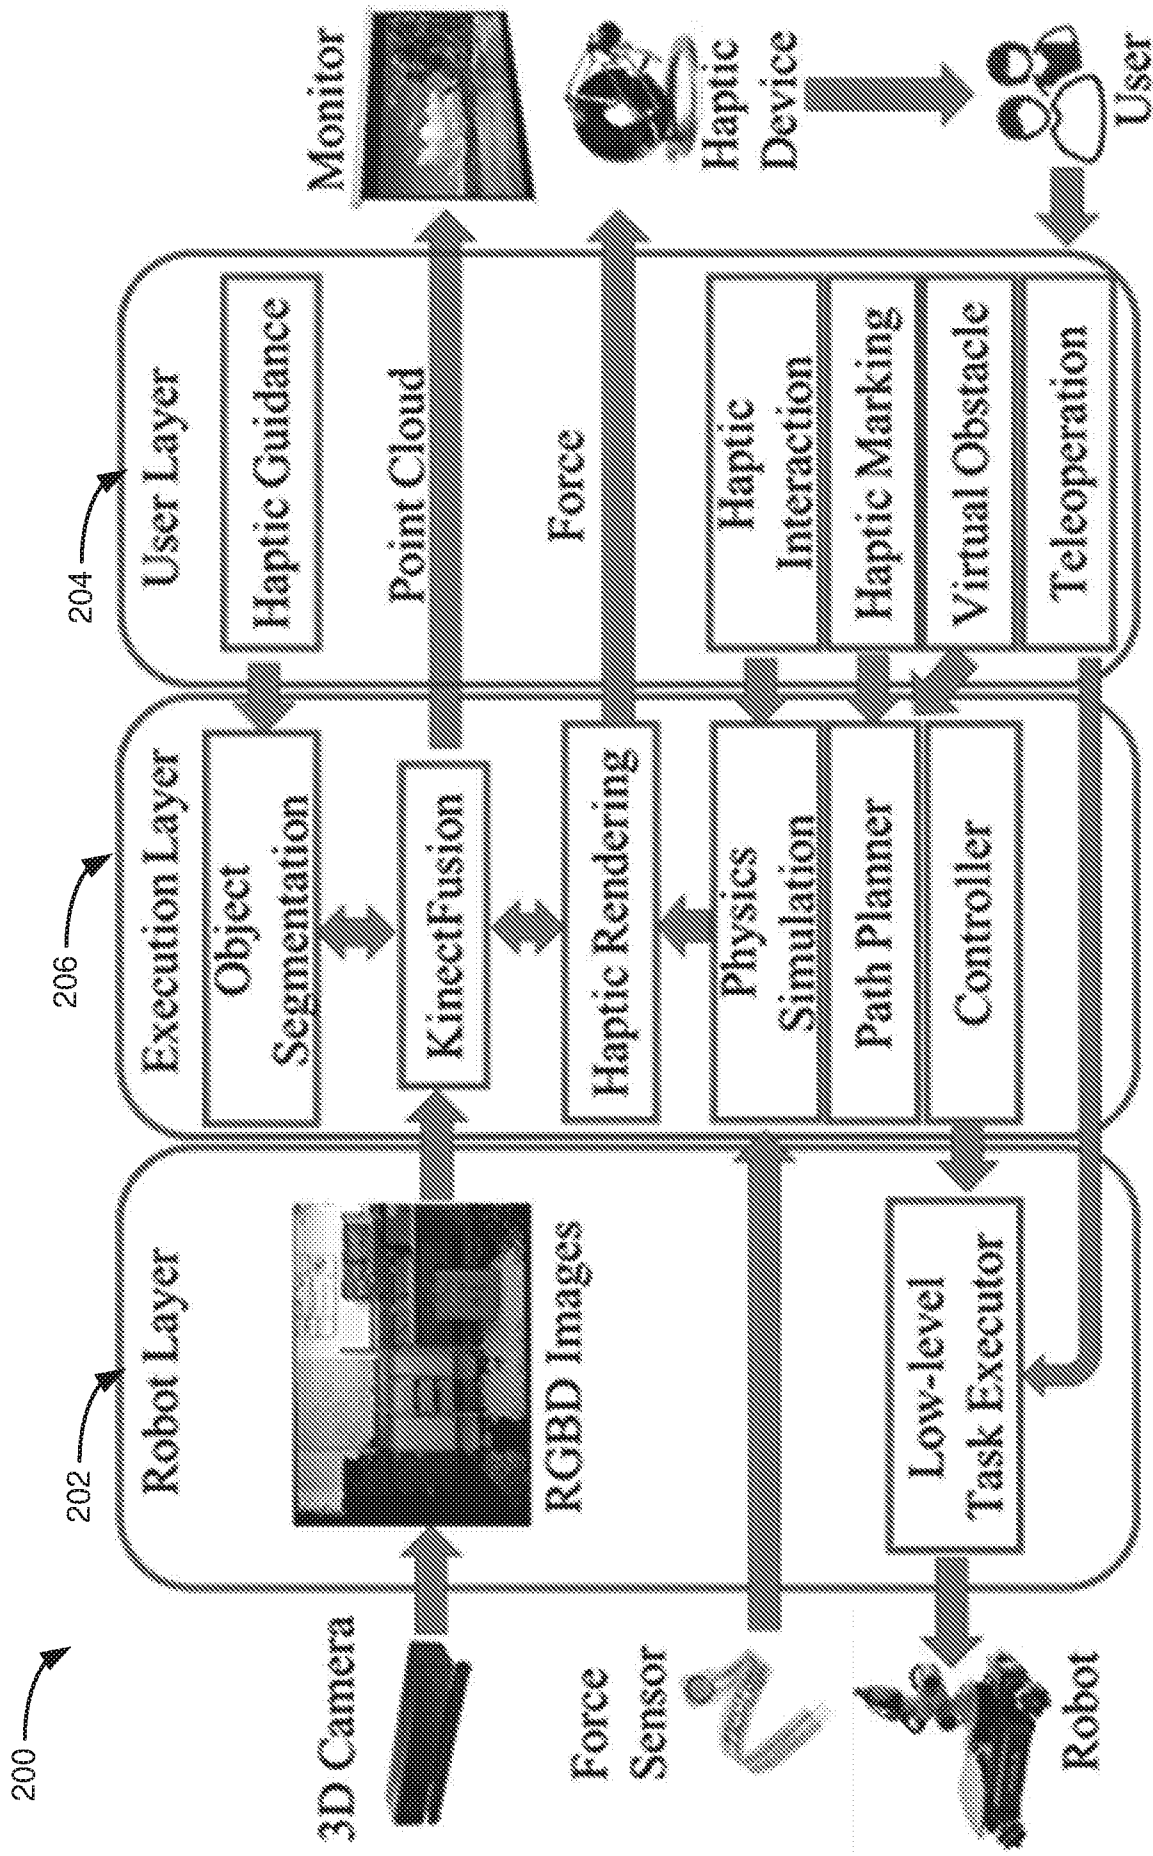


Figure 2






Figure 3



(2)

(1)

Figure 4



(2)

(1)

Figure 5




Figure 6

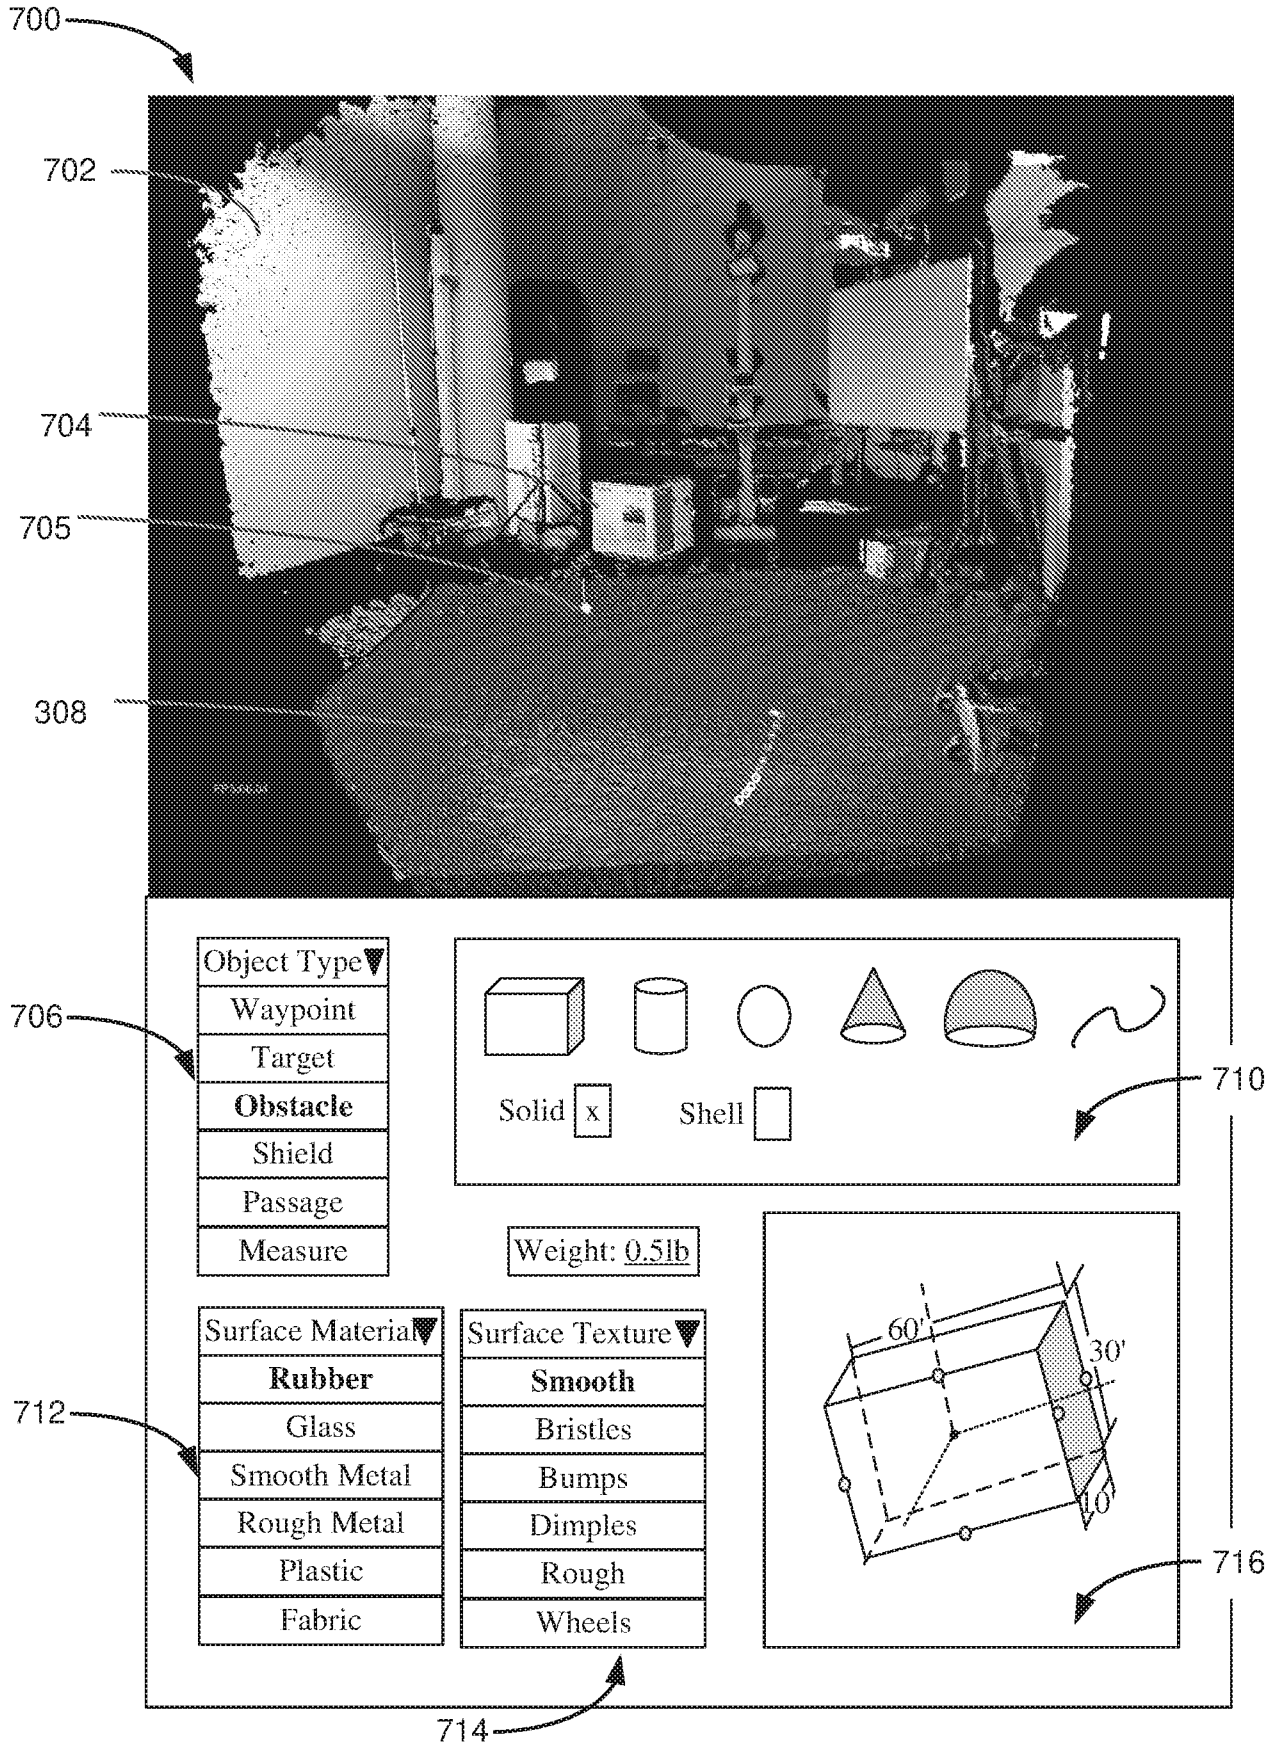


Figure 7A





Figure 7B




Figure 7C

800

802 At a computing device having one or more processors and memory, wherein the computing device is communicably coupled to a robot and is configured to generate a planned path for the robot in accordance with a first set of preprogrammed path-planning instructions, and the robot is configured to navigate within a physical environment in accordance with the planned path received from the computing device and locally-stored path-execution instructions:

804 Display a control user interface via a display generation component coupled to the computing device, including displaying a virtual pointer object within a virtualized environment corresponding to a first physical environment currently surrounding the robot, wherein the virtualized environment is generated and updated in accordance with streaming environment data received from a first set of sensors collocated with the robot

806 While displaying the virtual pointer object in the virtualized environment, detecting first user input via a haptic-enabled input device, including detecting a first movement input and one or more marking inputs at one or more locations along a movement path of the virtual pointer object in the virtualized environment

810 In response to detecting the first user input:

812 Move the virtual pointer object along a first movement path in the virtualized environment in accordance with the first movement input, wherein the first movement path is constrained by one or more simulated surfaces in the virtualized environment

814 Generate haptic feedback via the haptic-enabled input device in accordance with simulated material and/or structural characteristics of the first movement path in the virtualized environment

816 Modify at least a portion of the virtualized environment that corresponds to the locations of the one or more marking inputs in the virtualized environment, wherein the modified virtualized environment replaces the virtualized environment as environment basis for path planning for the robot within the first physical environment that is performed by the computing device in accordance with the first set of preprogrammed path-planning instructions

Figure 8

900

902 At a computing device having one or more processors and memory, wherein the computing device is communicably coupled, via a network, to a set of sensors and a robot that are collocated in a first physical environment:

904 Receive an environment data stream from the set of sensors, the environment data stream including correlated image and depth data capturing the first physical environment

906 Display, via a display generation component, a virtualized representation of the first physical environment, in a user interface region, wherein the virtualized representation of the first physical environment is generated and updated in accordance with the environment data stream

908 While displaying the virtualized representation of the first physical environment, receive a marking input from a haptic-enabled input device, wherein a position of the marking input in the virtualized representation of the first physical environment is determined in accordance with the marking input and is constrained by the a first virtualized surface in the virtualized representation of the first physical environment

910 In response to receiving the marking input, performing object segmentation from the position of the marking input on the correlated image and depth data, including, for a respective time-step in updating the virtualized representation of the first physical environment in accordance with the environment data stream:

912 Cluster pixels from a first position in the image data that corresponds to the position of the marking input to form a first cluster

914 Segment out a first labeled object based on the first cluster and the depth data corresponding to the first cluster

916 Fuse the first labeled object and the depth data into a first Truncated Signed Distance Function (TSDF) representation of the first physical environment corresponding to the respective time-step out a first labeled object based on the first cluster and the depth data corresponding to the first cluster

918 Obtain a first segmented point cloud from the first TSDF representation fused with the first labeled object and the depth data corresponding to the first cluster

Figure 9

1000

1002 At a computing device having one or more processors and memory, wherein the computing device is communicably coupled to a robot and is configured to generate a planned path for the robot in accordance with a first set of preprogrammed path-planning instructions, and the robot is configured to navigate within a physical environment in accordance with the planned path received from the computing device and locally-stored path-execution instructions:

1004 Display a control user interface via a display generation component coupled to the computing device, including displaying a virtualized environment corresponding to a first physical environment currently surrounding the robot, wherein the virtualized environment is generated and updated in accordance with streaming environment data received from a first set of sensors collocated with the robot

1006 While displaying the virtualized environment, detect a first user input inserting a first virtual object at a first location in the virtualized environment

1008 In response to detecting the first user input, modify the first virtualized environment in accordance with the insertion of the first virtual object at the first location, wherein the first virtual object at the first location causes the robot to execute a first navigation path in the physical environment that is generated in accordance with the first set of pre-programmed path-planning instructions

1010 While displaying the first virtual object at the first location in the virtualized environment and while the robot is executing the first navigation path in the physical environment, detect a second user input, including detecting a first movement input directed to the first virtual object via a haptic-enabled input device

1012 In response to detecting the second user input, move the first virtual object along a first movement path to a second location in the virtualized environment in accordance with the first movement input, wherein the first movement path is constrained by one or more simulated surfaces in the virtualized environment, wherein the first virtual object at the second location causes the robot to execute a modified navigation path in the physical environment that is generated in accordance with the first set of pre-programmed path-planning instructions

Figure 10

Robot 102

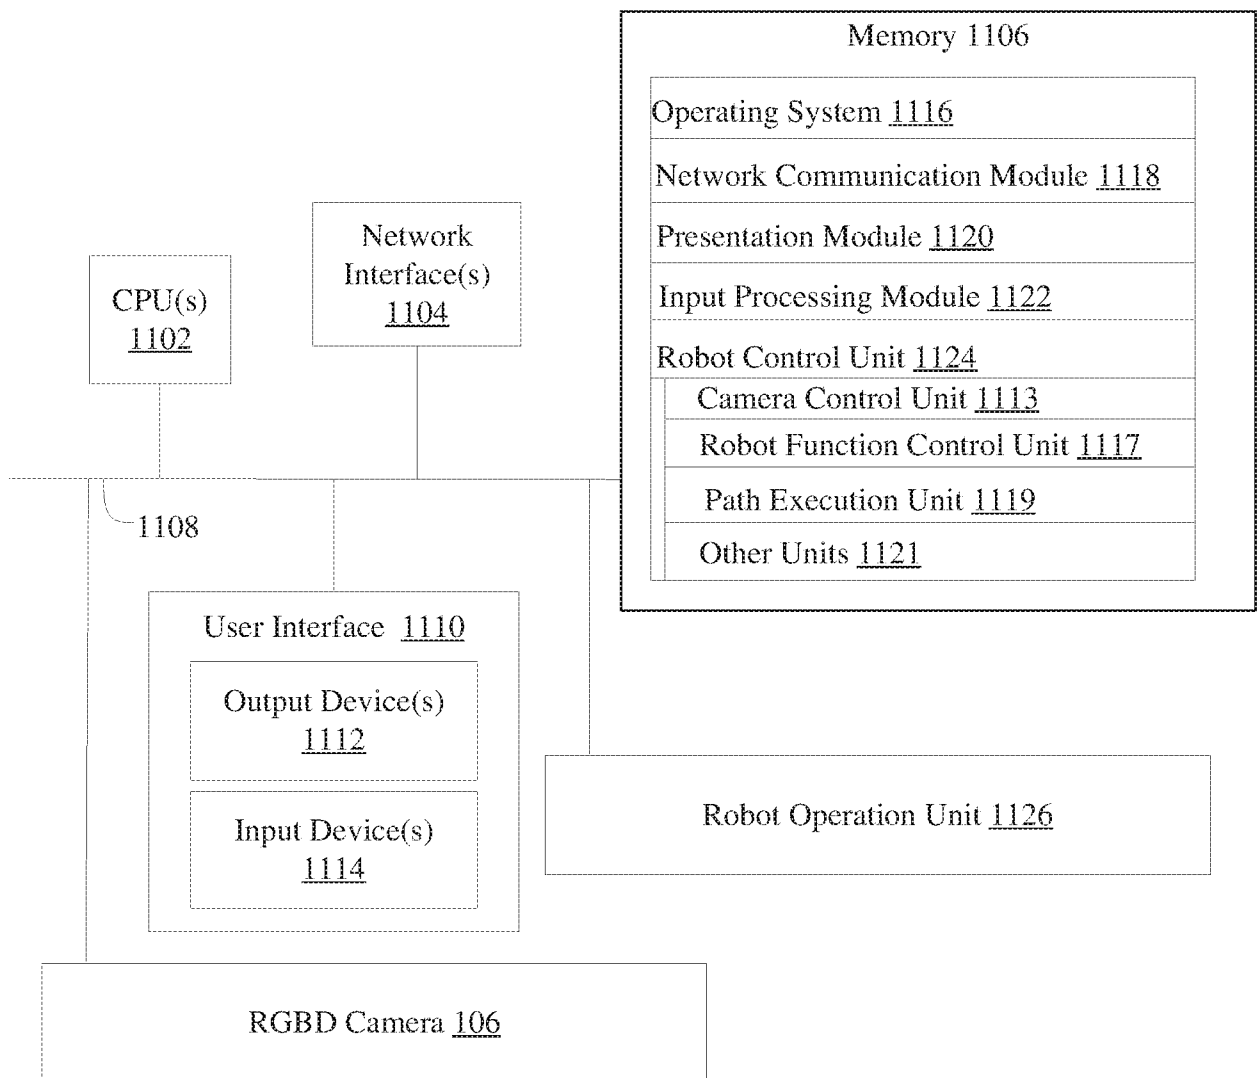


Figure 11

Server 110

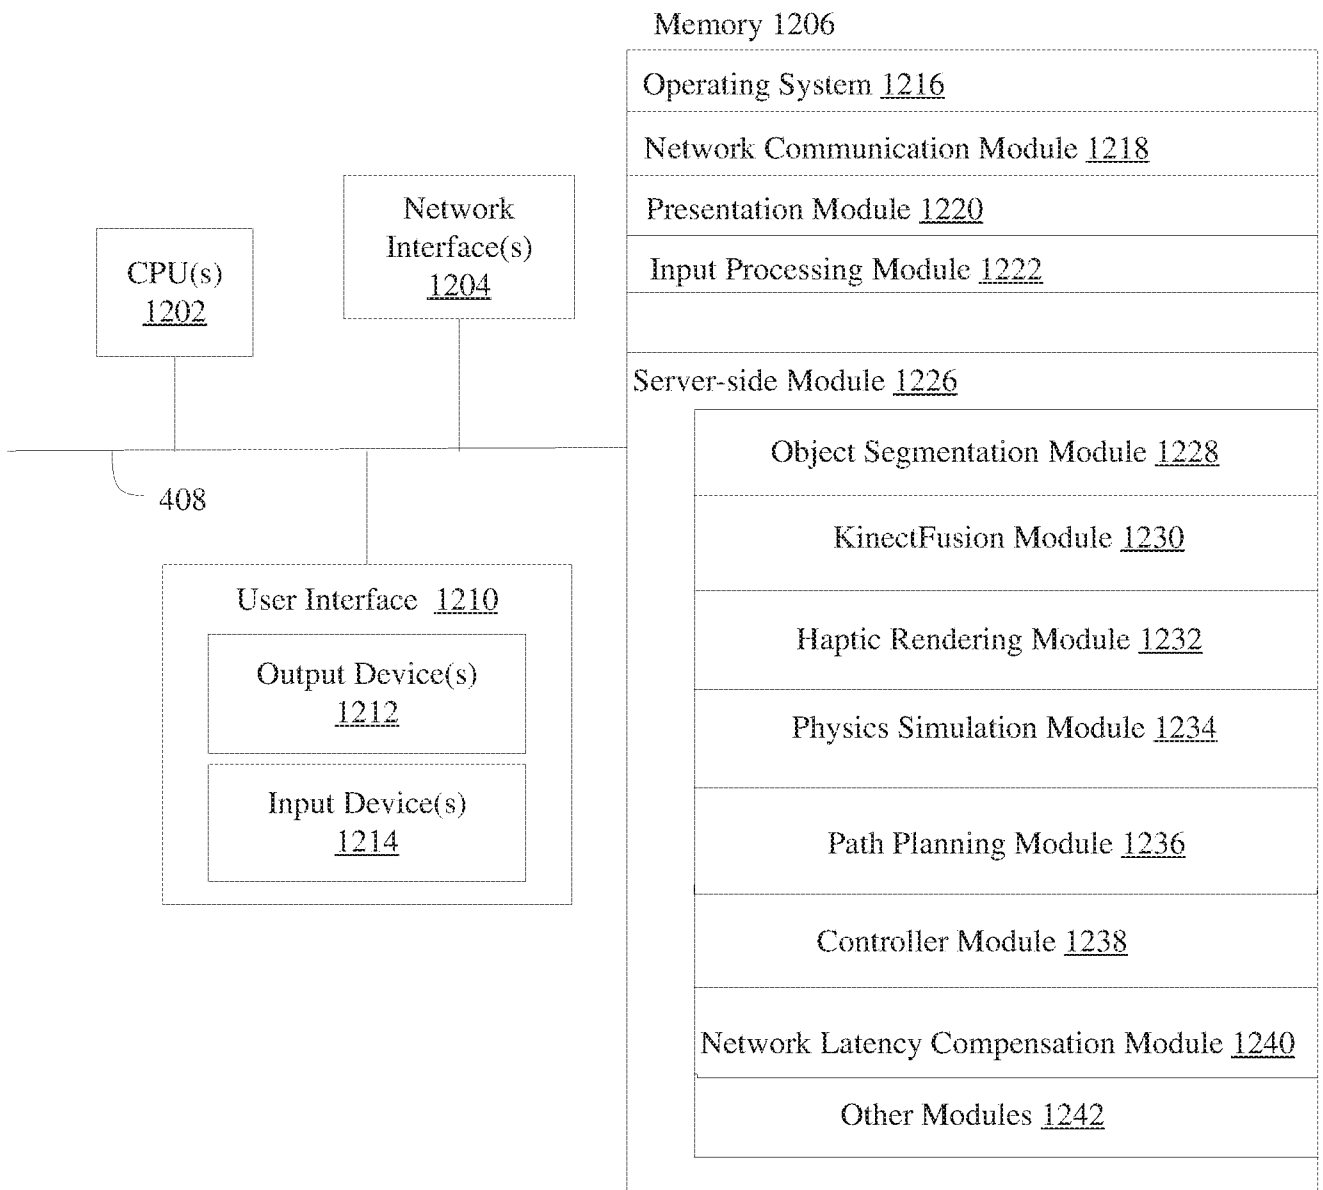


Figure 12

User-side device 1300

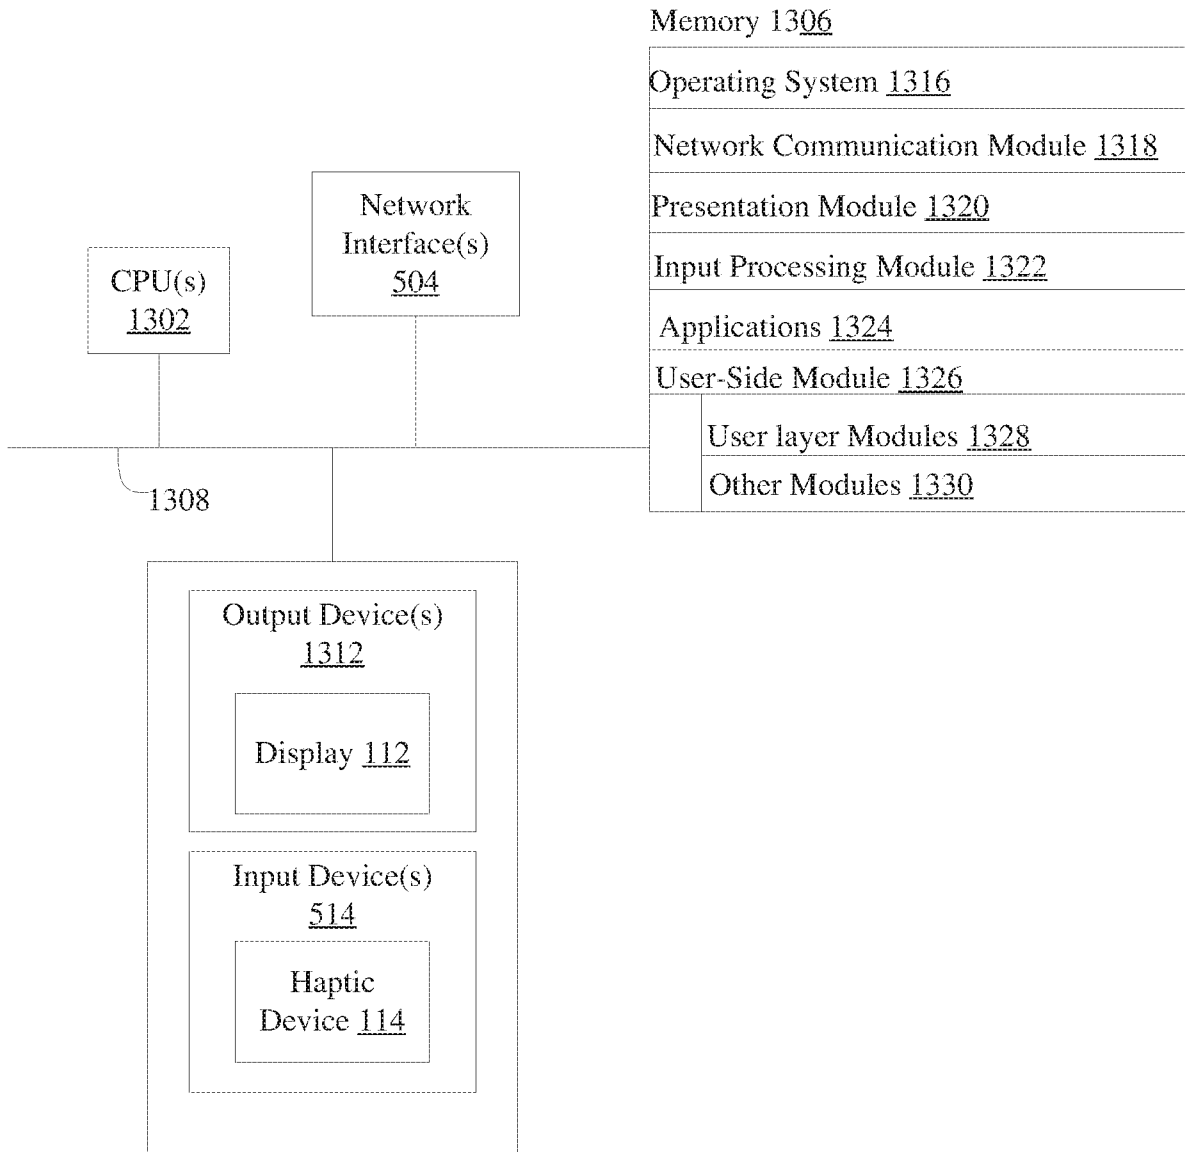


Figure 13

INTERNATIONAL SEARCH REPORT

International application No.

PCT/CN2019/092147

A. CLASSIFICATION OF SUBJECT MATTER		
G06F 3/048(2013.01)i		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED		
Minimum documentation searched (classification system followed by classification symbols) G06F; G05D; H04L		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) CNKI,CNABS,VEN,DWPI: robotic, physical, interface, virtual, haptic, simulated, feedback, movement, path		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	CN 102699914 A (UNIV ZHENGZHOU) 03 October 2012 (2012-10-03) description, pages 3-6 and figures 1-5	1-20
X	CN 103302668 A (UNIV SOUTHEAST) 18 September 2013 (2013-09-18) description, pages 2-5 and figures 1-5	1-20
X	US 2016217617 A1 (HEWLETT-PACKARD DEV COMPANY L P) 28 July 2016 (2016-07-28) description, pages 1-4 and figures 1-5	1-20
A	CN 101459857 A (SHENZHEN HUAWEI TELECOM TECHNOLOGIES CO) 17 June 2009 (2009-06-17) the whole document	1-20
A	JP 2013061937 A (PALO ALTO RES CT INC) 04 April 2013 (2013-04-04) the whole document	1-20
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family		
Date of the actual completion of the international search 03 September 2019		Date of mailing of the international search report 25 September 2019
Name and mailing address of the ISA/CN National Intellectual Property Administration, PRC 6, Xitucheng Rd., Jimen Bridge, Haidian District, Beijing 100088 China		Authorized officer TIAN,Dan
Facsimile No. (86-10)62019451		Telephone No. 62085521

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.

PCT/CN2019/092147

Patent document cited in search report			Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
CN	102699914	A	03 October 2012	CN	102699914	B	21 January 2015
CN	103302668	A	18 September 2013	CN	103302668	B	16 March 2016
US	2016217617	A1	28 July 2016	WO	2015030786	A1	05 March 2015
CN	101459857	A	17 June 2009	EP	2136602	A1	23 December 2009
				CN	101459857	B	05 September 2012
				WO	2009074110	A1	18 June 2009
				BR	PI0819332	A2	12 May 2015
				EP	2136602	A4	14 April 2010
				US	2010053307	A1	04 March 2010
JP	2013061937	A	04 April 2013	KR	20130028878	A	20 March 2013
				EP	2568355	A2	13 March 2013
				US	2013063560	A1	14 March 2013
				EP	2568355	A3	15 May 2013