



(51) International Patent Classification:
G06T 19/20 (2011.01)

(21) International Application Number:
PCT/US2024/031125

(22) International Filing Date:
24 May 2024 (24.05.2024)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
18/325,868 30 May 2023 (30.05.2023) US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US 18/325,868 (CON)
Filed on 30 May 2023 (30.05.2023)

(71) Applicant: SNAP INC. [US/US]; 3000 31st Street, Santa Monica, California 90405 (US).

(72) Inventors: ASSOULINE, Avihay; 3000 31st Street, Santa Monica, California 90405 (US). BERGER, Itamar; 3000 31st Street, Santa Monica, California 90405 (US). DU-DOVITCH, Gal; 3000 31st Street, Santa Monica, California 90405 (US). HAREL, Peleg; 3000 31st Street, Santa Monica, California 90405 (US).

(74) Agent: PERDOK, Monique M. et al.; Schwegman Lundberg & Woessner, P.A., P.O. Box 2938, Minneapolis, Minnesota 55402 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CV, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, MG, MK, MN, MU, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH,

(54) Title: MIRRORING 3D ASSETS FOR VIRTUAL EXPERIENCES




FIG. 6

(57) Abstract: Methods and systems are disclosed for generating mirrored 3D assets for an extended reality (XR) experience. The system receives a three-dimensional (3D) object comprising a target and analyzes the 3D object using one or more machine learning models to generate data associated with a mirrored version of the target of the 3D object. The system applies the mirrored version of the target to a mirrored version of the 3D object using the generated data and generates a new 3D object comprising the mirrored version of the 3D object and the mirrored version of the target.



TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS,
ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, CV, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SC, 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, ME, 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))*

MIRRORING 3D ASSETS FOR VIRTUAL EXPERIENCES**CLAIM OF PRIORITY**

[0001] This application claims the benefit of priority to U.S. Patent Application Serial No. 18/325,868, filed May 30, 2023, which is incorporated herein by reference in its entirety.

5

TECHNICAL FIELD

[0002] The present disclosure relates generally to generating 3D objects representing pairs of fashion items.

BACKGROUND

[0003] Augmented reality (AR) is a modification of a virtual environment. For example, in virtual reality (VR), a user is completely immersed in a virtual world, whereas in AR, the user is immersed in a world where virtual objects are combined or superimposed on the real world. An AR system aims to generate and present virtual objects that interact realistically with a real-world environment and with each other. Examples of AR applications can include single or multiple player video games, instant messaging systems, and the like. In general, these systems are referred to as extended reality (XR) systems.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0004] In the drawings, which are not necessarily drawn to scale, like numerals may describe similar components in different views. To easily identify the discussion of any particular element or act, the most significant digit or digits in a reference number refer to the figure number in which that element is first introduced. Some non-limiting examples are illustrated in the figures of the accompanying drawings in which:

[0005] FIG. 1 is a diagrammatic representation of a networked environment in which the present disclosure may be deployed, according to some examples.

[0006] FIG. 2 is a diagrammatic representation of a messaging system that has both client-side and server-side functionality, according to some examples.

[0007] FIG. 3 is a diagrammatic representation of a data structure as maintained in a database, according to some examples.

[0008] FIG. 4 is a diagrammatic representation of a message, according to some examples.

[0009] FIG. 5 is a diagrammatic representation of an object generation system, according to some examples.

[0010] FIG. 6 is a diagrammatic representation of example inputs and outputs of the object generation system, according to some examples.

[0011] FIG. 7 is a flowchart illustrating example operations and methods of the object generation system, according to some examples.

[0012] FIG. 8 is a diagrammatic representation of a machine in the form of a computer system within which a set of instructions may be executed to cause the machine to perform any one or more of the methodologies discussed herein, according to some examples.

[0013] FIG. 9 is a block diagram showing a software architecture within which examples may be implemented.

[0014] FIG. 10 illustrates a system in which a head-wearable apparatus may be implemented, according to some examples.

DETAILED DESCRIPTION

[0015] The description that follows includes systems, methods, techniques, instruction sequences, and computing machine program products that embody illustrative examples of the disclosure. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide an understanding of various examples. It will be evident, however, to those skilled in the art, that examples may be practiced without these specific details. In general, well-known instruction instances, protocols, structures, and techniques are not necessarily shown in detail.

[0016] Typically, various communication platforms allow users to share content and create images for transmission to other users. These images can be used to promote products or services and/or to simply represent different real-world objects in simulated or real environments. However, these systems require a user to use expensive equipment and technology to create high-quality, appealing images. Also, users may spend a great deal of effort meticulously placing objects in different environments and manually adjusting lighting and other image attributes to enhance the presentation of the objects in the images. All of these factors can add up to make the creation of high-quality images (e.g., for use in advertising) a significant expense and detract from the overall use and enjoyment of the system. In addition, because users may not have the resources needed to create high-quality images, opportunities to share and present objects in ideal settings are missed. Also, presenting lower quality images of such objects can cause other users to overlook the value of the objects, which wastes the resources used to create and display the objects.

[0017] The disclosed techniques seek to improve the efficiency of using an electronic device by intelligently and automatically generating 3D objects (e.g., including images of 3D objects) that depict pairs of objects (e.g., fashion items) in a simple and intuitive manner. The disclosed techniques create photorealistic images or videos or virtual 3D

objects that depict a pairs of fashion items in real-world scenes or simulated scenes very quickly and efficiently and with minimal user interaction or involvement. This can reduce the overall time and expense incurred to develop high-quality images that feature objects or products, such as shoes, shirts, or other fashion items.

5 **[0018]** For example, the disclosed techniques receive a 3D object (e.g., a real or virtual first fashion item of a pair of fashion items) that includes a target (e.g., an emblem, logo, and/or text). The disclosed techniques analyze the 3D object using one or more machine learning models to generate data associated with a mirrored version of the target of the 3D object. The mirrored version of the target can correspond to a second fashion item of the
10 pair of fashion items. The disclosed techniques apply the mirrored version of the target to a mirrored version of the 3D object and generate a new 3D object that includes the mirrored version of the 3D object that includes the mirrored version of the target. This new 3D object can be used in communication with other users and/or placed in an XR experience, such as a virtual try-on experience. Namely, the disclosed techniques can separately generate a
15 mirrored version of a logo or graphical element on a 3D object from generating the mirrored version of the 3D object and then overlay that mirrored version of the logo or graphical element on the mirrored version of the 3D object.

[0019] In this way, the disclosed techniques improve the overall experience of the user in using the electronic device and reduce the overall amount of resources needed to accomplish
20 a task of producing high-quality images and producing realistic XR experiences. As used herein, “article of clothing,” “fashion item,” and “garment” are used interchangeably and should be understood to have the same meaning. Article of clothing, garment, or fashion item can include a shirt, skirt, dress, shoes, purse, furniture item, household item, eyewear, eyeglasses, AR logo, AR emblem, pants, shorts, jacket, t-shirt, blouse, glasses, jewelry,
25 earrings, bunny ears, a hat, earmuffs, makeup, or any other suitable item or object.

NETWORKED COMPUTING ENVIRONMENT

[0020] FIG. 1 is a block diagram showing an example interaction system 100 for facilitating interactions (e.g., exchanging text messages, conducting text audio and video
30 calls, or playing games) over a network. The interaction system 100 includes multiple user systems 102, each of which hosts multiple applications, including an interaction client 104 and other applications 106. Each interaction client 104 is communicatively coupled, via one or more communication networks including a network 108 (e.g., the Internet), to other instances of the interaction client 104 (e.g., hosted on respective other user systems 102), an

interaction server system 110 and third-party servers 112). An interaction client 104 can also communicate with locally hosted applications 106 using Applications Program Interfaces (APIs).

[0021] Each user system 102 may include multiple user devices, such as a mobile device 114, head-wearable apparatus 116, and a computer client device 118 that are communicatively connected to exchange data and messages.

[0022] An interaction client 104 interacts with other interaction clients 104 and with the interaction server system 110 via the network 108. The data exchanged between the interaction clients 104 (e.g., interactions 120) and between the interaction clients 104 and the interaction server system 110 includes functions (e.g., commands to invoke functions) and payload data (e.g., text, audio, video, or other multimedia data).

[0023] The interaction server system 110 provides server-side functionality via the network 108 to the interaction clients 104. While certain functions of the interaction system 100 are described herein as being performed by either an interaction client 104 or by the interaction server system 110, the location of certain functionality either within the interaction client 104 or the interaction server system 110 may be a design choice. For example, it may be technically preferable to initially deploy particular technology and functionality within the interaction server system 110 but to later migrate this technology and functionality to the interaction client 104 where a user system 102 has sufficient processing capacity.

[0024] The interaction server system 110 supports various services and operations that are provided to the interaction clients 104. Such operations include transmitting data to, receiving data from, and processing data generated by the interaction clients 104. This data may include message content, client device information, geolocation information, media augmentation and overlays, message content persistence conditions, entity relationship information, and live event information. Data exchanges within the interaction system 100 are invoked and controlled through functions available via user interfaces (UIs) of the interaction clients 104.

[0025] Turning now specifically to the interaction server system 110, an API server 122 is coupled to and provides programmatic interfaces to interaction servers 124, making the functions of the interaction servers 124 accessible to interaction clients 104, other applications 106, and third-party server 112. The interaction servers 124 are communicatively coupled to a database server 126, facilitating access to a database 128 that stores data associated with interactions processed by the interaction servers 124. Similarly, a

web server 130 is coupled to the interaction servers 124 and provides web-based interfaces to the interaction servers 124. To this end, the web server 130 processes incoming network requests over Hypertext Transfer Protocol (HTTP) and several other related protocols.

[0026] The API server 122 receives and transmits interaction data (e.g., commands and message payloads) between the interaction servers 124 and the user systems 102 (and, for example, interaction clients 104 and other application 106) and the third-party server 112. Specifically, the API server 122 provides a set of interfaces (e.g., routines and protocols) that can be called or queried by the interaction client 104 and other applications 106 to invoke functionality of the interaction servers 124. The API server 122 exposes various functions supported by the interaction servers 124, including account registration; login functionality; the sending of interaction data, via the interaction servers 124, from a particular interaction client 104 to another interaction client 104; the communication of media files (e.g., images or video) from an interaction client 104 to the interaction servers 124; the settings of a collection of media data (e.g., a story); the retrieval of a list of friends of a user of a user system 102; the retrieval of messages and content; the addition and deletion of entities (e.g., friends) to an entity relationship graph (e.g., the entity graph 310); the location of friends within an entity relationship graph; and opening an application event (e.g., relating to the interaction client 104).

[0027] The interaction servers 124 host multiple systems and subsystems, described below with reference to FIG. 2.

LINKED APPLICATIONS

[0028] Returning to the interaction client 104, features and functions of an external resource (e.g., a linked application 106 or applet) are made available to a user via an interface of the interaction client 104. In this context, “external” refers to the fact that the application 106 or applet is external to the interaction client 104. The external resource is often provided by a third party but may also be provided by the creator or provider of the interaction client 104. The interaction client 104 receives a user selection of an option to launch or access features of such an external resource. The external resource may be the application 106 installed on the user system 102 (e.g., a “native app”), or a small-scale version of the application (e.g., an “applet”) that is hosted on the user system 102 or remote of the user system 102 (e.g., on third-party servers 112). The small-scale version of the application includes a subset of features and functions of the application (e.g., the full-scale, native version of the application) and is implemented using a markup-language document.

In some examples, the small-scale version of the application (e.g., an “applet”) is a web-based, markup-language version of the application and is embedded in the interaction client 104. In addition to using markup-language documents (e.g., a *.ml file), an applet may incorporate a scripting language (e.g., a *.js file or a .json file) and a style sheet (e.g., a *.ss file).

[0029] In response to receiving a user selection of the option to launch or access features of the external resource, the interaction client 104 determines whether the selected external resource is a web-based external resource or a locally-installed application 106. In some cases, applications 106 that are locally installed on the user system 102 can be launched independently of and separately from the interaction client 104, such as by selecting an icon corresponding to the application 106 on a home screen of the user system 102. Small-scale versions of such applications can be launched or accessed via the interaction client 104 and, in some examples, no or limited portions of the small-scale application can be accessed outside of the interaction client 104. The small-scale application can be launched by the interaction client 104 receiving, from a third-party server 112 for example, a markup-language document associated with the small-scale application and processing such a document.

[0030] In response to determining that the external resource is a locally-installed application 106, the interaction client 104 instructs the user system 102 to launch the external resource by executing locally-stored code corresponding to the external resource. In response to determining that the external resource is a web-based resource, the interaction client 104 communicates with the third-party servers 112 (for example) to obtain a markup-language document corresponding to the selected external resource. The interaction client 104 then processes the obtained markup-language document to present the web-based external resource within a user interface of the interaction client 104.

[0031] The interaction client 104 can notify a user of the user system 102, or other users related to such a user (e.g., “friends”), of activity taking place in one or more external resources. For example, the interaction client 104 can provide participants in a conversation (e.g., a chat session) in the interaction client 104 with notifications relating to the current or recent use of an external resource by one or more members of a group of users. One or more users can be invited to join in an active external resource or to launch a recently used but currently inactive (in the group of friends) external resource. The external resource can provide participants in a conversation, each using respective interaction clients 104, with the ability to share an item, status, state, or location in an external resource in a chat session with one or more members of a group of users. The shared item may be an interactive chat

card with which members of the chat can interact, for example, to launch the corresponding external resource, view specific information within the external resource, or take the member of the chat to a specific location or state within the external resource. Within a given external resource, response messages can be sent to users on the interaction client

104. The external resource can selectively include different media items in the responses, based on a current context of the external resource.

[0032] The interaction client 104 can present a list of the available external resources (e.g., applications 106 or applets) to a user to launch or access a given external resource. This list can be presented in a context-sensitive menu. For example, the icons representing different ones of the application 106 (or applets) can vary based on how the menu is launched by the user (e.g., from a conversation interface or from a non-conversation interface).

SYSTEM ARCHITECTURE

[0033] FIG. 2 is a block diagram illustrating further details regarding the interaction system 100, according to some examples. Specifically, the interaction system 100 is shown to comprise the interaction client 104 and the interaction servers 124. The interaction system 100 embodies multiple subsystems, which are supported on the client side by the interaction client 104 and on the server side by the interaction servers 124. Example subsystems are discussed below and can include an object generation system 500 that generates a mirror image of an object (e.g., a fashion item) depicted in an image including a mirror image of one or more targets (e.g., emblems, graphics, logos, text, and so forth) that are part of the object. The object generation system 500 can then modify the image to include the mirror image of the object including the mirror image of the target with the image depicting the object and the target. In some cases, the object generation system 500 can generate a new image that includes the mirror image of the object and the mirror image of the target to create a 3D virtual asset. Namely, the object generation system 500 can receive the object as a first virtual or real-world asset and can generate, using the disclosed techniques, a second virtual asset that is a mirror image of the first virtual or real-world asset including one or more targets depicted on the first virtual or real-world asset. An illustrative implementation of the object generation system 500 is shown and described in connection with FIG. 5 below.

[0034] In some examples, these subsystems are implemented as microservices. A microservice subsystem (e.g., a microservice application) may have components that enable

it to operate independently and communicate with other services. Example components of a microservice subsystem may include:

- Function logic: The function logic implements the functionality of the microservice subsystem, representing a specific capability or function that the microservice provides.
- API interface: Microservices may communicate with each other through well-defined APIs or interfaces, using lightweight protocols such as REST or messaging. The API interface defines the inputs and outputs of the microservice subsystem and how it interacts with other microservice subsystems of the interaction system 100.
- Data storage: A microservice subsystem may be responsible for its own data storage, which may be in the form of a database, cache, or other storage mechanism (e.g., using the database server 126 and database 128). This enables a microservice subsystem to operate independently of other microservices of the interaction system 100.
- Service discovery: Microservice subsystems may find and communicate with other microservice subsystems of the interaction system 100. Service discovery mechanisms enable microservice subsystems to locate and communicate with other microservice subsystems in a scalable and efficient way.
- Monitoring and logging: Microservice subsystems may need to be monitored and logged in order to ensure availability and performance. Monitoring and logging mechanisms enable the tracking of health and performance of a microservice subsystem.

[0035] In some examples, the interaction system 100 may employ a monolithic architecture, a service-oriented architecture (SOA), a function-as-a-service (FaaS) architecture, or a modular architecture. Example subsystems are discussed below.

[0036] An image processing system 202 provides various functions that enable a user to capture and augment (e.g., annotate or otherwise modify or edit) media content associated with a message.

[0037] A camera system 204 includes control software (e.g., in a camera application) that interacts with and controls camera hardware (e.g., directly or via operating system controls) of the user system 102 to modify and augment real-time images captured and displayed via the interaction client 104.

[0038] An augmentation system 206 provides functions related to the generation and publishing of augmentations (e.g., media overlays) for images captured in real-time by cameras of the user system 102 or retrieved from memory of the user system 102. For example, the augmentation system 206 operatively selects, presents, and displays media overlays (e.g., an image filter or an image lens) to the interaction client 104 for the

augmentation of real-time images received via the camera system 204 or stored images retrieved from memory 1002 (shown in FIG. 10) of a user system 102. These augmentations are selected by the augmentation system 206 and presented to a user of an interaction client 104, based on a number of inputs and data, such as, for example:

- 5 • Geolocation of the user system 102; and
- Entity relationship information of the user of the user system 102.

[0039] An augmentation may include audio and visual content and visual effects.

Examples of audio and visual content include pictures, texts, logos, animations, and sound effects. An example of a visual effect includes color overlaying. The audio and visual
10 content or the visual effects can be applied to a media content item (e.g., a photo or video) at user system 102 for communication in a message, or applied to video content, such as a video content stream or feed transmitted from an interaction client 104. As such, the image processing system 202 may interact with, and support, the various subsystems of the communication system 208, such as the messaging system 210 and the video
15 communication system 212.

[0040] A media overlay may include text or image data that can be overlaid on top of a photograph taken by the user system 102 or a video stream produced by the user system 102. In some examples, the media overlay may be a location overlay (e.g., Venice beach), a name of a live event, or a name of a merchant overlay (e.g., Beach Coffee House). In further
20 examples, the image processing system 202 uses the geolocation of the user system 102 to identify a media overlay that includes the name of a merchant at the geolocation of the user system 102. The media overlay may include other indicia associated with the merchant. The media overlays may be stored in the databases 128 and accessed through the database server 126.

[0041] The image processing system 202 provides a user-based publication platform that enables users to select a geolocation on a map and upload content associated with the selected geolocation. The user may also specify circumstances under which a particular media overlay should be offered to other users. The image processing system 202 generates a media overlay that includes the uploaded content and associates the uploaded content with
30 the selected geolocation.

[0042] An augmentation creation system 214 supports AR developer platforms and includes an application for content creators (e.g., artists and developers) to create and publish augmentations (e.g., AR experiences) of the interaction client 104. The

augmentation creation system 214 provides a library of built-in features and tools to content creators including, for example, custom shaders, tracking technology, and templates.

[0043] In some examples, the augmentation creation system 214 provides a merchant-based publication platform that enables merchants to select a particular augmentation

associated with a geolocation via a bidding process. For example, the augmentation creation system 214 associates a media overlay of the highest bidding merchant with a corresponding geolocation for a predefined amount of time.

[0044] A communication system 208 is responsible for enabling and processing multiple forms of communication and interaction within the interaction system 100 and includes a

messaging system 210, an audio communication system 216, and a video communication system 212. The messaging system 210 is responsible for enforcing the temporary or time-limited access to content by the interaction clients 104. The messaging system

210 incorporates multiple timers (e.g., within an ephemeral timer system) that, based on duration and display parameters associated with a message or collection of messages (e.g., a story), selectively enable access (e.g., for presentation and display) to messages and

associated content via the interaction client 104. The audio communication system 216 enables and supports audio communications (e.g., real-time audio chat) between multiple interaction clients 104. Similarly, the video communication system 212 enables and supports video communications (e.g., real-time video chat) between multiple interaction clients 104.

[0045] A user management system 218 is operationally responsible for the management of user data and profiles, and maintains entity information (e.g., stored in entity tables 308, entity graphs 310, and profile data 302 of FIG. 3) regarding users and relationships between users of the interaction system 100.

[0046] A collection management system 220 is operationally responsible for managing sets or collections of media (e.g., collections of text, image, video, and audio data). A collection of content (e.g., messages, including images, video, text, and audio) may be organized into an “event gallery” or an “event story.” Such a collection may be made available for a specified time period, such as the duration of an event to which the content relates. For example, content relating to a music concert may be made available as a “story” for the duration of that music concert. The collection management system 220 may also be responsible for publishing an icon that provides notification of a particular collection to the user interface of the interaction client 104. The collection management system 220 includes a curation function that allows a collection manager to manage and curate a particular

collection of content. For example, the curation interface enables an event organizer to curate a collection of content relating to a specific event (e.g., delete inappropriate content or redundant messages). Additionally, the collection management system 220 employs machine vision (or image recognition technology) and content rules to curate a content collection automatically. In certain examples, compensation may be paid to a user to include user-generated content into a collection. In such cases, the collection management system 220 operates to automatically make payments to such users to use their content.

[0047] A map system 222 provides various geographic location (e.g., geolocation) functions and supports the presentation of map-based media content and messages by the interaction client 104. For example, the map system 222 enables the display of user icons or avatars (e.g., stored in profile data 302) on a map to indicate a current or past location of "friends" of a user, as well as media content (e.g., collections of messages including photographs and videos) generated by such friends, within the context of a map. For example, a message posted by a user to the interaction system 100 from a specific geographic location may be displayed within the context of a map at that particular location to "friends" of a specific user on a map interface of the interaction client 104. A user can furthermore share his or her location and status information (e.g., using an appropriate status avatar) with other users of the interaction system 100 via the interaction client 104, with this location and status information being similarly displayed within the context of a map interface of the interaction client 104 to selected users.

[0048] A game system 224 provides various gaming functions within the context of the interaction client 104. The interaction client 104 provides a game interface providing a list of available games that can be launched by a user within the context of the interaction client 104 and played with other users of the interaction system 100. The interaction system 100 further enables a particular user to invite other users to participate in the play of a specific game by issuing invitations to such other users from the interaction client 104. The interaction client 104 also supports audio, video, and text messaging (e.g., chats) within the context of gameplay, provides a leaderboard for the games, and also supports the provision of in-game rewards (e.g., coins and items).

[0049] An external resource system 226 provides an interface for the interaction client 104 to communicate with remote servers (e.g., third-party servers 112) to launch or access external resources, e.g., applications or applets. Each third-party server 112 hosts, for example, a markup language (e.g., HTML5) based application or a small-scale version of an application (e.g., game, utility, payment, or ride-sharing application). The interaction client 104 may launch a web-based resource (e.g., application) by accessing the HTML5 file from

the third-party servers 112 associated with the web-based resource. Applications hosted by third-party servers 112 are programmed in JavaScript leveraging a Software Development Kit (SDK) provided by the interaction servers 124. The SDK includes APIs with functions that can be called or invoked by the web-based application. The interaction servers 124 host a JavaScript library that provides a given external resource access to specific user data of the interaction client 104. HTML5 is an example of technology for programming games, but applications and resources programmed based on other technologies can be used.

[0050] To integrate the functions of the SDK into the web-based resource, the SDK is downloaded by the third-party server 112 from the interaction servers 124 or is otherwise received by the third-party server 112. Once downloaded or received, the SDK is included as part of the application code of a web-based external resource. The code of the web-based resource can then call or invoke certain functions of the SDK to integrate features of the interaction client 104 into the web-based resource.

[0051] The SDK stored on the interaction server system 110 effectively provides the bridge between an external resource (e.g., applications 106 or applets) and the interaction client 104. This gives the user a seamless experience of communicating with other users on the interaction client 104 while also preserving the look and feel of the interaction client 104. To bridge communications between an external resource and an interaction client 104, the SDK facilitates communication between third-party servers 112 and the interaction client 104. A bridge script running on a user system 102 establishes two one-way communication channels between an external resource and the interaction client 104. Messages are sent between the external resource and the interaction client 104 via these communication channels asynchronously. Each SDK function invocation is sent as a message and callback. Each SDK function is implemented by constructing a unique callback identifier and sending a message with that callback identifier.

[0052] By using the SDK, not all information from the interaction client 104 is shared with third-party servers 112. The SDK limits which information is shared based on the needs of the external resource. Each third-party server 112 provides an HTML5 file corresponding to the web-based external resource to interaction servers 124. The interaction servers 124 can add a visual representation (such as a box art or other graphic) of the web-based external resource in the interaction client 104. Once the user selects the visual representation or instructs the interaction client 104 through a graphical user interface (GUI) of the interaction client 104 to access features of the web-based external resource, the interaction client 104 obtains the HTML5 file and instantiates the resources to access the features of the web-based external resource.

[0053] The interaction client 104 presents a GUI (e.g., a landing page or title screen) for an external resource. During, before, or after presenting the landing page or title screen, the interaction client 104 determines whether the launched external resource has been previously authorized to access user data of the interaction client 104. In response to determining that the launched external resource has been previously authorized to access user data of the interaction client 104, the interaction client 104 presents another GUI of the external resource that includes functions and features of the external resource. In response to determining that the launched external resource has not been previously authorized to access user data of the interaction client 104, after a threshold period of time (e.g., 3 seconds) of displaying the landing page or title screen of the external resource, the interaction client 104 slides up (e.g., animates a menu as surfacing from a bottom of the screen to a middle or other portion of the screen) a menu for authorizing the external resource to access the user data. The menu identifies the type of user data that the external resource will be authorized to use. In response to receiving a user selection of an accept option, the interaction client 104 adds the external resource to a list of authorized external resources and allows the external resource to access user data from the interaction client 104. The external resource is authorized by the interaction client 104 to access the user data under an OAuth 2 framework.

[0054] The interaction client 104 controls the type of user data that is shared with external resources based on the type of external resource being authorized. For example, external resources that include full-scale applications (e.g., an application 106) are provided with access to a first type of user data (e.g., two-dimensional (2D) avatars of users with or without different avatar characteristics). As another example, external resources that include small-scale versions of applications (e.g., web-based versions of applications) are provided with access to a second type of user data (e.g., payment information, 2D avatars of users, three-dimensional (3D) avatars of users, and avatars with various avatar characteristics). Avatar characteristics include different ways to customize a look and feel of an avatar, such as different poses, facial features, clothing, and so forth.

[0055] An advertisement system 228 operationally enables the purchasing of advertisements by third parties for presentation to end-users via the interaction clients 104 and also handles the delivery and presentation of these advertisements.

[0056] An artificial intelligence and machine learning system 230 provides a variety of services to different subsystems within the interaction system 100. For example, the artificial intelligence and machine learning system 230 operates with the image processing system 202 and the camera system 204 to analyze images and extract information such as

objects, text, or faces. This information can then be used by the image processing system 202 to enhance, filter, or manipulate images. The artificial intelligence and machine learning system 230 may be used by the augmentation system 206 to generate augmented content, XR experiences, and AR experiences, such as adding virtual objects or animations to real-world images and videos. The communication system 208 and messaging system 210 may use the artificial intelligence and machine learning system 230 to analyze communication patterns and provide insights into how users interact with each other and provide intelligent message classification and tagging, such as categorizing messages based on sentiment or topic.

[0057] The artificial intelligence and machine learning system 230 may also provide chatbot functionality to message interactions 120 between user systems 102 and between a user system 102 and the interaction server system 110. The artificial intelligence and machine learning system 230 may also work with the audio communication system 216 to provide speech recognition and natural language processing capabilities, allowing users to interact with the interaction system 100 using voice commands.

[0058] In some cases, the artificial intelligence and machine learning system 230 can implement one or more machine learning models that generate artificial 3D objects or images of a mirror image of a real-world or virtual 3D object (e.g., a fashion item), such as one depicted in a received image. The generated artificial 3D and the real-world or virtual 3D object can together form a pair of 3D objects corresponding to a fashion item pair, such as a pair of left and right shoes. In such cases, the artificial intelligence and machine learning system 230 can be implemented as one or more components shown in **FIG. 5**.

DATA ARCHITECTURE

[0059] **FIG. 3** is a schematic diagram illustrating data structures 300, which may be stored in a database 304 of the interaction server system 110, according to certain examples. While the content of the database 304 is shown to comprise multiple tables, it will be appreciated that the data could be stored in other types of data structures (e.g., as an object-oriented database).

[0060] The database 304 includes message data stored within a message table 306. This message data includes, for any particular message, at least message sender data, message recipient (or receiver) data, and a payload. Further details regarding information that may be included in a message, and included within the message data stored in the message table 306, are described below with reference to **FIG. 4**.

[0061] An entity table 308 stores entity data, and is linked (e.g., referentially) to an entity graph 310 and profile data 302. Entities for which records are maintained within the entity table 308 may include individuals, corporate entities, organizations, objects, places, events, and so forth. Regardless of entity type, any entity regarding which the interaction server system 110 stores data may be a recognized entity. Each entity is provided with a unique identifier, as well as an entity type identifier (not shown).

[0062] The entity graph 310 stores information regarding relationships and associations between entities. Such relationships may be social, professional (e.g., work at a common corporation or organization), interest-based, or activity-based, merely for example. Certain relationships between entities may be unidirectional, such as a subscription by an individual user to digital content of a commercial or publishing user (e.g., a newspaper or other digital media outlet, or a brand). Other relationships may be bidirectional, such as a “friend” relationship between individual users of the interaction system 100.

[0063] Certain permissions and relationships may be attached to each relationship, and also to each direction of a relationship. For example, a bidirectional relationship (e.g., a friend relationship between individual users) may include authorization for the publication of digital content items between the individual users, but may impose certain restrictions or filters on the publication of such digital content items (e.g., based on content characteristics, location data, or time of day data). Similarly, a subscription relationship between an individual user and a commercial user may impose different degrees of restrictions on the publication of digital content from the commercial user to the individual user, and may significantly restrict or block the publication of digital content from the individual user to the commercial user. A particular user, as an example of an entity, may record certain restrictions (e.g., by way of privacy settings) in a record for that entity within the entity table 308. Such privacy settings may be applied to all types of relationships within the context of the interaction system 100 or may selectively be applied to certain types of relationships.

[0064] The profile data 302 stores multiple types of profile data about a particular entity. The profile data 302 may be selectively used and presented to other users of the interaction system 100 based on privacy settings specified by a particular entity. Where the entity is an individual, the profile data 302 includes, for example, a user name, telephone number, address, settings (e.g., notification and privacy settings), as well as a user-selected avatar representation (or collection of such avatar representations). A particular user may then selectively include one or more of these avatar representations within the content of messages communicated via the interaction system 100 and on map interfaces displayed by

interaction clients 104 to other users. The collection of avatar representations may include “status avatars,” which present a graphical representation of a status or activity that the user may select to communicate at a particular time.

[0065] Where the entity is a group, the profile data 302 for the group may similarly include one or more avatar representations associated with the group, in addition to the group name, members, and various settings (e.g., notifications) for the relevant group.

[0066] The database 304 also stores augmentation data, such as overlays or filters, in an augmentation table 312. The augmentation data is associated with and applied to videos (for which data is stored in a video table 314) and images (for which data is stored in an image table 316).

[0067] Filters, in some examples, are overlays that are displayed as overlaid on an image or video during presentation to a recipient user. Filters may be of various types, including user-selected filters from a set of filters presented to a sending user by the interaction client 104 when the sending user is composing a message. Other types of filters include geolocation filters (also known as geo-filters), which may be presented to a sending user based on geographic location. For example, geolocation filters specific to a neighborhood or special location may be presented within a user interface by the interaction client 104, based on geolocation information determined by a Global Positioning System (GPS) unit of the user system 102.

[0068] Another type of filter is a data filter, which may be selectively presented to a sending user by the interaction client 104 based on other inputs or information gathered by the user system 102 during the message creation process. Examples of data filters include current temperature at a specific location, a current speed at which a sending user is traveling, battery life for a user system 102, or the current time.

[0069] Other augmentation data that may be stored within the image table 316 includes AR content items (e.g., corresponding to applying “lenses” or AR experiences). An AR content item may be a real-time special effect and sound that may be added to an image or a video.

[0070] A collections table 318 stores data regarding collections of messages and associated image, video, or audio data, which are compiled into a collection (e.g., a story or a gallery).

The creation of a particular collection may be initiated by a particular user (e.g., each user for which a record is maintained in the entity table 308). A user may create a “personal story” in the form of a collection of content that has been created and sent/broadcast by that user. To this end, the user interface of the interaction client 104 may include an icon that is user-selectable to enable a sending user to add specific content to his or her personal story.

[0071] A collection may also constitute a “live story,” which is a collection of content from multiple users that is created manually, automatically, or using a combination of manual and automatic techniques. For example, a “live story” may constitute a curated stream of user-submitted content from various locations and events. Users whose client devices have location services enabled and are at a common location event at a particular time may, for example, be presented with an option, via a user interface of the interaction client 104, to contribute content to a particular live story. The live story may be identified to the user by the interaction client 104, based on his or her location. The end result is a “live story” told from a community perspective.

[0072] A further type of content collection is known as a “location story,” which enables a user whose user system 102 is located within a specific geographic location (e.g., on a college or university campus) to contribute to a particular collection. In some examples, a contribution to a location story may employ a second degree of authentication to verify that the end-user belongs to a specific organization or other entity (e.g., is a student on the university campus).

[0073] As mentioned above, the video table 314 stores video data that, in some examples, is associated with messages for which records are maintained within the message table 306. Similarly, the image table 316 stores image data associated with messages for which message data is stored in the entity table 308. The entity table 308 may associate various augmentations from the augmentation table 312 with various images and videos stored in the image table 316 and the video table 314.

[0074] The databases 304 also include trained machine learning techniques 307 that store parameters of one or more machine learning models that have been trained during training of the object generation system 500. For example, trained machine learning techniques 307 store the trained parameters of one or more artificial neural network machine learning models or techniques.

DATA COMMUNICATIONS ARCHITECTURE

[0075] **FIG. 4** is a schematic diagram illustrating a structure of a message 400, according to some examples, generated by an interaction client 104 for communication to a further interaction client 104 via the interaction servers 124. The content of a particular message 400 is used to populate the message table 306 stored within the database 304, accessible by the interaction servers 124. Similarly, the content of a message 400 is stored in memory as

“in-transit” or “in-flight” data of the user system 102 or the interaction servers 124. A message 400 is shown to include the following example components:

- Message identifier 402: a unique identifier that identifies the message 400.
- Message text payload 404: text, to be generated by a user via a user interface of the user system 102, and that is included in the message 400.
- Message image payload 406: image data, captured by a camera component of a user system 102 or retrieved from a memory component of a user system 102, and that is included in the message 400. Image data for a sent or received message 400 may be stored in the image table 316.
- Message video payload 408: video data, captured by a camera component or retrieved from a memory component of the user system 102, and that is included in the message 400. Video data for a sent or received message 400 may be stored in the image table 316.
- Message audio payload 410: audio data, captured by a microphone or retrieved from a memory component of the user system 102, and that is included in the message 400.
- Message augmentation data 412: augmentation data (e.g., filters, stickers, or other annotations or enhancements) that represents augmentations to be applied to message image payload 406, message video payload 408, or message audio payload 410 of the message 400. Augmentation data for a sent or received message 400 may be stored in the augmentation table 312.
- Message duration parameter 414: parameter value indicating, in seconds, the amount of time for which content of the message (e.g., the message image payload 406, message video payload 408, message audio payload 410) is to be presented or made accessible to a user via the interaction client 104.
- Message geolocation parameter 416: geolocation data (e.g., latitudinal and longitudinal coordinates) associated with the content payload of the message. Multiple message geolocation parameter 416 values may be included in the payload, each of these parameter values being associated with respect to content items included in the content (e.g., a specific image within the message image payload 406, or a specific video in the message video payload 408).
- Message story identifier 418: identifier values identifying one or more content collections (e.g., “stories” identified in the collections table 318) with which a particular content item in the message image payload 406 of the message 400 is associated. For example, multiple images within the message image payload 406 may each be associated with multiple content collections using identifier values.

- Message tag 420: each message 400 may be tagged with multiple tags, each of which is indicative of the subject matter of content included in the message payload. For example, where a particular image included in the message image payload 406 depicts an animal (e.g., a lion), a tag value may be included within the message tag 420 that is indicative of the relevant animal. Tag values may be generated manually, based on user input, or may be automatically generated using, for example, image recognition.
- Message sender identifier 422: an identifier (e.g., a messaging system identifier, email address, or device identifier) indicative of a user of the user system 102 on which the message 400 was generated and from which the message 400 was sent.
- Message receiver identifier 424: an identifier (e.g., a messaging system identifier, email address, or device identifier) indicative of a user of the user system 102 to which the message 400 is addressed.

[0076] The contents (e.g., values) of the various components of message 400 may be pointers to locations in tables within which content data values are stored. For example, an image value in the message image payload 406 may be a pointer to (or address of) a location within an image table 316. Similarly, values within the message video payload 408 may point to data stored within an image table 316, values stored within the message augmentation data 412 may point to data stored in an augmentation table 312, values stored within the message story identifier 418 may point to data stored in a collections table 318, and values stored within the message sender identifier 422 and the message receiver identifier 424 may point to user records stored within an entity table 308.

OBJECT GENERATION SYSTEM

[0080] FIG. 5 is a block diagram showing an example object generation system 500, according to some examples. The object generation system 500 includes an image input component 510, an XR experience component 530, and a mirrored image generation component 522. Together, these components enable the object generation system 500 to receive a first image depicting or data describing a first virtual or real-world 3D object with a target (e.g., a first fashion item with a graphical element of a pair of fashion items) and generate a second image (or modify the first image) or generate a new 3D object (e.g., a second virtual 3D object) that is a mirror image of the 3D object and target (e.g., a second fashion item with the graphical element of the pair of fashion items). In some cases, the second 3D object can be used to generate or render the second virtual object as a virtual object that can be placed in an XR experience, such as a virtual try-on experience. This allows the user to interact independently with both the first

virtual or real-world object in the XR experience and the second virtual object, such as to try on a pair of shoes. Namely, the object generation system 500 can generate a second fashion item of a pair of fashion items by processing the first fashion item in the pair of fashion items while preserving any targets (e.g., emblems, logos, text, and so forth) that are depicted on the first fashion item.

[0081] The object generation system 500 can generate the mirror image of the 3D object (e.g., the second fashion item in the pair of fashion items) in any one or combination of two or more techniques. In some examples, the object generation system 500 can apply multiple machine learning models to the first 3D object (e.g., real-world or virtual object). In some cases, a first machine learning model can be trained to detect the first 3D object in a received image. A process can be applied to the first 3D object to generate a second 3D object that is a mirror image of the first 3D object. A second machine learning model can receive the output of the first machine learning model and can be trained to identify or generate a bounding box around respective targets (e.g., emblems, logos, text, and so forth) that are part of the detected first 3D object. The output of the second machine learning model can be provided to a third machine learning model that is trained to generate a segmentation of the respective targets to enable distinguishing pixels in the first 3D object of the bounding box of the target from other pixels in the bounding box.

[0082] A fourth machine learning model or other process can implement a flip transformation to flip or generate a mirrored image of the pixels of the first 3D object corresponding to the target. Mirroring a fashion item may also mirror the target and if the target includes text or an asymmetrical logo or graphical element, the text and/or logo or graphical element may appear reversed on the mirrored version. Because of this, the disclosed techniques intelligently generate the mirrored version of the text and/or logo or graphical element to ensure that it appears coherent and consistent with the original fashion item. This flipped or mirrored image of the pixels are then blended with the mirrored image of the first 3D object, such as with a specified portion of the mirrored image of the first 3D object that corresponds to or matches a portion of the first 3D object on which the target is identified. The pixels can be blended while preserving original texture features of the mirrored version of the 3D object. In some cases, inpainting techniques can be applied to populate unknown pixel values in the mirrored version of the 3D object. This results in a second 3D object that resembles a true mirror image of the first 3D object and can be combined with the first 3D object to form a pair of fashion items (e.g., a pair of virtual 3D objects).

[0083] In some examples, a single machine learning model can be trained to generate mirror images of 3D objects. This can avoid having to apply multiple machine learning models to

generate the second 3D object that resembles a true mirror image of the first 3D object and can be combined with the first 3D object to form a pair of fashion items. Namely, in such cases, the first 3D object can be processed or analyzed by the single machine learning model (e.g., a diffusion model and/or a convolutional neural network). The single machine learning model can generate an artificial image or a new 3D object that corresponds to or resembles a mirror image of the first 3D object including mirror images of one or more targets that are part of the first 3D object. The output of the single machine learning model can be used to generate a new image or modify an image that depicts the first 3D object to depict a pair of fashion items including the first 3D object and the second 3D object.

[0084] Specifically, the image input component 510 receives one or more input images. These input images can depict a user (a person) and can be received from a message, a communication, and/or from storage or an online database. These input images can also include or depict one or more target fashion items. The image input component 510, in some cases, activates a rear-facing or front-facing camera of the user system 102. The activated camera captures a first image or video that depicts a real-world object, such as a real-world person wearing a real-world fashion item within a certain real-world background. The image input component 510, in some cases, activates a rear-facing or front-facing camera of the user system 102. The activated camera can capture a second image or video that depicts a target fashion item. During training of the XR experience component 530, the image input component 510 retrieves one or more training images and corresponding ground truth images, as discussed below. In some cases, the image input component 510 receives a 3D object directly without receiving an image that depicts the 3D object. For example, the image input component 510 can receive data describing the 3D object.

[0085] In some examples, after the first and second 3D objects are available as part of the same fashion item pair (by way of the single machine learning model being applied to the first object and/or by way of multiple machine learning models being applied to the first object), a quality control machine learning model can be applied to the first and second 3D objects. The quality control machine learning model can be trained to generate an accuracy score representing how accurately the second 3D object resembles the first 3D object as a mirror image of the first 3D object. The quality control machine learning model can include a discriminator network that is trained to distinguish real fashion item pairs from fake fashion item pairs.

[0086] In some examples, the image input component 510 receives an instruction from a user to generate a second fashion item in a pair of fashion items when only the first fashion item pair is available. In some cases, the image input component 510 can receive an image that depicts only one fashion item in a pair of fashion items or can receive a 3D object that is only one

fashion item in the pair of fashion items. Namely, a left shoe in a pair of shoes can be available and can be presented as a first virtual or real-world object in an image. The instruction from the user can request that a right shoe in the pair of shoes be made available and generated so that the user can visualize and interact with both the left shoe and the right shoe in the pair of shoes.

- 5 While the disclosed techniques are discussed with respect to shoes as fashion item pairs, similar processes and techniques can be applied to any type of fashion item that involves pairs of fashion items including a pair of sleeves on a shirt, a pair of lenses of glasses, a pair of earrings, a pair of gloves, a pair of socks, and so forth.

[0087] In response to receiving the request to generate the pair of fashion items
10 (automatically or based on user input), the image input component 510 transmits the image that depicts the object or the 3D object itself (e.g., the first fashion item of a pair of fashion items) to the mirrored image generation component 522. In some cases, the image can correspond to a virtual object that represents the first fashion item and that includes one or more targets (e.g., logos, emblems, text, and so forth). For example, as shown in **FIG. 6**, the image input
15 component 510 receives and/or access an object 600 (e.g., a real or virtual left shoe). The object 600 can include one or more targets 610, such as a logo or emblem or text.

[0088] The object 600 can be provided to the mirrored image generation component 522. The mirrored image generation component 522 applies one or more machine learning models to generate a second virtual object 601 that is a mirrored version of the object 600 and that includes
20 a target 620 that is a mirrored version of the one or more targets 610.

[0089] For example, the mirrored image generation component 522 can implement a first machine learning model that is trained to detect certain types of objects in images. The first machine learning model can be specifically trained to detect fashion items of the type specified by the instructions, such as shoes. In such cases, the first machine learning model can output a
25 bounding box around the real-world or virtual object 601. A portion of the image is extracted corresponding to this bounding box. This portion is then applied and processed by a second machine learning model that is trained to detect or identify targets (e.g., emblems, logos, and/or text). The second machine learning model can be trained using training data that includes various samples of targets of different types along with ground truth labels indicating whether the targets
30 are present or not and their respective locations in the images. After iterating through the training data and updating parameters of the second machine learning model until a stopping criterion is reached, the second machine learning model is output as a trained second machine learning model. The trained second machine learning model can be applied to the portion of the image received from the first machine learning model and can generate additional bounding boxes
35 around each of the one or more targets 610 that are identified by the trained second machine

learning model. For example, the one or more targets 610 can include a logo around which a bounding box is generated.

[0090] One or more portions of the image corresponding to the one or more bounding boxes around the one or more targets 610 are then provided to a fourth machine learning model. The fourth machine learning model can be trained to generate one or more segmentations respectively for the one or more targets. The fourth machine learning model can be a specifically trained machine learning model that processes or analyzes input images and identifies targets (e.g., logos, text, emblems, graphical elements, and so forth) in the images. The output of the fourth machine learning model can include a border that defines pixels in the image that correspond to the target to distinguish pixels of the target from other pixels in the image. The segmentation can be applied to the image that depicts the targets 610 to extract only those pixels that correspond to the one or more targets 610. In some examples, the third and/or fourth machine learning model can include a neural network that is trained to analyze a mesh texture and identify any areas that contain text or logos using object detection and/or semantic segmentation techniques.

[0091] In some examples, the fourth machine learning model can include an artificial neural network that is trained on a set of training images that depicts targets of different types. The fourth machine learning model can select a first training image and can generate or predict an estimated segmentation of the target depicted in the first training image. The estimated segmentation of the target can be compared with a ground truth segmentation (that is also part of the training data and is paired with the first training image). A deviation can be computed based on the comparison and one or more parameters of the fourth machine learning model can be updated based on the deviation. Additional training images are processed in a similar manner until a stopping criterion is reached or satisfied, at which point parameters of the fourth machine learning model are stored and used to provide the trained fourth machine learning model. The fourth machine learning model can include an R-CNN model.

[0092] In some examples, another process or machine learning model is applied to the object 600. The output of the process or machine learning model can provide a virtual object 601 that is a mirror image of the object 600. The virtual object 601 can include a mirror image of the one or more targets 610. Namely, the process or machine learning model can flip the object 600 about an axis (e.g., a center axis) of a bounding box generated by the first machine learning model to provide an intermediate virtual object 601. The result of such flipping can provide mirror images (shown as one or more targets 620) of the one or more targets 610 that are not truly accurate or photorealistic. As such, the mirrored image generation component 522 can identify a particular region of the intermediate virtual object 601 that matches or corresponds to a region of the object

600 that includes a depiction of the one or more targets 610. The mirrored image generation component 522 can then obtain pixel values of the one or more targets 620 using the output of the fourth machine learning model. The mirrored image generation component 522 can flip those obtained pixel values, such as about an axis (e.g., a central axis) of the bounding box of the one or more targets 620. The mirrored image generation component 522 can then copy those pixel values to the particular region of the intermediate object to blend those pixel values with those present on the particular region.

[0093] In some examples, the blending of the pixels can be performed according to a Poisson blending process and/or using a blending map. This results in the generation of the second virtual object 602. In some examples, the mirrored image generation component 522 can detect portions of the second virtual object 602 that are missing colors and/or textures. The mirrored image generation component 522 can apply another machine learning model that performs inpainting for such portions to complete the missing colors and/or textures in a predictive manner.

[0094] In some examples, the object 600 and the second virtual object 602 can be paired to form a pair of fashion items and provided to the XR experience component 530. In some cases, prior to being provided to the XR experience component 530, the mirrored image generation component 522 can apply a quality control machine learning model to verify whether the object 600 and the second virtual object 602 accurately represent a pair of fashion items. If the score generated by the quality control machine learning model transgresses a minimum threshold, the pair of objects representing the pair of fashion items are output and can be used to provide an XR experience for a user. The XR experience component 530 can generate or modify a new image that depicts the object 600 and the second virtual object 602 on a real-world person to provide a virtual try-on experience.

[0095] The quality control machine learning model can include a discriminator that generates a score representing match accuracy between an object and a mirrored version of the object. The quality control machine learning model can be trained by accessing training data that includes a training pair of fashion items including a training fashion item and a corresponding mirrored version of the training fashion item. The quality control machine learning model can analyze the training pair of fashion items using the discriminator to generate a prediction indicative of whether the mirrored version of the training fashion item is a real mirrored version of the training fashion item or a fake mirrored version of the training fashion item. A fake mirrored version may represent a mirrored version of a fashion item with imperfections and may not correspond precisely to how the mirrored version should look (e.g., the mirrored version may have text that is reversed or a logo that is incoherent due to being mirrored). A real version of the mirrored fashion item can be imperceptibly similar to the original fashion item including the text

and/or logos or graphical elements. The quality control machine learning model accesses ground truth labels for the training pair of fashion items indicating whether the mirrored version of the training fashion item is real or fake and updates one or more parameters of the discriminator based on accuracy of the prediction in relation to the ground truth labels. The quality control machine learning model then repeats these operations for another set of training pairs of fashion items or batches of pairs of fashion items until a stopping criterion is reached.

[0096] In order to train the quality control machine learning model, the training data can be at least partially artificially generated. Namely, to generate the training data, a plurality of 3D objects resembling pairs of fashion items can be accessed, such as from a database or an online inventory of a store. A subset of the plurality of 3D objects can be selected. Noise can then be added to an individual fashion item of a pair of fashion items depicted in the subset of the plurality of 3D objects. For example, a first pair of fashion items can be identified in a first 3D object of the subset. A target can be identified in the first pair of fashion items and the target can be blurred or modified manually for only one fashion item in the first pair of fashion items. This creates at least a partial mismatch in the first pair of fashion items. The subset of the plurality of 3D objects is then associated with a ground truth label indicating that the subset corresponds to fake pairs of fashion items. The remaining images that are not in the subset of 3D objects in which noise was added are associated with a ground truth label indicating that the images depict real pairs of fashion items. These labels are used by the discriminator in training to update a loss and parameters of the discriminator.

[0097] In some examples, a single machine learning model is implemented by the mirrored image generation component 522 to generate a corresponding virtual object to form a pair of fashion items based on a given object. In such cases, the mirrored image generation component 522 can analyze an input 3D object including a target. The mirrored image generation component 522 can generatively create another virtual object that is a mirror image of the 3D object. The generated virtual object can be used as a virtual object in an XR experience. Any targets on the original given object are mirrored in a photorealistic and true manner to avoid having text or logos appear in the wrong direction resulting from simply mirroring the 3D object. For example, the single machine learning model can be applied to the 3D object 600 including the one or more targets 610. The single machine learning model can output the virtual object 602 including one or more targets 630 that are correctly aligned and appear in the right direction. The virtual object 602 can be a mirror image of the object 600 and any of the one or more targets 610 are mirrored in a realistic manner. Namely, rather than simply having the one or more targets 610 mirrored in an improper manner (where text runs in the wrong direction) as shown in object 601 and the one or more targets 620 that are improper or not photorealistic, the object 602 (e.g., a

right shoe) is a mirror image of the object 600 (e.g., a left shoe) including the targets 630 appearing in the correctly aligned direction.

[0098] To train this single machine learning model, various pairs of fashion items can be received as part of training data. The single machine learning model can be applied to a first fashion item in a pair of fashion items and can predict or estimate a second fashion item in the pair of fashion items. The ground truth second fashion item in the pair of fashion items can be retrieved and compared with the estimated second fashion item to generate a loss. This loss is used to update one or more parameters of the single machine learning model. The single machine learning model can then be applied similarly to another pair of fashion items in the training data until a stopping criterion is satisfied. The output of the single machine learning model can be processed by the quality control machine learning model to evaluate the accuracy of the pair of fashion items that are generated by the single machine learning model.

[0099] In some cases, the given 3D object and the generated virtual 3D object can form a pair of generated fashion items that are processed by the quality control machine learning model to verify or estimate their accuracy. If the score generated by the quality control machine learning model transgresses a minimum threshold, the pair of objects representing the pair of fashion items are output and can be used to provide an XR experience for a user. The XR experience component 530 can generate or modify an image that depicts the object 600 and the second virtual object 602 on a real-world person to provide a virtual try-on experience.

[0100] FIG. 7 is a flowchart of a process or method 700 performed by the object generation system 500, according to some examples. Although the flowchart can describe the operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be re-arranged. A process is terminated when its operations are completed. A process may correspond to a method, a procedure, and the like. The steps of methods may be performed in whole or in part, may be performed in conjunction with some or all of the steps in other methods, and may be performed by any number of different systems or any portion thereof, such as a processor included in any of the systems.

[0098] At operation 701, the object generation system 500 (e.g., a user system 102 or a server) receives a 3D object comprising a target, as discussed above.

[0099] At operation 702, the object generation system 500 analyzes the 3D object using one or more machine learning models to generate data associated with a mirrored version of the target of the 3D object, as discussed above.

[0100] At operation 704, the object generation system 500 applies the mirrored version of the target to a mirrored version of the 3D object using the generated data, as discussed above.

[0101] At operation 705, the object generation system 500 generates a new 3D object comprising the mirrored version of the 3D object and the mirrored version of the target, as discussed above.

EXAMPLES

[0102] Example 1. A method comprising: receiving a three-dimensional (3D) object comprising a target; analyzing the 3D object using one or more machine learning models to generate data associated with a mirrored version of the target of the 3D object; applying the mirrored version of the target to a mirrored version of the 3D object using the generated data; and generating a new 3D object comprising the mirrored version of the 3D object and the mirrored version of the target.

[0103] Example 2. The method of Example 1, wherein the 3D object and the mirrored version of the 3D object represent a pair of fashion items, and wherein the target and the mirrored version of the target comprise one or more graphical elements or text on the pair of fashion items.

[0104] Example 3. The method of any one of Examples 1-2, further comprising: extracting the 3D object from an image; and generating the mirrored version of the 3D object from the extracted 3D object, wherein the 3D object comprises a first virtual or real-world object and the mirrored version of the 3D object comprises a second virtual object, wherein the mirrored version of the 3D object is separately generated from the generated data associated with the mirrored version of the target.

[0105] Example 4. The method of Example 3, further comprising: identifying a first bounding box of the 3D object that includes the target; determining a second bounding box in the mirrored version of the 3D object that corresponds to the first bounding box; and replacing visual content in the second bounding box with the mirrored version of the 3D object to generate the mirrored version of the 3D object comprising the mirrored version of the target.

[0106] Example 5. The method of any one of Examples 1-4, further comprising: applying a first machine learning model to the 3D object to identify a bounding box comprising the target; and applying a second machine learning model to generate a segmentation of the target in the bounding box, the segmentation distinguishing pixels in the 3D object of the bounding box of the target from other pixels in the bounding box.

[0107] Example 6. The method of Example 5, further comprising: applying a flip transformation to the segmentation of the target in the bounding box, the flip transformation flipping the segmentation of the target around a flipping axis of the bounding box thereby to generate a flipped segmentation of the target.

5 **[0108]** Example 7. The method of Example 6, further comprising: blending pixels of the flipped segmentation of the target to a specified portion of the mirrored version of the 3D object.

[0109] Example 8. The method of Example 7, wherein the pixels of the flipped segmentation are blended to the specified portion while preserving original texture features
10 of the mirrored version of the 3D object.

[0110] Example 9. The method of any one of Examples 7-8, wherein the pixels are blended according to at least one of Poisson blending techniques or a blending map.

[0111] Example 10. The method of any one of Examples 7-9, further comprising: after blending the pixels, identifying an area in the mirrored version of the 3D object with
15 unknown pixel values; and applying inpainting techniques to populate the unknown pixel values with predicted values.

[0112] Example 11. The method of any one of Examples 1-10, further comprising: analyzing the 3D object comprising the target and the mirrored version of the 3D object comprising the mirrored version of the target with a quality control machine learning model
20 to generate a score representing match accuracy between the mirrored version of the object comprising the mirrored version of the target and the 3D object comprising the target.

[0113] Example 12. The method of Example 11, wherein the quality control machine learning model comprises a discriminator trained by performing training operations comprising: accessing training data comprising a training pair of fashion items including a
25 training fashion item and a corresponding mirrored version of the training fashion item; analyzing the training pair of fashion items using the discriminator to generate a prediction indicative of whether the mirrored version of the training fashion item is a real mirrored version of the training fashion item or a fake mirrored version of the training fashion item; accessing ground truth labels for the training pair of fashion items indicating whether the
30 mirrored version of the training fashion item is real or fake; and updating one or more parameters of the discriminator based on accuracy of the prediction in relation to the ground truth labels.

[0114] Example 13. The method of Example 12, further comprising: accessing a plurality of 3D objects representing pairs of fashion items; selecting a subset of the plurality of 3D

objects; adding noise to an individual fashion item of each pair of fashion items in the subset of the plurality of 3D objects; associating a first set of ground truth labels for each of the subset of the plurality of 3D objects indicating that the mirrored version of the fashion item in each of the subset of the plurality of 3D objects is fake; and associating a second set of ground truth labels for each of the plurality of 3D objects outside of the subset of the plurality of 3D objects indicating that the mirrored version of the fashion item in each of the plurality of 3D objects outside of the subset of the plurality of 3D objects is real.

[0115] Example 14. The method of any one of Examples 1-13, further comprising: generating an extended reality (XR) experience using the 3D object and the new 3D object.

[0116] Example 15. The method of any one of Examples 1-14, further comprising: generating an artificial image depicting the mirrored version of the 3D object, the artificial image being modified using at least a portion of the new 3D object.

[0117] Example 16. The method of Example 15, wherein the one or more machine learning models comprise a convolutional neural network trained by performing training operations comprising: accessing training data comprising a training pair of fashion items including a training fashion item and a ground truth mirrored version of the training fashion item; analyzing the training fashion item to estimate a mirrored version of the training fashion item; computing a deviation between the estimated mirrored version of the training fashion item and the ground truth mirrored version of the training fashion item; and updating one or more parameters of the convolutional neural network based on the computed deviation.

[0118] Example 17. The method of Example 16, further comprising repeating the training operations for additional training data until a stopping criterion is met.

[0119] Example 18. The method of any one of Examples 16-17, wherein the convolutional neural network is trained to establish a relationship between mapping of pixels in a first fashion item in a first pair of fashion items to pixels in a second fashion item in the first pair of fashion items.

[0120] Example 19. A system comprising: at least one processor; and at least one memory component having instructions stored thereon that, when executed by the at least one processor, cause the at least one processor to perform operations comprising: receiving a three-dimensional (3D) object comprising a target; analyzing the 3D object using one or more machine learning models to generate data associated with a mirrored version of the target of the 3D object; applying the mirrored version of the target to a mirrored version of the 3D object using the generated data; and generating a new 3D object comprising the mirrored version of the 3D object and the mirrored version of the target.

[0121] Example 20. A non-transitory computer-readable storage medium having stored thereon instructions that, when executed by at least one processor, cause the at least one processor to perform operations comprising: receiving a three-dimensional (3D) object comprising a target; analyzing the 3D object using one or more machine learning models to generate data associated with a mirrored version of the target of the 3D object; applying the mirrored version of the target to a mirrored version of the 3D object using the generated data; and generating a new 3D object comprising the mirrored version of the 3D object and the mirrored version of the target.

MACHINE ARCHITECTURE

[0122] FIG. 8 is a diagrammatic representation of a machine 800 within which instructions 802 (e.g., software, a program, an application, an applet, an app, or other executable code) for causing the machine 800 to perform any one or more of the methodologies discussed herein may be executed. For example, the instructions 802 may cause the machine 800 to execute any one or more of the methods described herein. The instructions 802 transform the general, non-programmed machine 800 into a particular machine 800 programmed to carry out the described and illustrated functions in the manner described. The machine 800 may operate as a standalone device or may be coupled (e.g., networked) to other machines. In a networked deployment, the machine 800 may operate in the capacity of a server machine or a client machine in a server-client network environment, or as a peer machine in a peer-to-peer (or distributed) network environment. The machine 800 may comprise, but not be limited to, a server computer, a client computer, a personal computer (PC), a tablet computer, a laptop computer, a netbook, a set-top box (STB), a personal digital assistant (PDA), an entertainment media system, a cellular telephone, a smartphone, a mobile device, a wearable device (e.g., a smartwatch), a smart home device (e.g., a smart appliance), other smart devices, a web appliance, a network router, a network switch, a network bridge, or any machine capable of executing the instructions 802, sequentially or otherwise, that specify actions to be taken by the machine 800. Further, while a single machine 800 is illustrated, the term “machine” shall also be taken to include a collection of machines that individually or jointly execute the instructions 802 to perform any one or more of the methodologies discussed herein. The machine 800, for example, may comprise the user system 102 or any one of multiple server devices forming part of the interaction server system 110. In some examples, the machine 800 may also comprise both client and server systems, with certain operations of a particular method or algorithm being performed on the server-side and with certain operations of the particular method or algorithm being performed on the client-side.

[0123] The machine 800 may include processors 804, memory 806, and input/output (I/O) components 808, which may be configured to communicate with each other via a bus 810. In an example, the processors 804 (e.g., a Central Processing Unit (CPU), a Reduced Instruction Set Computing (RISC) Processor, a Complex Instruction Set Computing (CISC) Processor, a Graphics Processing Unit (GPU), a Digital Signal Processor (DSP), an Application Specific Integrated Circuit (ASIC), a Radio-Frequency Integrated Circuit (RFIC), another processor, or any suitable combination thereof) may include, for example, a processor 812 and a processor 814 that execute the instructions 802. The term "processor" is intended to include multi-core processors that may comprise two or more independent processors (sometimes referred to as "cores") that may execute instructions contemporaneously. Although FIG. 8 shows multiple processors 804, the machine 800 may include a single processor with a single-core, a single processor with multiple cores (e.g., a multi-core processor), multiple processors with a single core, multiple processors with multiples cores, or any combination thereof.

[0124] The memory 806 includes a main memory 816, a static memory 818, and a storage unit 820, all accessible to the processors 804 via the bus 810. The main memory 806, the static memory 818, and storage unit 820 store the instructions 802 embodying any one or more of the methodologies or functions described herein. The instructions 802 may also reside, completely or partially, within the main memory 816, within the static memory 818, within machine-readable medium 822 within the storage unit 820, within at least one of the processors 804 (e.g., within the processor's cache memory), or any suitable combination thereof, during execution thereof by the machine 800.

[0125] The I/O components 808 may include a wide variety of components to receive input, provide output, produce output, transmit information, exchange information, capture measurements, and so on. The specific I/O components 808 that are included in a particular machine will depend on the type of machine. For example, portable machines such as mobile phones may include a touch input device or other such input mechanisms, while a headless server machine will likely not include such a touch input device. It will be appreciated that the I/O components 808 may include many other components that are not shown in FIG. 8. In various examples, the I/O components 808 may include user output components 824 and user input components 826. The user output components 824 may include visual components (e.g., a display such as a plasma display panel (PDP), a light-emitting diode (LED) display, a liquid crystal display (LCD), a projector, or a cathode ray tube (CRT)), acoustic components (e.g., speakers), haptic components (e.g., a vibratory motor, resistance mechanisms), other signal generators, and so forth. The user input

components 826 may include alphanumeric input components (e.g., a keyboard, a touch screen configured to receive alphanumeric input, a photo-optical keyboard, or other alphanumeric input components), point-based input components (e.g., a mouse, a touchpad, a trackball, a joystick, a motion sensor, or another pointing instrument), tactile input components (e.g., a physical button, a touch screen that provides location and force of touches or touch gestures, or other tactile input components), audio input components (e.g., a microphone), and the like. Any biometric collected by the biometric components is captured and stored with user approval and deleted on user request.

[0126] Further, such biometric data may be used for very limited purposes, such as identification verification. To ensure limited and authorized use of biometric information and other personally identifiable information (PII), access to this data is restricted to authorized personnel only, if allowed at all. Any use of biometric data may strictly be limited to identification verification purposes, and the data is not shared or sold to any third party without the explicit consent of the user. In addition, appropriate technical and organizational measures are implemented to ensure the security and confidentiality of this sensitive information.

[0127] In further examples, the I/O components 808 may include biometric components 828, motion components 830, environmental components 832, or position components 834, among a wide array of other components. For example, the biometric components 828 include components to detect expressions (e.g., hand expressions, facial expressions, vocal expressions, body gestures, or eye-tracking), measure biosignals (e.g., blood pressure, heart rate, body temperature, perspiration, or brain waves), identify a person (e.g., voice identification, retinal identification, facial identification, fingerprint identification, or electroencephalogram-based identification), and the like.

[0128] The biometric components may include a brain-machine interface (BMI) system that allows communication between the brain and an external device or machine. This may be achieved by recording brain activity data, translating this data into a format that can be understood by a computer, and then using the resulting signals to control the device or machine.

[0129] Example types of BMI technologies include:

- Electroencephalography (EEG) based BMIs, which record electrical activity in the brain using electrodes placed on the scalp.
- Invasive BMIs, which use electrodes that are surgically implanted into the brain.

- Optogenetics BMIs, which use light to control the activity of specific nerve cells in the brain.

[0130] The motion components 830 include acceleration sensor components (e.g., accelerometer), gravitation sensor components, rotation sensor components (e.g., gyroscope).

[0131] The environmental components 832 include, for example, one or cameras (with still image/photograph and video capabilities), illumination sensor components (e.g., photometer), temperature sensor components (e.g., one or more thermometers that detect ambient temperature), humidity sensor components, pressure sensor components (e.g., barometer), acoustic sensor components (e.g., one or more microphones that detect background noise), proximity sensor components (e.g., infrared sensors that detect nearby objects), gas sensors (e.g., gas detection sensors to detection concentrations of hazardous gases for safety or to measure pollutants in the atmosphere), or other components that may provide indications, measurements, or signals corresponding to a surrounding physical environment.

[0132] With respect to cameras, the user system 102 may have a camera system comprising, for example, front cameras on a front surface of the user system 102 and rear cameras on a rear surface of the user system 102. The front cameras may, for example, be used to capture still images and video of a user of the user system 102 (e.g., “selfies”), which may then be augmented with augmentation data (e.g., filters) described above. The rear cameras may, for example, be used to capture still images and videos in a more traditional camera mode, with these images similarly being augmented with augmentation data. In addition to front and rear cameras, the user system 102 may also include a 360° camera for capturing 360° photographs and videos.

[0133] Further, the camera system of the user system 102 may include dual rear cameras (e.g., a primary camera as well as a depth-sensing camera), or even triple, quad, or penta rear camera configurations on the front and rear sides of the user system 102. These multiple cameras systems may include a wide camera, an ultra-wide camera, a telephoto camera, a macro camera, and a depth sensor, for example.

[0134] The position components 834 include location sensor components (e.g., a GPS receiver component), altitude sensor components (e.g., altimeters or barometers that detect air pressure from which altitude may be derived), orientation sensor components (e.g., magnetometers), and the like.

[0135] Communication may be implemented using a wide variety of technologies. The I/O components 808 further include communication components 836 operable to couple the machine 800 to a network 838 or devices 840 via respective coupling or connections. For example, the communication components 836 may include a network interface component or another suitable device to interface with the network 838. In further examples, the communication components 836 may include wired communication components, wireless communication components, cellular communication components, Near Field Communication (NFC) components, Bluetooth® components (e.g., Bluetooth® Low Energy), Wi-Fi® components, and other communication components to provide communication via other modalities. The devices 840 may be another machine or any of a wide variety of peripheral devices (e.g., a peripheral device coupled via a universal serial bus (USB)).

[0136] Moreover, the communication components 836 may detect identifiers or include components operable to detect identifiers. For example, the communication components 836 may include Radio Frequency Identification (RFID) tag reader components, NFC smart tag detection components, optical reader components (e.g., an optical sensor to detect one-dimensional bar codes such as Universal Product Code (UPC) bar code, multi-dimensional bar codes such as Quick Response (QR) code, Aztec code, Data Matrix, Dataglyph™, MaxiCode, PDF417, Ultra Code, UCC RSS-2D bar code, and other optical codes), or acoustic detection components (e.g., microphones to identify tagged audio signals). In addition, a variety of information may be derived via the communication components 836, such as location via Internet Protocol (IP) geolocation, location via Wi-Fi® signal triangulation, location via detecting an NFC beacon signal that may indicate a particular location, and so forth.

[0137] The various memories (e.g., main memory 816, static memory 818, and memory of the processors 804) and storage unit 820 may store one or more sets of instructions and data structures (e.g., software) embodying or used by any one or more of the methodologies or functions described herein. These instructions (e.g., the instructions 802), when executed by processors 804, cause various operations to implement the disclosed examples.

[0138] The instructions 802 may be transmitted or received over the network 838, using a transmission medium, via a network interface device (e.g., a network interface component included in the communication components 836) and using any one of several well-known transfer protocols (e.g., HTTP). Similarly, the instructions 802 may be transmitted or received using a transmission medium via a coupling (e.g., a peer-to-peer coupling) to the devices 840.

SOFTWARE ARCHITECTURE

[0139] FIG. 9 is a block diagram 900 illustrating a software architecture 902, which can be installed on any one or more of the devices described herein. The software architecture 902 is supported by hardware such as a machine 904 that includes processors 906, memory 908, and I/O components 910. In this example, the software architecture 902 can be conceptualized as a stack of layers, where each layer provides a particular functionality. The software architecture 902 includes layers such as an operating system 912, libraries 914, frameworks 916, and applications 918. Operationally, the applications 918 invoke API calls 920 through the software stack and receive messages 922 in response to the API calls 920.

[0140] The operating system 912 manages hardware resources and provides common services. The operating system 912 includes, for example, a kernel 924, services 926, and drivers 928. The kernel 924 acts as an abstraction layer between the hardware and the other software layers. For example, the kernel 924 provides memory management, processor management (e.g., scheduling), component management, networking, and security settings, among other functionalities. The services 926 can provide other common services for the other software layers. The drivers 928 are responsible for controlling or interfacing with the underlying hardware. For instance, the drivers 928 can include display drivers, camera drivers, BLUETOOTH® or BLUETOOTH® Low Energy drivers, flash memory drivers, serial communication drivers (e.g., USB drivers), WI-FI® drivers, audio drivers, power management drivers, and so forth.

[0141] The libraries 914 provide a common low-level infrastructure used by the applications 918. The libraries 914 can include system libraries 930 (e.g., C standard library) that provide functions such as memory allocation functions, string manipulation functions, mathematic functions, and the like. In addition, the libraries 914 can include API libraries 932 such as media libraries (e.g., libraries to support presentation and manipulation of various media formats such as Moving Picture Experts Group-4 (MPEG4), Advanced Video Coding (H.264 or AVC), Moving Picture Experts Group Layer-3 (MP3), Advanced Audio Coding (AAC), Adaptive Multi-Rate (AMR) audio codec, Joint Photographic Experts Group (JPEG or JPG), or Portable Network Graphics (PNG)), graphics libraries (e.g., an OpenGL framework used to render in 2D and 3D in a graphic content on a display), database libraries (e.g., SQLite to provide various relational database functions), web libraries (e.g., WebKit to provide web browsing functionality), and the like. The libraries

914 can also include a wide variety of other libraries 934 to provide many other APIs to the applications 918.

[0142] The frameworks 916 provide a common high-level infrastructure that is used by the applications 918. For example, the frameworks 916 provide various GUI functions, high-level resource management, and high-level location services. The frameworks 916 can provide a broad spectrum of other APIs that can be used by the applications 918, some of which may be specific to a particular operating system or platform.

[0143] In an example, the applications 918 may include a home application 936, a contacts application 938, a browser application 940, a book reader application 942, a location application 944, a media application 946, a messaging application 948, a game application 950, and a broad assortment of other applications such as a third-party application 952. The applications 918 are programs that execute functions defined in the programs. Various programming languages can be employed to create one or more of the applications 918, structured in a variety of manners, such as object-oriented programming languages (e.g., Objective-C, Java, or C++) or procedural programming languages (e.g., C or assembly language). In a specific example, the third-party application 952 (e.g., an application developed using the ANDROID™ or IOS™ SDK by an entity other than the vendor of the particular platform) may be mobile software running on a mobile operating system such as IOS™, ANDROID™, WINDOWS® Phone, or another mobile operating system. In this example, the third-party application 952 can invoke the API calls 920 provided by the operating system 912 to facilitate functionalities described herein.

SYSTEM WITH HEAD-WEARABLE APPARATUS

[0144] FIG. 10 illustrates a system 1000 including a head-wearable apparatus 116 with a selector input device, according to some examples. FIG. 10 is a high-level functional block diagram of an example head-wearable apparatus 116 communicatively coupled to a mobile device 114 and various server systems 1004 (e.g., the interaction server system 110) via various networks 1016.

[0145] The head-wearable apparatus 116 includes one or more cameras, each of which may be, for example, a visible light camera 1006, an infrared emitter 1008, and an infrared camera 1010.

[0146] The mobile device 114 connects with head-wearable apparatus 116 using both a low-power wireless connection 1012 and a high-speed wireless connection 1014. The mobile device 114 is also connected to the server system 1004 and the network 1016.

[0147] The head-wearable apparatus 116 further includes two image displays of optical assembly 1018. The two image displays of optical assembly 1018 include one associated with the left lateral side and one associated with the right lateral side of the head-wearable apparatus 116. The head-wearable apparatus 116 also includes an image display driver 1020, an image processor 1022, low-power circuitry 1024, and high-speed circuitry 1026. The image display of optical assembly 1018 is for presenting images and videos, including an image that can include a GUI, to a user of the head-wearable apparatus 116.

[0148] The image display driver 1020 commands and controls the image display of optical assembly 1018. The image display driver 1020 may deliver image data directly to the image display of optical assembly 1018 for presentation or may convert the image data into a signal or data format suitable for delivery to the image display device. For example, the image data may be video data formatted according to compression formats, such as H.264 (MPEG-4 Part 10), HEVC, Theora, Dirac, RealVideo RV40, VP8, VP9, or the like, and still image data may be formatted according to compression formats such as PNG, JPEG, Tagged Image File Format (TIFF) or exchangeable image file format (EXIF) or the like.

[0149] The head-wearable apparatus 116 includes a frame and stems (or temples) extending from a lateral side of the frame. The head-wearable apparatus 116 further includes a user input device 1028 (e.g., touch sensor or push button), including an input surface on the head-wearable apparatus 116. The user input device 1028 (e.g., touch sensor or push button) is to receive from the user an input selection to manipulate the GUI of the presented image.

[0150] The components shown in FIG. 10 for the head-wearable apparatus 116 are located on one or more circuit boards, for example a PCB or flexible PCB, in the rims or temples. Alternatively, or additionally, the depicted components can be located in the chunks, frames, hinges, or bridge of the head-wearable apparatus 116. Left and right visible light cameras 1006 can include digital camera elements such as a complementary metal oxide-semiconductor (CMOS) image sensor, charge-coupled device, camera lenses, or any other respective visible or light-capturing elements that may be used to capture data, including images of scenes with unknown objects.

[0151] The head-wearable apparatus 116 includes a memory 1002, which stores instructions to perform a subset or all of the functions described herein. The memory 1002 can also include a storage device.

[0152] As shown in FIG. 10, the high-speed circuitry 1026 includes a high-speed processor 1030, a memory 1002, and high-speed wireless circuitry 1032. In some examples, the

image display driver 1020 is coupled to the high-speed circuitry 1026 and operated by the high-speed processor 1030 in order to drive the left and right image displays of the image display of optical assembly 1018. The high-speed processor 1030 may be any processor capable of managing high-speed communications and operation of any general computing system needed for the head-wearable apparatus 116. The high-speed processor 1030 includes processing resources needed for managing high-speed data transfers on a high-speed wireless connection 1014 to a wireless local area network (WLAN) using the high-speed wireless circuitry 1032. In certain examples, the high-speed processor 1030 executes an operating system such as a LINUX operating system or other such operating system of the head-wearable apparatus 116, and the operating system is stored in the memory 1002 for execution. In addition to any other responsibilities, the high-speed processor 1030 executing a software architecture for the head-wearable apparatus 116 is used to manage data transfers with high-speed wireless circuitry 1032. In certain examples, the high-speed wireless circuitry 1032 is configured to implement Institute of Electrical and Electronic Engineers (IEEE) 802.11 communication standards, also referred to herein as WiFi. In some examples, other high-speed communications standards may be implemented by the high-speed wireless circuitry 1032.

[0153] Low-power wireless circuitry 1034 and the high-speed wireless circuitry 1032 of the head-wearable apparatus 116 can include short-range transceivers (Bluetooth™) and wireless wide, local, or wide area network transceivers (e.g., cellular or WiFi). Mobile device 114, including the transceivers communicating via the low-power wireless connection 1012 and the high-speed wireless connection 1014, may be implemented using details of the architecture of the head-wearable apparatus 116, as can other elements of the network 1016.

[0154] The memory 1002 includes any storage device capable of storing various data and applications, including, among other things, camera data generated by the left and right visible light cameras 1006, the infrared camera 1010, and the image processor 1022, as well as images generated for display by the image display driver 1020 on the image displays of the image display of optical assembly 1018. While the memory 1002 is shown as integrated with high-speed circuitry 1026, in some examples, the memory 1002 may be an independent standalone element of the head-wearable apparatus 116. In certain such examples, electrical routing lines may provide a connection through a chip that includes the high-speed processor 1030 from the image processor 1022 or low-power processor 1036 to the memory 1002. In some examples, the high-speed processor 1030 may manage addressing of the

memory 1002 such that the low-power processor 1036 will boot the high-speed processor 1030 any time that a read or write operation involving memory 1002 is needed.

[0155] As shown in FIG. 10, the low-power processor 1036 or high-speed processor 1030 of the head-wearable apparatus 116 can be coupled to the camera (visible light camera 1006, infrared emitter 1008, or infrared camera 1010), the image display driver 1020, the user input device 1028 (e.g., touch sensor or push button), and the memory 1002.

[0156] The head-wearable apparatus 116 is connected to a host computer. For example, the head-wearable apparatus 116 is paired with the mobile device 114 via the high-speed wireless connection 1014 or connected to the server system 1004 via the network 1016. The server system 1004 may be one or more computing devices as part of a service or network computing system, for example, that includes a processor, a memory, and network communication interface to communicate over the network 1016 with the mobile device 114 and the head-wearable apparatus 116.

[0157] The mobile device 114 includes a processor and a network communication interface coupled to the processor. The network communication interface allows for communication over the network 1016, low-power wireless connection 1012, or high-speed wireless connection 1014. Mobile device 114 can further store at least portions of the instructions for generating binaural audio content in the memory of mobile device 114 to implement the functionality described herein.

[0158] Output components of the head-wearable apparatus 116 include visual components, such as a display such as a LCD, a PDP, a LED display, a projector, or a waveguide. The image displays of the optical assembly are driven by the image display driver 1020. The output components of the head-wearable apparatus 116 further include acoustic components (e.g., speakers), haptic components (e.g., a vibratory motor), other signal generators, and so forth. The input components of the head-wearable apparatus 116, the mobile device 114, and server system 1004, such as the user input device 1028, may include alphanumeric input components (e.g., a keyboard, a touch screen configured to receive alphanumeric input, a photo-optical keyboard, or other alphanumeric input components), point-based input components (e.g., a mouse, a touchpad, a trackball, a joystick, a motion sensor, or other pointing instruments), tactile input components (e.g., a physical button, a touch screen that provides location and force of touches or touch gestures, or other tactile input components), audio input components (e.g., a microphone), and the like.

[0159] The head-wearable apparatus 116 may also include additional peripheral device elements. Such peripheral device elements may include biometric sensors, additional

sensors, or display elements integrated with the head-wearable apparatus 116. For example, peripheral device elements may include any I/O components including output components, motion components, position components, or any other such elements described herein.

- 5 **[0160]** For example, the biometric components include components to detect expressions (e.g., hand expressions, facial expressions, vocal expressions, body gestures, or eye-tracking), measure biosignals (e.g., blood pressure, heart rate, body temperature, perspiration, or brain waves), identify a person (e.g., voice identification, retinal identification, facial identification, fingerprint identification, or electroencephalogram-based identification), and the like. The biometric components may include a BMI system that allows communication between the brain and an external device or machine. This may be achieved by recording brain activity data, translating this data into a format that can be understood by a computer, and then using the resulting signals to control the device or machine.
- 10
- 15 **[0161]** The motion components include acceleration sensor components (e.g., accelerometer), gravitation sensor components, rotation sensor components (e.g., gyroscope), and so forth. The position components include location sensor components to generate location coordinates (e.g., a GPS receiver component), Wi-Fi or Bluetooth™ transceivers to generate positioning system coordinates, altitude sensor components (e.g., altimeters or barometers that detect air pressure from which altitude may be derived),
- 20 orientation sensor components (e.g., magnetometers), and the like. Such positioning system coordinates can also be received over low-power wireless connections 1012 and high-speed wireless connection 1014 from the mobile device 114 via the low-power wireless circuitry 1034 or high-speed wireless circuitry 1032.

25

GLOSSARY

- [0162]** "Carrier signal" refers, for example, to any intangible medium that is capable of storing, encoding, or carrying instructions for execution by the machine and includes digital or analog communications signals or other intangible media to facilitate communication of such instructions. Instructions may be transmitted or received over a network using a transmission medium via a network interface device.
- 30

[0163] "Client device" refers, for example, to any machine that interfaces to a communications network to obtain resources from one or more server systems or other client devices. A client device may be, but is not limited to, a mobile phone, desktop

computer, laptop, portable digital assistant (PDA), smartphone, tablet, ultrabook, netbook, laptop, multi-processor system, microprocessor-based or programmable consumer electronics, game console, STB, or any other communication device that a user may use to access a network.

5 **[0164]** "Communication network" refers, for example, to one or more portions of a network that may be an ad hoc network, an intranet, an extranet, a virtual private network (VPN), a local area network (LAN), a WLAN, a wide area network (WAN), a wireless WAN (WWAN), a metropolitan area network (MAN), the Internet, a portion of the Internet, a portion of the Public Switched Telephone Network (PSTN), a plain old telephone service
10 (POTS) network, a cellular telephone network, a wireless network, a Wi-Fi® network, another type of network, or a combination of two or more such networks. For example, a network or a portion of a network may include a wireless or cellular network, and the coupling may be a Code Division Multiple Access (CDMA) connection, a Global System for Mobile communications (GSM) connection, or other types of cellular or wireless
15 coupling. In this example, the coupling may implement any of a variety of types of data transfer technology, such as Single Carrier Radio Transmission Technology (1xRTT), Evolution-Data Optimized (EVDO) technology, General Packet Radio Service (GPRS) technology, Enhanced Data rates for GSM Evolution (EDGE) technology, third Generation Partnership Project (3GPP) including 3G, fourth-generation wireless (4G) networks,
20 Universal Mobile Telecommunications System (UMTS), High Speed Packet Access (HSPA), Worldwide Interoperability for Microwave Access (WiMAX), Long Term Evolution (LTE) standard, others defined by various standard-setting organizations, other long-range protocols, or other data transfer technology.

[0165] "Component" refers, for example, to a device, physical entity, or logic having
25 boundaries defined by function or subroutine calls, branch points, APIs, or other technologies that provide for the partitioning or modularization of particular processing or control functions. Components may be combined via their interfaces with other components to carry out a machine process. A component may be a packaged functional hardware unit designed for use with other components and a part of a program that usually performs a
30 particular function of related functions. Components may constitute either software components (e.g., code embodied on a machine-readable medium) or hardware components.

[0166] A "hardware component" is a tangible unit capable of performing certain operations and may be configured or arranged in a certain physical manner. In various examples, one or more computer systems (e.g., a standalone computer system, a client computer system, or
35 a server computer system) or one or more hardware components of a computer system (e.g.,

a processor or a group of processors) may be configured by software (e.g., an application or application portion) as a hardware component that operates to perform certain operations as described herein.

[0167] A hardware component may also be implemented mechanically, electronically, or any suitable combination thereof. For example, a hardware component may include dedicated circuitry or logic that is permanently configured to perform certain operations. A hardware component may be a special-purpose processor, such as a field-programmable gate array (FPGA) or an ASIC. A hardware component may also include programmable logic or circuitry that is temporarily configured by software to perform certain operations. For example, a hardware component may include software executed by a general-purpose processor or other programmable processors. Once configured by such software, hardware components become specific machines (or specific components of a machine) uniquely tailored to perform the configured functions and are no longer general-purpose processors. It will be appreciated that the decision to implement a hardware component mechanically, in dedicated and permanently configured circuitry, or in temporarily configured circuitry (e.g., configured by software), may be driven by cost and time considerations. Accordingly, the phrase "hardware component" (or "hardware-implemented component") should be understood to encompass a tangible entity, be that an entity that is physically constructed, permanently configured (e.g., hardwired), or temporarily configured (e.g., programmed) to operate in a certain manner or to perform certain operations described herein.

[0168] Considering examples in which hardware components are temporarily configured (e.g., programmed), each of the hardware components need not be configured or instantiated at any one instance in time. For example, where a hardware component comprises a general-purpose processor configured by software to become a special-purpose processor, the general-purpose processor may be configured as respectively different special-purpose processors (e.g., comprising different hardware components) at different times. Software accordingly configures a particular processor or processors, for example, to constitute a particular hardware component at one instance of time and to constitute a different hardware component at a different instance of time. Hardware components can provide information to, and receive information from, other hardware components. Accordingly, the described hardware components may be regarded as being communicatively coupled. Where multiple hardware components exist contemporaneously, communications may be achieved through signal transmission (e.g., over appropriate circuits and buses) between or among two or more of the hardware components. In examples in which multiple hardware components are configured or instantiated at different times, communications between such hardware

components may be achieved, for example, through the storage and retrieval of information in memory structures to which the multiple hardware components have access. For example, one hardware component may perform an operation and store the output of that operation in a memory device to which it is communicatively coupled. A further hardware component may then, at a later time, access the memory device to retrieve and process the stored output. Hardware components may also initiate communications with input or output devices, and can operate on a resource (e.g., a collection of information). The various operations of example methods described herein may be performed, at least partially, by one or more processors that are temporarily configured (e.g., by software) or permanently configured to perform the relevant operations. Whether temporarily or permanently configured, such processors may constitute processor-implemented components that operate to perform one or more operations or functions described herein.

[0169] As used herein, "processor-implemented component" refers to a hardware component implemented using one or more processors. Similarly, the methods described herein may be at least partially processor-implemented, with a particular processor or processors being an example of hardware. For example, at least some of the operations of a method may be performed by one or more processors or processor-implemented components. Moreover, the one or more processors may also operate to support performance of the relevant operations in a "cloud computing" environment or as a "software as a service" (SaaS). For example, at least some of the operations may be performed by a group of computers (as examples of machines including processors), with these operations being accessible via a network (e.g., the Internet) and via one or more appropriate interfaces (e.g., an API). The performance of certain of the operations may be distributed among the processors, not only residing within a single machine, but deployed across a number of machines. In some examples, the processors or processor-implemented components may be located in a single geographic location (e.g., within a home environment, an office environment, or a server farm). In other examples, the processors or processor-implemented components may be distributed across a number of geographic locations.

[0170] "Computer-readable storage medium" refers, for example, to both machine-storage media and transmission media. Thus, the terms include both storage devices/media and carrier waves/modulated data signals. The terms "machine-readable medium," "computer-readable medium," and "device-readable medium" mean the same thing and may be used interchangeably in this disclosure. "Ephemeral message" refers, for example, to a message that is accessible for a time-limited duration. An ephemeral message may be a text, an

image, a video and the like. The access time for the ephemeral message may be set by the message sender. Alternatively, the access time may be a default setting or a setting specified by the recipient. Regardless of the setting technique, the message is transitory.

[0171] "Machine storage medium" refers, for example, to a single or multiple storage devices and media (e.g., a centralized or distributed database, and associated caches and servers) that store executable instructions, routines and data. The term shall accordingly be taken to include, but not be limited to, solid-state memories, and optical and magnetic media, including memory internal or external to processors. Specific examples of machine-storage media, computer-storage media and device-storage media include non-volatile memory, including by way of example semiconductor memory devices, e.g., erasable programmable read-only memory (EPROM), electrically erasable programmable read-only memory (EEPROM), FPGA, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The terms "machine-storage medium," "device-storage medium," "computer-storage medium" mean the same thing and may be used interchangeably in this disclosure.

[0172] The terms "machine-storage media," "computer-storage media," and "device-storage media" specifically exclude carrier waves, modulated data signals, and other such media, at least some of which are covered under the term "signal medium." "Non-transitory computer-readable storage medium" refers, for example, to a tangible medium that is capable of storing, encoding, or carrying the instructions for execution by a machine. "Signal medium" refers, for example, to any intangible medium that is capable of storing, encoding, or carrying the instructions for execution by a machine and includes digital or analog communications signals or other intangible media to facilitate communication of software or data. The term "signal medium" shall be taken to include any form of a modulated data signal, carrier wave, and so forth. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a matter as to encode information in the signal. The terms "transmission medium" and "signal medium" mean the same thing and may be used interchangeably in this disclosure.

[0173] "User device" refers, for example, to a device accessed, controlled or owned by a user and with which the user interacts perform an action, or interaction on the user device, including interaction with other users or computer systems. "Carrier signal" refers to any intangible medium that is capable of storing, encoding, or carrying instructions for execution by the machine and includes digital or analog communications signals or other intangible media to facilitate communication of such instructions. Instructions may be transmitted or received over a network using a transmission medium via a network interface

device. "Client device" refers to any machine that interfaces to a communications network to obtain resources from one or more server systems or other client devices. A client device may be, but is not limited to, a mobile phone, desktop computer, laptop, PDA, smartphone, tablet, ultrabook, netbook, laptop, multi-processor system, microprocessor-based or programmable consumer electronics, game console, STB, or any other communication device that a user may use to access a network.

[0174] "Communication network" refers to one or more portions of a network that may be an ad hoc network, an intranet, an extranet, a VPN, a LAN, a WLAN, a WAN, a WWAN, a MAN, the Internet, a portion of the Internet, a portion of the PSTN, a POTS network, a cellular telephone network, a wireless network, a Wi-Fi® network, another type of network, or a combination of two or more such networks. For example, a network or a portion of a network may include a wireless or cellular network, and the coupling may be a CDMA connection, a GSM connection, or other types of cellular or wireless coupling. In this example, the coupling may implement any of a variety of types of data transfer technology, such as 1xRTT, EVDO technology, GPRS technology, EDGE technology, 3GPP including 3G, 4G networks, UMTS, HSPA, WiMAX, LTE standard, others defined by various standard-setting organizations, other long-range protocols, or other data transfer technology.

[0175] Components may constitute either software components (e.g., code embodied on a machine-readable medium) or hardware components. A "hardware component" is a tangible unit capable of performing certain operations and may be configured or arranged in a certain physical manner. In various examples, one or more computer systems (e.g., a standalone computer system, a client computer system, or a server computer system) or one or more hardware components of a computer system (e.g., a processor or a group of processors) may be configured by software (e.g., an application or application portion) as a hardware component that operates to perform certain operations as described herein.

[0176] A hardware component may also be implemented mechanically, electronically, or any suitable combination thereof. For example, a hardware component may include dedicated circuitry or logic that is permanently configured to perform certain operations. A hardware component may be a special-purpose processor, such as a FPGA or an ASIC. A hardware component may also include programmable logic or circuitry that is temporarily configured by software to perform certain operations. For example, a hardware component may include software executed by a general-purpose processor or other programmable processor. Once configured by such software, hardware components become specific machines (or specific components of a machine) uniquely tailored to perform the configured functions and are no longer general-purpose processors. It will be appreciated that the

decision to implement a hardware component mechanically, in dedicated and permanently configured circuitry, or in temporarily configured circuitry (e.g., configured by software), may be driven by cost and time considerations. Accordingly, the phrase "hardware component"(or "hardware-implemented component") should be understood to encompass a tangible entity, be that an entity that is physically constructed, permanently configured (e.g., hardwired), or temporarily configured (e.g., programmed) to operate in a certain manner or to perform certain operations described herein.

[0177] The various operations of example methods described herein may be performed, at least partially, by one or more processors that are temporarily configured (e.g., by software) or permanently configured to perform the relevant operations. Whether temporarily or permanently configured, such processors may constitute processor-implemented components that operate to perform one or more operations or functions described herein.

[0178] Changes and modifications may be made to the disclosed examples without departing from the scope of the present disclosure. These and other changes or modifications are intended to be included within the scope of the present disclosure, as expressed in the following claims.

What is claimed is:

1. A method comprising:
receiving a three-dimensional (3D) object comprising a target;
analyzing the 3D object using one or more machine learning models to generate data associated with a mirrored version of the target of the 3D object;
applying the mirrored version of the target to a mirrored version of the 3D object using the generated data; and
generating a new 3D object comprising the mirrored version of the 3D object and the mirrored version of the target.
2. The method of claim 1, wherein the 3D object and the mirrored version of the 3D object represent a pair of fashion items, and wherein the target and the mirrored version of the target comprise one or more graphical elements or text on the pair of fashion items.
3. The method of any one of claims 1-2, further comprising:
extracting the 3D object from an image; and
generating the mirrored version of the 3D object from the extracted 3D object, wherein the 3D object comprises a first virtual or real-world object and the mirrored version of the 3D object comprises a second virtual object, wherein the mirrored version of the 3D object is separately generated from the generated data associated with the mirrored version of the target.
4. The method of claim 3, further comprising:
identifying a first bounding box of the 3D object that includes the target;
determining a second bounding box in the mirrored version of the 3D object that corresponds to the first bounding box; and
replacing visual content in the second bounding box with the mirrored version of the 3D object to generate the mirrored version of the 3D object comprising the mirrored version of the target.
5. The method of any one of claims 1-4, further comprising:
applying a first machine learning model to the 3D object to identify a bounding box comprising the target; and
applying a second machine learning model to generate a segmentation of the target in the bounding box, the segmentation distinguishing pixels in the 3D object of the bounding box of the target from other pixels in the bounding box.

6. The method of claim 5, further comprising:
applying a flip transformation to the segmentation of the target in the bounding box, the flip transformation flipping the segmentation of the target around a flipping axis of the bounding box thereby to generate a flipped segmentation of the target.
7. The method of any one of claims 1-6, further comprising:
blending pixels of the flipped segmentation of the target to a specified portion of the mirrored version of the 3D object.
8. The method of claim 7, wherein the pixels of the flipped segmentation are blended to the specified portion while preserving original texture features of the mirrored version of the 3D object.
9. The method of any one of claims 1-8, wherein the pixels are blended according to at least one of Poisson blending techniques or a blending map.
10. The method of any one of claims 1-9, further comprising:
after blending the pixels, identifying an area in the mirrored version of the 3D object with unknown pixel values; and
applying inpainting techniques to populate the unknown pixel values with predicted values.
11. The method of any one of claims 1-10, further comprising:
analyzing the 3D object comprising the target and the mirrored version of the 3D object comprising the mirrored version of the target with a quality control machine learning model to generate a score representing match accuracy between the mirrored version of the object comprising the mirrored version of the target and the 3D object comprising the target.
12. The method of claim 11, wherein the quality control machine learning model comprises a discriminator trained by performing training operations comprising:
accessing training data comprising a training pair of fashion items including a training fashion item and a corresponding mirrored version of the training fashion item;
analyzing the training pair of fashion items using the discriminator to generate a prediction indicative of whether the mirrored version of the training fashion item is a real

mirrored version of the training fashion item or a fake mirrored version of the training fashion item;

accessing ground truth labels for the training pair of fashion items indicating whether the mirrored version of the training fashion item is real or fake; and

updating one or more parameters of the discriminator based on accuracy of the prediction in relation to the ground truth labels.

13. The method of any one of claims 1-12, further comprising:

accessing a plurality of 3D objects representing pairs of fashion items;

selecting a subset of the plurality of 3D objects;

adding noise to an individual fashion item of each pair of fashion items in the subset of the plurality of 3D objects;

associating a first set of ground truth labels for each of the subset of the plurality of 3D objects indicating that the mirrored version of the fashion item in each of the subset of the plurality of 3D objects is fake; and

associating a second set of ground truth labels for each of the plurality of 3D objects outside of the subset of the plurality of 3D objects indicating that the mirrored version of the fashion item in each of the plurality of 3D objects outside of the subset of the plurality of 3D objects is real.

14. The method of any one of claims 1-13, further comprising:

generating an extended reality (XR) experience using the 3D object and the new 3D object.

15. The method of any one of claims 1-14, further comprising:

generating an artificial image depicting the mirrored version of the 3D object, the artificial image being modified using at least a portion of the new 3D object.

16. The method of claim 15, wherein the one or more machine learning models comprise a convolutional neural network trained by performing training operations comprising:

accessing training data comprising a training pair of fashion items including a training fashion item and a ground truth mirrored version of the training fashion item;

analyzing the training fashion item to estimate a mirrored version of the training fashion item;

computing a deviation between the estimated mirrored version of the training fashion item and the ground truth mirrored version of the training fashion item; and
updating one or more parameters of the convolutional neural network based on the computed deviation.

17. The method of any one of claims 1-16, further comprising repeating the training operations for additional training data until a stopping criterion is met.

18. The method of any one of claims 1-17, wherein the convolutional neural network is trained to establish a relationship between mapping of pixels in a first fashion item in a first pair of fashion items to pixels in a second fashion item in the first pair of fashion items.

19. A system comprising:
at least one processor; and
at least one memory component having instructions stored thereon that, when executed by the at least one processor, cause the at least one processor to perform operations comprising:
receiving a three-dimensional (3D) object comprising a target;
analyzing the 3D object using one or more machine learning models to generate data associated with a mirrored version of the target of the 3D object;
applying the mirrored version of the target to a mirrored version of the 3D object using the generated data; and
generating a new 3D object comprising the mirrored version of the 3D object and the mirrored version of the target.

20. A non-transitory computer-readable storage medium having stored thereon instructions that, when executed by at least one processor, cause the at least one processor to perform operations comprising:
receiving a three-dimensional (3D) object comprising a target;
analyzing the 3D object using one or more machine learning models to generate data associated with a mirrored version of the target of the 3D object;
applying the mirrored version of the target to a mirrored version of the 3D object using the generated data; and
generating a new 3D object comprising the mirrored version of the 3D object and the mirrored version of the target.




FIG. 1




FIG. 2




FIG. 3




FIG. 4

500 →




FIG. 5




FIG. 6




FIG. 7




FIG. 8





FIG. 9



INTERNATIONAL SEARCH REPORT

International application No
PCT/US2024/031125

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06T19/20
ADD.

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

B. FIELDS SEARCHED

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

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

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

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Anonymous: "How do I flip this texture? I've tried to mirror the UVs but it just makes it even more complicated. : blenderhelp", , 23 September 2022 (2022-09-23), XP093188241, Retrieved from the Internet: URL:https://web.archive.org/web/2022092305 0436/https://www.reddit.com/r/blenderhelp/ comments/xhhyd3/how_do_i_flip_this_texture _ive_tried_to_mirror/ the whole document - - - - - - / - -	1 - 20

☒ Further documents are listed in the continuation of Box C.

☐ See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

23 July 2024

Date of mailing of the international search report

01/08/2024

Name and mailing address of the ISA/
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Gao, Miao

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2024/031125

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	Learningbydoodling: "How to map a texture on a mirrored object - in Blender", , 12 April 2016 (2016-04-12), XP093188259, Retrieved from the Internet: URL:https://www.youtube.com/watch?v=BG9O_PDAVV0 the whole document -----	1-20
X	Anonymous: "Recognition, Object Detection, and Semantic Segmentation - MATLAB & Simulink - MathWorks Deutschland", , 29 July 2021 (2021-07-29), XP093188263, Retrieved from the Internet: URL:https://web.archive.org/web/20210729000209/https://de.mathworks.com/help/vision/recognition-object-detection-and-semantic-segmentation.html the whole document -----	1-20