



(19) **United States**

(12) **Patent Application Publication**
Holzer et al.

(10) **Pub. No.: US 2021/0227195 A1**

(43) **Pub. Date: Jul. 22, 2021**

(54) **CREATING CINEMATIC VIDEO FROM MULTI-VIEW CAPTURE DATA**

Publication Classification

(71) Applicant: **Fyusion, Inc.**, San Francisco, CA (US)

(72) Inventors: **Stefan Johannes Josef Holzer**, San Mateo, CA (US); **Rodrigo Ortiz Cayon**, San Francisco, CA (US); **Julius Santiago**, San Francisco, CA (US); **Milos Vlaski**, San Francisco, CA (US)

(51) **Int. Cl.**
H04N 13/282 (2006.01)
H04N 13/189 (2006.01)
H04N 5/232 (2006.01)
G11B 27/00 (2006.01)
(52) **U.S. Cl.**
CPC *H04N 13/282* (2018.05); *G11B 27/005* (2013.01); *H04N 5/23222* (2013.01); *H04N 13/189* (2018.05)

(73) Assignee: **Fyusion, Inc.**, San Francisco, CA (US)

(21) Appl. No.: **17/147,377**

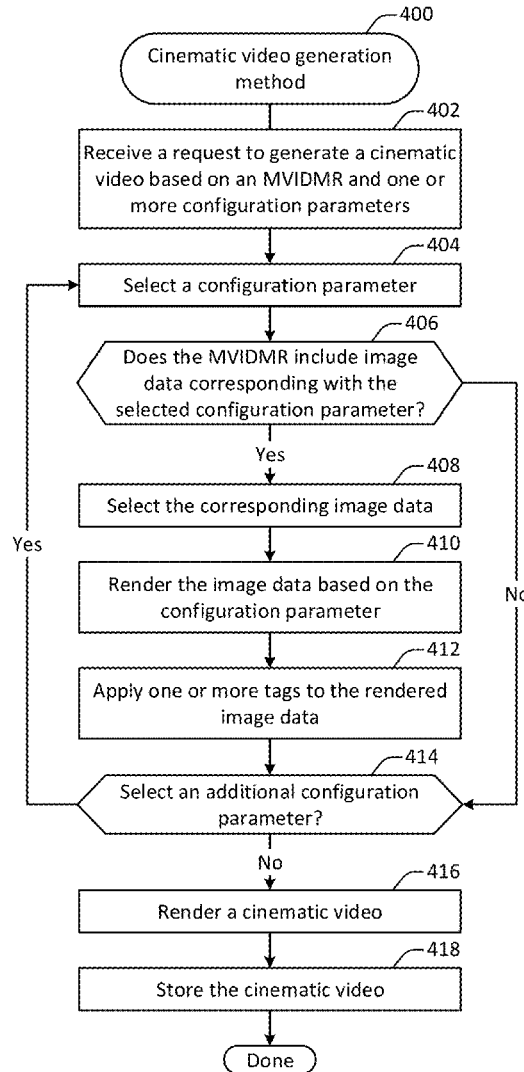
(22) Filed: **Jan. 12, 2021**

(57) **ABSTRACT**

Configuration parameters associated with the generation of a cinematic video may identify object components, an order in which to display the object components, and an object type. A multi-view representation that is navigable in one or more dimensions and that includes images of an object captured from different viewpoints may be identified. A cinematic video of the object may be generated based on a subset of the images, arranged in an order based on the configuration parameters.

Related U.S. Application Data

(60) Provisional application No. 62/961,813, filed on Jan. 16, 2020.



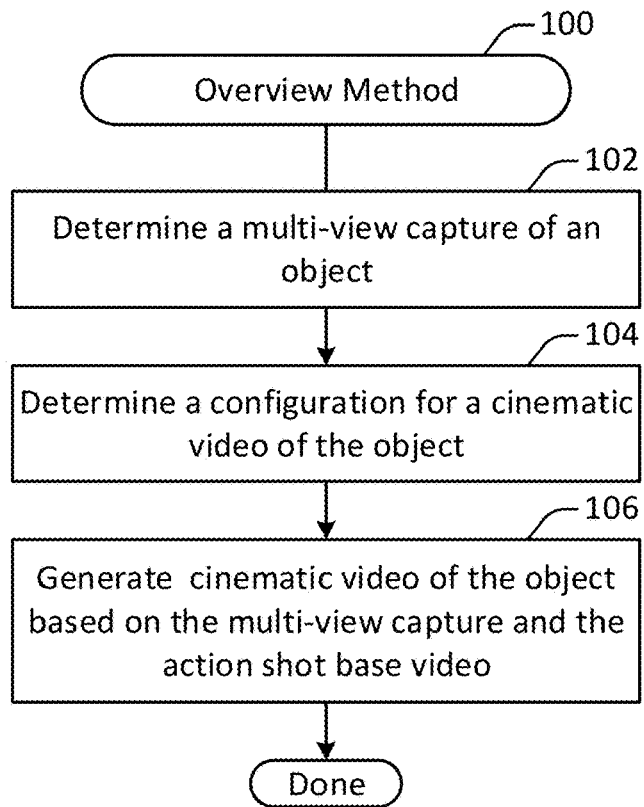


Figure 1

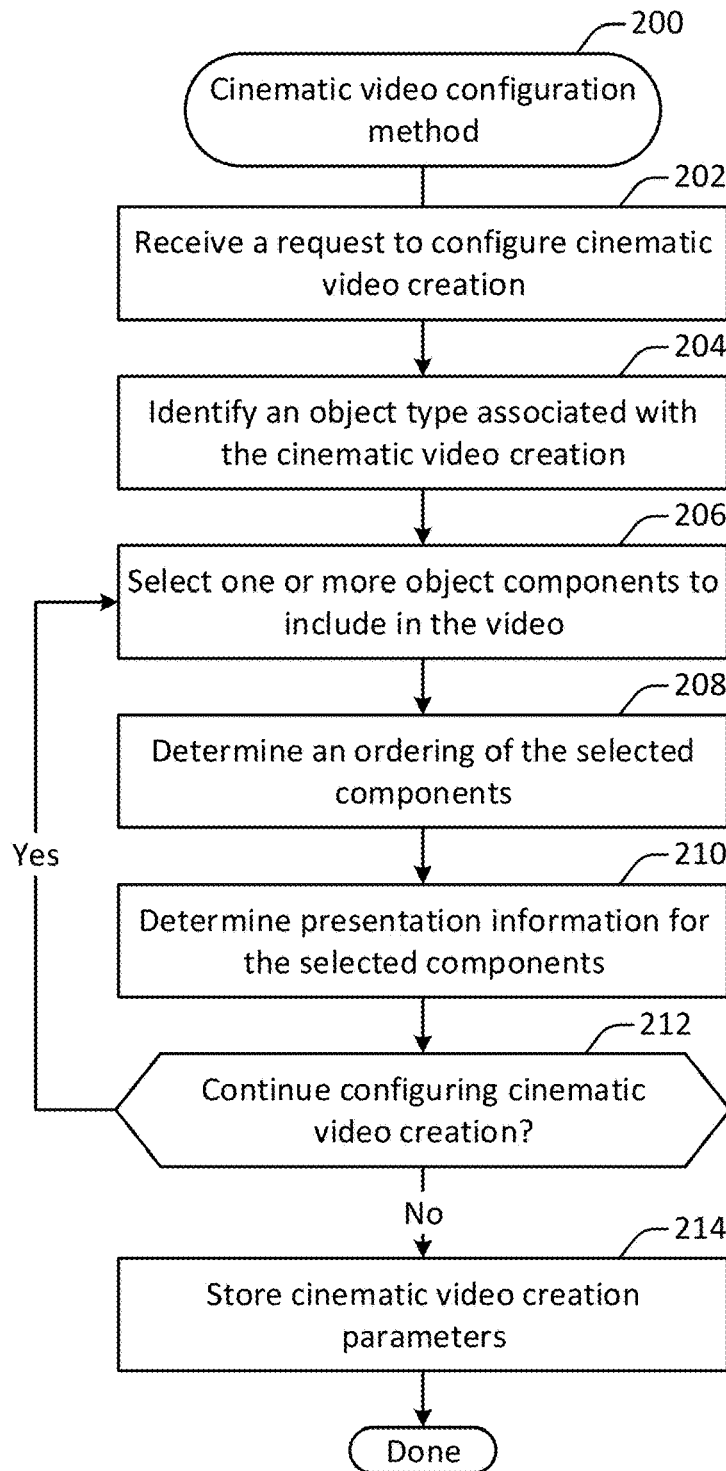


Figure 2

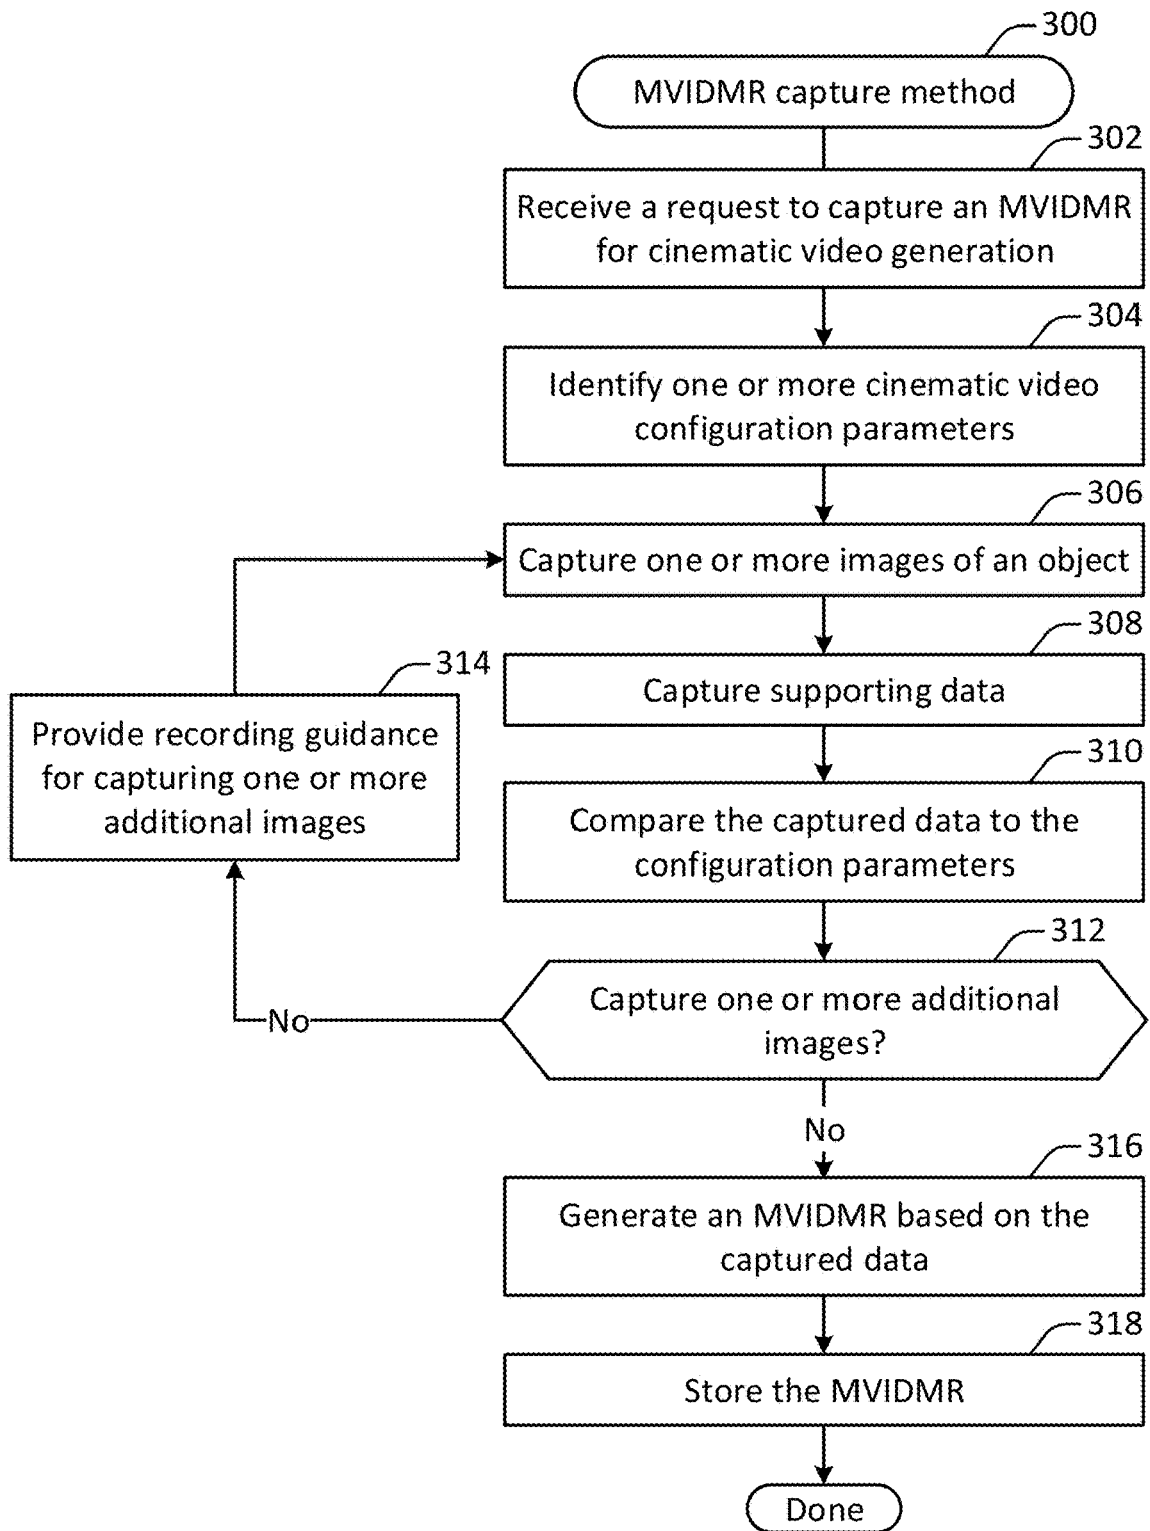


Figure 3

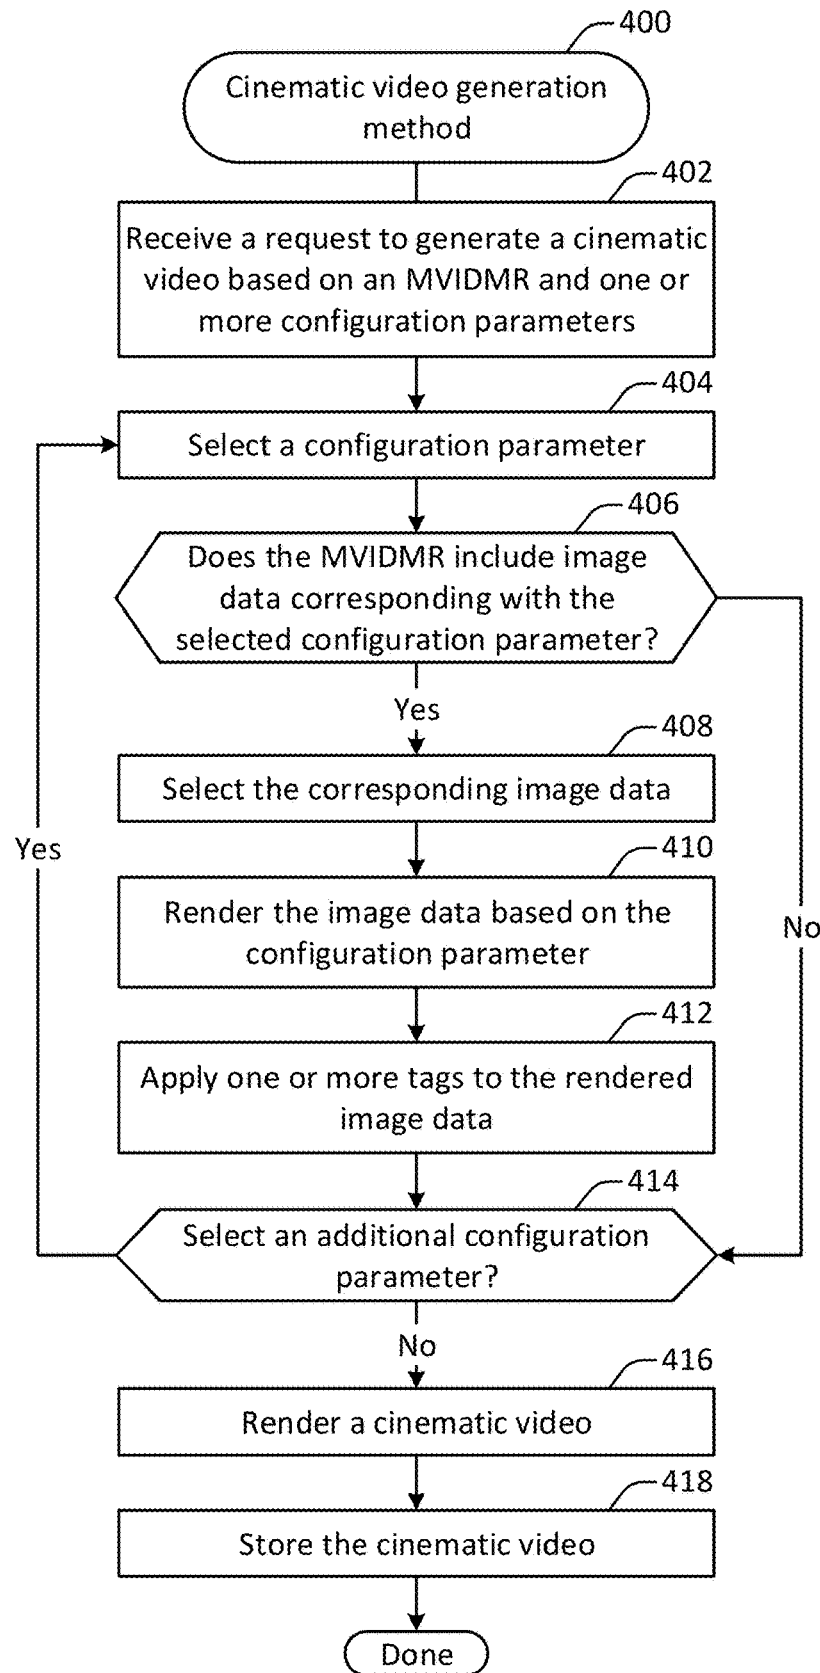


Figure 4

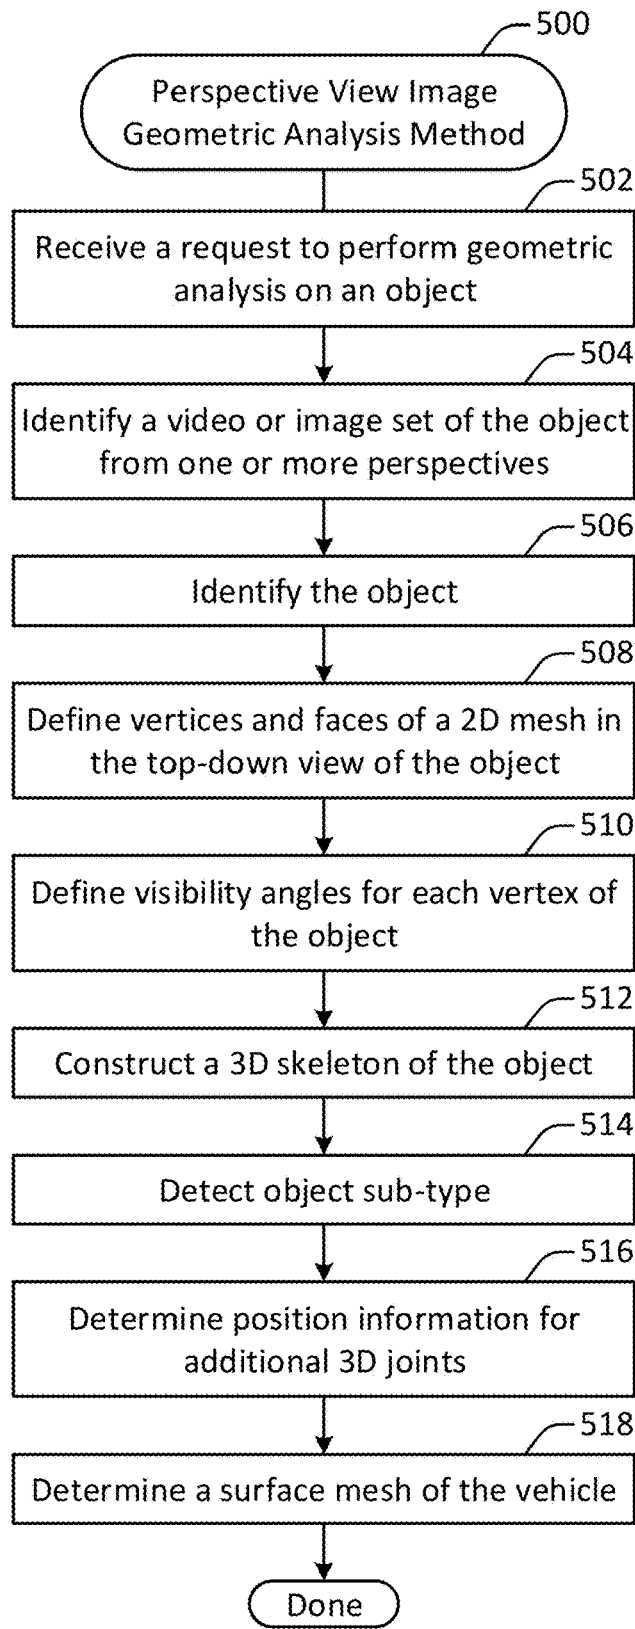


Figure 5

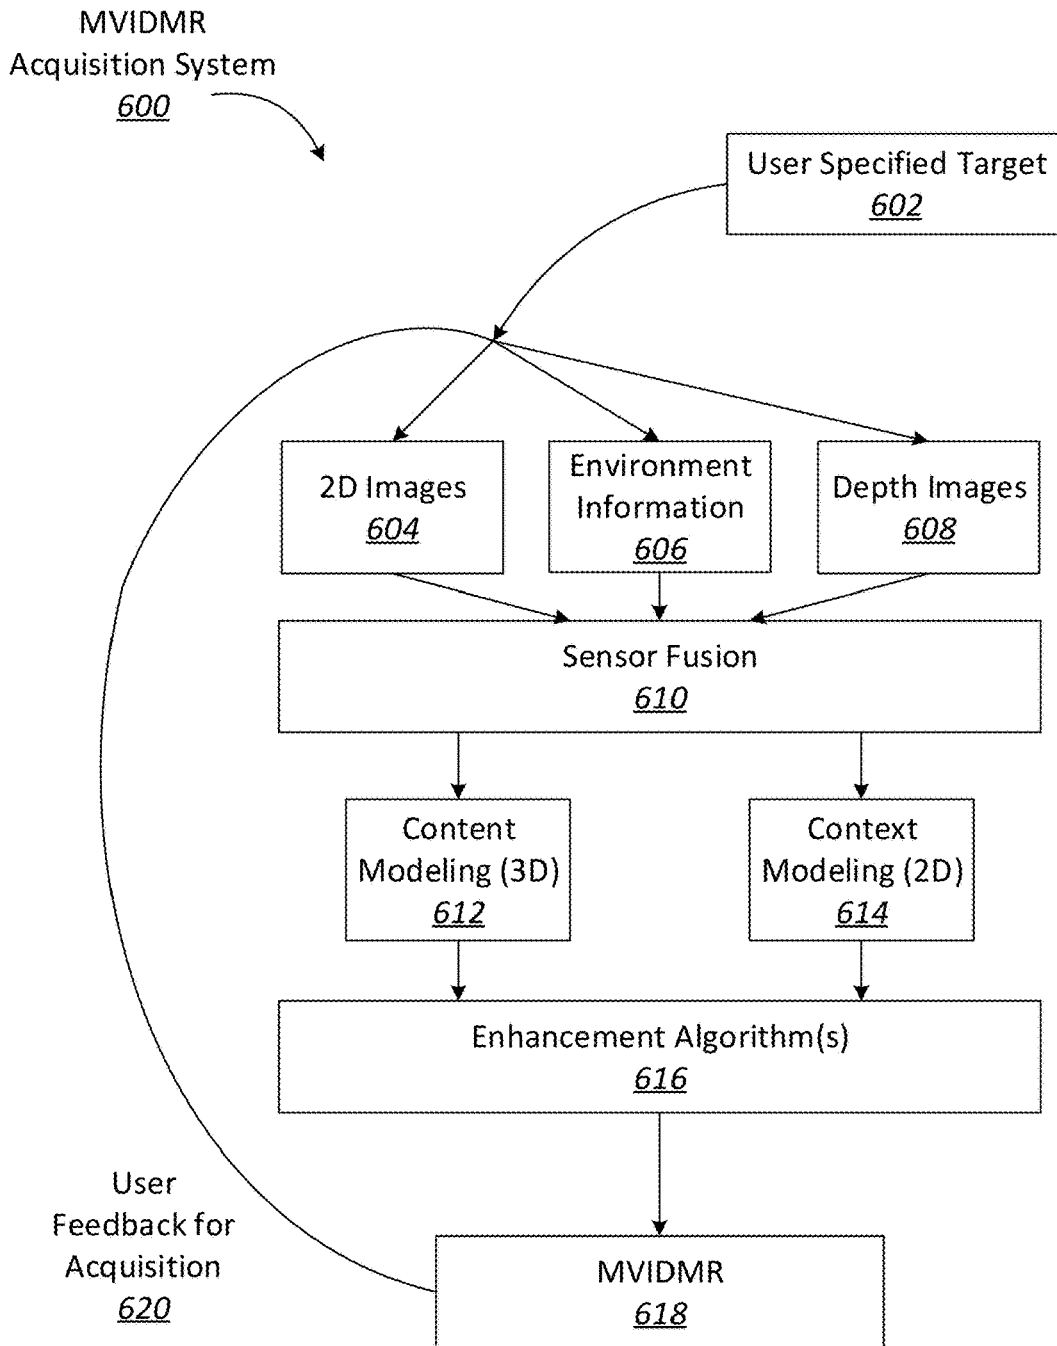


Figure 6

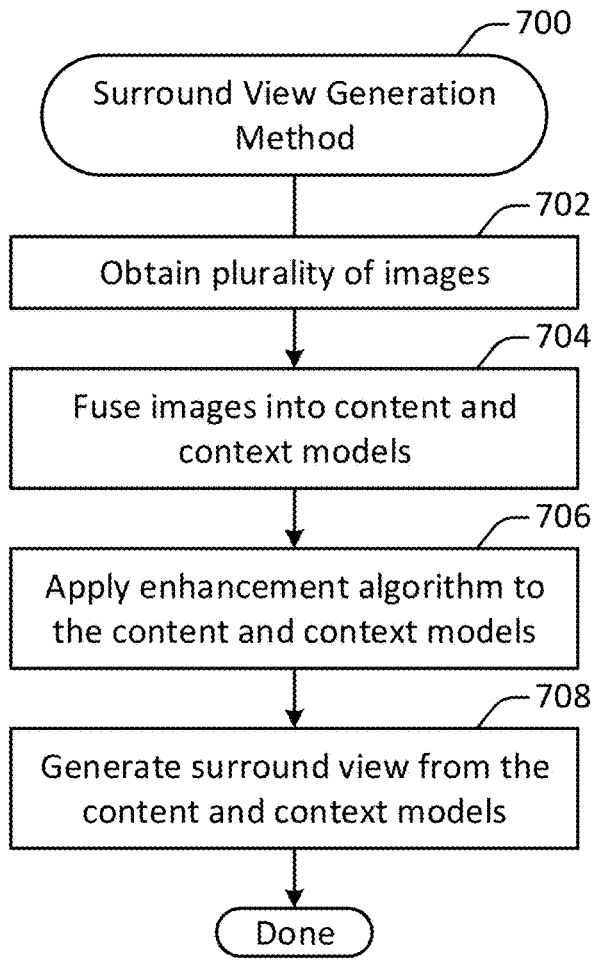


Figure 7




Figure 8




Figure 9




Figure 10A




Figure 10B




Figure 11A




Figure 11B




Figure 12A




Figure 12B

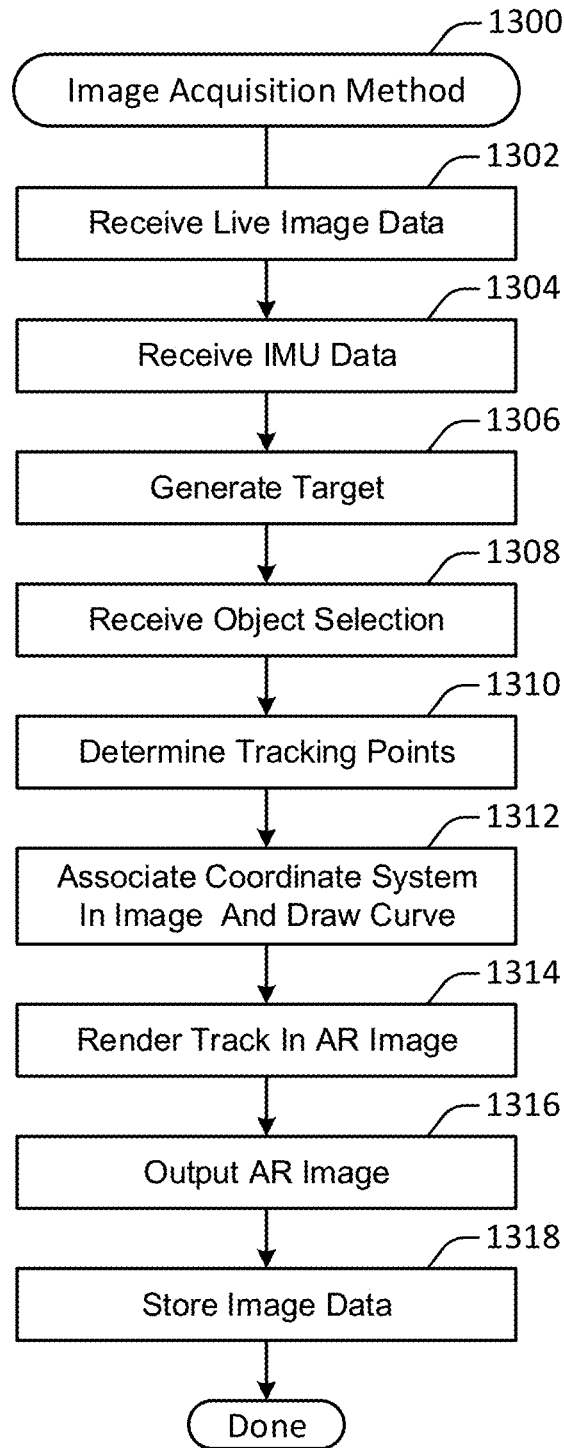


Figure 13




Figure 14




Figure 15A




Figure 15B




Figure 16




Figure 17A




Figure 17B

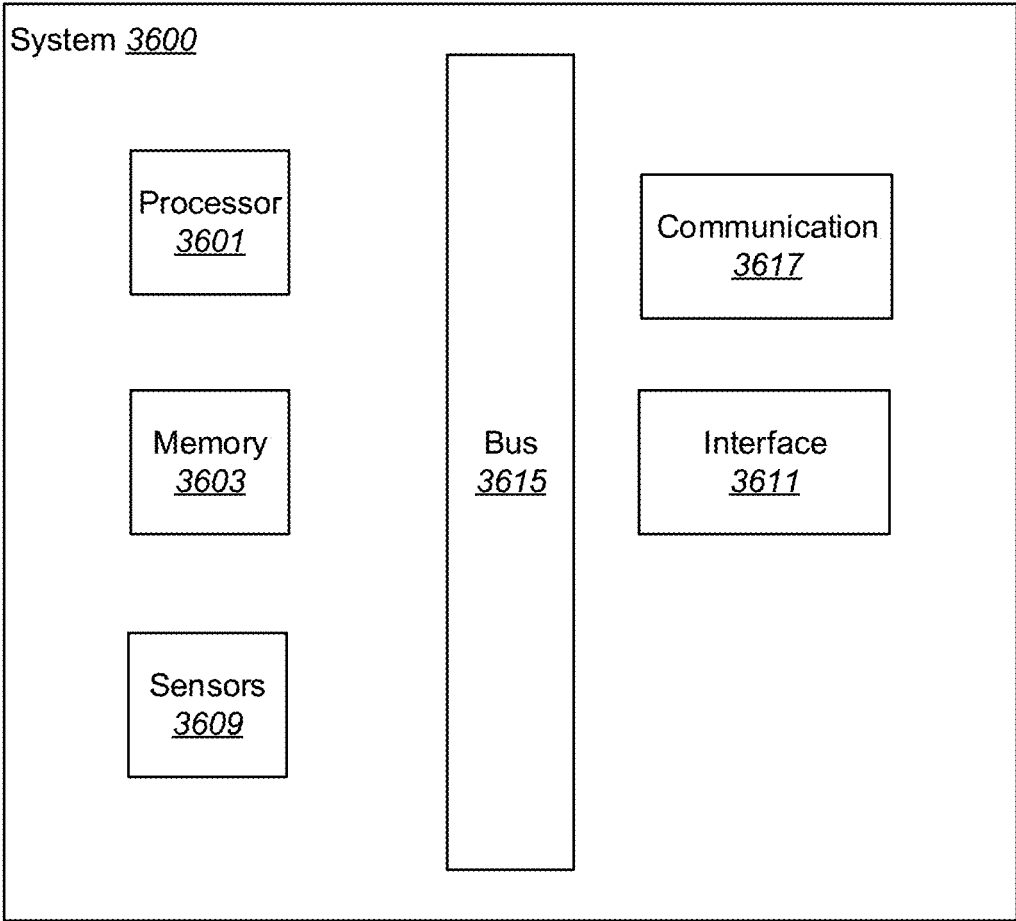


Figure 18





Figure 19

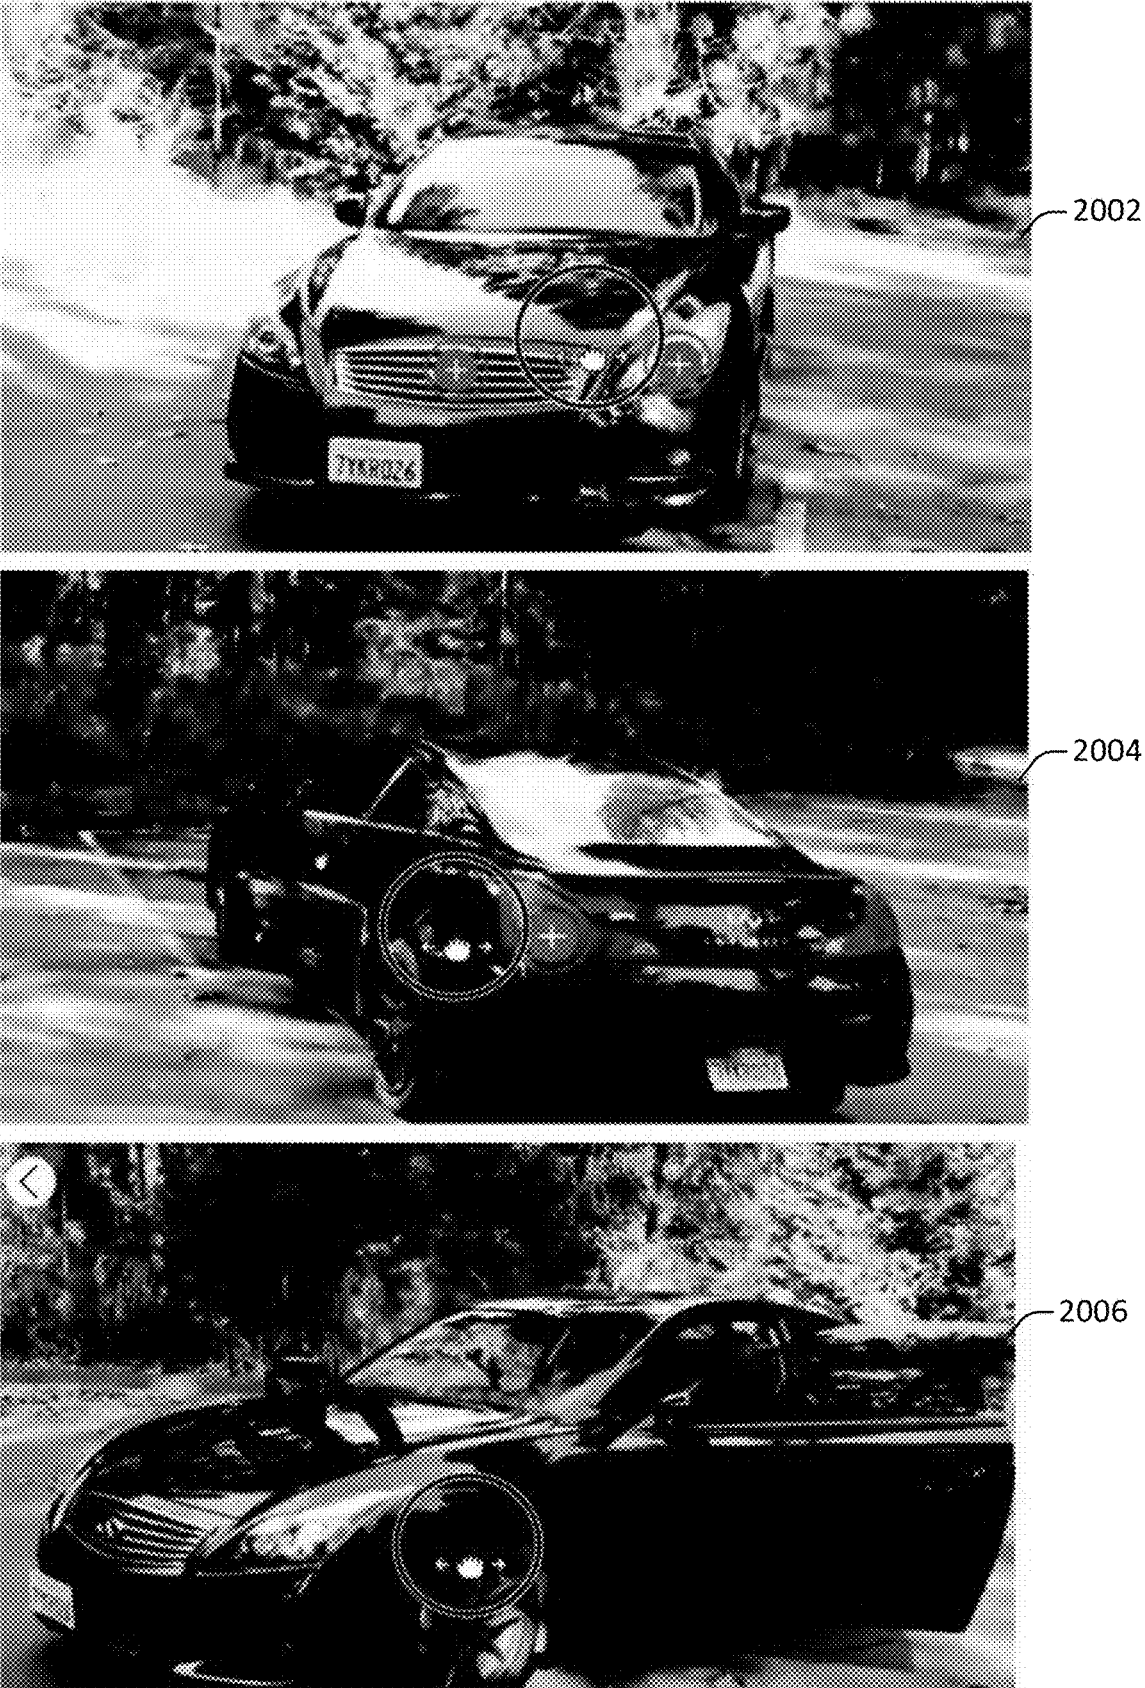


Figure 20




Figure 21



Figure 22







Figure 24







Figure 25

CREATING CINEMATIC VIDEO FROM MULTI-VIEW CAPTURE DATA

PRIORITY CLAIM

[0001] The present application claims priority under 35 U.S.C. 120 to U.S. Patent App. No. 62/961,813 (Atty Docket No. FYSNP065P), titled "Creating Cinematic Video from Multi-view Capture Data", filed Jan. 16, 2020 by Holzer, which is hereby incorporated by reference in its entirety and for all purposes.

COPYRIGHT NOTICE

[0002] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the United States Patent and Trademark Office patent file or records but otherwise reserves all copyright rights whatsoever.

TECHNICAL FIELD

[0003] The present disclosure relates generally to the processing of visual digital media content, and more specifically to the creation of video from image data.

DESCRIPTION OF RELATED ART

[0004] Visual data encompasses a wide variety of information, such as individual images, video, and multi-view captures. Images and multi-view captures are useful and informative tools for presenting information about objects, for instance in an online platform. However, they lack to the ability to create the same excitement and sense of presence as a video. Nevertheless, videos for many objects do not exist. Further, even when video of an object is available, the video may present the object in the wrong context.

[0005] Creating videos can be difficult and expensive. For instance, an object such as a vehicle must first be positioned in a suitable (e.g., attractive) context. A video of the object must then be captured in the context. Finally, the video must often be post-processed, such as with one or more editing and/or compression techniques, to render the video suitable for presentation.

Overview

[0006] Provided are various mechanisms and processes relating to the processing of visual media data. According to various embodiments, techniques and mechanisms described herein may facilitate cinematic video creation from multi-view capture data.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The patent or application file contains at least one drawing executed in color. Copies of this patent or patent application publication with color drawing(s) will be provided by the Office upon request and payment of the necessary fee.

[0008] The included drawings are for illustrative purposes and serve only to provide examples of possible structures and operations for the disclosed inventive systems, apparatus, methods and computer program products for processing visual data. These drawings in no way limit any changes in

form and detail that may be made by one skilled in the art without departing from the spirit and scope of the disclosed implementations.

[0009] FIG. 1 illustrates one example of an overview method, performed in accordance with one or more embodiments.

[0010] FIG. 2 illustrates an example of a cinematic video configuration method, performed in accordance with one or more embodiments.

[0011] FIG. 3 illustrates a method for capturing a multi-view image digital media representation (MVIDMR), performed in accordance with one or more embodiments.

[0012] FIG. 4 illustrates an example of a cinematic video generation method, performed in accordance with one or more embodiments.

[0013] FIG. 5 illustrates one example of a method for performing geometric analysis of a perspective view image, performed in accordance with one or more embodiments.

[0014] FIG. 6 illustrates an example of a multi-view image digital media representation (MVIDMR) acquisition system, configured in accordance with one or more embodiments.

[0015] FIG. 7 illustrates one example of a method for generating a MVIDMR, performed in accordance with one or more embodiments.

[0016] FIG. 8 illustrates one example of multiple camera views fused together into a three-dimensional (3D) model.

[0017] FIG. 9 illustrates one example of separation of content and context in a MVIDMR.

[0018] FIGS. 10A-10B illustrate examples of concave and convex views, where both views use a back-camera capture style.

[0019] FIGS. 11A-11B illustrates one example of a back-facing, concave MVIDMR, generated in accordance with one or more embodiments.

[0020] FIGS. 12A-12B illustrate examples of front-facing, concave and convex MVIDMRs generated in accordance with one or more embodiments.

[0021] FIG. 13 illustrates one example of a method for generating virtual data associated with a target using live image data, performed in accordance with one or more embodiments.

[0022] FIG. 14 illustrates one example of a method for generating MVIDMRs, performed in accordance with one or more embodiments.

[0023] FIGS. 15A and 15B illustrate some aspects of generating an Augmented Reality (AR) image capture track for capturing images used in a MVIDMR.

[0024] FIG. 16 illustrates one example of generating an Augmented Reality (AR) image capture track for capturing images used in a MVIDMR on a mobile device.

[0025] FIGS. 17A and 17B illustrate examples of generating an Augmented Reality (AR) image capture track including status indicators for capturing images used in a MVIDMR.

[0026] FIG. 18 illustrates a particular example of a computer system configured in accordance with various embodiments.

[0027] FIGS. 19, 20, 21, 22, 23, 24, and 25 illustrate examples of visual representations constructed or analyzed in accordance with various embodiments.

DETAILED DESCRIPTION

[0028] Techniques and mechanisms described herein are directed to creating cinematic videos from multi-view data

sets. According to various embodiments, one or more configuration parameters may specify operations for creating a cinematic video. Based on those parameters, sequences of images may be combined in an order designed to present a narrative stream.

[0029] For purposes of illustration, consider the example of commercials for vehicles. Dealerships often employ static images of vehicles to show prospective buyers the condition of a vehicle. However, according to various embodiments of techniques and mechanisms described herein, a vehicle dealership may be able to create a cinematic video of a vehicle without capturing new footage of the vehicle. Instead, existing media such as individual images may be used as input. The result may be a robust visualization that may appear as if a custom video was created for the specific vehicle. For instance, the visualization may include slow motion dramatic pans of key features that are automatically created based on input images. Such pans may be edited together into a single video that can be presented on a web page, creating an enhanced experience for users in the market for a new vehicle.

[0030] According to various embodiments, one or more examples are described herein with reference to vehicles and automotive contexts. However, techniques and mechanisms described herein are generally applicable to a wide range of contexts. For instance, cinematic videos may be created in the fashion context, where a video of an outfit is generated. Such a video may, for example, show a panning around of a person, and then close-up pans around specific fashion elements. Alternately, or additionally, such a video may involve pans around objects that are part of the outfit, separated from the rest of the scene.

[0031] In some implementations, techniques and mechanisms described herein may facilitate the generation of cinematic video from individual images with minimal effort from a user. One or more configuration parameters may guide the video creation process. The video may then be automatically created based on the configuration parameters and one or more images included in a multi-view capture of an object.

[0032] FIG. 1 illustrates one example of an overview method 100, performed in accordance with one or more embodiments. According to various embodiments, the method 100 may be performed on any suitable computing device. For example, the method 100 may be performed on a mobile computing device such as a mobile phone or laptop computer. As another example, the method 100 may be performed on a computing device such as a server.

[0033] In some implementations, the method 100 may be triggered in various ways. For instance, the user may click a playback button in a user interface on a mobile computing device such as a cell phone. Alternately, cinematic video generation may be automatically initiated, for instance if the user is idling or not interacting with the player and/or device for a designated period of time. In either case, the user may be able to take over control from the cinematic video once it is being created and/or playing. When the user regains control, the user may, for instance, explore the currently displayed visual data from a different viewpoint.

[0034] A multi-view capture of an object is determined at 102. According to various embodiments, a multi-view capture may include various multi-view capture data. Multi-view capture data may include two or more perspective view images of the object captured from different perspectives

relative to the object. Multi-view capture data may also include inertial measurement unit (IMU) data captured from an IMU, for instance within a mobile computing device.

[0035] In some implementations, a multi-view capture may be generated during the execution of the method 100. Alternately, a multi-view capture may be pre-generated and retrieved at 102. For instance, at 102 a user may provide user input to select a multi-view capture for cinematic video generation. Additional details regarding the creation of multi-view captures are discussed throughout the application, for example with respect to FIGS. 5-18.

[0036] A configuration for a cinematic video is determined at 104. According to various embodiments, a cinematic video configuration may involve one or more parameters identifying portions of an object to include in the video and/or orders in which those object portions are to be presented. For example, the cinematic video configuration may indicate that a camera is to pan around a vehicle in a 360-degree fashion, slowing down and zooming in on specific features. Additional details regarding the determination of a configuration for a cinematic video are discussed with respect to the method 200 shown in FIG. 2.

[0037] According to various embodiments, the configuration for the cinematic video may be generated during the execution of the method 100. Alternately, the configuration may be pre-generated and retrieved at 104.

[0038] A cinematic video of the object is generated at 106 based on the multi-view capture and the cinematic video configuration. According to various embodiments, various techniques may be used to generate the cinematic video. For instance, one or more images from a multi-view capture may be selected and arranged in a designated order.

[0039] FIG. 20 illustrates a method 200 for cinematic video configuration. According to various embodiments, the method 200 may be performed at any suitable computing device or combination of computing devices. For example, the method 200 may be performed at a mobile computing device such as a smart phone. As another example, the method 200 may be performed at a server. As still another example, the method 200 may be performed at a mobile computing device and a server in communication with one another.

[0040] A request to configure cinematic video creation is received at 202. According to various embodiments, the request may be received based on user input. For example, a user may provide user input via a user interface indicating a request to configure cinematic video creation.

[0041] An object type associated with the cinematic video creation is identified at 204. According to various embodiments, the object type may be identified via user input. For example, the user may select an object type from a list. As another example, the user may identify a 3D model or a representative image of the object type.

[0042] In some implementations, the object type may be generic. For example, the object type may be a vehicle or a clothing item. Alternately, or additionally, one or more specific features or object subtypes may be identified. For example, the object may be a truck, car, or sport utility vehicle. As another example, the object may be a dress, suit, or specific jewelry item.

[0043] One or more object components to include in the video are identified at 206. According to various embodiments, the one or more object components may be identified in any suitable fashion. For example, the object components

may be identified by receiving user input selecting them from a list or other user interface affordance. As another example, the object components may be identified by circling or otherwise identifying them in an image or images. As still another example, the object components may be identified by selecting them from a 3D model that includes various components.

[0044] In some implementations, one or more components may be suggested for inclusion in the video. For example, the system may include a default or base video configuration, for instance for a specific object type. The user may then add additional components to or remove components from those included in the default video. As another example, the system may include a default or base list of components that are specific to the object type. For instance, a default configuration for a vehicle may include a panoramic view around the entire vehicle, with specific components such as the vehicle headlights and door handles selected for inclusion.

[0045] An ordering of the selected components is determined at **208**. According to various embodiments, the ordering may be determined based on user input. For example, user input specifying an order in which the object components are to be presented, for instance by arranging them in a list, may be received.

[0046] In some embodiments, an ordering may be suggested based on an object type. For instance, in the context of a vehicle, a default ordering may involve a 360-degree pan around the vehicle. Different components of the vehicle may then be suggested for emphasis in a clockwise or counterclockwise ordering as they appear around the video.

[0047] Presentation information for the selected components is determined at **210**. According to various embodiments, the presentation information may include any information specifying an approach for including a component in the video. For example, the presentation information may specify a zoom level for the component. As another example, the presentation information may specify that the component is to be subjected to a filter. As still another example, the presentation information may specify a speed for panning over a component. As yet another example, the presentation information may specify a zoom-and-pan view in which the virtual camera zooms in on a single image and pans over it.

[0048] In particular embodiments, presentation information may be suggested based on an object type. For instance, in a context of a vehicle, default presentation information may involve panning around a vehicle in a 360-degree view, followed up by slow motion closeups of distinctive features such as headlights, and ending with a side view of the entire vehicle.

[0049] A determination is made at **212** as to whether to continue configuring cinematic video creation. According to various embodiments, the determination may be made at least in part based on user input. For example, the user may provide the information in a user interface and may indicate when configuration is complete.

[0050] In some embodiments, the determination made at **212** may be made at least in part based on an evaluation of the completeness of the configuration information. For example, a designated number and arrangement of object components may be necessary to create a complete video configuration.

[0051] One or more cinematic video creation parameters are stored at **214**. According to various embodiments, the one or more cinematic video creation parameters may include any of the information identified as part of the method **200**. Alternately, or additionally, any other configuration information for creating a cinematic video may be identified and stored.

[0052] FIG. 3 illustrates a method **300** for capturing a multi-view image digital media representation (MVIDMR), performed in accordance with one or more implementations. According to various embodiments, the method **300** may be performed at a mobile computing device such as a smartphone. The mobile computing device may include one or more cameras configured to capture image data.

[0053] In some implementations, a cinematic video may be created based on an existing MVIDMR. In such a configuration, the method **300** need not be performed. Instead, an existing MVIDMR may be identified and retrieved. However, the method **300** may be used to generate an MVIDMR that includes information particularly suitable for cinematic video generation.

[0054] A request to capture an MVIDMR for cinematic video generation is received at **302**. In some implementations, the request may be received at a user interface. For instance, a user may provide a request to generate an MVIDMR via a touch screen display. Alternately, or additionally, the system may automatically initiate a request to generate an MVIDMR when a request to generate a cinematic video is received and no suitable MVIDMR is available for video generation.

[0055] One or more cinematic video configuration parameters are identified at **304**. In some implementations, the one or more cinematic video parameters may be determined as discussed with respect to the method **200** shown in FIG. 2. The appropriate cinematic video configuration parameters may be identified at least in part based on user input. For example, the user may provide a request to create a particular cinematic video and identify a configuration for that video.

[0056] In some embodiments, the cinematic configuration parameters may be identified at least in part based on a determination made by the system. For instance, the system may analyze one or more preliminary image of an object and then identify the object from the images. Based on the identified object, the system may select configuration parameters suitable for the object. For example, if the object is a pickup truck, the system may identify configuration parameters suitable for generating a cinematic video from an MVIDMR of a pickup truck.

[0057] One or more images of an object are captured at **306**. According to various embodiments, the one or more images may be captured via a camera at the mobile computing device. For instance, a user may hold a cell phone and move it along a path through space around an object to be captured.

[0058] In some implementations, the one or more images of the object may be captured from multiple cameras at the same time. For example, a mobile computing device such as a smart phone may include two or more cameras such as a wide-angle camera, a normal lens camera, and/or a telephoto lens camera.

[0059] Supporting data is captured at **308**. In some implementations, the supporting data may include any data captured from one or more other sensors at the mobile com-

puting device. For instance, inertial measurement unit (IMU) data may be captured from an IMU at the mobile computing device.

[0060] According to various embodiments, the supporting data may be linked with corresponding image data. For instance, IMU data may be captured at the same time one or more images are captured, and then linked with that image data to facilitate subsequent analysis of the image.

[0061] The captured data is compared to the configuration parameters at **310**. In some embodiments, the comparison may be used to determine whether the captured data is sufficient for generating a cinematic video according to the one or more parameters identified at **304**. For instance, if the one or more cinematic video configuration parameters identified at **304** indicate that the object is a vehicle and that the cinematic video is to include a closeup pan around the headlight of the vehicle, then the comparison at **310** may seek to determine whether the captured data includes sufficient image data to generate such a video.

[0062] A determination is made at **312** as to whether to capture one or more additional images. In some implementations, the determination may be made at least in part on the comparison performed at **310**. For instance, if the captured data is insufficient for generating a cinematic video as specified by the configuration parameters identified at **304**, then one or more additional images may be captured.

[0063] In some implementations, multi-view capture data of a vehicle for example may include data such as one or more 360-degree captures of the vehicle. Such a capture may include a walkaround or turntable capture. For instance, one 360-degree capture may be completed with the doors closed, while another may be completed with the doors open. Alternately, or additionally, close-up captures of important details may be captured. Interior image data may be captured as one or more multi-view captures. Alternately, or additionally, image data may be captured via a 360-degree image such as an image captured from a Ricoh Theta available from Ricoh Company of Chūō, Tokyo. Internal data may include one or more close-up captures of the vehicle's interior.

[0064] When it is determined to capture one or more additional images, recording guidance for capturing one or more additional images is provided at **314**. In some implementations, the recording guidance may be provided as an overlay on a live camera feed. For instance, a user may be provided with a target, a guide box, or other visual affordance to assist in positioning the camera. The guidance may be selected so as to assist the user in positioning the camera to capture one or more images to expand the captured data to include information sufficient for generating the cinematic video.

[0065] When it is determined to not capture one or more additional images, an MVIDMR is generated based on the captured data at **316**. Techniques for generating MVIDMRs are discussed throughout the application, such as with respect to the FIGS. **5-18**.

[0066] The generated MVIDMR is stored at **318**. In some implementations, storing the MVIDMR may involve sending the MVIDMR to a storage medium such as a hard drive or network attached storage (NAS) system, transmitting the MVIDMR to a remote computing device, or storing the MVIDMR in temporary memory such as random-access memory (RAM).

[0067] FIG. **4** illustrates an example of a cinematic video generation method **400**, performed in accordance with one or more embodiments. According to various embodiments, the method **400** may be performed on a computing device such as a local machine such as a mobile computing device. Alternately, or additionally, the method **400** may be performed on a remote server, or a local machine in communication with a remote server.

[0068] A request to generate a cinematic video based on an MVIDMR and one or more configuration parameters is received at **402**. According to various embodiments, the request may be received as part of the method **100** shown in FIG. **1**. In some instances, the request may be generated automatically. For example, a system may automatically create a cinematic video of an object when an MVIDMR of the object is generated and a suitable set of configuration parameters is available. Alternately, or additionally, a request may be generated based on user input. For instance, a user may provide user input indicating a request to generate a cinematic video.

[0069] A configuration parameter is selected at **404**. According to various embodiments, configuration parameters may be selected in any suitable order. For instance, as discussed with respect to the method **200** shown in FIG. **2**, configuration parameters may specify portions of an object to include in a video, as well as an order in which to display such portions. Accordingly, configuration parameters may be selected in order of their appearance in the video.

[0070] A determination is made at **406** as to whether the MVIDMR includes image data corresponding with the selected configuration parameter. In some implementations, the determination may be made at **406** based on an analysis of the MVIDMR data. For instance, object component identification, pose determination, coverage analysis, and/or other such processes may be performed on one or more portions of the MVIDMR data.

[0071] According to various embodiments, the camera poses may be extracted using a computer vision algorithm. Additional details regarding pose detection, which is also referred to herein as object orientation detection, are discussed with respect to FIGS. **19** and **20**. As used herein, the term camera pose may refer to the viewpoint or viewing angle of the camera relative to the object captured in the image.

[0072] According to various embodiments, camera pose information may be estimated for a viewpoint in order to associate each image in a sequence of images with a viewpoint. Camera-poses may be estimated by an algorithm that uses the images and, optionally, additional information received from the capturing device. Such information may include angle, location, acceleration, velocity, drift correction, or other such data collected or determined based on information from a gyroscope, an accelerometer, magnetometer, or other such sensor associated with the camera. Such input data may be used to facilitate the determination of camera pose information using a bundle adjustment procedure. The estimated poses provide information about the relative position of the camera towards the object of interest.

[0073] If the MVIDMR does include corresponding image data, then the corresponding image data may be selected for analysis at **408**. At **410** the image data may then be rendered based on the configuration parameter selected at **404**. If the MVIDMR does not include image data corresponding the

selected configuration parameter, then the corresponding portion of the object may be omitted from the cinematic video. Alternately, in some embodiments, if the MVIDMR does not include image data precisely corresponding with the selected configuration parameter, then suitable image data may be generated.

[0074] In some embodiments, a view of an object from a novel viewpoint may be generated. For instance, the multi-view capture may include a number of static images of the object each captured from a respective perspective viewpoint. However, the cinematic base video may include a camera that pans continuously around a point in space. Accordingly, rendering the object into the scene may involve generating an image of the object that does not exist in the underlying multi-view capture. In this case, image data corresponding to viewpoints close to the object portion associated with the configuration parameter may be selected for analysis.

[0075] In some implementations, a novel viewpoint of an object may be generated from a set of perspective view images using light field reconstruction. Techniques and mechanisms for rendering a novel viewpoint of an object are discussed in additional detail in co-pending and commonly assigned U.S. patent application Ser. No. 16/574,622, titled FREE-VIEWPOINT PHOTOREALISTIC VIEW SYNTHESIS FROM CASUALLY CAPTURED VIDEO, by Kar et al., filed Sep. 18, 2019, which is hereby incorporated by reference in its entirety and for all purposes.

[0076] In some embodiments, a parameter may specify that a close-up of a particular portion of the object is to be included. If such a close-up is not included in the MVIDMR, then a wider-view image may be cropped to focus on the identified portion. For instance, if a closeup of the front left headlight of a vehicle is to be included, then one or more images of the front of the vehicle may be cropped to focus on the front left headlight.

[0077] One or more tags are applied to the rendered image data at **412**. In some implementations, tags may be added to highlight features of the object during playback of the cinematic video. Such tags can be rendered into the cinematic video. Tags may be generated in any of various ways.

[0078] In some embodiments, tags may be placed by a user, for instance during MVIDMR capture. For example, a user may identify a portion of an object and then provide a tag corresponding to the identified object portion.

[0079] In some embodiments, a tag may be placed automatically. For example, one or more pre-set locations may be automatically populated from a database. For instance, in the vehicle context, textual information about a vehicle such as performance, mileage, component name, warranty information, or other such data may be retrieved and added to the cinematic video. As another example, one or more pre-set locations may be automatically defined through the order of the capture process. For instance, during the capture process a user may be guided to capture specific parts of an object. Those objects may then be automatically populated with visual information at the corresponding location. As still another example, a tag may be automatically placed based on an estimate of the location of an object component from visual data. For instance, one or more images of an object may be analyzed to identify components of the object. Then one or more tags may be selected that correspond to those components.

[0080] In some implementations, applying the one or more tags may involve rendering information into the video itself. For example, textual information may be added directly to one or more video frames. As still another example, additional detailed visual portion may be added. For instance, if the front left light of a vehicle is captured, then a close-up of it may be added as part of the vehicle when the tag is selected.

[0081] According to various embodiments, a tag location can be used for any of various techniques. For example, text may be placed at a tag location. As another example, visual data in the video may be referenced based on the tag location. As still another example, one or more object components in the video may be highlighted based on the tag location. As yet another example, a tag location may be used to queue a transition effect. For instance, when a tagged location in the video is reached, the video may zoom into the tagged object and then blend to a close-up capture of the object.

[0082] A determination is made at **414** as to whether to select an additional configuration parameter. In some implementations, the determination may be made at least in part based on whether each of the configuration parameters identified at **402** have been selected for analysis at **404**.

[0083] When no additional configuration parameters are selected for analysis, the cinematic video is rendered at **416**. According to various embodiments, rendering the cinematic video may involve combining the image data rendered at **410** into a unified whole. For example, individual images modified in accordance with the configuration parameters and tags may be combined in sequence to generate the video.

[0084] In some embodiments, one or more effects may be applied to the cinematic video based on the configuration parameters identified at **402**. For instance, one such technique may be to artificially destabilize the video to make it appear as if the video was captured by a human operating a handheld camera, even if the MVIDMR was instead captured in a more fixed manner, such as via a fixed camera pointing at an object rotated on a turntable.

[0085] The rendered video is stored at **418**. In some implementations, storing the rendered video may involve sending the rendered video to a storage medium such as a hard drive or network attached storage (NAS) system, transmitting the rendered video to a remote computing device, or storing the rendered video in temporary memory such as random-access memory (RAM).

[0086] According to various embodiments, one or more of the operations shown in FIG. 4 may be omitted. For example, the operation **412** need not necessarily be performed for any or all frames.

[0087] FIGS. 19, 20, 21, and 22 illustrate examples of visual representations constructed or analyzed in accordance with various embodiments. FIGS. 19 and 20 show images in an MVIDMR within a configuration interface for creating a cinematic video. The user may configure a cinematic video at least in part by navigating the images in a particular order. For example, the images **1902**, **1904**, **2002**, **2004**, and **2006** show a circle and hand affordance that a user may employ to transition from one image to the next, counterclockwise or clockwise around the vehicle. As another example, the images **1902**, **1904**, **1906**, **2002**, and **2004** show a plus symbol that a user may employ to add additional context to the cinematic video. For instance, the user may place a tag

at such a location, zoom in on the location, and/or add a zoomed multi-view capture of the location at a point in the cinematic video.

[0088] In some embodiments, a user may employ the user interface to configure the cinematic video to transition between views. For example, the image 1906 shows a picture-in-picture of a similar view of the vehicle with the doors opened. The user could transition to such a view in the cinematic video by clicking on the picture-in-picture. The images 2102, 2104, and 2202 are examples of the images presented when such a transition has been made.

[0089] In some embodiments, the user may employ the user interface to configure the cinematic video to transition to a sub-MVIDMR. For example, the user may click on one of the plus buttons. The image 2106 shows an example of an MVIDMR of the area under the vehicle hood after the user has made such a selection. In the image 2106, the area under the hood is navigable by using the circle with hand affordance. The image 2204 shows a different view of the same area, after the user has navigated to a nearby image of the area.

[0090] FIGS. 23-25 show images from a cinematic video rendered in accordance with one or more embodiments. At 2302, an initial image is provided that only includes a tag that describes the object represented in the video—in this case a 2012 IPL G37 Coupe vehicle. At 2304, a subsequent frame begins a transition to an image of the vehicle. At 2306, a subsequent frame has mostly transitioned to an image of the vehicle. At 2308, a subsequent frame has transitioned entirely to an image of the vehicle, and includes another tag characterizing the vehicle. At 2402, a subsequent frame includes an image of the vehicle from a rotated perspective. At 2404, a subsequent frame from the video has transitioned to an interior view of the steering wheel and instrument panel. At 2406, a subsequent frame from the video shows the same components from a rotated perspective. At 2408, a subsequent video frame shows a closeup of the engine. At 2502, a subsequent video frame shows a closeup of the engine from a rotated perspective. At 2504, a subsequent video frame shows a tag that identifies information about the engine that was just presented. At 2506, a subsequent video frame shows a transition to an open-door view of the vehicle, with a tag that describes characteristics of the interior. At 2508, a subsequent video frame shows the same content from a rotated perspective.

[0091] According to various embodiments, the images shown in FIGS. 23-25 are frames taken from a cinematic video of a vehicle rendered in accordance with techniques and mechanisms described herein. The actual video includes many frames not shown in FIGS. 23-25. Although a vehicle is shown for the purpose of illustration, a cinematic video may be generated for any suitable object in accordance with techniques and mechanisms described herein. The frames included in the video may be selected from, and/or interpolated between, images including in an MVIDMR. Accordingly, the specific viewpoints presented in the video may depend on the availability of images in the MVIDMR. Further, not all images in the MVIDMR may be included. As discussed herein, specific images may be automatically selected as video frames based on the cinematic video configuration information.

[0092] In some embodiments, rendering cinematic video may involve identifying a viewpoint of an object that aligns with a viewpoint of the cinematic video configuration. For

instance, the multi-view capture may include a number of images of the object each captured from a respective perspective viewpoint, while the cinematic video configuration may include a virtual camera that pans around a point in space. For successive frames of the cinematic video configuration, a suitable image may be selected from the perspective view images. For example, estimated camera and object pose information may be compared with configuration information to select an image of the object from the multi-view capture that aligns with the camera perspective at a particular point in the cinematic video configuration.

[0093] FIG. 5 illustrates one example of a method 500 for performing geometric analysis of a perspective view image, performed in accordance with one or more embodiments. The method 500 may be performed on any suitable computing device. For example, the method 500 may be performed on a mobile computing device such as a smart phone. Alternately, or additionally, the method 500 may be performed on a remote server in communication with a mobile computing device.

[0094] A request to perform geometric analysis on an object is received at 502. According to various embodiments, the request may be received as part of a procedure for generating a cinematic video. For example, the method 500 may be used to determine object pose and/or other information for an object in order to include the object within a cinematic video

[0095] At 504, a video or image set of the object captured from one or more perspectives is identified. The video or image set is referred to herein as “source data”. According to various embodiments, the source data may include a 360-degree view of the object. Alternately, the source data may include a view that has less than 360-degree coverage.

[0096] In some embodiments, the source data may include data captured from a camera. For example, the camera may be located on a mobile computing device such a mobile phone. As another example, one or more traditional cameras may be used to capture such information.

[0097] In some implementations, the source data may include data collected from an inertial measurement unit (IMU). IMU data may include information such as camera location, camera angle, device velocity, device acceleration, or any of a wide variety of data collected from accelerometers or other such sensors.

[0098] The object is identified at 506. In some implementations, the object may be identified based on user input. For example, a user may identify the object as a vehicle or person via a user interface component such as a drop-down menu.

[0099] In some embodiments, the object may be identified based on image recognition. For example, the source data may be analyzed to determine that the subject of the source data is a vehicle, a person, or another such object. The source data may include a variety of image data. However, in case of a multi-view capture the source data focuses in a particular object from different viewpoints, the image recognition procedure may identify commonalities between the different perspective views to isolate the object that is the subject of the source data from other objects that are present in some portion of the source data but not in other portions of the source data.

[0100] At 508, vertices and faces of a 2D mesh are defined in the top-down view of the object. According to various embodiments, each face may represent a part of the object

surface that could be approximated as being planar. For example, when a vehicle is captured in the source data, the vehicle's door panel or roof may be represented as a face in a 2D mesh because the door and roof are approximately planar despite being slightly curved.

[0101] In some embodiments, vertices and faces of a 2D mesh may be identified by analyzing the source data. Alternately, or additionally, the identification of the object at **206** may allow for the retrieval of a predetermined 2D mesh. For example, a vehicle object may be associated with a default 2D mesh that may be retrieved upon request.

[0102] Visibility angles are determined for each vertex of the object at **510**. According to various embodiments, a visibility angle indicates the range of object angles with respect to the camera for which the vertex is visible. In some embodiments, visibility angles of a 2D mesh may be identified by analyzing the source data. Alternately, or additionally, the identification of the object at **506** may allow for the retrieval of predetermined visibility angle along with a predetermined 2D mesh. For example, a vehicle object may be associated with a default 2D mesh with associated visibility angle that may be retrieved upon request.

[0103] A 3D skeleton of the object is constructed at **512**. According to various embodiments, constructing a 3D skeleton may involve any of a variety of operations. For example, 2D skeleton detection may be performed on every frame using a machine learning procedure. As another example, 3D camera pose estimation may be performed to determine a location and angle of the camera with respect to the object for a particular frame. As yet another example, a 3D skeleton may be reconstructed from 2D skeletons and or poses.

[0104] Various embodiments described herein relate generally to systems and methods for analyzing the spatial relationship between multiple images and video together with location information data, for the purpose of creating a single representation, a multi-view image digital media representation (MVIDMR, also referred to as a multi-view capture), which eliminates redundancy in the data, and presents a user with an interactive and immersive active viewing experience. According to various embodiments, active is described in the context of providing a user with the ability to control the viewpoint of the visual information displayed on a screen.

[0105] In particular example embodiments, augmented reality (AR) is used to aid a user in capturing the multiple images used in a MVIDMR. For example, a virtual guide can be inserted into live image data from a mobile. The virtual guide can help the user guide the mobile device along a desirable path useful for creating the MVIDMR. The virtual guide in the AR images can respond to movements of the mobile device. The movement of mobile device can be determined from a number of different sources, including but not limited to an Inertial Measurement Unit and image data.

[0106] Various aspects also relate generally to systems and methods for providing feedback when generating a MVIDMR. For example, object recognition may be used to recognize an object present in a MVIDMR. Then, feedback such as one or more visual indicators may be provided to guide the user in collecting additional MVIDMR data to collect a high-quality MVIDMR of the object. As another example, a target view may be determined for a MVIDMR, such as the terminal point when capturing a 360-degree

MVIDMR. Then, feedback such as one or more visual indicators may be provided to guide the user in collecting additional MVIDMR data to reach the target view.

[0107] FIG. 6 shows an example of a MVIDMR acquisition system **600**, configured in accordance with one or more embodiments. The MVIDMR acquisition system **600** is depicted in a flow sequence that can be used to generate a MVIDMR. According to various embodiments, the data used to generate a MVIDMR can come from a variety of sources.

[0108] In particular, data such as, but not limited to two-dimensional (2D) images **604** can be used to generate a MVIDMR. These 2D images can include color image data streams such as multiple image sequences, video data, etc., or multiple images in any of various formats for images, depending on the application. As will be described in more detail below with respect to FIGS. 7A-11B, during an image capture process, an AR system can be used. The AR system can receive and augment live image data with virtual data. In particular, the virtual data can include guides for helping a user direct the motion of an image capture device.

[0109] Another source of data that can be used to generate a MVIDMR includes environment information **606**. This environment information **606** can be obtained from sources such as accelerometers, gyroscopes, magnetometers, GPS, WiFi, IMU-like systems (Inertial Measurement Unit systems), and the like. Yet another source of data that can be used to generate a MVIDMR can include depth images **608**. These depth images can include depth, 3D, or disparity image data streams, and the like, and can be captured by devices such as, but not limited to, stereo cameras, time-of-flight cameras, three-dimensional cameras, and the like.

[0110] In some embodiments, the data can then be fused together at sensor fusion block **610**. In some embodiments, a MVIDMR can be generated a combination of data that includes both 2D images **604** and environment information **606**, without any depth images **608** provided. In other embodiments, depth images **608** and environment information **606** can be used together at sensor fusion block **610**. Various combinations of image data can be used with environment information at **606**, depending on the application and available data.

[0111] In some embodiments, the data that has been fused together at sensor fusion block **610** is then used for content modeling **612** and context modeling **614**. As described in more detail with regard to FIG. 9, the subject matter featured in the images can be separated into content and context. The content can be delineated as the object of interest and the context can be delineated as the scenery surrounding the object of interest. According to various embodiments, the content can be a three-dimensional model, depicting an object of interest, although the content can be a two-dimensional image in some embodiments, as described in more detail below with regard to FIG. 9. Furthermore, in some embodiments, the context can be a two-dimensional model depicting the scenery surrounding the object of interest. Although in many examples the context can provide two-dimensional views of the scenery surrounding the object of interest, the context can also include three-dimensional aspects in some embodiments. For instance, the context can be depicted as a "flat" image along a cylindrical "canvas," such that the "flat" image appears on the surface of a cylinder. In addition, some examples may include three-dimensional context models, such as when some

objects are identified in the surrounding scenery as three-dimensional objects. According to various embodiments, the models provided by content modeling 612 and context modeling 614 can be generated by combining the image and location information data, as described in more detail with regard to FIG. 8.

[0112] According to various embodiments, context and content of a MVIDMR are determined based on a specified object of interest. In some embodiments, an object of interest is automatically chosen based on processing of the image and location information data. For instance, if a dominant object is detected in a series of images, this object can be selected as the content. In other examples, a user specified target 602 can be chosen, as shown in FIG. 6. It should be noted, however, that a MVIDMR can be generated without a user-specified target in some applications.

[0113] In some embodiments, one or more enhancement algorithms can be applied at enhancement algorithm(s) block 616. In particular example embodiments, various algorithms can be employed during capture of MVIDMR data, regardless of the type of capture mode employed. These algorithms can be used to enhance the user experience. For instance, automatic frame selection, stabilization, view interpolation, filters, and/or compression can be used during capture of MVIDMR data. In some embodiments, these enhancement algorithms can be applied to image data after acquisition of the data. In other examples, these enhancement algorithms can be applied to image data during capture of MVIDMR data.

[0114] According to various embodiments, automatic frame selection can be used to create a more enjoyable MVIDMR. Specifically, frames are automatically selected so that the transition between them will be smoother or more even. This automatic frame selection can incorporate blur- and overexposure-detection in some applications, as well as more uniformly sampling poses such that they are more evenly distributed.

[0115] In some embodiments, stabilization can be used for a MVIDMR in a manner similar to that used for video. In particular, keyframes in a MVIDMR can be stabilized for to produce improvements such as smoother transitions, improved/enhanced focus on the content, etc. However, unlike video, there are many additional sources of stabilization for a MVIDMR, such as by using IMU information, depth information, computer vision techniques, direct selection of an area to be stabilized, face detection, and the like.

[0116] For instance, IMU information can be very helpful for stabilization. In particular, IMU information provides an estimate, although sometimes a rough or noisy estimate, of the camera tremor that may occur during image capture. This estimate can be used to remove, cancel, and/or reduce the effects of such camera tremor.

[0117] In some embodiments, depth information, if available, can be used to provide stabilization for a MVIDMR. Because points of interest in a MVIDMR are three-dimensional, rather than two-dimensional, these points of interest are more constrained and tracking/matching of these points is simplified as the search space reduces. Furthermore, descriptors for points of interest can use both color and depth information and therefore, become more discriminative. In addition, automatic or semi-automatic content selection can be easier to provide with depth information. For instance, when a user selects a particular pixel of an image, this selection can be expanded to fill the entire surface that

touches it. Furthermore, content can also be selected automatically by using a foreground/background differentiation based on depth. According to various embodiments, the content can stay relatively stable/visible even when the context changes.

[0118] According to various embodiments, computer vision techniques can also be used to provide stabilization for MVIDMRs. For instance, keypoints can be detected and tracked. However, in certain scenes, such as a dynamic scene or static scene with parallax, no simple warp exists that can stabilize everything. Consequently, there is a trade-off in which certain aspects of the scene receive more attention to stabilization and other aspects of the scene receive less attention. Because a MVIDMR is often focused on a particular object of interest, a MVIDMR can be content-weighted so that the object of interest is maximally stabilized in some examples.

[0119] Another way to improve stabilization in a MVIDMR includes direct selection of a region of a screen. For instance, if a user taps to focus on a region of a screen, then records a convex MVIDMR, the area that was tapped can be maximally stabilized. This allows stabilization algorithms to be focused on a particular area or object of interest.

[0120] In some embodiments, face detection can be used to provide stabilization. For instance, when recording with a front-facing camera, it is often likely that the user is the object of interest in the scene. Thus, face detection can be used to weight stabilization about that region. When face detection is precise enough, facial features themselves (such as eyes, nose, and mouth) can be used as areas to stabilize, rather than using generic keypoints. In another example, a user can select an area of image to use as a source for keypoints.

[0121] According to various embodiments, view interpolation can be used to improve the viewing experience. In particular, to avoid sudden “jumps” between stabilized frames, synthetic, intermediate views can be rendered on the fly. This can be informed by content-weighted keypoint tracks and IMU information as described above, as well as by denser pixel-to-pixel matches. If depth information is available, fewer artifacts resulting from mismatched pixels may occur, thereby simplifying the process. As described above, view interpolation can be applied during capture of a MVIDMR in some embodiments. In other embodiments, view interpolation can be applied during MVIDMR generation.

[0122] In some embodiments, filters can also be used during capture or generation of a MVIDMR to enhance the viewing experience. Just as many popular photo sharing services provide aesthetic filters that can be applied to static, two-dimensional images, aesthetic filters can similarly be applied to surround images. However, because a MVIDMR representation is more expressive than a two-dimensional image, and three-dimensional information is available in a MVIDMR, these filters can be extended to include effects that are ill-defined in two dimensional photos. For instance, in a MVIDMR, motion blur can be added to the background (i.e. context) while the content remains crisp. In another example, a drop-shadow can be added to the object of interest in a MVIDMR.

[0123] According to various embodiments, compression can also be used as an enhancement algorithm 616. In particular, compression can be used to enhance user-experience by reducing data upload and download costs. Because

MVIDMRs use spatial information, far less data can be sent for a MVIDMR than a typical video, while maintaining desired qualities of the MVIDMR. Specifically, the IMU, keypoint tracks, and user input, combined with the view interpolation described above, can all reduce the amount of data that must be transferred to and from a device during upload or download of a MVIDMR. For instance, if an object of interest can be properly identified, a variable compression style can be chosen for the content and context. This variable compression style can include lower quality resolution for background information (i.e. context) and higher quality resolution for foreground information (i.e. content) in some examples. In such examples, the amount of data transmitted can be reduced by sacrificing some of the context quality, while maintaining a desired level of quality for the content.

[0124] In the present embodiment, a MVIDMR 618 is generated after any enhancement algorithms are applied. The MVIDMR can provide a multi-view interactive digital media representation. According to various embodiments, the MVIDMR can include three-dimensional model of the content and a two-dimensional model of the context. However, in some examples, the context can represent a “flat” view of the scenery or background as projected along a surface, such as a cylindrical or other-shaped surface, such that the context is not purely two-dimensional. In yet other examples, the context can include three-dimensional aspects.

[0125] According to various embodiments, MVIDMRs provide numerous advantages over traditional two-dimensional images or videos. Some of these advantages include: the ability to cope with moving scenery, a moving acquisition device, or both; the ability to model parts of the scene in three-dimensions; the ability to remove unnecessary, redundant information and reduce the memory footprint of the output dataset; the ability to distinguish between content and context; the ability to use the distinction between content and context for improvements in the user-experience; the ability to use the distinction between content and context for improvements in memory footprint (an example would be high quality compression of content and low quality compression of context); the ability to associate special feature descriptors with MVIDMRs that allow the MVIDMRs to be indexed with a high degree of efficiency and accuracy; and the ability of the user to interact and change the viewpoint of the MVIDMR. In particular example embodiments, the characteristics described above can be incorporated natively in the MVIDMR representation, and provide the capability for use in various applications. For instance, MVIDMRs can be used to enhance various fields such as e-commerce, visual search, 3D printing, file sharing, user interaction, and entertainment.

[0126] According to various example embodiments, once a MVIDMR 618 is generated, user feedback for acquisition 620 of additional image data can be provided. In particular, if a MVIDMR is determined to need additional views to provide a more accurate model of the content or context, a user may be prompted to provide additional views. Once these additional views are received by the MVIDMR acquisition system 600, these additional views can be processed by the system 600 and incorporated into the MVIDMR.

[0127] FIG. 7 shows an example of a process flow diagram for generating a MVIDMR 700. In the present example, a plurality of images is obtained at 702. According to various

embodiments, the plurality of images can include two-dimensional (2D) images or data streams. These 2D images can include location information that can be used to generate a MVIDMR. In some embodiments, the plurality of images can include depth images 608, as also described above with regard to FIG. 6. The depth images can also include location information in various examples.

[0128] In some embodiments, when the plurality of images is captured, images output to the user can be augmented with the virtual data. For example, the plurality of images can be captured using a camera system on a mobile device. The live image data, which is output to a display on the mobile device, can include virtual data, such as guides and status indicators, rendered into the live image data. The guides can help a user guide a motion of the mobile device. The status indicators can indicate what portion of images needed for generating a MVIDMR have been captured. The virtual data may not be included in the image data captured for the purposes of generating the MVIDMR.

[0129] According to various embodiments, the plurality of images obtained at 702 can include a variety of sources and characteristics. For instance, the plurality of images can be obtained from a plurality of users. These images can be a collection of images gathered from the internet from different users of the same event, such as 2D images or video obtained at a concert, etc. In some embodiments, the plurality of images can include images with different temporal information. In particular, the images can be taken at different times of the same object of interest. For instance, multiple images of a particular statue can be obtained at different times of day, different seasons, etc. In other examples, the plurality of images can represent moving objects. For instance, the images may include an object of interest moving through scenery, such as a vehicle traveling along a road or a plane traveling through the sky. In other instances, the images may include an object of interest that is also moving, such as a person dancing, running, twirling, etc.

[0130] In some embodiments, the plurality of images is fused into content and context models at 704. According to various embodiments, the subject matter featured in the images can be separated into content and context. The content can be delineated as the object of interest and the context can be delineated as the scenery surrounding the object of interest. According to various embodiments, the content can be a three-dimensional model, depicting an object of interest, and the content can be a two-dimensional image in some embodiments.

[0131] According to the present example embodiment, one or more enhancement algorithms can be applied to the content and context models at 706. These algorithms can be used to enhance the user experience. For instance, enhancement algorithms such as automatic frame selection, stabilization, view interpolation, filters, and/or compression can be used. In some embodiments, these enhancement algorithms can be applied to image data during capture of the images. In other examples, these enhancement algorithms can be applied to image data after acquisition of the data.

[0132] In the present embodiment, a MVIDMR is generated from the content and context models at 708. The MVIDMR can provide a multi-view interactive digital media representation. According to various embodiments, the MVIDMR can include a three-dimensional model of the content and a two-dimensional model of the context.

According to various embodiments, depending on the mode of capture and the viewpoints of the images, the MVIDMR model can include certain characteristics. For instance, some examples of different styles of MVIDMRs include a locally concave MVIDMR, a locally convex MVIDMR, and a locally flat MVIDMR. However, it should be noted that MVIDMRs can include combinations of views and characteristics, depending on the application.

[0133] FIG. 8 shows an example of multiple camera views that can be fused together into a three-dimensional (3D) model to create an immersive experience. According to various embodiments, multiple images can be captured from various viewpoints and fused together to provide a MVIDMR. In some embodiments, three cameras **812**, **814**, and **816** are positioned at locations **822**, **824**, and **826**, respectively, in proximity to an object of interest **808**. Scenery can surround the object of interest **808** such as object **810**. Views **802**, **804**, and **806** from their respective cameras **812**, **814**, and **816** include overlapping subject matter. Specifically, each view **802**, **804**, and **806** includes the object of interest **808** and varying degrees of visibility of the scenery surrounding the object **810**. For instance, view **802** includes a view of the object of interest **808** in front of the cylinder that is part of the scenery surrounding the object **810**. View **806** shows the object of interest **808** to one side of the cylinder, and view **804** shows the object of interest without any view of the cylinder.

[0134] In some embodiments, the various views **802**, **804**, and **816** along with their associated locations **822**, **824**, and **826**, respectively, provide a rich source of information about object of interest **808** and the surrounding context that can be used to produce a MVIDMR. For instance, when analyzed together, the various views **802**, **804**, and **826** provide information about different sides of the object of interest and the relationship between the object of interest and the scenery. According to various embodiments, this information can be used to parse out the object of interest **808** into content and the scenery as the context. Furthermore, various algorithms can be applied to images produced by these viewpoints to create an immersive, interactive experience when viewing a MVIDMR.

[0135] FIG. 9 illustrates one example of separation of content and context in a MVIDMR. According to various embodiments, a MVIDMR is a multi-view interactive digital media representation of a scene **900**. With reference to FIG. 9, shown is a user **902** located in a scene **900**. The user **902** is capturing images of an object of interest, such as a statue. The images captured by the user constitute digital visual data that can be used to generate a MVIDMR.

[0136] According to various embodiments of the present disclosure, the digital visual data included in a MVIDMR can be, semantically and/or practically, separated into content **904** and context **906**. According to particular embodiments, content **904** can include the object(s), person(s), or scene(s) of interest while the context **906** represents the remaining elements of the scene surrounding the content **904**. In some embodiments, a MVIDMR may represent the content **904** as three-dimensional data, and the context **906** as a two-dimensional panoramic background. In other examples, a MVIDMR may represent both the content **904** and context **906** as two-dimensional panoramic scenes. In yet other examples, content **904** and context **906** may include three-dimensional components or aspects. In par-

ticular embodiments, the way that the MVIDMR depicts content **904** and context **906** depends on the capture mode used to acquire the images.

[0137] In some embodiments, such as but not limited to: recordings of objects, persons, or parts of objects or persons, where only the object, person, or parts of them are visible, recordings of large flat areas, and recordings of scenes where the data captured appears to be at infinity (i.e., there are no subjects close to the camera), the content **904** and the context **906** may be the same. In these examples, the MVIDMR produced may have some characteristics that are similar to other types of digital media such as panoramas. However, according to various embodiments, MVIDMRs include additional features that distinguish them from these existing types of digital media. For instance, a MVIDMR can represent moving data. Additionally, a MVIDMR is not limited to a specific cylindrical, spherical or translational movement. Various motions can be used to capture image data with a camera or other capture device. Furthermore, unlike a stitched panorama, a MVIDMR can display different sides of the same object.

[0138] FIGS. 10A-10B illustrate examples of concave and convex views, respectively, where both views use a back-camera capture style. In particular, if a camera phone is used, these views use the camera on the back of the phone, facing away from the user. In particular embodiments, concave and convex views can affect how the content and context are designated in a MVIDMR.

[0139] With reference to FIG. 10A, shown is one example of a concave view **1000** in which a user is standing along a vertical axis **1008**. In this example, the user is holding a camera, such that camera location **1002** does not leave axis **1008** during image capture. However, as the user pivots about axis **1008**, the camera captures a panoramic view of the scene around the user, forming a concave view. In this embodiment, the object of interest **1004** and the distant scenery **1006** are all viewed similarly because of the way in which the images are captured. In this example, all objects in the concave view appear at infinity, so the content is equal to the context according to this view.

[0140] With reference to FIG. 10B, shown is one example of a convex view **1020** in which a user changes position when capturing images of an object of interest **1024**. In this example, the user moves around the object of interest **1024**, taking pictures from different sides of the object of interest from camera locations **1028**, **1030**, and **1032**. Each of the images obtained includes a view of the object of interest, and a background of the distant scenery **1026**. In the present example, the object of interest **1024** represents the content, and the distant scenery **1026** represents the context in this convex view.

[0141] FIGS. 11A-11B illustrate examples of various capture modes for MVIDMRs. Although various motions can be used to capture a MVIDMR and are not constrained to any particular type of motion, three general types of motion can be used to capture particular features or views described in conjunction MVIDMRs. These three types of motion, respectively, can yield a locally concave MVIDMR, a locally convex MVIDMR, and a locally flat MVIDMR. In some embodiments, a MVIDMR can include various types of motions within the same MVIDMR.

[0142] With reference to FIG. 11A, shown is an example of a back-facing, concave MVIDMR being captured. According to various embodiments, a locally concave

MVIDMR is one in which the viewing angles of the camera or other capture device diverge. In one dimension this can be likened to the motion required to capture a spherical 360 panorama (pure rotation), although the motion can be generalized to any curved sweeping motion in which the view faces outward. In the present example, the experience is that of a stationary viewer looking out at a (possibly dynamic) context.

[0143] In some embodiments, a user 1102 is using a back-facing camera 1106 to capture images towards world 1100, and away from user 1102. As described in various examples, a back-facing camera refers to a device with a camera that faces away from the user, such as the camera on the back of a smart phone. The camera is moved in a concave motion 1108, such that views 1104a, 1104b, and 1104c capture various parts of capture area 1109.

[0144] With reference to FIG. 11B, shown is an example of a back-facing, convex MVIDMR being captured. According to various embodiments, a locally convex MVIDMR is one in which viewing angles converge toward a single object of interest. In some embodiments, a locally convex MVIDMR can provide the experience of orbiting about a point, such that a viewer can see multiple sides of the same object. This object, which may be an “object of interest,” can be segmented from the MVIDMR to become the content, and any surrounding data can be segmented to become the context. Previous technologies fail to recognize this type of viewing angle in the media-sharing landscape.

[0145] In some embodiments, a user 1102 is using a back-facing camera 1114 to capture images towards world 1100, and away from user 1102. The camera is moved in a convex motion 1110, such that views 1112a, 1112b, and 1112c capture various parts of capture area 1111. As described above, world 1100 can include an object of interest in some examples, and the convex motion 1110 can orbit around this object. Views 1112a, 1112b, and 1112c can include views of different sides of this object in these examples.

[0146] With reference to FIG. 12A, shown is an example of a front-facing, concave MVIDMR being captured. As described in various examples, a front-facing camera refers to a device with a camera that faces towards the user, such as the camera on the front of a smart phone. For instance, front-facing cameras are commonly used to take “selfies” (i.e., self-portraits of the user).

[0147] In some embodiments, camera 1220 is facing user 1202. The camera follows a concave motion 1206 such that the views 1218a, 1218b, and 1218c diverge from each other in an angular sense. The capture area 1217 follows a concave shape that includes the user at a perimeter.

[0148] With reference to FIG. 12B, shown is an example of a front-facing, convex MVIDMR being captured. In some embodiments, camera 1226 is facing user 1202. The camera follows a convex motion 1222 such that the views 1224a, 1224b, and 1224c converge towards the user 1202. As described above, various modes can be used to capture images for a MVIDMR. These modes, including locally concave, locally convex, and locally linear motions, can be used during capture of separate images or during continuous recording of a scene. Such recording can capture a series of images during a single session.

[0149] In some embodiments, the augmented reality system can be implemented on a mobile device, such as a cell phone. In particular, the live camera data, which is output to

a display on the mobile device, can be augmented with virtual objects. The virtual objects can be rendered into the live camera data. In some embodiments, the virtual objects can provide a user feedback when images are being captured for a MVIDMR.

[0150] FIGS. 13 and 14 illustrate an example of a process flow for capturing images in a MVIDMR using augmented reality. In 1302, live image data can be received from a camera system. For example, live image data can be received from one or more cameras on a handheld mobile device, such as a smartphone. The image data can include pixel data captured from a camera sensor. The pixel data varies from frame to frame. In some embodiments, the pixel data can be 2-D. In other embodiments, depth data can be included with the pixel data.

[0151] In 1304, sensor data can be received. For example, the mobile device can include an IMU with accelerometers and gyroscopes. The sensor data can be used to determine an orientation of the mobile device, such as a tilt orientation of the device relative to the gravity vector. Thus, the orientation of the live 2-D image data relative to the gravity vector can also be determined. In addition, when the user applied accelerations can be separated from the acceleration due to gravity, it may be possible to determine changes in position of the mobile device as a function of time.

[0152] In particular embodiments, a camera reference frame can be determined. In the camera reference frame, one axis is aligned with a line perpendicular to the camera lens. Using an accelerometer on the phone, the camera reference frame can be related to an Earth reference frame. The earth reference frame can provide a 3-D coordinate system where one of the axes is aligned with the Earths’ gravitational vector. The relationship between the camera frame and Earth reference frame can be indicated as yaw, roll and tilt/pitch. Typically, at least two of the three of yaw, roll and pitch are available typically from sensors available on a mobile device, such as smart phone’s gyroscopes and accelerometers.

[0153] The combination of yaw-roll-tilt information from the sensors, such as a smart phone or tablets accelerometers and the data from the camera including the pixel data can be used to relate the 2-D pixel arrangement in the camera field of view to the 3-D reference frame in the real world. In some embodiments, the 2-D pixel data for each picture can be translated to a reference frame as if the camera were resting on a horizontal plane perpendicular to an axis through the gravitational center of the Earth where a line drawn through the center of lens perpendicular to the surface of lens is mapped to a center of the pixel data. This reference frame can be referred as an Earth reference frame. Using this calibration of the pixel data, a curve or object defined in 3-D space in the earth reference frame can be mapped to a plane associated with the pixel data (2-D pixel data). If depth data is available, i.e., the distance of the camera to a pixel. Then, this information can also be utilized in a transformation.

[0154] In alternate embodiments, the 3-D reference frame in which an object is defined doesn’t have to be an Earth reference frame. In some embodiments, a 3-D reference in which an object is drawn and then rendered into the 2-D pixel frame of reference can be defined relative to the Earth reference frame. In another embodiment, a 3-D reference frame can be defined relative to an object or surface identified in the pixel data and then the pixel data can be calibrated to this 3-D reference frame.

[0155] As an example, the object or surface can be defined by a number of tracking points identified in the pixel data. Then, as the camera moves, using the sensor data and a new position of the tracking points, a change in the orientation of the 3-D reference frame can be determined from frame to frame. This information can be used to render virtual data in a live image data and/or virtual data into a MVIDMR.

[0156] Returning to FIG. 13, in 1306, virtual data associated with a target can be generated in the live image data. For example, the target can be cross hairs. In general, the target can be rendered as any shape or combinations of shapes. In some embodiments, via an input interface, a user may be able to adjust a position of the target. For example, using a touch screen over a display on which the live image data is output, the user may be able to place the target at a particular location in the synthetic image. The synthetic image can include a combination of live image data rendered with one or more virtual objects.

[0157] For example, the target can be placed over an object that appears in the image, such as a face or a person. Then, the user can provide an additional input via an interface that indicates the target is in a desired location. For example, the user can tap the touch screen proximate to the location where the target appears on the display. Then, an object in the image below the target can be selected. As another example, a microphone in the interface can be used to receive voice commands which direct a position of the target in the image (e.g., move left, move right, etc.) and then confirm when the target is in a desired location (e.g., select target).

[0158] In some instances, object recognition can be available. Object recognition can identify possible objects in the image. Then, the live images can be augmented with a number of indicators, such as targets, which mark identified objects. For example, objects, such as people, parts of people (e.g., faces), cars, wheels, can be marked in the image. Via an interface, the person may be able to select one of the marked objects, such as via the touch screen interface. In another embodiment, the person may be able to provide a voice command to select an object. For example, the person may be to say something like “select face,” or “select car.”

[0159] In 1308, the object selection can be received. The object selection can be used to determine an area within the image data to identify tracking points. When the area in the image data is over a target, the tracking points can be associated with an object appearing in the live image data.

[0160] In 1310, tracking points can be identified which are related to the selected object. Once an object is selected, the tracking points on the object can be identified on a frame to frame basis. Thus, if the camera translates or changes orientation, the location of the tracking points in the new frame can be identified and the target can be rendered in the live images so that it appears to stay over the tracked object in the image. This feature is discussed in more detail below. In particular embodiments, object detection and/or recognition may be used for each or most frames, for instance to facilitate identifying the location of tracking points.

[0161] In some embodiments, tracking an object can refer to tracking one or more points from frame to frame in the 2-D image space. The one or more points can be associated with a region in the image. The one or more points or regions can be associated with an object. However, the object doesn't have to be identified in the image. For example, the

boundaries of the object in 2-D image space don't have to be known. Further, the type of object doesn't have to be identified. For example, a determination doesn't have to be made as to whether the object is a car, a person or something else appearing in the pixel data. Instead, the one or more points may be tracked based on other image characteristics that appear in successive frames. For instance, edge tracking, corner tracking, or shape tracking may be used to track one or more points from frame to frame.

[0162] One advantage of tracking objects in the manner described in the 2-D image space is that a 3-D reconstruction of an object or objects appearing in an image don't have to be performed. The 3-D reconstruction step may involve operations such as “structure from motion (SFM)” and/or “simultaneous localization and mapping (SLAM).” The 3-D reconstruction can involve measuring points in multiple images, and the optimizing for the camera poses and the point locations. When this process is avoided, significant computation time is saved. For example, avoiding the SLAM/SFM computations can enable the methods to be applied when objects in the images are moving. Typically, SLAM/SFM computations assume static environments.

[0163] In 1312, a 3-D coordinate system in the physical world can be associated with the image, such as the Earth reference frame, which as described above can be related to camera reference frame associated with the 2-D pixel data. In some embodiments, the 2-D image data can be calibrated so that the associated 3-D coordinate system is anchored to the selected target such that the target is at the origin of the 3-D coordinate system.

[0164] Then, in 1314, a 2-D or 3-D trajectory or path can be defined in the 3-D coordinate system. For example, a trajectory or path, such as an arc or a parabola can be mapped to a drawing plane which is perpendicular to the gravity vector in the Earth reference frame. As described above, based upon the orientation of the camera, such as information provided from an IMU, the camera reference frame including the 2-D pixel data can be mapped to the Earth reference frame. The mapping can be used to render the curve defined in the 3-D coordinate system into the 2-D pixel data from the live image data. Then, a synthetic image including the live image data and the virtual object, which is the trajectory or path, can be output to a display.

[0165] In general, virtual objects, such as curves or surfaces can be defined in a 3-D coordinate system, such as the Earth reference frame or some other coordinate system related to an orientation of the camera. Then, the virtual objects can be rendered into the 2-D pixel data associated with the live image data to create a synthetic image. The synthetic image can be output to a display.

[0166] In some embodiments, the curves or surfaces can be associated with a 3-D model of an object, such as person or a car. In another embodiment, the curves or surfaces can be associated with text. Thus, a text message can be rendered into the live image data. In other embodiments, textures can be assigned to the surfaces in the 3-D model. When a synthetic image is created, these textures can be rendered into the 2-D pixel data associated with the live image data.

[0167] When a curve is rendered on a drawing plane in the 3-D coordinate system, such as the Earth reference frame, one or more of the determined tracking points can be projected onto the drawing plane. As another example, a centroid associated with the tracked points can be projected onto the drawing plane. Then, the curve can be defined

relative to one or more points projected onto the drawing plane. For example, based upon the target location, a point can be determined on the drawing plane. Then, the point can be used as the center of a circle or arc of some radius drawn in the drawing plane.

[0168] In **1314**, based upon the associated coordinate system, a curve can be rendered into the live image data as part of the AR system. In general, one or more virtual objects including plurality of curves, lines or surfaces can be rendered into the live image data. Then, the synthetic image including the live image data and the virtual objects can be output to a display in real-time.

[0169] In some embodiments, the one or more virtual object rendered into the live image data can be used to help a user capture images used to create a MVIDMR. For example, the user can indicate a desire to create a MVIDMR of a real object identified in the live image data. The desired MVIDMR can span some angle range, such as forty-five, ninety, one hundred eighty degrees or three hundred sixty degrees. Then, a virtual object can be rendered as a guide where the guide is inserted into the live image data. The guide can indicate a path along which to move the camera and the progress along the path. The insertion of the guide can involve modifying the pixel data in the live image data in accordance with coordinate system in **1312**.

[0170] In the example above, the real object can be some object which appears in the live image data. For the real object, a 3-D model may not be constructed. Instead, pixel locations or pixel areas can be associated with the real object in the 2-D pixel data. This definition of the real object is much less computationally expensive than attempting to construct a 3-D model of the real object in physical space.

[0171] The virtual objects, such as lines or surfaces can be modeled in the 3-D space. The virtual objects can be defined a priori. Thus, the shape of the virtual object doesn't have to be constructed in real-time, which is computationally expensive. The real objects which may appear in an image are not known a priori. Hence, 3-D models of the real object are not typically available. Therefore, the synthetic image can include "real" objects which are only defined in the 2-D image space via assigning tracking points or areas to the real object and virtual objects which are modeled in a 3-D coordinate system and then rendered into the live image data.

[0172] Returning to FIG. 13, in **1316**, AR image with one or more virtual objects can be output. The pixel data in the live image data can be received at a particular frame rate. In particular embodiments, the augmented frames can be output at the same frame rate as it received. In other embodiments, it can be output at a reduced frame rate. The reduced frame rate can lessen computation requirements. For example, live data received at 30 frames per second can be output at 15 frames per second. In another embodiment, the AR images can be output at a reduced resolution, such as 240p instead of 480p. The reduced resolution can also be used to reduce computational requirements.

[0173] In **1318**, one or more images can be selected from the live image data and stored for use in a MVIDMR. In some embodiments, the stored images can include one or more virtual objects. Thus, the virtual objects can be become part of the MVIDMR. In other embodiments, the virtual objects are only output as part of the AR system. But, the image data which is stored for use in the MVIDMR may not include the virtual objects.

[0174] In yet other embodiments, a portion of the virtual objects output to the display as part of the AR system can be stored. For example, the AR system can be used to render a guide during the MVIDMR image capture process and render a label associated with the MVIDMR. The label may be stored in the image data for the MVIDMR. However, the guide may not be stored. To store the images without the added virtual objects, a copy may have to be made. The copy can be modified with the virtual data and then output to a display and the original stored or the original can be stored prior to its modification.

[0175] In FIG. 14, the method in FIG. 13 is continued. In **1422**, new image data can be received. In **1424**, new IMU data (or, in general sensor data) can be received. The IMU data can represent a current orientation of the camera. In **1426**, the location of the tracking points identified in previous image data can be identified in the new image data.

[0176] The camera may have tilted and/or moved. Hence, the tracking points may appear at a different location in the pixel data. As described above, the tracking points can be used to define a real object appearing in the live image data. Thus, identifying the location of the tracking points in the new image data allows the real object to be tracked from image to image. The differences in IMU data from frame to frame and knowledge of the rate at which the frames are recorded can be used to help to determine a change in location of tracking points in the live image data from frame to frame.

[0177] The tracking points associated with a real object appearing in the live image data may change over time. As a camera moves around the real object, some tracking points identified on the real object may go out of view as new portions of the real object come into view and other portions of the real object are occluded. Thus, in **1426**, a determination may be made whether a tracking point is still visible in an image. In addition, a determination may be made as to whether a new portion of the targeted object has come into view. New tracking points can be added to the new portion to allow for continued tracking of the real object from frame to frame.

[0178] In **1428**, a coordinate system can be associated with the image. For example, using an orientation of the camera determined from the sensor data, the pixel data can be calibrated to an Earth reference frame as previously described. In **1430**, based upon the tracking points currently placed on the object and the coordinate system a target location can be determined. The target can be placed over the real object which is tracked in live image data. As described above, a number and a location of the tracking points identified in an image can vary with time as the position of the camera changes relative to the camera. Thus, the location of the target in the 2-D pixel data can change. A virtual object representing the target can be rendered into the live image data. In particular embodiments, a coordinate system may be defined based on identifying a position from the tracking data and an orientation from the IMU (or other) data.

[0179] In **1432**, a track location in the live image data can be determined. The track can be used to provide feedback associated with a position and orientation of a camera in physical space during the image capture process for a MVIDMR. As an example, as described above, the track can be rendered in a drawing plane which is perpendicular to the gravity vector, such as parallel to the ground. Further, the

track can be rendered relative to a position of the target, which is a virtual object, placed over a real object appearing in the live image data. Thus, the track can appear to surround or partially surround the object. As described above, the position of the target can be determined from the current set of tracking points associated with the real object appearing in the image. The position of the target can be projected onto the selected drawing plane.

[0180] In **1434**, a capture indicator status can be determined. The capture indicator can be used to provide feedback in regards to what portion of the image data used in a MVIDMR has been captured. For example, the status indicator may indicate that half of angle range of images for use in a MVIDMR has been captured. In another embodiment, the status indicator may be used to provide feedback in regards to whether the camera is following a desired path and maintaining a desired orientation in physical space. Thus, the status indicator may indicate the current path or orientation of the camera is desirable or not desirable. When the current path or orientation of the camera is not desirable, the status indicator may be configured to indicate what type of correction which is needed, such as but not limited to moving the camera more slowly, starting the capture process over, tilting the camera in a certain direction and/or translating the camera in a particular direction.

[0181] In **1436**, a capture indicator location can be determined. The location can be used to render the capture indicator into the live image and generate the synthetic image. In some embodiments, the position of the capture indicator can be determined relative to a position of the real object in the image as indicated by the current set of tracking points, such as above and to left of the real object. In **1438**, a synthetic image, i.e., a live image augmented with virtual objects, can be generated. The synthetic image can include the target, the track and one or more status indicators at their determined locations, respectively. In **1440**, image data captured for the purposes of use in a MVIDMR can be captured. As described above, the stored image data can be raw image data without virtual objects or may include virtual objects.

[0182] In **1442**, a check can be made as to whether images needed to generate a MVIDMR have been captured in accordance with the selected parameters, such as a MVIDMR spanning a desired angle range. When the capture is not complete, new image data may be received and the method may return to **1422**. When the capture is complete, a virtual object can be rendered into the live image data indicating the completion of the capture process for the MVIDMR and a MVIDMR can be created. Some virtual objects associated with the capture process may cease to be rendered. For example, once the needed images have been captured the track used to help guide the camera during the capture process may no longer be generated in the live image data.

[0183] FIGS. **15A** and **15B** illustrate aspects of generating an Augmented Reality (AR) image capture track for capturing images used in a MVIDMR. In FIG. **15A**, a mobile device **1514** with a display **1516** is shown. The mobile device can include at least one camera (not shown) with a field of view **1500**. A real object **1502**, which is a person, is selected in the field of view **1500** of the camera. A virtual object, which is a target (not shown), may have been used to help select the real object. For example, the target on a touch

screen display of the mobile device **1514** may have been placed over the object **1502** and then selected.

[0184] The camera can include an image sensor which captures light in the field of view **1500**. The data from the image sensor can be converted to pixel data. The pixel data can be modified prior to its output on display **1516** to generate a synthetic image. The modifications can include rendering virtual objects in the pixel data as part of an augmented reality (AR) system.

[0185] Using the pixel data and a selection of the object **1502**, tracking points on the object can be determined. The tracking points can define the object in image space. Locations of a current set of tracking points, such as **1505**, **1506** and **1508**, which can be attached to the object **1502** are shown. As a position and orientation of the camera on the mobile device **1514**, the shape and position of the object **1502** in the captured pixel data can change. Thus, the location of the tracking points in the pixel data can change. Thus, a previously defined tracking point can move from a first location in the image data to a second location. Also, a tracking point can disappear from the image as portions of the object are occluded.

[0186] Using sensor data from the mobile device **1514**, an Earth reference frame 3-D coordinate system **1504** can be associated with the image data. The direction of the gravity vector is indicated by arrow **1510**. As described above, in a particular embodiment, the 2-D image data can be calibrated relative to the Earth reference frame. The arrow representing the gravity vector is not rendered into the live image data. However, if desired, an indicator representative of the gravity could be rendered into the synthetic image.

[0187] A plane which is perpendicular to the gravity vector can be determined. The location of the plane can be determined using the tracking points in the image, such as **1505**, **1506** and **1508**. Using this information, a curve, which is a circle, is drawn in the plane. The circle can be rendered into to the 2-D image data and output as part of the AR system. As is shown on display **1516**, the circle appears to surround the object **1502**. In some embodiments, the circle can be used as a guide for capturing images used in a MVIDMR.

[0188] If the camera on the mobile device **1514** is rotated in some way, such as tilted, the shape of the object will change on display **1516**. However, the new orientation of the camera can be determined in space including a direction of the gravity vector. Hence, a plane perpendicular to the gravity vector can be determined. The position of the plane and hence, a position of the curve in the image can be based upon a centroid of the object determined from the tracking points associated with the object **1502**. Thus, the curve can appear to remain parallel to the ground, i.e., perpendicular to the gravity vector, as the camera **1514** moves. However, the position of the curve can move from location to location in the image as the position of the object and its apparent shape in the live images changes.

[0189] In FIG. **15B**, a mobile device **1534** including a camera (not shown) and a display **1536** for outputting the image data from the camera is shown. A cup **1522** is shown in the field of view of camera **1520** of the camera. Tracking points, such as **1524** and **1526**, have been associated with the object **1522**. These tracking points can define the object **1522** in image space. Using the IMU data from the mobile device **1534**, a reference frame has been associated with the image data. As described above, In some embodiments, the

pixel data can be calibrated to the reference frame. The reference frame is indicated by the 3-D axes **1524** and the direction of the gravity vector is indicated by arrow **1528**.

[0190] As described above, a plane relative to the reference frame can be determined. In this example, the plane is parallel to the direction of the axis associated with the gravity vector as opposed to perpendicular to the frame. This plane is used to proscribe a path for the MVIDMR which goes over the top of the object **1530**. In general, any plane can be determined in the reference frame and then a curve, which is used as a guide, can be rendered into the selected plane.

[0191] Using the locations of the tracking points, in some embodiments, a centroid of the object **1522** on the selected plane in the reference can be determined. A curve **1530**, such as a circle, can be rendered relative to the centroid. In this example, a circle is rendered around the object **1522** in the selected plane.

[0192] The curve **1530** can serve as a track for guiding the camera along a particular path where the images captured along the path can be converted into a MVIDMR. In some embodiments, a position of the camera along the path can be determined. Then, an indicator can be generated which indicates a current location of the camera along the path. In this example, current location is indicated by arrow **1532**.

[0193] The position of the camera along the path may not directly map to physical space, i.e., the actual position of the camera in physical space doesn't have to be necessarily determined. For example, an angular change can be estimated from the IMU data and optionally the frame rate of the camera. The angular change can be mapped to a distance moved along the curve where the ratio of the distance moved along the path **1530** is not a one to one ratio with the distance moved in physical space. In another example, a total time to traverse the path **1530** can be estimated and then the length of time during which images have been recorded can be tracked. The ratio of the recording time to the total time can be used to indicate progress along the path **1530**.

[0194] The path **1530**, which is an arc, and arrow **1532** are rendered into the live image data as virtual objects in accordance with their positions in the 3-D coordinate system associated with the live 2-D image data. The cup **1522**, the circle **1530** and the **1532** arrow are shown output to display **1536**. The orientation of the curve **1530** and the arrow **1532** shown on display **1536** relative to the cup **1522** can change if the orientation of the camera is changed, such as if the camera is tilted.

[0195] In particular embodiments, a size of the object **1522** in the image data can be changed. For example, the size of the object can be made bigger or smaller by using a digital zoom. In another example, the size of the object can be made bigger or smaller by moving the camera, such as on mobile device **1534**, closer or farther away from the object **1522**.

[0196] When the size of the object changes, the distances between the tracking points can change, i.e., the pixel distances between the tracking points can increase or can decrease. The distance changes can be used to provide a scaling factor. In some embodiments, as the size of the object changes, the AR system can be configured to scale a size of the curve **1530** and/or arrow **1532**. Thus, a size of the curve relative to the object can be maintained.

[0197] In another embodiment, a size of the curve can remain fixed. For example, a diameter of the curve can be

related to a pixel height or width of the image, such as 150 percent of the pixel height or width. Thus, the object **1522** can appear to grow or shrink as a zoom is used or a position of the camera is changed. However, the size of curve **1530** in the image can remain relatively fixed.

[0198] FIG. 16 illustrates a second example of generating an Augmented Reality (AR) image capture track for capturing images used in a MVIDMR on a mobile device. FIG. 16 includes a mobile device at three times **1600a**, **1600b** and **1600c**. The device can include at least one camera, a display, an IMU, a processor (CPU), memory, microphone, audio output devices, communication interfaces, a power supply, graphic processor (GPU), graphical memory and combinations thereof. The display is shown with images at three times **1606a**, **1606b** and **1606c**. The display can be overlaid with a touch screen.

[0199] In **1606a**, an image of an object **1608** is output to the display in state **1606a**. The object is a rectangular box. The image data output to the display can be live image data from a camera on the mobile device. The camera could also be a remote camera.

[0200] In some embodiments, a target, such as **1610**, can be rendered to the display. The target can be combined with the live image data to create a synthetic image. Via the input interface on the phone, a user may be able to adjust a position of the target on the display. The target can be placed on an object and then an additional input can be made to select the object. For example, the touch screen can be tapped at the location of the target.

[0201] In another embodiment, object recognition can be applied to the live image data. Various markers can be rendered to the display, which indicate the position of the identified objects in the live image data. To select an object, the touchscreen can be tapped at a location of one of markers appearing in the image or another input device can be used to select the recognized object.

[0202] After an object is selected, a number of initial tracking points can be identified on the object, such as **1612**, **1614** and **1616**. In some embodiments, the tracking points may not appear on the display. In another embodiment, the tracking points may be rendered to the display. In some embodiments, if the tracking point is not located on the object of interest, the user may be able to select the tracking point and delete it or move it so that the tracking point lies on the object.

[0203] Next, an orientation of the mobile device can change. The orientation can include a rotation through one or more angles and translational motion as shown in **1604**. The orientation change and current orientation of the device can be captured via the IMU data from IMU **1602** on the device.

[0204] As the orientation of the device is changed, one or more of the tracking points, such as **1612**, **1614** and **1616**, can be occluded. In addition, the shape of surfaces currently appearing in the image can change. Based on changes between frames, movement at various pixel locations can be determined. Using the IMU data and the determined movement at the various pixel locations, surfaces associated with the object **1608** can be predicted. The new surfaces can be appearing in the image as the position of the camera changes. New tracking points can be added to these surfaces.

[0205] As described above, the mobile device can be used to capture images used in a MVIDMR. To aid in the capture, the live image data can be augmented with a track or other

guides to help the user move the mobile device correctly. The track can include indicators that provide feedback to a user while images associated with a MVIDMR are being recorded. In **1606c**, the live image data is augmented with a path **1622**. The beginning and end of the path is indicated by the text, “start” and “finish.” The distance along the path is indicated by shaded region **1618**.

[0206] The circle with the arrow **1620** is used to indicate a location on the path. In some embodiments, the position of the arrow relative to the path can change. For example, the arrow can move above or below the path or point in a direction which is not aligned with the path. The arrow can be rendered in this way when it is determined the orientation of the camera relative to the object or position of the camera diverges from a path that is desirable for generating the MVIDMR. Colors or other indicators can be used to indicate the status. For example, the arrow and/or circle can be rendered green when the mobile device is properly following the path and red when the position/orientation of the camera relative to the object is less than optimal.

[0207] FIGS. **17A** and **17B** illustrate yet another example of generating an Augmented Reality (AR) image capture track including status indicators for capturing images used in a MVIDMR. The synthetic image generated by the AR system can consist of live image data from a camera augmented with one or more virtual objects. For example, as described above, the live image data can be from a camera on a mobile device.

[0208] In FIG. **17A**, an object **1700a**, which is a statue, is shown in an image **1715** from a camera at a first position and orientation. The object **1700a** can be selected via the cross hairs **1704a**. Once the cross hairs are placed on the object and the object is selected, the cross hairs can move and remain on the object as the object **1700a** moves in the image data. As described above, as an object’s position/orientation changes in an image, a location to place the cross hairs in an image can be determined. In some embodiments, the position of the cross hairs can be determined via tracking the movements of points in the image, i.e., the tracking points.

[0209] In particular embodiments, if another object is moved in front of a tracked object, it may not be possible to associate the target **1704a** with the object. For example, if a person moves in front of the camera, a hand is passed in front of the camera or the camera is moved so the object no longer appears in the camera field of view, then the object which is being tracked will no longer be visible. Hence, it may not be possible to determine a location for the target associated with the tracked object. In the instance where the object reappears in the image, such as if a person that blocked the view of the object moved into and out of the view, then the system can be configured to reacquire the tracking points and reposition the target.

[0210] A first virtual object is rendered as indicator **1702a**. Indicator **1702a** can be used to indicate the progress in capturing images for a MVIDMR. A second virtual object is rendered as curve **1710**. Third and fourth virtual objects are rendered as lines **1706** and **1708**. A fifth virtual object is rendered as curve **1712**.

[0211] The curve **1710** can be used to depict a path of a camera. Whereas lines **1706** and **1708** and curve **1712** can be used to indicate an angle range for the MVIDMR. In this example, the angle range is about ninety degrees.

[0212] In FIG. **17B**, the position of the camera is different as compared to FIG. **17A**. Hence, a different view of object

1700b is presented in image **1725**. In particular, the camera view shows more of the front of the object as compared to the view in FIG. **17A**. The target **1704b** is still affixed to the object **1700b**. However, the target is fixed in a different location on the object, i.e., on a front surface as opposed to an arm.

[0213] The curve **1716** with arrow **1720** at the end is used to indicate the progress of the image capture along curve **1710**. The circle **1718** around the arrow **1720** further highlights the current position of the arrow. As described above, a position and a direction of the arrow **1720** can be used to provide feedback to a user on a deviation of the camera position and/or orientation from curve **1710**. Based upon this information, the user may adjust a position and/or orientation of the camera while it is capturing the image data.

[0214] Lines **1706** and **1708** still appear in the image but are positioned differently relative to object **1700b**. The lines again indicate an angle range. In **1720**, the arrow is about half way between lines **1706** and **1708**. Hence, an angle of about 45 degrees has been captured around the object **1700b**.

[0215] The indicator **1702b** now includes a shaded region **1722**. The shaded region can indicate a portion of a MVIDMR angle range currently captured. In some embodiments, lines **1706** and **1708** can only indicate a portion of the angle range in a MVIDMR that is being captured and the total angle range can be shown via indicator **1702b**. In this example, the angle range shown by indicator **1702b** is three hundred sixty degrees while lines **1706** and **1708** show a portion of this range which ninety degrees.

[0216] With reference to FIG. **18**, shown is a particular example of a computer system that can be used to implement particular examples. For instance, the computer system **1800** can be used to provide MVIDMRs according to various embodiments described above. According to various embodiments, a system **1800** suitable for implementing particular embodiments includes a processor **1801**, a memory **1803**, an interface **1811**, and a bus **1815** (e.g., a PCI bus). The system **1800** can include one or more sensors, such as light sensors, accelerometers, gyroscopes, microphones, cameras including stereoscopic or structured light cameras. As described above, the accelerometers and gyroscopes may be incorporated in an IMU. The sensors can be used to detect movement of a device and determine a position of the device. Further, the sensors can be used to provide inputs into the system. For example, a microphone can be used to detect a sound or input a voice command.

[0217] In the instance of the sensors including one or more cameras, the camera system can be configured to output native video data as a live video feed. The live video feed can be augmented and then output to a display, such as a display on a mobile device. The native video can include a series of frames as a function of time. The frame rate is often described as frames per second (fps). Each video frame can be an array of pixels with color or gray scale values for each pixel. For example, a pixel array size can be 512 by 512 pixels with three color values (red, green and blue) per pixel. The three-color values can be represented by varying amounts of bits, such as 24, 30, 36, 40 bits, etc. per pixel. When more bits are assigned to representing the RGB color values for each pixel, a larger number of colors values are possible. However, the data associated with each image also increases. The number of possible colors can be referred to as the color depth.

[0218] The video frames in the live video feed can be communicated to an image processing system that includes hardware and software components. The image processing system can include non-persistent memory, such as random-access memory (RAM) and video RAM (VRAM). In addition, processors, such as central processing units (CPUs) and graphical processing units (GPUs) for operating on video data and communication busses and interfaces for transporting video data can be provided. Further, hardware and/or software for performing transformations on the video data in a live video feed can be provided.

[0219] In particular embodiments, the video transformation components can include specialized hardware elements configured to perform functions necessary to generate a synthetic image derived from the native video data and then augmented with virtual data. In data encryption, specialized hardware elements can be used to perform a specific data transformation, i.e., data encryption associated with a specific algorithm. In a similar manner, specialized hardware elements can be provided to perform all or a portion of a specific video data transformation. These video transformation components can be separate from the GPU(s), which are specialized hardware elements configured to perform graphical operations. All or a portion of the specific transformation on a video frame can also be performed using software executed by the CPU.

[0220] The processing system can be configured to receive a video frame with first RGB values at each pixel location and apply operation to determine second RGB values at each pixel location. The second RGB values can be associated with a transformed video frame which includes synthetic data. After the synthetic image is generated, the native video frame and/or the synthetic image can be sent to a persistent memory, such as a flash memory or a hard drive, for storage. In addition, the synthetic image and/or native video data can be sent to a frame buffer for output on a display or displays associated with an output interface. For example, the display can be the display on a mobile device or a view finder on a camera.

[0221] In general, the video transformations used to generate synthetic images can be applied to the native video data at its native resolution or at a different resolution. For example, the native video data can be a 512 by 512 array with RGB values represented by 24 bits and at frame rate of 24 fps. In some embodiments, the video transformation can involve operating on the video data in its native resolution and outputting the transformed video data at the native frame rate at its native resolution.

[0222] In other embodiments, to speed up the process, the video transformations may involve operating on video data and outputting transformed video data at resolutions, color depths and/or frame rates different than the native resolutions. For example, the native video data can be at a first video frame rate, such as 24 fps. But, the video transformations can be performed on every other frame and synthetic images can be output at a frame rate of 12 fps. Alternatively, the transformed video data can be interpolated from the 12 fps rate to 24 fps rate by interpolating between two of the transformed video frames.

[0223] In another example, prior to performing the video transformations, the resolution of the native video data can be reduced. For example, when the native resolution is 512 by 512 pixels, it can be interpolated to a 256 by 256 pixel array using a method such as pixel averaging and then the

transformation can be applied to the 256 by 256 array. The transformed video data can output and/or stored at the lower 256 by 256 resolution. Alternatively, the transformed video data, such as with a 256 by 256 resolution, can be interpolated to a higher resolution, such as its native resolution of 512 by 512, prior to output to the display and/or storage. The coarsening of the native video data prior to applying the video transformation can be used alone or in conjunction with a coarser frame rate.

[0224] As mentioned above, the native video data can also have a color depth. The color depth can also be coarsened prior to applying the transformations to the video data. For example, the color depth might be reduced from 40 bits to 24 bits prior to applying the transformation.

[0225] As described above, native video data from a live video can be augmented with virtual data to create synthetic images and then output in real-time. In particular embodiments, real-time can be associated with a certain amount of latency, i.e., the time between when the native video data is captured and the time when the synthetic images including portions of the native video data and virtual data are output. In particular, the latency can be less than 100 milliseconds. In other embodiments, the latency can be less than 50 milliseconds. In other embodiments, the latency can be less than 30 milliseconds. In yet other embodiments, the latency can be less than 20 milliseconds. In yet other embodiments, the latency can be less than 10 milliseconds.

[0226] The interface **1811** may include separate input and output interfaces, or may be a unified interface supporting both operations. Examples of input and output interfaces can include displays, audio devices, cameras, touch screens, buttons and microphones. When acting under the control of appropriate software or firmware, the processor **1801** is responsible for such tasks such as optimization. Various specially configured devices can also be used in place of a processor **1801** or in addition to processor **1801**, such as graphical processor units (GPUs). The complete implementation can also be done in custom hardware. The interface **1811** is typically configured to send and receive data packets or data segments over a network via one or more communication interfaces, such as wireless or wired communication interfaces. Particular examples of interfaces the device supports include Ethernet interfaces, frame relay interfaces, cable interfaces, DSL interfaces, token ring interfaces, and the like.

[0227] In addition, various very high-speed interfaces may be provided such as fast Ethernet interfaces, Gigabit Ethernet interfaces, ATM interfaces, HSSI interfaces, POS interfaces, FDDI interfaces and the like. Generally, these interfaces may include ports appropriate for communication with the appropriate media. In some cases, they may also include an independent processor and, in some instances, volatile RAM. The independent processors may control such communications intensive tasks as packet switching, media control and management.

[0228] According to various embodiments, the system **1800** uses memory **1803** to store data and program instructions and maintained a local side cache. The program instructions may control the operation of an operating system and/or one or more applications, for example. The memory or memories may also be configured to store received metadata and batch requested metadata.

[0229] The system **1800** can be integrated into a single device with a common housing. For example, system **1800**

can include a camera system, processing system, frame buffer, persistent memory, output interface, input interface and communication interface. In various embodiments, the single device can be a mobile device like a smart phone, an augmented reality and wearable device like Google Glass™ or a virtual reality head set that includes a multiple cameras, like a Microsoft HoloLens™. In other embodiments, the system 1800 can be partially integrated. For example, the camera system can be a remote camera system. As another example, the display can be separate from the rest of the components like on a desktop PC.

[0230] In the case of a wearable system, like a head-mounted display, as described above, a virtual guide can be provided to help a user record a MVIDMR. In addition, a virtual guide can be provided to help teach a user how to view a MVIDMR in the wearable system. For example, the virtual guide can be provided in synthetic images output to head mounted display which indicate that the MVIDMR can be viewed from different angles in response to the user moving some manner in physical space, such as walking around the projected image. As another example, the virtual guide can be used to indicate a head motion of the user can allow for different viewing functions. In yet another example, a virtual guide might indicate a path that a hand could travel in front of the display to instantiate different viewing functions.

[0231] FIG. 19 illustrates a method 1900 of determining an orientation of an object, performed in accordance with one or more embodiments. The method 1900 may be performed on any suitable computing device. For example, the method 1900 may be performed on a mobile computing device such as a smartphone. Alternately or additionally, the method 1900 may be performed on a remote server in communication with a mobile computing device.

[0232] A request to determine an orientation of an object is received at 1902. According to various embodiments, the request may be received at a user interface. For example, the request may be received via an application at a mobile computing device. As another example, the request may be received at a user interface at a laptop.

[0233] At 1904, an image of the object is selected for analysis. According to various embodiments, the image may be a single two-dimensional image of the object. Alternately or additionally, the image may be one of a series of two-dimensional images of the object captured from different viewpoints. In some implementations, the image may be a single frame of a video stream of the object. In some embodiments, the image may be a 360-degree multi-view capture of the object. Alternately, the image may include a view that has less than 360-degree coverage.

[0234] Optionally, a bounding box for the object is determined at 1906. According to various embodiments, a bounding box for the object may include coordinates of a rectangular or non-rectangular border that encloses the object in the selected image. The bounding box may be determined by a neural network. For example, a neural network may be trained to determine the coordinates of the visible object pixels in an image of the object. The neural network may facilitate the determination of the pixel coordinates and the width, height, or other characteristics of the bounding box enclosing the object. Alternately or additionally, the bounding box may be determined by a different type of algorithm. In some configurations, the bounding box may be determined at least in part based on user input. For example, the

bounding box may be determined by a user super-imposing a border on the image to enclose the object. In another example, the bounding box may be determined by the user tapping on the object with an algorithm automatically estimating the extent of the bounding based on the selection. In yet another example, the user may select points or regions in the foreground and background, and an algorithm automatically separates both to estimate the extent of a corresponding bounding box.

[0235] At 1908, an estimated object orientation is determined. According to various embodiments, the estimated object orientation may be determined by a neural network. For example, a neural network may be trained to determine the estimated object orientation through an analysis of the object from one or more images of an object. The estimated object orientation may include estimates about an object's roll, elevation, angular position, attitude, and azimuth angle.

[0236] An estimated camera orientation is identified at 1910. According to various embodiments, camera orientation may be estimated from data collected from an inertial measurement unit (IMU). In some implementations, the IMU may be associated with a mobile computing device, such as a smartphone. For example, a mobile computing device may include sensors such as a camera capable of capturing visual data such as an image or video stream. A mobile computing device may also include an accelerometer, a gyroscope, and other such sensors capable of capturing IMU data. IMU data may include information such as camera location, camera angle, device velocity, device acceleration, or any of a wide variety of data collected from accelerometers or other such sensors.

[0237] In some implementations, IMU orientation may identify orientation information associated with the camera. The IMU orientation may then be used to infer orientation information about the object. For example, the IMU orientation may indicate that the image was captured with a camera located at a particular height and angle relative to the object and/or ground. Alternatively, or additionally, the orientation estimation might be done using a pose estimation algorithm based on image data. In some configurations, the pose estimation method might also consider the available IMU data.

[0238] At 1912, an offset is determined between the estimated object orientation and the estimated camera orientation. According to various embodiments, the determination may be made at least in part by identifying a difference between the two values. Alternately, or additionally, the determination may be made at least in part by refining the object orientation based on additional information characterizing the position and angle of the camera.

[0239] A determination is made at 1914 as to whether to select an additional image of the object for analysis. In some implementations, additional images may be selected until all available images have been analyzed. For example, each image associated with a multi-view capture may be analyzed. If multiple images have been processed, then a combined offset may be to be calculated. For example, the offset can be averaged from the results obtained from multiple images. In another example, a robust method for estimating the final offset from multiple computed offsets might be used, for instance to handle outliers such as an inaccurate orientation estimation for certain images.

[0240] In some embodiments, the determination made at 1914 may be made at least in part automatically. For

example, the estimated object orientation procedure at **1908** may attempt to estimate an orientation for the object, but may fail due to incomplete image information or other information. A determination may be automatically made to prompt the user to capture or upload additional images of the object.

[**0241**] In some implementations, the determination made at **1914** may be made at least in part as a request by a user. For example, a user may be prompted to select an additional image for analysis. Alternately or additionally, a user may be prompted to review an estimated object orientation, and a user may review the estimated object orientation and may elect to analyze additional images of the object in order to improve the accuracy or quality of the estimated object orientation. In some implementations, every N-th frame of the multi-view capture might be processed.

[**0242**] At **1916**, the orientation estimation information is stored. According to various embodiments, storing the orientation estimation information may involve storing one or more orientation values for each of the images analyzed in FIG. **19**. For example, orientation information may be stored to supplement a multi-view capture.

[**0243**] In some implementations, the orientation estimation may be stored on a storage device. Alternately, or additionally, the orientation estimation may be transmitted to a remote location via a network.

[**0244**] According to various embodiments, one or more operations shown in FIG. **19** may be omitted. For example, any or all of operations **1910** and **1912** may be omitted.

[**0245**] FIG. **20** illustrates a method for determining an object orientation based on 3D data. The method **2000** may be performed on any suitable computing device. For example, the method **2000** may be performed on a mobile computing device such as a smartphone. Alternately or additionally, the method **2000** may be performed on a remote server in communication with a mobile computing device.

[**0246**] A request to determine an orientation of an object is received at **2002**. In some implementations, the request may be received at a user interface. At **2004**, image information for the object is captured. In some implementations, the image information may include a video or image set of the object is captured from multiple perspectives. The image information may include a 360-degree view of the object. Alternately, the image information may include a view that has less than 360-degree coverage.

[**0247**] An object is identified for analysis at **2006**. In some embodiments, the object may be identified via an algorithm. For example, the object may be identified via an image recognition algorithm. As another example, the object may be identified via a bounding box algorithm.

[**0248**] In some implementations, the object may be identified based at least in part on user input. For instance, a user may identify the object as a vehicle via a user interface component such as a drop-down menu.

[**0249**] At **2008**, one or more points of reference are identified on the object. According to various embodiments, points of reference may include one or more fiducial points on the object. Fiducial points are identifiable locations on objects of a particular type that in general should be present on all or nearly all objects of that type and that provide an indication as to the object's orientation. For example, in the

case of a vehicle, fiducial points may include but are not limited to the vehicle's headlights, rearview mirror, and wheels.

[**0250**] In some implementations, the reference points may be identified via an algorithm, such as an image recognition algorithm. Alternately or additionally, reference points may be identified in part through user input.

[**0251**] A sparse 3D mesh of an object is determined at **2010**. According to various embodiments, the sparse 3D mesh may be determined by using the fiducial points identified in operation **2008** as vertices of the mesh. For example, each face of the mesh may approximate the object shape with a planar surface.

[**0252**] At **2012**, a 3D representation of the object is determined. According to various embodiments, the 3D representation may be constructed by elevating the fiducial points to 3D space and constructing a sparse 3D mesh from their locations. Alternately, or additionally, the object's orientation may be inferred based on cues such as surface normal of the faces of the mesh, visibility states of the fiducial points in a plurality of image, or other such features.

[**0253**] In some embodiments, the elevation into 3D space might be done using triangulation. Alternately, or additionally, a depth map might be computed for each frame which is then used to elevate the 2D points into 3D.

[**0254**] According to various embodiments, the 3D representation may be a 3D skeleton. For example, 2D skeleton detection may be performed on every frame using a machine learning procedure. As another example, 3D camera pose estimation may be performed to determine a location and angle of the camera with respect to the object for a particular frame. As yet another example, a 3D skeleton may be reconstructed from one or more 2D skeletons and/or one or more 3D models.

[**0255**] According to various embodiments, the 3D reconstruction of the object surface may include one or more of a variety of reconstruction procedures. For example, dominant axes may be computed to estimate orientation. As another example, an object model such as a vehicle model that has a reference orientation may be fit into a 3D reconstructed model based on the input data. As another example, one or more parts of a 3D reconstructed model may be classified to estimate an orientation of the object.

[**0256**] Estimated orientation for the object is determined at **2014**. According to various embodiments, the estimated object orientation may be determined by a neural network. For example, a neural network may be trained to determine the estimated object orientation through an analysis of the object from one or more images of an object. The estimated object orientation may include estimates about an object's dominant axes, roll, elevation, angular position, attitude, and azimuth angle.

[**0257**] According to various embodiments, the object orientation may be estimated based on the sparse 3D mesh determined at operation **2010** and/or the 3D skeleton determined at operation **2012**. For example, the fiducial points identified in operation **2008** facilitate the triangulation of fiducial points of the object in 3D space to help determine the azimuth, roll, elevation, and axes of the designated object. The 3D skeleton may facilitate the inference of the object's orientation based on cues such as, but not limited to, the surface normals of the faces of the mesh, and the visibility states of the fiducial points in the images of the

object, and information about the three axes of the 3D skeleton model determined at operation 2012.

[0258] In some embodiments, the estimated object orientation may be determined based on a 3D reconstruction of the object. For example, the object orientation may be estimated based on the dominant axes of the 3D reconstruction of an object. As another example, the orientation for a designated object may be estimated by incorporating the orientation of a known reference object with reference orientation with a 3D reconstructed model of the designated object. As yet another example, parts of the 3D reconstructed model may be identified to determine orientation of the object. For example, in the case of a vehicle, the license plate may be identified as a reference point relative to the rearview mirror to estimate the orientation of the vehicle.

[0259] At 2016, the orientation estimation information is stored. In some implementations, the orientation estimation may be stored on a storage device. Alternately, or additionally, the orientation estimation may be transmitted to a remote location via a network.

[0260] According to various embodiments, the object orientation may be estimated based on the determination of specific reference points on the object. For example, in the case of a vehicle, a specific reference point such as the license plate may be detected in one or more images of a multi-view image set. Based on the location of the license plate in the image relative to other reference points, such as the vehicle's headlights, the orientation of the object may be inferred.

[0261] In some implementations, this determination may be performed on a 3D model of the object. For example, for a 3D model of a vehicle, the location of the license plate may be used to estimate global angle for the vehicle. Alternately, or additionally, the trajectory of a feature such as a license plate may be fit to an ellipse in image space, which corresponds to a sphere around the vehicle in 3D space, and the location on the ellipse is used to estimate the vehicle's orientation.

[0262] In some embodiments, the estimated object orientation may be determined from a single image. In some implementations, the object may not be visible from all angles or perspectives from a single image, or an image may not supply a minimum number of reference points visible to determine a sparse 3D mesh capable of facilitating the determination of a 3D skeleton of the object. In such situations, a neural network may be trained to predict coordinates of the visible object pixels in a top-down view of the object.

[0263] In particular embodiments, once the locations are mapped, the points may be lifted to 3D based on a predefined correspondence between the top-down view and a 3D mesh. Then, the transformation between the image points and the 3D mesh may be used to obtain the 3D orientation of the vehicle.

[0264] According to various embodiments, the estimated object orientation may be determined from a single image via top-down view mapping. For example, a network may be trained that takes an image of an object such as a vehicle as input and maps every point of the image onto a top-down view. Based on the distribution of the points on the top-down view, the system may then draw an inference about the orientation of the object with respect to the camera.

[0265] According to various embodiments, orientation may be estimated from a single image using a component

classifier. For example, a network may be trained to identify object components such as car parts. For instance, each pixel of the image may be classified as belonging to an object component or as part of the background. The visible object components may then be used to estimate an orientation of the object.

[0266] According to various embodiments, the object orientation may be directly estimated from the detected 2D skeleton by considering which joints of the 2D skeleton are detected (i.e. are visible).

[0267] In some implementations, the estimated object orientation may be determined from a single image via using a neural network trained directly to determine an N degree-of-freedom pose. For example, a 3 degree-of-freedom pose may include azimuth, roll, and elevation. As another example, a 6 degree-of-freedom pose may include 3 degrees of freedom for orientation and 3 degrees of freedom for translation. As another example, a combination of 3 degrees of freedom for translation and 3 degrees of freedom for position may be used. For instance, a 2 degree-of-freedom pose may include 1 degree-of-freedom for orientation and 1 degree-of-freedom for the distance between the camera and the object.

[0268] In particular embodiments, any of the approaches involving a single image may be applied on a multi-view data set and the results combined (e.g., averaged) to improve the overall results. In general, any combination of the methods described herein may be used to increase robustness and/or provide a failsafe. For example, orientation may be independent estimated using three different methods, and then the results cross-checked.

[0269] While the present disclosure has been particularly shown and described with reference to specific embodiments thereof, it will be understood by those skilled in the art that changes in the form and details of the disclosed embodiments may be made without departing from the spirit or scope of the invention. Specifically, there are many alternative ways of implementing the processes, systems, and apparatuses described. It is therefore intended that the invention be interpreted to include all variations and equivalents that fall within the true spirit and scope of the present invention.

1. A method comprising:

retrieving from a storage medium a plurality of configuration parameters associated with the generation of a cinematic video, the configuration parameters identifying a plurality of object components, the configuration parameters further identifying an order in which to display the object components, the configuration parameters further identifying an object type;

identifying a multi-view representation of an object corresponding with the identified object type, the multi-view representation including a first plurality of images of the object captured from different viewpoints, the multi-view representation being navigable in one or more dimensions;

generating via a processor a cinematic video of the object based on a second plurality of images selected from the first plurality of images and arranged in an order based on the retrieved configuration parameters; and

storing the cinematic video on the storage medium.

2. The method recited in claim 1, the method further comprising:

analyzing the first plurality of images to identify a component of the object;

determining a tag that includes text that describes the identified component, wherein the cinematic video includes one or more video frames in which the tag is superimposed on the object.

3. The method recited in claim 1, wherein identifying the multi-view representation comprises providing recording guidance via a user interface to facilitate the capture of one or more of the first plurality of images.

4. The method recited in claim 3, wherein the recording guidance is provided via an overlay presented on a live camera feed at a smartphone.

5. The method recited in claim 3, wherein the recording guidance is determined based on the retrieved configuration parameters.

6. The method recited in claim 1, wherein generating the cinematic video comprises applying a slow-motion effect to a subset of the second plurality of images, the slow-motion effect being defined by the configuration parameters.

7. The method recited in claim 1, wherein a subset of the second plurality of images include closeup views of a component of the object, and wherein generating the cinematic video comprises shifting from a comprehensive view of the object to the second plurality of images.

8. The method recited in claim 1, wherein the object is a vehicle, and wherein the cinematic video includes a designated portion in which a viewpoint pans around the vehicle.

9. The method recited in claim 1, wherein the multi-view representation includes a 3D model of the object determined based on analyzing the first plurality of images.

10. The method recited in claim 9, wherein determining the 3D model of the object comprises applying a neural network to one or more of the first plurality of images.

11. The method recited in claim 1, wherein the multi-view representation of the object is generated on a mobile computing device comprising a camera, and wherein each of the plurality of images of the object are captured by the camera.

12. The method recited in claim 11, wherein the mobile computing device includes an inertial measurement unit configured to capture inertial measurement data, and wherein determining the representation of the object involves analyzing the inertial measurement data.

13. A computing device comprising:

a storage medium operable to store a plurality of configuration parameters associated with the generation of a cinematic video, the configuration parameters identifying a plurality of object components, the configuration parameters further identifying an order in which to display the object components, the configuration parameters further identifying an object type;

a processor operable to:

identify a multi-view representation of an object corresponding with the identified object type, the multi-view representation including a first plurality of images of the object captured from different viewpoints, the multi-view representation being navigable in one or more dimensions, and

generate via a processor a cinematic video of the object based on a second plurality of images selected from

the first plurality of images and arranged in an order based on the configuration parameters; and

a display screen operable to present the cinematic video.

14. The computing device recited in claim 13, wherein the processor is further operable to:

analyze the first plurality of images to identify a component of the object;

determine a tag that includes text that describes the identified component, wherein the cinematic video includes one or more video frames in which the tag is superimposed on the object.

15. The computing device recited in claim 13, wherein identifying the multi-view representation comprises providing recording guidance via a user interface to facilitate the capture of one or more of the first plurality of images, wherein the recording guidance is provided via an overlay presented on a live camera feed at a smartphone, and wherein the recording guidance is determined based on the configuration parameters.

16. The computing device recited in claim 13, wherein generating the cinematic video comprises applying a slow-motion effect to a subset of the second plurality of images, the slow-motion effect being defined by the configuration parameters.

17. The computing device recited in claim 13, wherein a subset of the second plurality of images include closeup views of a component of the object, and wherein generating the cinematic video comprises shifting from a comprehensive view of the object to the second plurality of images.

18. The computing device recited in claim 13, wherein the multi-view representation includes a 3D model of the object determined based on analyzing the first plurality of images.

19. The computing device recited in claim 18, wherein determining the 3D model of the object comprises applying a neural network to one or more of the first plurality of images.

20. One or more non-transitory machine-readable media having instructions stored thereon for performing a method, the method comprising:

retrieving from a storage medium a plurality of configuration parameters associated with the generation of a cinematic video, the configuration parameters identifying a plurality of object components, the configuration parameters further identifying an order in which to display the object components, the configuration parameters further identifying an object type;

identifying a multi-view representation of an object corresponding with the identified object type, the multi-view representation including a first plurality of images of the object captured from different viewpoints, the multi-view representation being navigable in one or more dimensions;

generating via a processor a cinematic video of the object based on a second plurality of images selected from the first plurality of images and arranged in an order based on the retrieved configuration parameters; and

storing the cinematic video on the storage medium.

* * * * *