



- (51) **International Patent Classification:**
G06K 9/00 (2006.01) **G06K 9/62** (2006.01)
- (21) **International Application Number:**
PCT/US2018/027684
- (22) **International Filing Date:**
14 April 2018 (14.04.2018)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/485,983 16 April 2017 (16.04.2017) US
15/953,107 13 April 2018 (13.04.2018) US
- (71) **Applicant: FACEBOOK, INC.** [US/US]; 1601 Willow Road, Menlo Park, California 94025 (US).
- (72) **Inventors: KUZYAKOV, Evgeny V.**; c/o Facebook, Inc., 1601 Willow Road, Menlo Park, California 94025 (US). **PENG, Renbin**; c/o Facebook, Inc., 1601 Willow Road, Menlo Park, California 94025 (US). **CHEN, Chien-Nan**; c/o Facebook, Inc., 1601 Willow Road, Menlo Park, California 94025 (US).

(74) **Agent: SOONG, James W.** et al.; Sheppard, Mullin, Richter & Hampton LLP, 379 Lytton Avenue, Palo Alto, California 94301 (US).

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

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

(54) **Title:** SYSTEMS AND METHODS FOR PROVISIONING CONTENT




FIGURE 1

(57) **Abstract:** Systems, methods, and non-transitory computer-readable media can provide at least one frame of a content item to a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items. Information describing at least a first salient point of interest that appears in the at least one frame can be obtained from the saliency prediction model. The first salient point of interest can be predicted to be of interest to users accessing the content item. A view-based projection can be applied to a region corresponding to the first salient point of interest, wherein the view-based projection enhances a quality in which the region is presented.

Published:

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

SYSTEMS AND METHODS FOR PROVISIONING CONTENT

FIELD OF THE INVENTION

[0001] The present technology relates to the field of content provisioning. More particularly, the present technology relates to techniques for evaluating content to be presented through computing devices.

BACKGROUND

[0002] Today, people often utilize computing devices (or systems) for a wide variety of purposes. Users can operate their computing devices to, for example, interact with one another, create content, share content, and access information. Under conventional approaches, content items (e.g., images, videos, audio files, etc.) can be made available through a content sharing platform. Users can operate their computing devices to access the content items through the platform. Typically, the content items can be provided, or uploaded, by various entities including, for example, content publishers and also users of the content sharing platform.

SUMMARY

[0003] Various embodiments of the present disclosure can include systems, methods, and non-transitory computer readable media are configured to provide at least one frame of a content item to a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items; obtain information describing at least a first salient point of interest that appears in the at least one frame from the saliency prediction model, wherein the first salient point of interest is predicted to be of interest to users accessing the content item; and apply a view-based projection to a region corresponding to the first salient point of interest, wherein the view-based projection enhances a quality in which the region is presented.

[0004] In some embodiments, the systems, methods, and non-transitory computer readable media are configured to obtain a saliency map for the at least one frame; determining a vector-based representation of the saliency map; determine an offset corresponding to the at least one frame based at least in part on the vector-based representation; and enhance the region corresponding to the first salient point of interest based at least in part on the determined offset.

[0005] In some embodiments, the vector-based representation includes a set of yaw bins and a set of pitch bins, and wherein each bin is associated with a pre-defined vector and a corresponding magnitude.

[0006] In some embodiments, the systems, methods, and non-transitory computer readable media are configured to determine a spherical representation of the at least one frame, the spherical representation including a virtual camera positioned at a center of the spherical representation; determine a displacement for the virtual camera based at least in part on the offset; and determine an amount of pixels to allocate for the region based at least in part on the displacement of the virtual camera.

[0007] In some embodiments, the region is enhanced by increasing a pixel density corresponding to the region.

[0008] In some embodiments, respective pixel densities for one or more other regions of the at least one frame are decreased in proportion to the displacement of the virtual camera.

[0009] In some embodiments, a total pixel count associated with the at least one frame remains unchanged.

[0010] Various embodiments of the present disclosure can include systems, methods, and non-transitory computer readable media are configured to determine at least one salient point of interest in a frame of a content item based at least in part on a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items; determine a barrel projection representation for the frame; and apply a view-based projection to the barrel projection representation for the frame, wherein the view-based projection enhances a quality in which a region corresponding to the at least one salient point of interest is presented.

[0011] In some embodiments, the barrel projection includes separate faces corresponding to a top portion of the frame, a bottom portion of the frame, and a middle portion of the frame.

[0012] In some embodiments, the middle portion of the frame represents a middle 90 degrees of a scene represented in the frame.

[0013] In some embodiments, the systems, methods, and non-transitory computer readable media are configured to bias a region in the middle portion of the frame that corresponds to the at least one salient point of interest, wherein the region is biased to increase a pixel density associated with the region.

[0014] In some embodiments, the region is stretched horizontally by a threshold amount.

[0015] In some embodiments, the systems, methods, and non-transitory computer readable media are configured to bias one or more other regions in the middle portion of the frame, wherein the one or more regions are biased to decrease respective pixel densities associated with the one or more regions.

[0016] In some embodiments, the one or more regions are shortened horizontally by a threshold amount.

[0017] It should be appreciated that many other features, applications, embodiments, and/or variations of the disclosed technology will be apparent from the accompanying drawings and from the following detailed description. Additional and/or alternative implementations of the structures, systems, non-transitory

computer readable media, and methods described herein can be employed without departing from the principles of the disclosed technology.

[0018] Embodiments according to the invention are in particular disclosed in the attached claims directed to a method, a system, and a storage medium, wherein any feature mentioned in one claim category, e.g. method, can be claimed in another claim category, e.g. system, storage medium, computer program product as well. The dependencies or references back in the attached claims are chosen for formal reasons only. However any subject matter resulting from a deliberate reference back to any previous claims (in particular multiple dependencies) can be claimed as well, so that any combination of claims and the features thereof is disclosed and can be claimed regardless of the dependencies chosen in the attached claims. The subject-matter which can be claimed comprises not only the combinations of features as set out in the attached claims but also any other combination of features in the claims, wherein each feature mentioned in the claims can be combined with any other feature or combination of other features in the claims. Furthermore, any of the embodiments and features described or depicted herein can be claimed in a separate claim and/or in any combination with any embodiment or feature described or depicted herein or with any of the features of the attached claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0019] **FIGURE 1** illustrates an example system including an example content provider module, according to an embodiment of the present disclosure.

[0020] **FIGURE 2** illustrates an example of a content features module, according to an embodiment of the present disclosure.

[0021] **FIGURE 3** illustrates an example of a view-based projection module, according to an embodiment of the present disclosure.

[0022] **FIGURES 4A-4H** illustrate examples diagrams, according to an embodiment of the present disclosure.

[0023] **FIGURES 5A-5B** illustrate example methods, according to an embodiment of the present disclosure.

[0024] **FIGURE 6** illustrates a network diagram of an example system including an example social networking system that can be utilized in various scenarios, according to an embodiment of the present disclosure.

[0025] **FIGURE 7** illustrates an example of a computer system or computing device that can be utilized in various scenarios, according to an embodiment of the present disclosure.

[0026] The figures depict various embodiments of the disclosed technology for purposes of illustration only, wherein the figures use like reference numerals to identify like elements. One skilled in the art will readily recognize from the following discussion that alternative embodiments of the structures and methods illustrated in the figures can be employed without departing from the principles of the disclosed technology described herein.

DETAILED DESCRIPTION

APPROACHES FOR PROVISIONING CONTENT

[0027] People use computing devices (or systems) for a wide variety of purposes. As mentioned, under conventional approaches, a user can utilize a computing device to share content items (e.g., documents, images, videos, audio, etc.) with other users. Such content items can be made available through a content sharing platform. Users can operate their computing devices to access the content items through the platform. Typically, the content items can be provided, or uploaded, by various entities including, for example, content publishers and also users of the content sharing platform.

[0028] In some instances, a user can access virtual reality content through a content provider. Such virtual reality content can be presented, for example, in a viewport that is accessible through a computing device (e.g., a virtual reality device, headset, or any computing device capable of presenting virtual reality content). In general, a virtual reality content item (or immersive video) corresponds to any virtual reality media that encompasses (or surrounds) a viewer (or user). Some examples of virtual reality content items include spherical videos, half sphere videos (e.g., 180 degree videos), arbitrary partial spheres, 225 degree videos, and 3D 360 videos. Such virtual reality content items need not be limited to videos that are formatted using a spherical shape but may also be applied to immersive videos formatted using other shapes including, for example, cubes, pyramids, and other shape representations of a video recorded three-dimensional world. In some embodiments, a virtual reality content item can be created by stitching together various video streams (or feeds) that were captured by cameras that are placed at particular locations and/or positions to capture a view of the scene (e.g., 180 degree view, 225 degree view, 360 degree view, etc.). Once stitched together, a user can access, or present (e.g., playback), the virtual reality content item. Generally, while accessing the virtual reality content item, the user can zoom and change the direction (e.g., pitch, yaw, roll) of the viewport to access different portions of the scene in the virtual reality content item. The direction of the viewport can be used to determine which stream of the virtual reality content item is presented. In general, a content item (e.g., virtual reality content item, immersive video, spherical video, etc.) may capture scenes that include various points of interest (e.g., persons, objects, landscapes,

etc.). In some instances, conventional models (e.g., neural network) can be trained to evaluate the content item to identify points of interest appearing in scenes (e.g., frames) during presentation (e.g., playback) of the content item. Although conventional approaches can be used to identify a number of different points of interest in a given content item, these approaches are typically unable to indicate which of these identified points of interest are likely to be relevant (or interesting) to a given user or a group of users. Further, conventional approaches may lack the ability to emphasize relevant (or interesting) content in content items. Accordingly, such conventional approaches may not be effective in addressing these and other problems arising in computer technology.

[0029] An improved approach overcomes the foregoing and other disadvantages associated with conventional approaches. In various embodiments, a saliency prediction model can be trained to identify content that is likely to be of interest to users (e.g., salient points of interest) during presentation of a given content item. In some embodiments, the content predicted by the saliency prediction model is expected to be more relevant, or interesting, to a given user or group of users (e.g., users sharing one or more demographic attributes). In some embodiments, these salient points of interest can be used to improve the delivery (or streaming) of the content item. For example, in some embodiments, saliency information outputted by the saliency prediction model can be used to implement dynamic streaming. For example, in dynamic streaming, a content item may be associated with a number of different streams in which different parts of frames are enhanced (or emphasized). In this example, the stream to be presented depends on a view direction of a viewer's computing device (e.g., virtual reality device). In some embodiments, the saliency information can be used to implement content-dependent streaming. In such embodiments, regions in frames that include salient points of interest can be emphasized and distributed as a single stream. In various embodiments, enhancement (or emphasis) of content can be achieved using view-based projection. For example, in some embodiments, the saliency information can be used to implement a view-based projection that emphasizes one view (e.g., a most relevant, interesting view) among many potential views in a content item (or frame) without changing (e.g., increasing) a total pixel count associated with the content item (or frame). As a result, users can enjoy an immersive and interactive

virtual experience that is visually pleasing without having to experience drawbacks that may result from increased resource requirements (e.g., processing, bandwidth, etc.). More details relating to the disclosed technology are provided below.

[0030] **FIGURE 1** illustrates an example system 100 including an example content provider module 102, according to an embodiment of the present disclosure. As shown in the example of **FIGURE 1**, the content provider module 102 can include a content module 104, a streaming module 106, and a content features module 108. In some instances, the example system 100 can include at least one data store 112. A client module 114 can interact with the content provider module 102 over one or more networks 150 (e.g., the Internet, a local area network, etc.). The client module 114 can be implemented in a software application running on a computing device (e.g., a virtual reality device, headset, or any computing device capable of presenting virtual reality content). In various embodiments, the network 150 can be any wired or wireless computer network through which devices can exchange data. For example, the network 150 can be a personal area network, a local area network, or a wide area network, to name some examples. The components (e.g., modules, elements, etc.) shown in this figure and all figures herein are exemplary only, and other implementations may include additional, fewer, integrated, or different components. Some components may not be shown so as not to obscure relevant details.

[0031] In some embodiments, the content provider module 102 can be implemented, in part or in whole, as software, hardware, or any combination thereof. In general, a module, as discussed herein, can be associated with software, hardware, or any combination thereof. In some implementations, one or more functions, tasks, and/or operations of modules can be carried out or performed by software routines, software processes, hardware, and/or any combination thereof. In some cases, the content provider module 102 can be implemented, in part or in whole, as software running on one or more computing devices or systems, such as on a user computing device or client computing system. For example, the content provider module 102, or at least a portion thereof, can be implemented as or within an application (e.g., app), a program, or an applet, etc., running on a user computing device or a client computing system, such as the user device 610 of **FIGURE 6**. Further, the content provider module 102, or at least a portion thereof, can be implemented using one or more computing devices or systems that include one or

more servers, such as network servers or cloud servers. In some instances, the content provider module 102 can, in part or in whole, be implemented within or configured to operate in conjunction with a social networking system (or service), such as the social networking system 630 of **FIGURE 6**. It should be understood that there can be many variations or other possibilities.

[0032] In some embodiments, the content provider module 102 can be configured to communicate and/or operate with the at least one data store 112 in the example system 100. In various embodiments, the at least one data store 112 can store data relevant to the function and operation of the content provider module 102. One example of such data can be content items (e.g., virtual reality content items) that are available for access (e.g., streaming). In some implementations, the at least one data store 112 can store information associated with the social networking system (e.g., the social networking system 630 of **FIGURE 6**). The information associated with the social networking system can include data about users, social connections, social interactions, locations, geo-fenced areas, maps, places, events, pages, groups, posts, communications, content, feeds, account settings, privacy settings, a social graph, and various other types of data. In some implementations, the at least one data store 112 can store information associated with users, such as user identifiers, user information, profile information, user specified settings, content produced or posted by users, and various other types of user data. It should be appreciated that there can be many variations or other possibilities.

[0033] In various embodiments, the content module 104 can provide access to various types of content items (e.g., virtual reality content items, immersive videos, etc.) to be presented through a viewport. This viewport may be provided through a display of a computing device (e.g., a virtual reality computing device) in which the client module 114 is implemented, for example. In some instances, the computing device may be running a software application (e.g., social networking application) that is configured to present content items. Some examples of virtual reality content can include videos composed using monoscopic 360 degree views or videos composed using stereoscopic 180 degree views, to name some examples. In various embodiments, virtual reality content items can capture views (e.g., 180 degree views, 225 degree views, 360 degree views, etc.) of one or more scenes over some duration of time. Such scenes may be captured from the real world and/or be

computer generated. In some instances, a virtual reality content item can be created by stitching together various video streams (or feeds) that were captured by cameras that are placed at particular locations and/or positions to capture a view of the scene. Such streams may be pre-determined for various directions, e.g., angles (e.g., 0 degree, 30 degrees, 60 degrees, etc.), accessible in a virtual reality content item. Once stitched together, a user can access, or present, the virtual reality content item to view a portion of the virtual reality content item along some direction (or angle). The portion of the virtual reality content item (e.g., stream) shown to the user can be determined based on techniques for dynamic streaming. In some instances, the virtual reality content item may include only one corresponding stream in which different portions are emphasized based on techniques for content-dependent streaming. Many variations are possible.

[0034] In one example, the computing device in which the client module 114 is implemented can request presentation of a virtual reality content item (e.g., spherical video). In this example, the streaming module 106 can provide one or more streams of the virtual reality content item to be presented through the computing device. In dynamic streaming, the stream(s) provided will typically correspond to a direction of the viewport in the virtual reality content item being accessed. As presentation of the virtual reality content item progresses, the client module 114 can continually provide the content provider module 102 with information describing the direction at which the viewport is facing. The streaming module 106 can use this information to determine which stream to provide the client module 114. In contrast, in content-dependent streaming, the streaming module 106 can provide a stream associated with the virtual reality content item to be presented through the computing device. As mentioned, different regions (or parts) of this stream can be emphasized based on saliency information.

[0035] In some embodiments, the content features module 108 provides a number of different features for enhancing the presentation of content items. For example, in some embodiments, the content features module 108 can generate a saliency prediction model that can be used to identify salient points of interest in a given content item. The content features module 108 can use the identified salient points of interest to improve the presentation of the content item. More details

describing the content features module 108 will be provided below in reference to **FIGURE 2**.

[0036] **FIGURE 2** illustrates an example of a content features module 202, according to an embodiment of the present disclosure. In some embodiments, the content features module 108 of **FIGURE 1** can be implemented with the content features module 202. As shown in the example of **FIGURE 2**, the content features module 202 can include a training content module 204, a view tracking data module 206, a heat map data module 208, a saliency module 210, and a view-based projection module 212.

[0037] In various embodiments, the training content module 204 can be configured to obtain content items to be used for training one or more models (e.g., saliency prediction models). Such content items may include videos (e.g., virtual reality content items, immersive videos, etc.). In general, a virtual reality content item (or immersive video) corresponds to any virtual reality media that encompasses (or surrounds) a viewer (or user). Some examples of virtual reality content items include spherical videos, half sphere videos (e.g., 180 degree videos), arbitrary partial spheres, 225 degree videos, and 3D 360 videos. Such virtual reality content items need not be limited to videos that are formatted using a spherical shape but may also be applied to immersive videos formatted using other shapes including, for example, cubes, pyramids, and other shape representations of a video recorded three-dimensional world.

[0038] The content items obtained by the training content module 204 can vary depending on the type of model being trained. For example, in some embodiments, a general saliency prediction model may be trained using various unrelated content items that were created by various publishers and corresponding heat map data for those content items. In some embodiments, such heat map data for a given content item may be generated based on view tracking data for the content item, as described below. This general saliency prediction model can be used to determine salient points of interest in various types of content items. In some embodiments, a publisher-specific saliency prediction model may be trained using content items that were posted by a given publisher (e.g., content creator) and corresponding heat map data for those content items. This publisher-specific saliency prediction model can be used to determine salient points of interest in content that is subsequently posted by

that publisher in which salient points of interest are not initially known. In some embodiments, a category-specific saliency prediction model may be trained using content items that all correspond to a given category (e.g., genre, topic, interest, etc.) and corresponding heat map data for those content items. This category-specific saliency prediction model can be used to determine salient points of interest in new content items that correspond to the given category.

[0039] In some embodiments, the view tracking data module 206 can be configured to obtain respective view tracking data for each of the content items being used to train the models. For example, view tracking data for a given content item may be collected for each user (or viewer) that has accessed the content item. The view tracking data for a user may identify regions that were accessed through the user's viewport during presentation of the content item. Such view tracking data may be collected for each frame corresponding to the content item. In some embodiments, a user's view tracking data for a content item can be determined based on changes to the user's viewport during presentation of the content item. Such changes to the viewport may be measured using various approaches that can be used either alone or in combination. For example, changes to the viewport may be measured using sensor data (e.g., gyroscope data, inertial measurement unit data, etc.) that describes movement of the computing device being used to present the content item. In another example, changes to the viewport can be measured using gesture data describing the types of gestures (e.g., panning, zooming, etc.) that were performed during presentation of the content item. Some other examples for measuring changes to the viewport include using input device data that describes input operations (e.g., mouse movement, dragging, etc.) performed during presentation of the content item, headset movement data that describes changes in the viewport direction during presentation of the content item, and eye tracking data collected during presentation of the content item, to name some examples.

[0040] In some embodiments, the heat map data module 208 can be configured to generate (or obtain) heat maps for each of the content items being used to train the models. In some embodiments, heat maps for a given content item may be generated based on view tracking data for the content item. As mentioned, the view tracking data module 206 can obtain respective view tracking data for users that viewed a content item. Each user's view tracking data can indicate which regions of a

given frame (or set of frames) were accessed using a user's viewport during presentation of a content item. That is, for any given frame in the content item, the heat map data module 208 can generate (or obtain) user-specific heat maps that graphically represent regions in the frame that were of interest to a given user. In some embodiments, heat maps can be generated for a set of frames that correspond to some interval of time. For example, a respective heat map can be generated for every second of the content item. In some embodiments, user-specific heat maps for a given content item can be combined to generate aggregated heat maps that represent aggregated regions of interest in frames corresponding to the content item. Thus, for example, the respective user-specific heat maps can be aggregated on a frame-by-frame basis so that each frame of the content item is associated with its own aggregated heat map that identifies the regions of interest in the frame. These regions of interest can correspond to various points of interest that appear in frames and were determined to be of interest to some, or all, of the users that viewed the content item. In some embodiments, these regions of interest can correspond to various points of interest that appear in frames and were determined to be of interest to users sharing one or more common characteristics with the user who is to view the content item.

[0041] In some embodiments, the saliency module 210 can be configured to train a saliency prediction model. In such embodiments, the saliency prediction model can be used to identify content (e.g., points of interest) that is likely to be of interest to a given user accessing a content item in which the identified content appears. For example, the saliency prediction model can determine that a first point of interest which appears in a given frame of a content item is likely to be of interest to a user over a second point of interest that also appears in the frame. In some embodiments, the saliency prediction model is trained using the content items that were obtained by the training content module 204 and their respective aggregated heat maps. For example, in some embodiments, each frame of a content item and its corresponding aggregated heat map can be provided as a training example to the saliency prediction model. In some embodiments, the saliency prediction model is trained using aggregated heat map data that has been labeled to identify points of interest. The aggregated heat map can be used to identify regions of the frame that were viewed more than others. Such view activity can be represented in the

aggregated heat map using various shapes that describe the size of the view region and/or colors that indicate concentrations of view activity in any given region of the frame. Based on this information, the saliency prediction model can learn which pixels in the frame were interesting (or relevant) to users in the aggregate. In some embodiments, pixels in the frame that fall within the shapes and/or colors represented in the aggregated heat map can be identified as being interesting (or relevant) to users in the aggregate. In some embodiments, these pixels correlate to points of interest that appear in frames. As a result, the saliency prediction model can learn which points of interest appearing in a frame were of interest to users in the aggregate with respect to other points of interest that also appear in the frame. Once trained, the saliency prediction model can be used to identify content (e.g., points of interest) that is likely to be of interest in new content items. In some embodiments, the saliency prediction model can be used to predict salient points of interest for stored content items (e.g., video on-demand). In some embodiments, the saliency prediction model can be used to predict salient points of interest (e.g., points of interest that are likely to be of interest) for live content items (e.g., live video broadcasts). In some embodiments, the saliency prediction model can be trained to output corresponding saliency maps for content items. For example, in some embodiments, the saliency prediction model can output a corresponding saliency map for each frame of a content item. In some embodiments, a saliency map for a given frame can assign a respective saliency value for each pixel in the frame. The saliency value for a pixel can provide a measure of saliency associated with that pixel. In various embodiments, heat map data used to generate the saliency prediction model, aggregated or otherwise, need not be actual heat maps that are represented graphically but may instead be some representation of view tracking data. For example, in some embodiments, the heat map data may identify clusters of view activity within individual frames of content items. In some embodiments, the clusters of view activity that are identified from heat map data can be used independently to identify salient points of interest in various content items. For example, in some embodiments, heat map data identifying clusters of view activity in frames during a live video broadcast (e.g., over the past n seconds of the broadcast) can be used to identify salient points of interest that appear in subsequent frames. May variations are possible.

[0042] The ability to predict salient content (e.g., points of interest) in content items provides a number of advantages. For example, in some embodiments, the view-based projection module 212 can use saliency information to selectively enhance (or emphasize) content (or portions of content) during presentation. More details describing the view-based projection module 212 will be provided below in reference to **FIGURE 3**.

[0043] **FIGURE 3** illustrates an example of a view-based projection module 300, according to an embodiment of the present disclosure. In some embodiments, the view-based projection module 212 of **FIGURE 2** can be implemented with the view-based projection module 300. As shown in the example of **FIGURE 3**, the view-based projection module 300 can include a frame module 302, a saliency map module 304, a vectorization module 306, and a projection module 308.

[0044] The frame module 302 can be configured to obtain a frame of a content item being enhanced based on view-based projection. In various embodiments, the view-based projection module 300 can enhance (or emphasize) regions in frames of the content item using the view-based projection techniques described herein.

[0045] The saliency map module 304 can be configured to obtain a saliency map for the frame of the content item being enhanced. As mentioned, such saliency maps can be determined based on information generated by the saliency prediction model. In some embodiments, the saliency map for the frame can indicate a respective saliency value for each pixel in the frame. In some embodiments, the saliency value for a pixel provides a measure of saliency associated with that pixel.

[0046] The vectorization module 306 can determine a respective vector-based representation of the frame. For example, in some embodiments, the vectorization module 306 can segment the saliency map corresponding to the frame into a set of bins. For example, the saliency map can be segmented into a set of yaw bins (e.g., 32 yaw bins) and a set of pitch bins (e.g., 16 pitch bins). In such embodiments, each bin can be associated with its own corresponding vector and magnitude. For example, in some embodiments, the set of yaw bins can each be associated with a pre-defined vector direction that ranges from negative 180 degrees to positive 180 degrees. Similarly, the set of pitch bins can each be associated with a pre-defined vector direction that ranges from negative 180 degrees to positive 180 degrees. In various embodiments, the vectorization module 306 can determine a net vector and

magnitude based on the vector-based representation of the saliency map. For example, in some embodiments, the vectorization module 306 can determine the net vector and magnitude for the frame by adding corresponding vectors and magnitudes for each bin (e.g., yaw bins, pitch bins) in the vector-based representation of the frame. In some embodiments, the net magnitude is normalized, for example, as a number between 0 and 1. In some embodiments, the net vector and normalized magnitude can reflect an amount of displacement, or offset, that is expected from a viewer (e.g., camera, computing device, etc.) when viewing the frame.

[0047] The projection module 308 can be configured to enhance the frame. In various embodiments, the projection module 308 implements view-based projection using an offset technique. In some embodiments, the offset technique can be applied to enhance different regions of the frame based on the amount of displacement that is expected in view of the net vector and normalized magnitude determined by the vectorization module 306. For example, in such embodiments, a spherical representation of the frame being enhanced can be projected. The projection module 308 can then correlate the net vector and normalized magnitude to an amount (or proportion) of pixels to allocate to a region in the frame that corresponds to the net vector and normalized. More details describing pixel allocation are described below in reference to **FIGURES 4D-4E**. In some embodiments, the projection module 308 can apply such content enhancement techniques to output an enhanced version of the frame 412 in a barrel layout (or barrel projection), as described below in reference to **FIGURES 4F-4G**. Many variations are possible.

[0048] **FIGURE 4A** illustrates an example frame 412 of a content item which includes a first point of interest 414 and a second point of interest 416. The frame 412 can be provided to the saliency prediction model to determine salient points of interest. In this example, the saliency prediction model may determine that the second point of interest 416 is a salient point of interest that is likely to be of interest to users viewing the content item. In some embodiments, the second point of interest 416 can be enhanced visually during presentation of the content item. For example, in some embodiments, a saliency map corresponding to the frame 412 can be determined by the saliency prediction model. In some embodiments, the saliency map indicates a respective saliency value for each pixel in the frame 412. As mentioned, the saliency value for a pixel can provide a measure of saliency

associated with that pixel. Next, a vector-based representation 422 of the frame 412 can be determined from the saliency map, as illustrated in the example of **FIGURE 4B**. For example, in some embodiments, the saliency map can be segmented into a set of yaw bins 424 and a set of pitch bins 426. In such embodiments, each bin 428 can be associated with its own corresponding vector 430 and magnitude 432 (e.g., saliency value). For example, in some embodiments, the set of yaw bins can each be associated with a pre-defined vector direction that ranges from negative 180 degrees to positive 180 degrees. Similarly, the set of pitch bins can each be associated with a pre-defined vector direction that ranges from negative 180 degrees to positive 180 degrees. In various embodiments, a net vector and magnitude 434 can be determined based on the vector-based representation 422, as illustrated in the example of **FIGURE 4C**. For example, the net vector and magnitude 434 for the frame 412 can be determined by adding corresponding vectors and magnitudes for each bin (e.g., yaw bins 424, pitch bins 426) in the vector-based representation 422. In some embodiments, the net magnitude is normalized.

[0049] As mentioned, in some embodiments, an offset technique can be applied to enhance salient regions of the frame 412 based on the net vector and normalized magnitude. For example, in some embodiments, a spherical representation 440 of the frame 412 being enhanced can be determined, as illustrated in the example of **FIGURE 4D**. The spherical representation 440 can be divided into regions based on a pre-defined number of rays 444 that emanate from a virtual camera 442 positioned at the center of the spherical representation 440. In some embodiments, an amount (or proportion) of pixels to be allocated to a region (e.g., the first point of interest 414 or the second point of interest 416) in the spherical representation 440 is determined based at least in part on a number of rays needed to capture (or fully view) that region. For example, when the virtual camera 442 is positioned at the center of the spherical representation 440, three rays are needed to capture the first point of interest 414 and three rays are also needed to capture the second point of interest 416, as illustrated in the example of **FIGURE 4D**. As a result, both the first point of interest 414 and the second point of interest 416 are allocated the same (or similar) amount of pixels. In some embodiments, the net vector and normalized magnitude for the frame 412 can be used to determine an offset for the virtual camera 442. In other words, the offset can represent a direction and

magnitude the virtual camera 442 is expected to move to view some region in the spherical representation 440. For example, in **FIGURE 4E**, the offset (i.e., net vector and normalized magnitude) determined for the frame 412 is used to reposition the virtual camera 442 within the spherical representation 440. In this example, the number of rays needed to capture the first point of interest 414 is now 1 while the number of rays needed to capture the second point of interest 416 is now 5. Since the number of rays needed to capture the first point of interest 414 decreased from 3 to 1, a lower density (or number) of pixels can be allocated to a region in the frame 412 that corresponds to the first point of interest 414. Similarly, since the number of rays needed to capture the second point of interest 416 increased from 3 to 5, a higher density (or number) of pixels can be allocated to a region in the frame 412 that corresponds to the second point of interest 416. In various embodiments, allocation (or re-allocation) of pixels to different regions of a frame can be performed dynamically without having to increase a total pixel count associated with the frame.

[0050] **FIGURE 4F** illustrates an example barrel layout (or barrel projection) 450 representation of the frame 412. In some embodiments, the frame 412 can be represented as a barrel layout in which a top portion 452 of the frame 412 (e.g., top 25 percent) and a bottom portion 454 of the frame 412 (e.g., bottom 25 percent) is formatted in an equirectangular layout. The remaining middle portion 456 of the frame 412 (e.g., the middle 90 degrees of the scene) then corresponds to an image that contains a quite uniform angular sample distribution. In some embodiments, regions in the portion 456 can be biased (e.g., stretched) to increase pixel density in those regions. Similarly, regions in the portion 456 can be biased (e.g., reduced, shortened) to decrease pixel density in those regions. As mentioned, in various embodiments, the offset (i.e., net vector and normalized magnitude) determined for the frame 412 can be used to enhance the portion 456. In this example, the offset indicates that a pixel density for a region corresponding to the second point of interest 416 should increase while a pixel density for a region corresponding to the first point of interest 414 should decrease. Accordingly, the region corresponding to the first point of interest 414 can be stretched horizontally, as illustrated in the example of **FIGURE 4G**. Further, the region corresponding to the first point of interest 414 can be shortened horizontally, as also illustrated in the example of **FIGURE 4G**. Many variations are possible. In various embodiments, the frame 412 can be

reconstructed from its barrel projection 450 using generally known environmental mapping techniques. The barrel projection 450 provides a number of advantages. For example, for the same surface area, the barrel projection 450 has higher pixel density in the middle 90 degrees when compared with a cube map. The barrel projection 450 also compresses better than a cube map for the same texture resolution, since the barrel projection 450 reuses a lot of data for movements in the horizontal plane.

[0051] **FIGURE 4H** illustrates an example activity diagram 480, according to an embodiment of the present disclosure. It should be appreciated that there can be additional, fewer, or alternative steps performed in similar or alternative orders, or in parallel, within the scope of the various embodiments discussed herein unless otherwise stated.

[0052] At block 482, content items to be used for training one or more models (e.g., a saliency prediction model) are obtained. As mentioned, such content items may include videos (e.g., virtual reality content items, immersive videos, etc.). In general, a virtual reality content item (or immersive video) corresponds to any virtual reality media that encompasses (or surrounds) a viewer (or user). Some examples of virtual reality content items include spherical videos, half sphere videos (e.g., 180 degree videos), arbitrary partial spheres, 225 degree videos, and 3D 360 videos. Such immersive videos need not be limited to videos that are formatted using a spherical shape but may also be applied to immersive videos formatted using other shapes including, for example, cubes, pyramids, and other shape representations of a video recorded three-dimensional world.

[0053] At block 484, respective view tracking data for each of the content items being used to train the models is obtained. For example, view tracking data for a given content item may be collected for each user (or viewer) that accessed the content item. The view tracking data for a user may identify regions that were accessed through the user's viewport during presentation of the content item. Such view tracking data may be collected for each frame corresponding to the content item. In some embodiments, heat maps for a given content item may be generated based on view tracking data for the content item, as described above. For example, user-specific heat maps can be generated for a given content item based on the respective view activity of those users. In some embodiments, such user-specific

heat maps can be combined to generate an aggregated heat map for the content item. This aggregated heat map can indicate which regions in a given frame of the content item were most popular, or interesting, to users that viewed the content item.

[0054] At block 486, one or more saliency prediction models can be generated. For example, in some embodiments, a general saliency prediction model can be trained using aggregated heat maps that describe user view tracking data for various content items, as described above. In some embodiments, this saliency prediction model can be used to predict content (e.g., points of interest) that is likely to be of interest to users during presentation of a content item.

[0055] At block 488, content items can be enhanced using view-based projection techniques, as described above. For example, regions in frames of a given content item that correspond to salient points of interest can be allocated more pixels while the remaining portions of the frame can be allocated fewer pixels. Many variations are possible.

[0056] At block 490, respective barrel projections can be determined for frames of content items. In some embodiments, a barrel projection is determined for each frame of a content item being enhanced. Many variations are possible.

[0057] **FIGURE 5A** illustrates an example method 500, according to an embodiment of the present disclosure. It should be appreciated that there can be additional, fewer, or alternative steps performed in similar or alternative orders, or in parallel, within the scope of the various embodiments discussed herein unless otherwise stated.

[0058] At block 502, at least one frame of a content item is provided to a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items. At block 504, information describing at least a first salient point of interest that appears in the at least one frame is obtained from the saliency prediction model, wherein the first salient point of interest is predicted to be of interest to users accessing the content item. At block 506, a view-based projection is applied to a region corresponding to the first salient point of interest, wherein the view-based projection enhances a quality in which the region is presented.

[0059] **FIGURE 5B** illustrates an example method 550, according to an embodiment of the present disclosure. It should be appreciated that there can be

additional, fewer, or alternative steps performed in similar or alternative orders, or in parallel, within the scope of the various embodiments discussed herein unless otherwise stated.

[0060] At block 552, at least one salient point of interest in a frame of a content item is determined based at least in part on a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items. At block 554, a barrel projection representation for the frame is determined. At block 556, a view-based projection is applied to the barrel projection representation for the frame. The view-based projection enhances a quality in which a region corresponding to the at least one salient point of interest is presented.

[0061] It is contemplated that there can be many other uses, applications, and/or variations associated with the various embodiments of the present disclosure. For example, in some cases, user can choose whether or not to opt-in to utilize the disclosed technology. The disclosed technology can also ensure that various privacy settings and preferences are maintained and can prevent private information from being divulged. In another example, various embodiments of the present disclosure can learn, improve, and/or be refined over time.

SOCIAL NETWORKING SYSTEM – EXAMPLE IMPLEMENTATION

[0062] **FIGURE 6** illustrates a network diagram of an example system 600 that can be utilized in various scenarios, in accordance with an embodiment of the present disclosure. The system 600 includes one or more user devices 610, one or more external systems 620, a social networking system (or service) 630, and a network 650. In an embodiment, the social networking service, provider, and/or system discussed in connection with the embodiments described above may be implemented as the social networking system 630. For purposes of illustration, the embodiment of the system 600, shown by **FIGURE 6**, includes a single external system 620 and a single user device 610. However, in other embodiments, the system 600 may include more user devices 610 and/or more external systems 620. In certain embodiments, the social networking system 630 is operated by a social network provider, whereas the external systems 620 are separate from the social networking system 630 in that they may be operated by different entities. In various embodiments, however, the social networking system 630 and the external systems

620 operate in conjunction to provide social networking services to users (or members) of the social networking system 630. In this sense, the social networking system 630 provides a platform or backbone, which other systems, such as external systems 620, may use to provide social networking services and functionalities to users across the Internet.

[0063] The user device 610 comprises one or more computing devices (or systems) that can receive input from a user and transmit and receive data via the network 650. In one embodiment, the user device 610 is a conventional computer system executing, for example, a Microsoft Windows compatible operating system (OS), Apple OS X, and/or a Linux distribution. In another embodiment, the user device 610 can be a computing device or a device having computer functionality, such as a smart-phone, a tablet, a personal digital assistant (PDA), a mobile telephone, a laptop computer, a wearable device (e.g., a pair of glasses, a watch, a bracelet, etc.), a camera, an appliance, etc. The user device 610 is configured to communicate via the network 650. The user device 610 can execute an application, for example, a browser application that allows a user of the user device 610 to interact with the social networking system 630. In another embodiment, the user device 610 interacts with the social networking system 630 through an application programming interface (API) provided by the native operating system of the user device 610, such as iOS and ANDROID. The user device 610 is configured to communicate with the external system 620 and the social networking system 630 via the network 650, which may comprise any combination of local area and/or wide area networks, using wired and/or wireless communication systems.

[0064] In one embodiment, the network 650 uses standard communications technologies and protocols. Thus, the network 650 can include links using technologies such as Ethernet, 802.11, worldwide interoperability for microwave access (WiMAX), 3G, 4G, CDMA, GSM, LTE, digital subscriber line (DSL), etc. Similarly, the networking protocols used on the network 650 can include multiprotocol label switching (MPLS), transmission control protocol/Internet protocol (TCP/IP), User Datagram Protocol (UDP), hypertext transport protocol (HTTP), simple mail transfer protocol (SMTP), file transfer protocol (FTP), and the like. The data exchanged over the network 650 can be represented using technologies and/or formats including hypertext markup language (HTML) and extensible markup

language (XML). In addition, all or some links can be encrypted using conventional encryption technologies such as secure sockets layer (SSL), transport layer security (TLS), and Internet Protocol security (IPsec).

[0065] In one embodiment, the user device 610 may display content from the external system 620 and/or from the social networking system 630 by processing a markup language document 614 received from the external system 620 and from the social networking system 630 using a browser application 612. The markup language document 614 identifies content and one or more instructions describing formatting or presentation of the content. By executing the instructions included in the markup language document 614, the browser application 612 displays the identified content using the format or presentation described by the markup language document 614. For example, the markup language document 614 includes instructions for generating and displaying a web page having multiple frames that include text and/or image data retrieved from the external system 620 and the social networking system 630. In various embodiments, the markup language document 614 comprises a data file including extensible markup language (XML) data, extensible hypertext markup language (XHTML) data, or other markup language data. Additionally, the markup language document 614 may include JavaScript Object Notation (JSON) data, JSON with padding (JSONP), and JavaScript data to facilitate data-interchange between the external system 620 and the user device 610. The browser application 612 on the user device 610 may use a JavaScript compiler to decode the markup language document 614.

[0066] The markup language document 614 may also include, or link to, applications or application frameworks such as FLASH™ or Unity™ applications, the Silverlight™ application framework, etc.

[0067] In one embodiment, the user device 610 also includes one or more cookies 616 including data indicating whether a user of the user device 610 is logged into the social networking system 630, which may enable modification of the data communicated from the social networking system 630 to the user device 610.

[0068] The external system 620 includes one or more web servers that include one or more web pages 622a, 622b, which are communicated to the user device 610 using the network 650. The external system 620 is separate from the social networking system 630. For example, the external system 620 is associated with a

first domain, while the social networking system 630 is associated with a separate social networking domain. Web pages 622a, 622b, included in the external system 620, comprise markup language documents 614 identifying content and including instructions specifying formatting or presentation of the identified content. As discussed previously, it should be appreciated that there can be many variations or other possibilities.

[0069] The social networking system 630 includes one or more computing devices for a social network, including a plurality of users, and providing users of the social network with the ability to communicate and interact with other users of the social network. In some instances, the social network can be represented by a graph, i.e., a data structure including edges and nodes. Other data structures can also be used to represent the social network, including but not limited to databases, objects, classes, meta elements, files, or any other data structure. The social networking system 630 may be administered, managed, or controlled by an operator. The operator of the social networking system 630 may be a human being, an automated application, or a series of applications for managing content, regulating policies, and collecting usage metrics within the social networking system 630. Any type of operator may be used.

[0070] Users may join the social networking system 630 and then add connections to any number of other users of the social networking system 630 to whom they desire to be connected. As used herein, the term “friend” refers to any other user of the social networking system 630 to whom a user has formed a connection, association, or relationship via the social networking system 630. For example, in an embodiment, if users in the social networking system 630 are represented as nodes in the social graph, the term “friend” can refer to an edge formed between and directly connecting two user nodes.

[0071] Connections may be added explicitly by a user or may be automatically created by the social networking system 630 based on common characteristics of the users (e.g., users who are alumni of the same educational institution). For example, a first user specifically selects a particular other user to be a friend. Connections in the social networking system 630 are usually in both directions, but need not be, so the terms “user” and “friend” depend on the frame of reference. Connections between users of the social networking system 630 are usually bilateral (“two-way”), or

“mutual,” but connections may also be unilateral, or “one-way.” For example, if Bob and Joe are both users of the social networking system 630 and connected to each other, Bob and Joe are each other’s connections. If, on the other hand, Bob wishes to connect to Joe to view data communicated to the social networking system 630 by Joe, but Joe does not wish to form a mutual connection, a unilateral connection may be established. The connection between users may be a direct connection; however, some embodiments of the social networking system 630 allow the connection to be indirect via one or more levels of connections or degrees of separation.

[0072] In addition to establishing and maintaining connections between users and allowing interactions between users, the social networking system 630 provides users with the ability to take actions on various types of items supported by the social networking system 630. These items may include groups or networks (i.e., social networks of people, entities, and concepts) to which users of the social networking system 630 may belong, events or calendar entries in which a user might be interested, computer-based applications that a user may use via the social networking system 630, transactions that allow users to buy or sell items via services provided by or through the social networking system 630, and interactions with advertisements that a user may perform on or off the social networking system 630. These are just a few examples of the items upon which a user may act on the social networking system 630, and many others are possible. A user may interact with anything that is capable of being represented in the social networking system 630 or in the external system 620, separate from the social networking system 630, or coupled to the social networking system 630 via the network 650.

[0073] The social networking system 630 is also capable of linking a variety of entities. For example, the social networking system 630 enables users to interact with each other as well as external systems 620 or other entities through an API, a web service, or other communication channels. The social networking system 630 generates and maintains the “social graph” comprising a plurality of nodes interconnected by a plurality of edges. Each node in the social graph may represent an entity that can act on another node and/or that can be acted on by another node. The social graph may include various types of nodes. Examples of types of nodes include users, non-person entities, content items, web pages, groups, activities, messages, concepts, and any other things that can be represented by an object in

the social networking system 630. An edge between two nodes in the social graph may represent a particular kind of connection, or association, between the two nodes, which may result from node relationships or from an action that was performed by one of the nodes on the other node. In some cases, the edges between nodes can be weighted. The weight of an edge can represent an attribute associated with the edge, such as a strength of the connection or association between nodes. Different types of edges can be provided with different weights. For example, an edge created when one user “likes” another user may be given one weight, while an edge created when a user befriends another user may be given a different weight.

[0074] As an example, when a first user identifies a second user as a friend, an edge in the social graph is generated connecting a node representing the first user and a second node representing the second user. As various nodes relate or interact with each other, the social networking system 630 modifies edges connecting the various nodes to reflect the relationships and interactions.

[0075] The social networking system 630 also includes user-generated content, which enhances a user’s interactions with the social networking system 630. User-generated content may include anything a user can add, upload, send, or “post” to the social networking system 630. For example, a user communicates posts to the social networking system 630 from a user device 610. Posts may include data such as status updates or other textual data, location information, images such as photos, videos, links, music or other similar data and/or media. Content may also be added to the social networking system 630 by a third party. Content “items” are represented as objects in the social networking system 630. In this way, users of the social networking system 630 are encouraged to communicate with each other by posting text and content items of various types of media through various communication channels. Such communication increases the interaction of users with each other and increases the frequency with which users interact with the social networking system 630.

[0076] The social networking system 630 includes a web server 632, an API request server 634, a user profile store 636, a connection store 638, an action logger 640, an activity log 642, and an authorization server 644. In an embodiment of the invention, the social networking system 630 may include additional, fewer, or

different components for various applications. Other components, such as network interfaces, security mechanisms, load balancers, failover servers, management and network operations consoles, and the like are not shown so as to not obscure the details of the system.

[0077] The user profile store 636 maintains information about user accounts, including biographic, demographic, and other types of descriptive information, such as work experience, educational history, hobbies or preferences, location, and the like that has been declared by users or inferred by the social networking system 630. This information is stored in the user profile store 636 such that each user is uniquely identified. The social networking system 630 also stores data describing one or more connections between different users in the connection store 638. The connection information may indicate users who have similar or common work experience, group memberships, hobbies, or educational history. Additionally, the social networking system 630 includes user-defined connections between different users, allowing users to specify their relationships with other users. For example, user-defined connections allow users to generate relationships with other users that parallel the users' real-life relationships, such as friends, co-workers, partners, and so forth. Users may select from predefined types of connections, or define their own connection types as needed. Connections with other nodes in the social networking system 630, such as non-person entities, buckets, cluster centers, images, interests, pages, external systems, concepts, and the like are also stored in the connection store 638.

[0078] The social networking system 630 maintains data about objects with which a user may interact. To maintain this data, the user profile store 636 and the connection store 638 store instances of the corresponding type of objects maintained by the social networking system 630. Each object type has information fields that are suitable for storing information appropriate to the type of object. For example, the user profile store 636 contains data structures with fields suitable for describing a user's account and information related to a user's account. When a new object of a particular type is created, the social networking system 630 initializes a new data structure of the corresponding type, assigns a unique object identifier to it, and begins to add data to the object as needed. This might occur, for example, when a user becomes a user of the social networking system 630, the social networking

system 630 generates a new instance of a user profile in the user profile store 636, assigns a unique identifier to the user account, and begins to populate the fields of the user account with information provided by the user.

[0079] The connection store 638 includes data structures suitable for describing a user's connections to other users, connections to external systems 620 or connections to other entities. The connection store 638 may also associate a connection type with a user's connections, which may be used in conjunction with the user's privacy setting to regulate access to information about the user. In an embodiment of the invention, the user profile store 636 and the connection store 638 may be implemented as a federated database.

[0080] Data stored in the connection store 638, the user profile store 636, and the activity log 642 enables the social networking system 630 to generate the social graph that uses nodes to identify various objects and edges connecting nodes to identify relationships between different objects. For example, if a first user establishes a connection with a second user in the social networking system 630, user accounts of the first user and the second user from the user profile store 636 may act as nodes in the social graph. The connection between the first user and the second user stored by the connection store 638 is an edge between the nodes associated with the first user and the second user. Continuing this example, the second user may then send the first user a message within the social networking system 630. The action of sending the message, which may be stored, is another edge between the two nodes in the social graph representing the first user and the second user. Additionally, the message itself may be identified and included in the social graph as another node connected to the nodes representing the first user and the second user.

[0081] In another example, a first user may tag a second user in an image that is maintained by the social networking system 630 (or, alternatively, in an image maintained by another system outside of the social networking system 630). The image may itself be represented as a node in the social networking system 630. This tagging action may create edges between the first user and the second user as well as create an edge between each of the users and the image, which is also a node in the social graph. In yet another example, if a user confirms attending an event, the user and the event are nodes obtained from the user profile store 636, where the

attendance of the event is an edge between the nodes that may be retrieved from the activity log 642. By generating and maintaining the social graph, the social networking system 630 includes data describing many different types of objects and the interactions and connections among those objects, providing a rich source of socially relevant information.

[0082] The web server 632 links the social networking system 630 to one or more user devices 610 and/or one or more external systems 620 via the network 650. The web server 632 serves web pages, as well as other web-related content, such as Java, JavaScript, Flash, XML, and so forth. The web server 632 may include a mail server or other messaging functionality for receiving and routing messages between the social networking system 630 and one or more user devices 610. The messages can be instant messages, queued messages (e.g., email), text and SMS messages, or any other suitable messaging format.

[0083] The API request server 634 allows one or more external systems 620 and user devices 610 to call access information from the social networking system 630 by calling one or more API functions. The API request server 634 may also allow external systems 620 to send information to the social networking system 630 by calling APIs. The external system 620, in one embodiment, sends an API request to the social networking system 630 via the network 650, and the API request server 634 receives the API request. The API request server 634 processes the request by calling an API associated with the API request to generate an appropriate response, which the API request server 634 communicates to the external system 620 via the network 650. For example, responsive to an API request, the API request server 634 collects data associated with a user, such as the user's connections that have logged into the external system 620, and communicates the collected data to the external system 620. In another embodiment, the user device 610 communicates with the social networking system 630 via APIs in the same manner as external systems 620.

[0084] The action logger 640 is capable of receiving communications from the web server 632 about user actions on and/or off the social networking system 630. The action logger 640 populates the activity log 642 with information about user actions, enabling the social networking system 630 to discover various actions taken by its users within the social networking system 630 and outside of the social networking system 630. Any action that a particular user takes with respect to

another node on the social networking system 630 may be associated with each user's account, through information maintained in the activity log 642 or in a similar database or other data repository. Examples of actions taken by a user within the social networking system 630 that are identified and stored may include, for example, adding a connection to another user, sending a message to another user, reading a message from another user, viewing content associated with another user, attending an event posted by another user, posting an image, attempting to post an image, or other actions interacting with another user or another object. When a user takes an action within the social networking system 630, the action is recorded in the activity log 642. In one embodiment, the social networking system 630 maintains the activity log 642 as a database of entries. When an action is taken within the social networking system 630, an entry for the action is added to the activity log 642. The activity log 642 may be referred to as an action log.

[0085] Additionally, user actions may be associated with concepts and actions that occur within an entity outside of the social networking system 630, such as an external system 620 that is separate from the social networking system 630. For example, the action logger 640 may receive data describing a user's interaction with an external system 620 from the web server 632. In this example, the external system 620 reports a user's interaction according to structured actions and objects in the social graph.

[0086] Other examples of actions where a user interacts with an external system 620 include a user expressing an interest in an external system 620 or another entity, a user posting a comment to the social networking system 630 that discusses an external system 620 or a web page 622a within the external system 620, a user posting to the social networking system 630 a Uniform Resource Locator (URL) or other identifier associated with an external system 620, a user attending an event associated with an external system 620, or any other action by a user that is related to an external system 620. Thus, the activity log 642 may include actions describing interactions between a user of the social networking system 630 and an external system 620 that is separate from the social networking system 630.

[0087] The authorization server 644 enforces one or more privacy settings of the users of the social networking system 630. A privacy setting of a user determines how particular information associated with a user can be shared. The privacy setting

comprises the specification of particular information associated with a user and the specification of the entity or entities with whom the information can be shared. Examples of entities with which information can be shared may include other users, applications, external systems 620, or any entity that can potentially access the information. The information that can be shared by a user comprises user account information, such as profile photos, phone numbers associated with the user, user's connections, actions taken by the user such as adding a connection, changing user profile information, and the like.

[0088] The privacy setting specification may be provided at different levels of granularity. For example, the privacy setting may identify specific information to be shared with other users; the privacy setting identifies a work phone number or a specific set of related information, such as, personal information including profile photo, home phone number, and status. Alternatively, the privacy setting may apply to all the information associated with the user. The specification of the set of entities that can access particular information can also be specified at various levels of granularity. Various sets of entities with which information can be shared may include, for example, all friends of the user, all friends of friends, all applications, or all external systems 620. One embodiment allows the specification of the set of entities to comprise an enumeration of entities. For example, the user may provide a list of external systems 620 that are allowed to access certain information. Another embodiment allows the specification to comprise a set of entities along with exceptions that are not allowed to access the information. For example, a user may allow all external systems 620 to access the user's work information, but specify a list of external systems 620 that are not allowed to access the work information. Certain embodiments call the list of exceptions that are not allowed to access certain information a "block list". External systems 620 belonging to a block list specified by a user are blocked from accessing the information specified in the privacy setting. Various combinations of granularity of specification of information, and granularity of specification of entities, with which information is shared are possible. For example, all personal information may be shared with friends whereas all work information may be shared with friends of friends.

[0089] The authorization server 644 contains logic to determine if certain information associated with a user can be accessed by a user's friends, external

systems 620, and/or other applications and entities. The external system 620 may need authorization from the authorization server 644 to access the user's more private and sensitive information, such as the user's work phone number. Based on the user's privacy settings, the authorization server 644 determines if another user, the external system 620, an application, or another entity is allowed to access information associated with the user, including information about actions taken by the user.

[0090] In some embodiments, the social networking system 630 can include a content provider module 646. The content provider module 646 can, for example, be implemented as the content provider module 102 of **FIGURE 1**. In some embodiments, the content provider module 646, or some of its features, can be implemented in a computing device, e.g., the user device 610. In some embodiments, the user device 610 can include a client module 618. The client module 618 can, for example, be implemented as the client module 114 of **FIGURE 1**. The network 650 can, for example, be implemented as the network 150 of **FIGURE 1**. As discussed previously, it should be appreciated that there can be many variations or other possibilities.

HARDWARE IMPLEMENTATION

[0091] The foregoing processes and features can be implemented by a wide variety of machine and computer system architectures and in a wide variety of network and computing environments. **FIGURE 7** illustrates an example of a computer system 700 that may be used to implement one or more of the embodiments described herein in accordance with an embodiment of the invention. The computer system 700 includes sets of instructions for causing the computer system 700 to perform the processes and features discussed herein. The computer system 700 may be connected (e.g., networked) to other machines. In a networked deployment, the computer system 700 may operate in the capacity of a server machine or a client machine in a client-server network environment, or as a peer machine in a peer-to-peer (or distributed) network environment. In an embodiment of the invention, the computer system 700 may be the social networking system 630, the user device 610, and the external system 720, or a component thereof. In an

embodiment of the invention, the computer system 700 may be one server among many that constitutes all or part of the social networking system 630.

[0092] The computer system 700 includes a processor 702, a cache 704, and one or more executable modules and drivers, stored on a computer-readable medium, directed to the processes and features described herein. Additionally, the computer system 700 includes a high performance input/output (I/O) bus 706 and a standard I/O bus 708. A host bridge 710 couples processor 702 to high performance I/O bus 706, whereas I/O bus bridge 712 couples the two buses 706 and 708 to each other. A system memory 714 and one or more network interfaces 716 couple to high performance I/O bus 706. The computer system 700 may further include video memory and a display device coupled to the video memory (not shown). Mass storage 718 and I/O ports 720 couple to the standard I/O bus 708. The computer system 700 may optionally include a keyboard and pointing device, a display device, or other input/output devices (not shown) coupled to the standard I/O bus 708. Collectively, these elements are intended to represent a broad category of computer hardware systems, including but not limited to computer systems based on the x86-compatible processors manufactured by Intel Corporation of Santa Clara, California, and the x86-compatible processors manufactured by Advanced Micro Devices (AMD), Inc., of Sunnyvale, California, as well as any other suitable processor.

[0093] An operating system manages and controls the operation of the computer system 700, including the input and output of data to and from software applications (not shown). The operating system provides an interface between the software applications being executed on the system and the hardware components of the system. Any suitable operating system may be used, such as the LINUX Operating System, the Apple Macintosh Operating System, available from Apple Computer Inc. of Cupertino, California, UNIX operating systems, Microsoft® Windows® operating systems, BSD operating systems, and the like. Other implementations are possible.

[0094] The elements of the computer system 700 are described in greater detail below. In particular, the network interface 716 provides communication between the computer system 700 and any of a wide range of networks, such as an Ethernet (e.g., IEEE 802.3) network, a backplane, etc. The mass storage 718 provides permanent storage for the data and programming instructions to perform the above-

described processes and features implemented by the respective computing systems identified above, whereas the system memory 714 (e.g., DRAM) provides temporary storage for the data and programming instructions when executed by the processor 702. The I/O ports 720 may be one or more serial and/or parallel communication ports that provide communication between additional peripheral devices, which may be coupled to the computer system 700.

[0095] The computer system 700 may include a variety of system architectures, and various components of the computer system 700 may be rearranged. For example, the cache 704 may be on-chip with processor 702. Alternatively, the cache 704 and the processor 702 may be packed together as a “processor module”, with processor 702 being referred to as the “processor core”. Furthermore, certain embodiments of the invention may neither require nor include all of the above components. For example, peripheral devices coupled to the standard I/O bus 708 may couple to the high performance I/O bus 706. In addition, in some embodiments, only a single bus may exist, with the components of the computer system 700 being coupled to the single bus. Moreover, the computer system 700 may include additional components, such as additional processors, storage devices, or memories.

[0096] In general, the processes and features described herein may be implemented as part of an operating system or a specific application, component, program, object, module, or series of instructions referred to as “programs”. For example, one or more programs may be used to execute specific processes described herein. The programs typically comprise one or more instructions in various memory and storage devices in the computer system 700 that, when read and executed by one or more processors, cause the computer system 700 to perform operations to execute the processes and features described herein. The processes and features described herein may be implemented in software, firmware, hardware (e.g., an application specific integrated circuit), or any combination thereof.

[0097] In one implementation, the processes and features described herein are implemented as a series of executable modules run by the computer system 700, individually or collectively in a distributed computing environment. The foregoing modules may be realized by hardware, executable modules stored on a computer-readable medium (or machine-readable medium), or a combination of both. For example, the modules may comprise a plurality or series of instructions to be

executed by a processor in a hardware system, such as the processor 702. Initially, the series of instructions may be stored on a storage device, such as the mass storage 718. However, the series of instructions can be stored on any suitable computer readable storage medium. Furthermore, the series of instructions need not be stored locally, and could be received from a remote storage device, such as a server on a network, via the network interface 716. The instructions are copied from the storage device, such as the mass storage 718, into the system memory 714 and then accessed and executed by the processor 702. In various implementations, a module or modules can be executed by a processor or multiple processors in one or multiple locations, such as multiple servers in a parallel processing environment.

[0098] Examples of computer-readable media include, but are not limited to, recordable type media such as volatile and non-volatile memory devices; solid state memories; floppy and other removable disks; hard disk drives; magnetic media; optical disks (e.g., Compact Disk Read-Only Memory (CD ROMs), Digital Versatile Disks (DVDs)); other similar non-transitory (or transitory), tangible (or non-tangible) storage medium; or any type of medium suitable for storing, encoding, or carrying a series of instructions for execution by the computer system 700 to perform any one or more of the processes and features described herein.

[0099] For purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the description. It will be apparent, however, to one skilled in the art that embodiments of the disclosure can be practiced without these specific details. In some instances, modules, structures, processes, features, and devices are shown in block diagram form in order to avoid obscuring the description. In other instances, functional block diagrams and flow diagrams are shown to represent data and logic flows. The components of block diagrams and flow diagrams (e.g., modules, blocks, structures, devices, features, etc.) may be variously combined, separated, removed, reordered, and replaced in a manner other than as expressly described and depicted herein.

[00100] Reference in this specification to “one embodiment”, “an embodiment”, “other embodiments”, “one series of embodiments”, “some embodiments”, “various embodiments”, or the like means that a particular feature, design, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the disclosure. The appearances of, for example, the phrase “in one

embodiment” or “in an embodiment” in various places in the specification are not necessarily all referring to the same embodiment, nor are separate or alternative embodiments mutually exclusive of other embodiments. Moreover, whether or not there is express reference to an “embodiment” or the like, various features are described, which may be variously combined and included in some embodiments, but also variously omitted in other embodiments. Similarly, various features are described that may be preferences or requirements for some embodiments, but not other embodiments.

[00101] The language used herein has been principally selected for readability and instructional purposes, and it may not have been selected to delineate or circumscribe the inventive subject matter. It is therefore intended that the scope of the invention be limited not by this detailed description, but rather by any claims that issue on an application based hereon. Accordingly, the disclosure of the embodiments of the invention is intended to be illustrative, but not limiting, of the scope of the invention, which is set forth in the following claims.

WHAT IS CLAIMED IS:

1. A computer-implemented method comprising:
 - providing, by a computing system, at least one frame of a content item to a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items;
 - obtaining, by the computing system, information describing at least a first salient point of interest that appears in the at least one frame from the saliency prediction model, wherein the first salient point of interest is predicted to be of interest to users accessing the content item; and
 - applying, by the computing system, a view-based projection to a region corresponding to the first salient point of interest, wherein the view-based projection enhances a quality in which the region is presented.
2. The computer-implemented method of claim 1, wherein applying the view-based projection further comprises:
 - obtaining, by the computing system, a saliency map for the at least one frame;
 - determining, by the computing system, a vector-based representation of the saliency map;
 - determining, by the computing system, an offset corresponding to the at least one frame based at least in part on the vector-based representation; and
 - enhancing, by the computing system, the region corresponding to the first salient point of interest based at least in part on the determined offset.
3. The computer-implemented method of claim 2, wherein the vector-based representation includes a set of yaw bins and a set of pitch bins, and wherein each bin is associated with a pre-defined vector and a corresponding magnitude.
4. The computer-implemented method of claim 2, wherein enhancing the region corresponding to the first salient point of interest based at least in part on the determined offset further comprises:

determining, by the computing system, a spherical representation of the at least one frame, the spherical representation including a virtual camera positioned at a center of the spherical representation;

determining, by the computing system, a displacement for the virtual camera based at least in part on the offset; and

determining, by the computing system, an amount of pixels to allocate for the region based at least in part on the displacement of the virtual camera.

5. The computer-implemented method of claim 4, wherein the region is enhanced by increasing a pixel density corresponding to the region.

6. The computer-implemented method of claim 4, wherein respective pixel densities for one or more other regions of the at least one frame are decreased in proportion to the displacement of the virtual camera.

7. The computer-implemented method of claim 4, wherein a total pixel count associated with the at least one frame remains unchanged.

8. A system comprising:

at least one processor; and

a memory storing instructions that, when executed by the at least one processor, cause the system to perform:

providing at least one frame of a content item to a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items;

obtaining information describing at least a first salient point of interest that appears in the at least one frame from the saliency prediction model, wherein the first salient point of interest is predicted to be of interest to users accessing the content item; and

applying a view-based projection to a region corresponding to the first salient point of interest, wherein the view-based projection enhances a quality in which the region is presented.

9. The system of claim 8, wherein applying the view-based projection further causes the system to perform:

- obtaining a saliency map for the at least one frame;
- determining a vector-based representation of the saliency map;
- determining an offset corresponding to the at least one frame based at least in part on the vector-based representation; and
- enhancing the region corresponding to the first salient point of interest based at least in part on the determined offset.

10. The system of claim 9, wherein the vector-based representation includes a set of yaw bins and a set of pitch bins, and wherein each bin is associated with a pre-defined vector and a corresponding magnitude.

11. The system of claim 9, wherein enhancing the region corresponding to the first salient point of interest based at least in part on the determined offset further causes the system to perform:

- determining a spherical representation of the at least one frame, the spherical representation including a virtual camera positioned at a center of the spherical representation;
- determining a displacement for the virtual camera based at least in part on the offset; and
- determining an amount of pixels to allocate for the region based at least in part on the displacement of the virtual camera.

12. The system of claim 11, wherein the region is enhanced by increasing a pixel density corresponding to the region.

13. The system of claim 11, wherein respective pixel densities for one or more other regions of the at least one frame are decreased in proportion to the displacement of the virtual camera.

14. The system of claim 11, wherein a total pixel count associated with the at least one frame remains unchanged.

15. A non-transitory computer-readable storage medium including instructions that, when executed by at least one processor of a computing system, cause the computing system to perform a method comprising:

providing at least one frame of a content item to a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items;

obtaining information describing at least a first salient point of interest that appears in the at least one frame from the saliency prediction model, wherein the first salient point of interest is predicted to be of interest to users accessing the content item; and

applying a view-based projection to a region corresponding to the first salient point of interest, wherein the view-based projection enhances a quality in which the region is presented.

16. The non-transitory computer-readable storage medium of claim 15, wherein applying the view-based projection further causes the computing system to perform:

obtaining a saliency map for the at least one frame;

determining a vector-based representation of the saliency map;

determining an offset corresponding to the at least one frame based at least in part on the vector-based representation; and

enhancing the region corresponding to the first salient point of interest based at least in part on the determined offset.

17. The non-transitory computer-readable storage medium of claim 16, wherein the vector-based representation includes a set of yaw bins and a set of pitch bins, and wherein each bin is associated with a pre-defined vector and a corresponding magnitude.

18. The non-transitory computer-readable storage medium of claim 16, wherein enhancing the region corresponding to the first salient point of interest based

at least in part on the determined offset further causes the computing system to perform:

- determining a spherical representation of the at least one frame, the spherical representation including a virtual camera positioned at a center of the spherical representation;

- determining a displacement for the virtual camera based at least in part on the offset; and

- determining an amount of pixels to allocate for the region based at least in part on the displacement of the virtual camera.

19. The non-transitory computer-readable storage medium of claim 18, wherein the region is enhanced by increasing a pixel density corresponding to the region.

20. The non-transitory computer-readable storage medium of claim 18, wherein respective pixel densities for one or more other regions of the at least one frame are decreased in proportion to the displacement of the virtual camera.

21. A computer-implemented method comprising:

- providing, by a computing system, at least one frame of a content item to a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items;

- obtaining, by the computing system, information describing at least a first salient point of interest that appears in the at least one frame from the saliency prediction model, wherein the first salient point of interest is predicted to be of interest to users accessing the content item; and

- applying, by the computing system, a view-based projection to a region corresponding to the first salient point of interest, wherein the view-based projection enhances a quality in which the region is presented.

22. The computer-implemented method of claim 21, wherein applying the view-based projection further comprises:

- obtaining, by the computing system, a saliency map for the at least one frame;

determining, by the computing system, a vector-based representation of the saliency map;

determining, by the computing system, an offset corresponding to the at least one frame based at least in part on the vector-based representation; and

enhancing, by the computing system, the region corresponding to the first salient point of interest based at least in part on the determined offset.

23. The computer-implemented method of claim 22, wherein the vector-based representation includes a set of yaw bins and a set of pitch bins, and wherein each bin is associated with a pre-defined vector and a corresponding magnitude.

24. The computer-implemented method of claim 22 or 23, wherein enhancing the region corresponding to the first salient point of interest based at least in part on the determined offset further comprises:

determining, by the computing system, a spherical representation of the at least one frame, the spherical representation including a virtual camera positioned at a center of the spherical representation;

determining, by the computing system, a displacement for the virtual camera based at least in part on the offset; and

determining, by the computing system, an amount of pixels to allocate for the region based at least in part on the displacement of the virtual camera.

25. The computer-implemented method of claim 24, wherein the region is enhanced by increasing a pixel density corresponding to the region; and/or

wherein respective pixel densities for one or more other regions of the at least one frame are decreased in proportion to the displacement of the virtual camera; and/or

wherein a total pixel count associated with the at least one frame remains unchanged.

26. A system comprising:

at least one processor; and

a memory storing instructions that, when executed by the at least one processor, cause the system to perform:

providing at least one frame of a content item to a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items;

obtaining information describing at least a first salient point of interest that appears in the at least one frame from the saliency prediction model, wherein the first salient point of interest is predicted to be of interest to users accessing the content item; and

applying a view-based projection to a region corresponding to the first salient point of interest, wherein the view-based projection enhances a quality in which the region is presented.

27. The system of claim 26, wherein applying the view-based projection further causes the system to perform:

obtaining a saliency map for the at least one frame;

determining a vector-based representation of the saliency map;

determining an offset corresponding to the at least one frame based at least in part on the vector-based representation; and

enhancing the region corresponding to the first salient point of interest based at least in part on the determined offset.

28. The system of claim 27, wherein the vector-based representation includes a set of yaw bins and a set of pitch bins, and wherein each bin is associated with a pre-defined vector and a corresponding magnitude.

29. The system of claim 27 or 28, wherein enhancing the region corresponding to the first salient point of interest based at least in part on the determined offset further causes the system to perform:

determining a spherical representation of the at least one frame, the spherical representation including a virtual camera positioned at a center of the spherical representation;

determining a displacement for the virtual camera based at least in part on the offset; and

determining an amount of pixels to allocate for the region based at least in part on the displacement of the virtual camera.

30. The system of claim 29, wherein the region is enhanced by increasing a pixel density corresponding to the region; and/or

wherein respective pixel densities for one or more other regions of the at least one frame are decreased in proportion to the displacement of the virtual camera; and/or

wherein a total pixel count associated with the at least one frame remains unchanged.

31. A non-transitory computer-readable storage medium including instructions that, when executed by at least one processor of a computing system, cause the computing system to perform a method comprising:

providing at least one frame of a content item to a saliency prediction model, the saliency prediction model being trained to identify salient points of interest that appear in content items;

obtaining information describing at least a first salient point of interest that appears in the at least one frame from the saliency prediction model, wherein the first salient point of interest is predicted to be of interest to users accessing the content item; and

applying a view-based projection to a region corresponding to the first salient point of interest, wherein the view-based projection enhances a quality in which the region is presented.

32. The non-transitory computer-readable storage medium of claim 31, wherein applying the view-based projection further causes the computing system to perform:

obtaining a saliency map for the at least one frame;

determining a vector-based representation of the saliency map;

determining an offset corresponding to the at least one frame based at least in part on the vector-based representation; and

enhancing the region corresponding to the first salient point of interest based at least in part on the determined offset.

33. The non-transitory computer-readable storage medium of claim 32, wherein the vector-based representation includes a set of yaw bins and a set of pitch bins, and wherein each bin is associated with a pre-defined vector and a corresponding magnitude.

34. The non-transitory computer-readable storage medium of claim 32 or 33, wherein enhancing the region corresponding to the first salient point of interest based at least in part on the determined offset further causes the computing system to perform:

determining a spherical representation of the at least one frame, the spherical representation including a virtual camera positioned at a center of the spherical representation;

determining a displacement for the virtual camera based at least in part on the offset; and

determining an amount of pixels to allocate for the region based at least in part on the displacement of the virtual camera.

35. The non-transitory computer-readable storage medium of claim 34, wherein the region is enhanced by increasing a pixel density corresponding to the region; and/or
wherein respective pixel densities for one or more other regions of the at least one frame are decreased in proportion to the displacement of the virtual camera.

1/15

100 ↗






FIGURE 1

2/15

**FIGURE 2**

3/15

**FIGURE 3**




FIGURE 4A




FIGURE 4B




FIGURE 4C

7/15




FIGURE 4D




FIGURE 4E

9/15




FIGURE 4F

10/15




FIGURE 4G

11/15

480 ↘




FIGURE 4H

12/15

500 ↘




FIGURE 5A

13/15

550 ↘




FIGURE 5B

14/15




FIGURE 6

15/15




FIGURE 7

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2018/027684**A. CLASSIFICATION OF SUBJECT MATTER****G06K 9/00(2006.01)i, G06K 9/62(2006.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06K 9/00; G06Q 30/02; G06K 9/46; G06F 17/30; G06T 3/40; G06K 9/62; G06T 7/00Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: content, salient, model, interest, projection**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2015-0213020 A1 (FUJITSU LIMITED) 30 July 2015 See claims 1-4 and figures 6-10.	1-35
A	US 2015-0371367 A1 (XIAOMI INC.) 24 December 2015 See paragraphs [0077]-[0098] and figures 1-2A,4.	1-35
A	US 2011-0229025 A1 (QI ZHAO et al.) 22 September 2011 See paragraphs [0095]-[0103] and figures 3A-3B.	1-35
A	US 2017-0068872 A1 (ACCENTURE GLOBAL SOLUTIONS LIMITED) 09 March 2017 See claims 1-3 and figures 1-4.	1-35
A	US 2010-0086221 A1 (BRIAN J. STANKIEWICZ et al.) 08 April 2010 See claims 21-28 and figure 1.	1-35



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

08 August 2018 (08.08.2018)

Date of mailing of the international search report

08 August 2018 (08.08.2018)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

KANG, Sung Chul

Telephone No. +82-42-481-8405



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2018/027684

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015-0213020 A1	30/07/2015	US 9946795 B2	17/04/2018
US 2015-0371367 A1	24/12/2015	BR 112015003438 A2	04/07/2017
		CN 104123720 A	29/10/2014
		CN 104123720 B	04/07/2017
		EP 2975574 A2	20/01/2016
		EP 2975574 A3	02/03/2016
		JP 2016-532947 A	20/10/2016
		JP 6170626 B2	26/07/2017
		KR 10-1641175 B1	20/07/2016
		KR 10-2016-0010389 A	27/01/2016
		MX 2015002442 A	03/03/2016
		RU 2015106257 A	10/09/2016
		RU 2614541 C2	28/03/2017
		US 9665925 B2	30/05/2017
		WO 2015-196715 A1	30/12/2015
US 2011-0229025 A1	22/09/2011	US 8649606 B2	11/02/2014
		WO 2011-152893 A1	08/12/2011
US 2017-0068872 A1	09/03/2017	US 2017-0068909 A1	09/03/2017
US 2010-0086221 A1	08/04/2010	AU 2009-298416 A1	08/04/2010
		AU 2009-298416 A8	29/08/2013
		AU 2009-298416 B2	15/08/2013
		AU 2009-298416 B8	29/08/2013
		AU 2009-298428 A1	08/04/2010
		AU 2009-298428 B2	14/03/2013
		AU 2009-298438 A1	08/04/2010
		AU 2009-298438 B2	24/10/2013
		CA 2739023 A1	08/04/2010
		CA 2739023 C	04/07/2017
		CA 2739260 A1	08/04/2010
		CA 2739273 A1	08/04/2010
		CN 102227748 A	26/10/2011
		CN 102227753 A	26/10/2011
		CN 102227753 B	19/01/2018
		CN 102239505 A	09/11/2011
		CN 102239505 B	22/10/2014
		EP 2347387 A1	27/07/2011
		EP 2356631 A1	17/08/2011
		EP 2359336 A1	24/08/2011
		JP 2012-504827 A	23/02/2012
		JP 2012-504828 A	23/02/2012
		JP 2012-504830 A	23/02/2012
		JP 5486006 B2	07/05/2014
		JP 5563580 B2	30/07/2014
		JP 5667061 B2	12/02/2015
		KR 10-2011-0081240 A	13/07/2011

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2018/027684Patent document
cited in search reportPublication
datePatent family
member(s)Publication
date

KR 10-2011-0082018 A	15/07/2011
KR 10-2011-0086010 A	27/07/2011
US 2010-0086200 A1	08/04/2010
US 2010-0086278 A1	08/04/2010
US 2013-0230253 A1	05/09/2013
US 2013-0259448 A1	03/10/2013
US 8442328 B2	14/05/2013
US 8478111 B2	02/07/2013
US 8577156 B2	05/11/2013
US 8811753 B2	19/08/2014
US 9131204 B2	08/09/2015
WO 2010-039954 A1	08/04/2010
WO 2010-039966 A1	08/04/2010
WO 2010-039976 A1	08/04/2010