

(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).

Declarations under Rule 4.17:

- *as to the identity of the inventor (Rule 4.17(i))*
- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

SCENE GRAPH GENERATION FOR UNLABELED DATA

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This PCT application claims priority to U.S. Provisional Patent Application Serial Number 63/030,734, filed May 27, 2020, and U.S. Patent Application No. 17/226,561, filed on April 9, 2021, both of which are hereby incorporated herein in their entirety and for all intents and purposes.

BACKGROUND

[0002] An increasing number of applications and services rely on techniques such as object recognition and computer vision. In order to provide accurate results, these techniques often involve the generation of scene graphs and the learning of visual relationships for objects represented in image or video content. Conventional scene graph generation techniques rely on the availability of expensive annotations or labels. Attempting to generate scene graphs without such annotations, such as by using unlabeled real data, is a daunting task. Generating synthetic data is a viable alternative and has emerged as a promising direction for scaling supervised learning since generated labels can be obtained from the data synthesis process. It is still challenging, however, for networks or models trained using synthetic data to perform reliably for complex tasks such as scene graph generation when evaluated on real data due, due to issues relating to a performance lag or domain gap.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] Various embodiments in accordance with the present disclosure will be described with reference to the drawings, in which:

[0004] FIGS. 1A and 1B illustrate images that can be utilized to generate scene graphs, according to at least one embodiment;

[0005] FIG. 2 illustrate a scene graph generation system, according to at least one embodiment;

[0006] FIGS. 3A and 3B illustrate an image and corresponding scene graph, according to at least one embodiment;

[0007] FIGS. 4A and 4B illustrate an image and corresponding scene graph, according to at least one embodiment;

[0008] FIG. 5 illustrates a process for training a scene graph generator, according to at least one embodiment;

[0009] FIG. 6 illustrates components of a system for generating a scene graph, according to at least one embodiment;

[0010] FIG. 7A illustrates inference and/or training logic, according to at least one embodiment;

[0011] FIG. 7B illustrates inference and/or training logic, according to at least one embodiment;

[0012] FIG. 8 illustrates an example data center system, according to at least one embodiment;

[0013] FIG. 9 illustrates a computer system, according to at least one embodiment;

[0014] FIG. 10 illustrates a computer system, according to at least one embodiment;

[0015] FIG. 11 illustrates at least portions of a graphics processor, according to one or more embodiments;

[0016] FIG. 12 illustrates at least portions of a graphics processor, according to one or more embodiments;

[0017] FIG. 13 is an example data flow diagram for an advanced computing pipeline, in accordance with at least one embodiment;

[0018] FIG. 14 is a system diagram for an example system for training, adapting, instantiating and deploying machine learning models in an advanced computing pipeline, in accordance with at least one embodiment; and

[0019] FIGS. 15A and 15B illustrate a data flow diagram for a process to train a machine learning model, as well as client-server architecture to enhance annotation tools with pre-trained annotation models, in accordance with at least one embodiment.

DETAILED DESCRIPTION

[0020] Approaches in accordance with various embodiments can provide for the generation and use of scene graphs or other such data representations. In particular, various embodiments relate to scene graph generation for transfer learning. Since conventional scene graph training approaches utilize labeled data that can be costly or difficult to obtain, it can be advantageous to utilize synthetic data that can provide labels at little to no extra cost. Unfortunately, models trained using synthetic data generally do not perform well when running inferences on real data due, at least in part, to domain gap. In order to overcome this gap problem, a scene generation technique can be utilized that decomposes the domain gap into individual types of discrepancies, such as may relate to appearance, label, and prediction discrepancies. This can include minimizing the appearance and prediction discrepancies by aligning the corresponding latent and output distributions using gradient reversal layers. Label discrepancies can be addressed using principles of self-learning, but using pseudo-statistics collected from target data and leveraging a synthetic data generator to produce accurate labels. Pseudo-statistic based self-learning and adversarial techniques are used to manage these discrepancies without the need for costly supervision from a real-world dataset.

[0021] When training a model or network to generate scene graphs, it can be desirable to be able to generate scene graphs for input that may correspond to different domains. In this context, domains can refer to different settings, categories, styles, or spaces, among other such options. For example, a first set of images may include images taken in a first location, such as metropolitan areas in the United States. As illustrated in the example image 100 of FIG. 1A, these images may include objects of specific types that are of interest for a particular application, as may include representations of vehicles 102 and pedestrians 104, among other such object types that may be present in such locations. These images may include representations of other

types of objects as well, as may include buildings, road signs, billboards, and so on. For labeled data, this image may also include, or have associated, a set of bounding boxes that illustrate the general locations or regions of this image that correspond to individual instances of these object types of interest. These bounding boxes, or other such location identifiers, can be used to help train a model or network to identify these types of objects in similar images. The network can then learn to differentiate these objects with high accuracy when surrounded by the buildings and other types of objects represented in this general domain, or in metropolitan areas in the United States.

[0022] It is possible, however, that it may be desirable to use such a trained model to perform similar classification or inferencing on data from a different domain. This might include, for example, a domain corresponding to the European alpine region, which may include images such as the image 150 illustrated in FIG. 1B. As illustrated, these images can include very different types of background and other objects, as may relate to mountains, farm animals, mountain cabins, and other objects that may have very different appearances than what was present in images from the first domain. While there may be similar objects of interest in these images from the second domain, the model or network may not be as accurate in identifying these objects and differentiating them from the surrounding image regions, as the model was not trained using images with such representations. In order to improve the accuracy of such a network, such as a scene graph generation network, it may be desirable to train that network using the data from the first domain, but also data from the second domain. In this example, however, the data from the second domain is not labeled. As mentioned, going through and manually annotating these images can be expensive and time consuming, particularly for a large number of images that may be needed to accurately train the model.

[0023] In some instances, synthetic images can be generated that include representations of objects of the types of interest. Since synthetic images can be generated using scene graphs, for example, these synthetic images can include the labels that are already known for the types of objects included in these synthesized images. It may also be possible to synthesize images targeted to a specific domain, such as European alpine regions. Unfortunately, this still results in

different domains as there will be images in a first domain (here a synthetic domain), and images in a second domain (here a “real” image domain). Thus, even if a model is trained on labeled synthetic images for a target domain, the performance may still not be sufficiently accurate for use on real images captured or otherwise obtained for that target domain. As mentioned, this is due in large part to the domain gap between the domains, which can include both content and appearance gap portions. An appearance gap can represent differences in visual appearance between real and synthetic images, or images from different domains. A content gap can represent differences in placement, number, selection, and type of object in an image, among other such options. In at least some embodiments, an appearance gap can be addressed by making scenes photo-realistic, such as by using one or more image translations, using feature alignment, or learning robust representations based on domain randomization. Approaches can be utilized to address the content gap for image classification as well.

[0024] Accordingly, approaches in accordance with various embodiments can provide scalable methods for training scene graph generation models, networks, or other such approaches that addresses such a domain gap. These approaches can further utilize both labeled synthetic data and unlabeled real data. Approaches in accordance with various embodiments can provide tractable ways of reducing gaps of appearance and content that may occur between domains, such as between synthetic and real data. Such approaches can also provide for the learning of neural networks to generate scene graphs for unlabeled real images. This can include using synthetic data for learning to generate scene graphs (SGs) on real images, or images from a different domain. Such approaches can utilize a generation pipeline including pseudo statistics to manage the domain gap. In at least one embodiment, scene graph generation can be learned for various applications or scenarios, such as for computer vision, object recognition, or navigation assistance, among other such options. Applications can be extended to where synthetic data can be generated with only a raw estimate of the target dataset distribution.

[0025] As mentioned, scene graphs (SGs) can be used in applications such as computer vision and computer graphics to provide an interpretable and structural representation of scenes. A scene graph can be used to identify or summarize entities in a scene, as well as plausible relationships

between or among them. Scene graphs can be presented as a manifestation of vision through inverse graphics. Scene graphs have proven to be beneficial for a variety of applications, including but not limited to image captioning, visual question answering, high level reasoning tasks, image retrieval, image generation, and so on. Much of this prior work on scene graph generation relies on the availability of expensive and limited labeled datasets, such as the Visual Genome and Visual Relationship Dataset (VRD). A significant limitation for at least some machine learning applications is the general lack of sufficient labeled data for supervised learning tasks. Synthetic data can provide a viable alternative to this problem since annotations are essentially free, as being part of a synthesizing process. Synthetic data has been used for a variety of tasks, such as image classification, object detection, semantic segmentation, optical flow modeling, 3D key point extraction, object pose estimation, and 3D reconstruction, among others. Synthetic data has also shown to be effective in initializing task networks and for data augmentation. Synthetic data has not yet been utilized, at least fully utilized, for scene graph generation and representation of visual relationships.

[0026] In at least one embodiment, gaps can be assumed to exist in both appearance and content, and those gaps can be expanded into various addressable sub-components. Such an approach can help to reduce the domain gap for scene graph generation, and these techniques can also be applied to other vision-related tasks, such as image classification, image segmentation, and object detection, among others. In at least one embodiment, a simulation-to-scene graph (SimS2G) approach can be utilized, wherein a model learns sim-to-real scene graph generation leveraging labeled synthetic data and unlabeled real data. A Sim2SG approach can address the domain gap by bounding error for the task, such as scene graph generation, on real data through appearance, prediction, and label (*e.g.*, ground truth) discrepancies between the two domains, as well as task error on synthetic data. Such an approach can minimize the appearance and prediction discrepancies by aligning the corresponding latent and output distributions via one or more gradient reversal layers (GRLs). In at least one embodiment, the discrepancy in labeling can be addressed using principles of self-learning. However, self-learning based on pseudo labels may generate inaccurate labels that result performance issues. Accordingly, approaches in accordance

with various embodiments can collect pseudo-statistics from target data and leverage the synthetic data generator to produce accurate labels.

[0027] An example pipeline 200 that can be utilized for scene graph generation is illustrated in FIG. 2. As illustrated, this pipeline can receive input from at least two domains, such as may relate to labeled synthetic images 202, as well as unlabeled “real” or captured images 204. This example pipeline can function in three stages, although other pipelines or implementations might use a single stage or other number of stages. In a first stage, this pipeline accepts the labeled synthetic images 202, or images from a first domain, and processes these images using an encoder 206 to determine relevant features of these images and encode those features into a latent space or representation 208. This pipeline can then process this latent space with a scene graph predictor 210 to attempt to generate an appropriate scene graph for each of those input images. In this example, both the encoder 206 and the scene graph predictor 210 can be neural networks, such as a convolutional neural network (CNN) for the encoder 206 and a generative neural network for the scene graph predictor 210. In this first state, this pipeline can learn to generate scene graphs for the labeled synthetic images.

[0028] In a second stage, this pipeline 200 can process the unlabeled real images 204. The unlabeled real images 204 are processed using the encoder 206 to determine relevant features and encode those features into a latent space or representation, which here is a shared representation Z 208 that serves as a single latent space for features of both the labeled synthetic images 202 and the unlabeled real images 204. In at least one embodiment, the encoder 206 maps the labeled synthetic images and unlabeled real images to this shared representation Z 208. An encoder in at least one embodiment can be a ResNet-101, ResNet-50 or VGGNet (*e.g.*, VGG-16) network, among others, and different encoders may be utilized for different types or categories of input data in at least some embodiments.

[0029] In a third stage, the scene graph predictor model 210 generates one or more scene graphs from the unlabeled real data in the shared latent space 208. One or more pseudo statistics can then be calculated for these scene graphs. In at least one embodiment, these pseudo statistics can

represent a high-level statistical description of the scene, which can then be fed back into the generator, here the scene graph predictor 210. These pseudo statistics can then be used to attempt to generate aligned scenes where synthetic data more closely mimics real data, such as where synthetic images have an appearance that is closer to that of real, captured images. This aligned data, such as may include generated scene graph data 216, can then be fed to a renderer (for image data) and used as additional data for training the encoder 206 and the generator or scene graph predictor 210, such as a graph convolution network. This data can be used to train these networks again for another epoch. After this additional training, the networks can be used to generate another round of aligned data, and the process can continue until an end criterion is satisfied, such as the network converging, a maximum number of training iterations or epochs being reached, or all training data being used for training, among other such options.

[0030] In such a process, a scene graph prediction network P can be trained on a shared representation Z using synthetic data. The domain gap can be bridged by aligning features in the representation space Z (for the appearance gap) using a first gradient reversal layer (GRL) 212 and in the output space (for the content gap) using a second GRL 214, along with a domain discriminator. The label discrepancies (for the content gap) can be aligned between the two domains using pseudo statistics-based self-learning. It should be noted that the arrows in the figure correspond to flows of synthetic data, real data, and/or back propagation. A GRL can be a feature layer before the final scene graph prediction layer, where a type of discriminator will attempt to determine a domain for the features, such as whether they correspond to a real domain or a synthesized domain.

[0031] Such an approach can be beneficial at least because an initial scene graph prediction based on synthetic pre-training is not going to be accurate due, at least in part, to the domain gap. After an initial scene graph prediction, pseudo statistics can be determined, such as may relate to a number or placement of objects in a scene, and this can be used to help generate slightly more aligned data in a next iteration or epoch. Better aligned data can help improve performance of the synthetic prediction network, as well as to improve performance in predicting or obtaining pseudo-statistics of real data. With time, such a self-learning based approach can predict data

that is increasingly aligned with real data, such that the network will converge and the synthetic data will be a very accurate representation of real data.

[0032] As mentioned, in a third state this synthetic data that now accurately resembles real data can have the appearance and prediction features aligned between the two domains. GRLs for the appearance loss and prediction loss can be utilized, and the distribution of feature matches at the encoder level and at the output level. Using such a process, appearance alignment will actually generate features that are invariant, or appearance invariant, to both domains. The prediction loss can also be used to teach the network to be consistent in terms of both domains. Once trained, a certain set of features categorized as “chair” in domain A should also be categorized as “chair” in domain B, which can help to be consistent with respect to prediction loss.

[0033] In at least one embodiment, a SG predictor module can be trained using labels from a source domain. A scene graph in at least one embodiment can effectively be a collection of bounding boxes that serve as nodes of the graph, with the edges of the graph representing the relationships between those objects. A scene graph of a given image I can be noted as a graph G with nodes o and edges r . The node can be a tuple $o_i = \langle b_i, c_i \rangle$ of bounding box $b_i = \{xmin_i, ymin_i, w_i, h_i\}$ and category c_i . Relationships r are a triplet of $\langle o_i, p, o_j \rangle$ where p is a predicate. In at least one embodiment, scene graph prediction has at least two key components, including extractor ϕ and graph predictor P . Extractor ϕ maps input space x to a latent space z , and predictor P maps from latent space z to output space $y(G)$. The predicted SG is $G = P(\phi(x))$. A network such as Resnet-101 can be used to implement ϕ and an architecture such as GraphRCNN used to implement P . These networks ϕ and P can be trained using the following task loss: cross entropy loss for object classification and relationship classification, and l_1 loss for bounding boxes, such as discussed with respect to FIG. 2. In at least some embodiments, relationships among objects in the scene are predicted but not their attributes.

[0034] In at least one embodiment, the domain gap can be analyzed between source and target domains to formulate task error (e.g., for SG generation) on a real domain as a function of the

task error on a synthetic domain, as well as an appearance and content gap between the domains. Each sub-component of this domain gap can then be addressed.

[0035] In an example representing a domain gap, two domains will be discussed including a synthetic domain $\langle x_s, y_s \rangle \sim p(x, y)$ and real domain $\langle x_r, y_r \rangle \sim q(x, y)$. In this example y_r is not known. This example also assumes that the synthetic and real domains contain the same categories of objects. Both domains also share similar scenarios (e.g. both have driving scenes). However, joint distributions of scenes (e.g., images) are different in both domains (i.e. $p(x) \neq q(x)$). The label (e.g., ground truth) distributions are also different in the two domains (i.e. $p(y) \neq q(y)$). This task (e.g., SG generation) error can then be expressed in the synthetic domain as a function of the latent space z :

$$\epsilon_s(\phi, \gamma) = \int p(z) e_s dz$$

where e_s is the risk defined as $e_s = |p(y | z) - y_s|$. Here, $p(z)$ represents the distribution of features, $p(y)$ is the distribution of labels, and $p(y | z)$ is the output distribution. Since there can be access to the label y_s , the task error can be minimized. There may not, however, be access to the label distribution $q(y)$ for the real domain. Therefore, the task error or risk can be bound in the real domain $\epsilon_r(\phi, P)$ by tractable components as follows:

$$\begin{aligned} \epsilon_r(\phi, \gamma) &= \int q(z) e_r dz = \epsilon_s(\phi, \gamma) + \int (q(z) - p(z)) e_s dz + \int q(z) (e_r - e_s) dz \\ &= \epsilon_s(\phi, \gamma) + \epsilon^a(\phi, \gamma) + \epsilon^c(\phi, \gamma) \end{aligned}$$

[0036] Using simplification in the above equation, it can be shown that task error on the real domain is a function of three terms: the task error on the synthetic domain $\epsilon_s(\phi, P)$, feature discrepancy between the two domains $\epsilon^a(\phi, P)$, and risk discrepancy between the domains $\epsilon^c(\phi, P)$. For simplicity, terms ϕ and P will be dropped from the discussion going forward. As features frequently account for appearance, feature discrepancy ϵ^a as can be referred to as an appearance gap, and empirically shown through ablations. Similarly analyzing risk discrepancy ϵ^c and ablations can relate to the content gap. Nonetheless, an approach in accordance with at least

one embodiment can attempt to minimize ϵ_s , ϵ^a , and ϵ^c to reduce the task error on the real domain. If the error ϵ_r reduces to zero on the target domain, then the approach has effectively closed the domain gap. The first term ϵ_s may be is straightforward to minimize discussed herein.

[0037] With respect to the second term, the appearance gap refers to the discrepancy in the appearance of the two domains. This can include differences in, for example, the texture, color, light, and reflectance of objects in a scene. To address the appearance gap, one approach can attempt to avoid photo-realism in synthetic data as may require high quality assets and a tremendous amount of effort from artists. Accordingly, an appearance invariant representation z can be learned to avoid appearance bias in the model. However, the features $z : \varphi(x)$ have both content and appearance components, and aligning them may be detrimental. To overcome this challenge, the appearance gap ϵ^a can be aligned only after aligning the content gap ϵ^c of the two domains as discussed herein. This appearance gap ϵ^a can be minimized by aligning the feature distributions $p(z)$ and $q(z)$ as follows:

$$\theta^* = \arg_{\theta} \min \int (q(z) - p(z)) e_s dz$$

where the fact that $p(z) \equiv q(z)$ is a sufficient condition for ϵ^a to be zero is exploited. A gradient reversal layer (GRL) can be used to align the distributions $p(z)$ and $q(z)$, along with a domain classifier D to classify them. The loss in D can be minimized with respect to its own parameters, while maximizing with respect to the network parameters of $\varphi(\theta)$. This can be performed using a GRL that acts as an identity function during forward propagation and flips the sign of the gradients during back propagation from D to φ .

[0038] With respect to the third term, the content gap refers to discrepancies between the two domains, including the difference in distribution of number of objects and their class, placement, pose, and scale. This can lead to a problem of label shift, or discrepancy in the ground truth distribution. It can be desirable to understand this content discrepancy in at least some embodiments. Since minimizing ϵ^c in the current form is not tractable, risk discrepancy ($e_r - e_s$) can be reduced as follows:

$$\begin{aligned}\arg_{\theta} \min \epsilon^c &= \arg_{\theta} \min \int q(z)(y_s - y_r)dz + \int q(z)(q(y|z) - p(y|z))dz \\ &= \arg_{\theta} \min \epsilon^{c,label} + \epsilon^{c,pred}\end{aligned}$$

The risk can be assumed to be positive, and the risk dependency split further into label discrepancy $\epsilon^{c,pred}$.

[0039] To address the prediction discrepancy $\epsilon^{c,pred}$, the output distribution can be aligned for latent code belonging to the same visual category. This is reasonable since the output of the scene graph generation model should be the same for same categories in different domains. The term $\epsilon^{c,pred}$ can be handled by aligning the output distributions $p(y | z)$ and $q(y | z)$ using the same GRL-based technique discussed previously.

[0040] On the other hand, minimizing $\epsilon^{c,label}$ can be challenging because there may not be access to the label (y_r) of the target (e.g., real) domain. In at least one embodiment, self-learning using pseudo statistics can be utilized to obtain to get an estimate of the label distribution of real data. Self-learning based on pseudo labels may generate inaccurate labels because of prediction inaccuracies when evaluated on a different domain. However, approaches presented herein can produce accurate labels using a synthetic data generator. It can be noted that the pseudo statistics method can be orthogonal to pseudo labels based self-learning, and the latter can potentially be applied in conjunction.

[0041] In at least one embodiment, label alignment $\epsilon^{c,label}$ works at least in part by generating generate N scene graphs (SGs) for all N scenes (x_r) of the target domain $\{x_r\}_{r=1}^N$. For each predicted SG $\hat{G}_k$, where $k \in \{1, \dots, N\}$, the pseudo statistic ps_k can be extracted. The statistic ps_k includes a list of objects $\{o_1, o_2, \dots, o_j \dots\}$ and relationships $\{r_1, r_2, \dots, r_j \dots\}$ among them. Each object o_j has a bounding box b_j and a category c_j associated with it, and each relationship triplet r_j has two objects and a predicate as described herein. This pseudo statistic can then be used to create a scene using the synthetic data generator. From ps_k , the first object can be selected according to its category c_j . An asset can then be randomly selected from this category, and its two-dimensional bounding box b_j projected into a three-dimensional scene using camera

parameters. The camera parameters of the target data can be used if available, or diligent assumptions can be made as may relate to height or pitch, etc.

[0042] Pose can be chosen intelligently or at random. For instance, for a driving simulation environment the aspect ratio of a car's bounding box b_j can be used to either align it with lanes (0° or 180° rotation) or rotate it 90° . The same steps can be repeated for every object in the scene. These relationships can then be used for fine adjustments and correction of object positions and pose.

[0043] For instance, in a dining simulation environment there may be various objects visible in a scene, such as a table 302, chairs 304, 306, 308, and laptop computers 310, 312 visible in an image 300, as illustrated in FIG. 3A. In at least one embodiment, an "on" relationship can be used to specify that the laptops are to be placed on the table. Various assumptions can be made, such as a flat ground plane, all objects remaining on the ground, minimal context (*e.g.*, walls, ground, road, or other objects) and real world scale. A view of a three-dimensional scene can then be rendered using a synthetic data generator to obtain a synthetic image 300 and ground truth x_s, y_s . Synthetic data can be generated for all statistics p_{sk} to train the encoding and generation networks in the next epoch. This can be referred to as self-learning, as predictions on the unlabeled target are used to generate synthetic data similar to the target, as well as to train the model with the new data in next epoch. FIG. 3B illustrates an example scene graph 350 that could be used to generate such a scene, here indicating the positional relationships of the chairs and laptops to the table 302. Various other scene graphs could be generated for such a scene as well within the scope of the various embodiments.

[0044] FIGS. 4A and 4B illustrate another example image and scene graph, respectively, that can be generated in accordance with various embodiments. In the example image 400 of FIG. 4A, there are three objects represented, including a cylinder 402, a cube 404, and a sphere 406. The example scene graph 450 corresponding to this scene, as illustrated in FIG. 4B, can include positional relationship information for each pair of objects in the scene, instead of with respect to a single object in the scene. Such a scene graph may contain redundant information, such as the cylinder being to the right of the cube and the cube being to the left of the cylinder, but such

an approach can provide many paths that can be used to generate, and check, the relationships used to generate a scene from this scene graph.

[0045] In at least one embodiment, a process for training one or more networks, as may be part of a SimS2G pipeline, can assume an encoder and scene graph predictor, as well as the availability of synthetic images, synthetic labels, a data generator, and real images. For a given number of iterations or epochs, an attempt can be made to minimize a loss function that includes loss terms for the appearance gap, the content prediction gap, and the scene graph generation task. After the appropriate network parameters have been adjusted to minimize the loss, the appropriate pseudo statistics can be extracted from the predictions, and synthetic data can be generated with labels aligned to real data.

[0046] FIG. 5 illustrates an example process 500 that can be used to train a scene graph generation model in accordance with at least one embodiment. This process can utilize the loss function with terms for appearance, content, and SG generation terms discussed above. In this example, a set of labeled synthetic data is obtained 502, along with a set of unlabeled “real,” or non-synthetic, data. This data can be input to at least one encoder network, which can extract and encode 504 features of these data sets into a shared representation, such as a shared latent space. This latent space can be provided 506 as input for use in training a scene graph prediction network, which can predict scene graphs based on features from this latent space. In this example, features in this latent space can be aligned with an output space of the scene graph prediction network. This can result in an alignment of real and synthetic features such that synthetic data (*e.g.*, images) generated using inferred scene graphs may be indistinguishable from real data. As mentioned, this can also be utilized to minimize the gap between other domains as well. The labels for objects in these scene graphs can also be aligned 510 using pseudo statistic-based self-learning as discussed herein, so that synthetic data can be generated with labels aligned to those for real data. The scene graph generation model can be trained 512 until an end criterion is satisfied, such as the network converging or the training proceeding for a determined number of epochs (*e.g.*, three epochs). Once trained, the pipeline can be used for various purposes, such as to generate synthetic data for simulation or

other such purposes. Such an approach can also be used to generate scene graphs for unlabeled data from a different domain. Given data for a real scene without data labels, a scene graph can be generated that can be used to synthesize a labeled version of that scene.

[0047] Such an approach may be beneficial for any content creation application, such as for online gaming, where it may be desirable to generate realistic synthetic data. Such an approach can also be used to automatically build simulation environments, as may be useful for robotics or navigation simulation. Approaches in accordance with various embodiments are also not limited to scene graphs, but can be used with any relevant representation of two or more domains, particularly where insufficient labeled training data is available for one of those domains. Applications that can benefit from such functionality can include area/room scanning, security, space analysis, anomaly detection, and realism simulation, among many others.

[0048] As an example, FIG. 6 illustrates an example network configuration 600 that can be used to provide or generate content. In at least one embodiment, a client device 602 can generate content for a session using components of a content application 604 on client device 602 and data stored locally on that client device. In at least one embodiment, a content application 624 (*e.g.*, an image generation or editing application) executing on content server 620 (*e.g.*, a cloud server or edge server) may initiate a session associated with at least client device 602, as may utilize a session manager and user data stored in a user database 634, and can cause content 632 to be determined by a content manager 626 and rendered using a rendering engine, if needed for this type of content or platform, and transmitted to client device 602 using an appropriate transmission manager 622 to send by download, streaming, or another such transmission channel. In at least one embodiment, this content 632 can include assets that can be used by a rendering engine to render a scene based on a determined scene graph. The content server 620 can also include one or more training modules 630 for training a scene graph generation component, network, or pipeline 628. In at least one embodiment, client device 602 receiving this content can provide this content to a corresponding content application 604, which may also or alternatively include a rendering engine (if necessary) for rendering at least some of this content for presentation via client device 602, such as image or video content through a display

606 and audio, such as sounds and music, through at least one audio playback device 608, such as speakers or headphones. For live video content captured by one or more cameras, for example, such a rendering engine may not be needed, unless used to augment that video content in some way.

[0049] In at least one embodiment, at least some of this content may already be stored on, rendered on, or accessible to client device 602 such that transmission over network 640 is not required for at least that portion of content, such as where that content may have been previously downloaded or stored locally on a hard drive or optical disk. In at least one embodiment, a transmission mechanism such as data streaming can be used to transfer this content from server 620, or content database 634, to client device 602. In at least one embodiment, at least a portion of this content can be obtained or streamed from another source, such as a third party content service 660 that may also include a content application 662 for generating or providing content. In at least one embodiment, portions of this functionality can be performed using multiple computing devices, or multiple processors within one or more computing devices, such as may include a combination of CPUs and GPUs.

[0050] In at least one embodiment, content application 624 includes a content manager 626 that can determine or analyze content before this content is transmitted to client device 602. In at least one embodiment, content manager 626 can also include, or work with, other components that are able to generate, modify, or enhance content to be provided. In at least one embodiment, this can include a rendering engine for rendering image or video content. In at least one embodiment, an image, video, or scene graph generation component 628 can be used to generate a scene graph, which can be used by the content application 624 on the server or the content application 604 on the client device to generate image, video, or other media content. In at least one embodiment, an enhancement component 630, which can also include a neural network, can perform one or more enhancements on this content, as discussed and suggested herein. In at least one embodiment, content manager 626 can cause this content to be transmitted to client device 602. In at least one embodiment, a content application 604 on client device 602 may also include components such as a rendering engine, scene graph generator 612, and content

generation module, such that any or all of this functionality can additionally, or alternatively, be performed on client device 602. In at least one embodiment, a content application 662 on a third party content service system 660 can also include such functionality. In at least one embodiment, locations where at least some of this functionality is performed may be configurable, or may depend upon factors such as a type of client device 602 or availability of a network connection with appropriate bandwidth, among other such factors. In at least one embodiment, a system for content generation can include any appropriate combination of hardware and software in one or more locations. In at least one embodiment, generated image or video content of one or more resolutions can also be provided, or made available, to other client devices 650, such as for download or streaming from a media source storing a copy of that image or video content. In at least one embodiment, this may include transmitting images of game content for a multiplayer game, where different client devices may display that content at different resolutions, including one or more super-resolutions.

[0051] In this example, these client devices can include any appropriate computing devices, as may include a desktop computer, notebook computer, edge device, set-top box, streaming device, gaming console, smartphone, tablet computer, VR headset, AR goggles, wearable computer, or a smart television. Each client device can submit a request across at least one wired or wireless network, as may include the Internet, an Ethernet, a local area network (LAN), or a cellular network, among other such options. In this example, these requests can be submitted to an address associated with a cloud provider, who may operate or control one or more electronic resources in a cloud provider environment, such as may include a data center or server farm. In at least one embodiment, the request may be received or processed by at least one edge server, that sits on a network edge and is outside at least one security layer associated with the cloud provider environment. In this way, latency can be reduced by enabling the client devices to interact with servers that are in closer proximity, while also improving security of resources in the cloud provider environment.

[0052] In at least one embodiment, such a system can be used for performing graphical rendering operations. In other embodiments, such a system can be used for other purposes, such

as for providing image or video content to test or validate autonomous machine applications, or for performing deep learning operations. In at least one embodiment, such a system can be implemented using an edge device, or may incorporate one or more Virtual Machines (VMs). In at least one embodiment, such a system can be implemented at least partially in a data center or at least partially using cloud computing resources.

INFERENCE AND TRAINING LOGIC

[0053] FIG. 7A illustrates inference and/or training logic 715 used to perform inferencing and/or training operations associated with one or more embodiments. Details regarding inference and/or training logic 715 are provided below in conjunction with FIGS. 7A and/or 7B.

[0054] In at least one embodiment, inference and/or training logic 715 may include, without limitation, code and/or data storage 701 to store forward and/or output weight and/or input/output data, and/or other parameters to configure neurons or layers of a neural network trained and/or used for inferencing in aspects of one or more embodiments. In at least one embodiment, training logic 715 may include, or be coupled to code and/or data storage 701 to store graph code or other software to control timing and/or order, in which weight and/or other parameter information is to be loaded to configure, logic, including integer and/or floating point units (collectively, arithmetic logic units (ALUs)). In at least one embodiment, code, such as graph code, loads weight or other parameter information into processor ALUs based on an architecture of a neural network to which the code corresponds. In at least one embodiment, code and/or data storage 701 stores weight parameters and/or input/output data of each layer of a neural network trained or used in conjunction with one or more embodiments during forward propagation of input/output data and/or weight parameters during training and/or inferencing using aspects of one or more embodiments. In at least one embodiment, any portion of code and/or data storage 701 may be included with other on-chip or off-chip data storage, including a processor's L1, L2, or L3 cache or system memory.

[0055] In at least one embodiment, any portion of code and/or data storage 701 may be internal or external to one or more processors or other hardware logic devices or circuits. In at least one

embodiment, code and/or code and/or data storage 701 may be cache memory, dynamic randomly addressable memory (“DRAM”), static randomly addressable memory (“SRAM”), non-volatile memory (e.g., Flash memory), or other storage. In at least one embodiment, choice of whether code and/or code and/or data storage 701 is internal or external to a processor, for example, or comprised of DRAM, SRAM, Flash or some other storage type may depend on available storage on-chip versus off-chip, latency requirements of training and/or inferencing functions being performed, batch size of data used in inferencing and/or training of a neural network, or some combination of these factors.

[0056] In at least one embodiment, inference and/or training logic 715 may include, without limitation, a code and/or data storage 705 to store backward and/or output weight and/or input/output data corresponding to neurons or layers of a neural network trained and/or used for inferencing in aspects of one or more embodiments. In at least one embodiment, code and/or data storage 705 stores weight parameters and/or input/output data of each layer of a neural network trained or used in conjunction with one or more embodiments during backward propagation of input/output data and/or weight parameters during training and/or inferencing using aspects of one or more embodiments. In at least one embodiment, training logic 715 may include, or be coupled to code and/or data storage 705 to store graph code or other software to control timing and/or order, in which weight and/or other parameter information is to be loaded to configure, logic, including integer and/or floating point units (collectively, arithmetic logic units (ALUs)). In at least one embodiment, code, such as graph code, loads weight or other parameter information into processor ALUs based on an architecture of a neural network to which the code corresponds. In at least one embodiment, any portion of code and/or data storage 705 may be included with other on-chip or off-chip data storage, including a processor’s L1, L2, or L3 cache or system memory. In at least one embodiment, any portion of code and/or data storage 705 may be internal or external to on one or more processors or other hardware logic devices or circuits. In at least one embodiment, code and/or data storage 705 may be cache memory, DRAM, SRAM, non-volatile memory (e.g., Flash memory), or other storage. In at least one embodiment, choice of whether code and/or data storage 705 is internal or external to a

processor, for example, or comprised of DRAM, SRAM, Flash or some other storage type may depend on available storage on-chip versus off-chip, latency requirements of training and/or inferencing functions being performed, batch size of data used in inferencing and/or training of a neural network, or some combination of these factors.

[0057] In at least one embodiment, code and/or data storage 701 and code and/or data storage 705 may be separate storage structures. In at least one embodiment, code and/or data storage 701 and code and/or data storage 705 may be same storage structure. In at least one embodiment, code and/or data storage 701 and code and/or data storage 705 may be partially same storage structure and partially separate storage structures. In at least one embodiment, any portion of code and/or data storage 701 and code and/or data storage 705 may be included with other on-chip or off-chip data storage, including a processor's L1, L2, or L3 cache or system memory.

[0058] In at least one embodiment, inference and/or training logic 715 may include, without limitation, one or more arithmetic logic unit(s) ("ALU(s)") 710, including integer and/or floating point units, to perform logical and/or mathematical operations based, at least in part on, or indicated by, training and/or inference code (e.g., graph code), a result of which may produce activations (e.g., output values from layers or neurons within a neural network) stored in an activation storage 720 that are functions of input/output and/or weight parameter data stored in code and/or data storage 701 and/or code and/or data storage 705. In at least one embodiment, activations stored in activation storage 720 are generated according to linear algebraic and or matrix-based mathematics performed by ALU(s) 710 in response to performing instructions or other code, wherein weight values stored in code and/or data storage 705 and/or code and/or data storage 701 are used as operands along with other values, such as bias values, gradient information, momentum values, or other parameters or hyperparameters, any or all of which may be stored in code and/or data storage 705 or code and/or data storage 701 or another storage on or off-chip.

[0059] In at least one embodiment, ALU(s) 710 are included within one or more processors or other hardware logic devices or circuits, whereas in another embodiment, ALU(s) 710 may be

external to a processor or other hardware logic device or circuit that uses them (e.g., a co-processor). In at least one embodiment, ALUs 710 may be included within a processor's execution units or otherwise within a bank of ALUs accessible by a processor's execution units either within same processor or distributed between different processors of different types (e.g., central processing units, graphics processing units, fixed function units, etc.). In at least one embodiment, code and/or data storage 701, code and/or data storage 705, and activation storage 720 may be on same processor or other hardware logic device or circuit, whereas in another embodiment, they may be in different processors or other hardware logic devices or circuits, or some combination of same and different processors or other hardware logic devices or circuits. In at least one embodiment, any portion of activation storage 720 may be included with other on-chip or off-chip data storage, including a processor's L1, L2, or L3 cache or system memory. Furthermore, inferencing and/or training code may be stored with other code accessible to a processor or other hardware logic or circuit and fetched and/or processed using a processor's fetch, decode, scheduling, execution, retirement and/or other logical circuits.

[0060] In at least one embodiment, activation storage 720 may be cache memory, DRAM, SRAM, non-volatile memory (e.g., Flash memory), or other storage. In at least one embodiment, activation storage 720 may be completely or partially within or external to one or more processors or other logical circuits. In at least one embodiment, choice of whether activation storage 720 is internal or external to a processor, for example, or comprised of DRAM, SRAM, Flash or some other storage type may depend on available storage on-chip versus off-chip, latency requirements of training and/or inferencing functions being performed, batch size of data used in inferencing and/or training of a neural network, or some combination of these factors. In at least one embodiment, inference and/or training logic 715 illustrated in FIG. 7a may be used in conjunction with an application-specific integrated circuit ("ASIC"), such as Tensorflow® Processing Unit from Google, an inference processing unit (IPU) from Graphcore™, or a Nervana® (e.g., "Lake Crest") processor from Intel Corp. In at least one embodiment, inference and/or training logic 715 illustrated in FIG. 7a may be used in conjunction with central

processing unit (“CPU”) hardware, graphics processing unit (“GPU”) hardware or other hardware, such as field programmable gate arrays (“FPGAs”).

[0061] FIG. 7b illustrates inference and/or training logic 715, according to at least one or more embodiments. In at least one embodiment, inference and/or training logic 715 may include, without limitation, hardware logic in which computational resources are dedicated or otherwise exclusively used in conjunction with weight values or other information corresponding to one or more layers of neurons within a neural network. In at least one embodiment, inference and/or training logic 715 illustrated in FIG. 7b may be used in conjunction with an application-specific integrated circuit (ASIC), such as Tensorflow® Processing Unit from Google, an inference processing unit (IPU) from Graphcore™, or a Nervana® (e.g., “Lake Crest”) processor from Intel Corp. In at least one embodiment, inference and/or training logic 715 illustrated in FIG. 7b may be used in conjunction with central processing unit (CPU) hardware, graphics processing unit (GPU) hardware or other hardware, such as field programmable gate arrays (FPGAs). In at least one embodiment, inference and/or training logic 715 includes, without limitation, code and/or data storage 701 and code and/or data storage 705, which may be used to store code (e.g., graph code), weight values and/or other information, including bias values, gradient information, momentum values, and/or other parameter or hyperparameter information. In at least one embodiment illustrated in FIG. 7b, each of code and/or data storage 701 and code and/or data storage 705 is associated with a dedicated computational resource, such as computational hardware 702 and computational hardware 706, respectively. In at least one embodiment, each of computational hardware 702 and computational hardware 706 comprises one or more ALUs that perform mathematical functions, such as linear algebraic functions, only on information stored in code and/or data storage 701 and code and/or data storage 705, respectively, result of which is stored in activation storage 720.

[0062] In at least one embodiment, each of code and/or data storage 701 and 705 and corresponding computational hardware 702 and 706, respectively, correspond to different layers of a neural network, such that resulting activation from one “storage/computational pair 701/702” of code and/or data storage 701 and computational hardware 702 is provided as an

input to “storage/computational pair 705/706” of code and/or data storage 705 and computational hardware 706, in order to mirror conceptual organization of a neural network. In at least one embodiment, each of storage/computational pairs 701/702 and 705/706 may correspond to more than one neural network layer. In at least one embodiment, additional storage/computation pairs (not shown) subsequent to or in parallel with storage computation pairs 701/702 and 705/706 may be included in inference and/or training logic 715.

DATA CENTER

[0063] FIG. 8 illustrates an example data center 800, in which at least one embodiment may be used. In at least one embodiment, data center 800 includes a data center infrastructure layer 810, a framework layer 820, a software layer 830, and an application layer 840.

[0064] In at least one embodiment, as shown in FIG. 8, data center infrastructure layer 810 may include a resource orchestrator 812, grouped computing resources 814, and node computing resources (“node C.R.s”) 816(1)-816(N), where “N” represents any whole, positive integer. In at least one embodiment, node C.R.s 816(1)-816(N) may include, but are not limited to, any number of central processing units (“CPUs”) or other processors (including accelerators, field programmable gate arrays (FPGAs), graphics processors, etc.), memory devices (e.g., dynamic read-only memory), storage devices (e.g., solid state or disk drives), network input/output (“NW I/O”) devices, network switches, virtual machines (“VMs”), power modules, and cooling modules, etc. In at least one embodiment, one or more node C.R.s from among node C.R.s 816(1)-816(N) may be a server having one or more of above-mentioned computing resources.

[0065] In at least one embodiment, grouped computing resources 814 may include separate groupings of node C.R.s housed within one or more racks (not shown), or many racks housed in data centers at various geographical locations (also not shown). Separate groupings of node C.R.s within grouped computing resources 814 may include grouped compute, network, memory or storage resources that may be configured or allocated to support one or more workloads. In at least one embodiment, several node C.R.s including CPUs or processors may grouped within one or more racks to provide compute resources to support one or more workloads. In at least one

embodiment, one or more racks may also include any number of power modules, cooling modules, and network switches, in any combination.

[0066] In at least one embodiment, resource orchestrator 812 may configure or otherwise control one or more node C.R.s 816(1)-816(N) and/or grouped computing resources 814. In at least one embodiment, resource orchestrator 812 may include a software design infrastructure (“SDI”) management entity for data center 800. In at least one embodiment, resource orchestrator may include hardware, software or some combination thereof.

[0067] In at least one embodiment, as shown in FIG. 8, framework layer 820 includes a job scheduler 822, a configuration manager 824, a resource manager 826 and a distributed file system 828. In at least one embodiment, framework layer 820 may include a framework to support software 832 of software layer 830 and/or one or more application(s) 842 of application layer 840. In at least one embodiment, software 832 or application(s) 842 may respectively include web-based service software or applications, such as those provided by Amazon Web Services, Google Cloud and Microsoft Azure. In at least one embodiment, framework layer 820 may be, but is not limited to, a type of free and open-source software web application framework such as Apache SparkTM (hereinafter “Spark”) that may utilize distributed file system 828 for large-scale data processing (*e.g.*, “big data”). In at least one embodiment, job scheduler 822 may include a Spark driver to facilitate scheduling of workloads supported by various layers of data center 800. In at least one embodiment, configuration manager 824 may be capable of configuring different layers such as software layer 830 and framework layer 820 including Spark and distributed file system 828 for supporting large-scale data processing. In at least one embodiment, resource manager 826 may be capable of managing clustered or grouped computing resources mapped to or allocated for support of distributed file system 828 and job scheduler 822. In at least one embodiment, clustered or grouped computing resources may include grouped computing resource 814 at data center infrastructure layer 810. In at least one embodiment, resource manager 826 may coordinate with resource orchestrator 812 to manage these mapped or allocated computing resources.

[0068] In at least one embodiment, software 832 included in software layer 830 may include software used by at least portions of node C.R.s 816(1)-816(N), grouped computing resources 814, and/or distributed file system 828 of framework layer 820. The one or more types of software may include, but are not limited to, Internet web page search software, e-mail virus scan software, database software, and streaming video content software.

[0069] In at least one embodiment, application(s) 842 included in application layer 840 may include one or more types of applications used by at least portions of node C.R.s 816(1)-816(N), grouped computing resources 814, and/or distributed file system 828 of framework layer 820. One or more types of applications may include, but are not limited to, any number of a genomics application, a cognitive compute, and a machine learning application, including training or inferencing software, machine learning framework software (*e.g.*, PyTorch, TensorFlow, Caffe, etc.) or other machine learning applications used in conjunction with one or more embodiments.

[0070] In at least one embodiment, any of configuration manager 824, resource manager 826, and resource orchestrator 812 may implement any number and type of self-modifying actions based on any amount and type of data acquired in any technically feasible fashion. In at least one embodiment, self-modifying actions may relieve a data center operator of data center 800 from making possibly bad configuration decisions and possibly avoiding underutilized and/or poor performing portions of a data center.

[0071] In at least one embodiment, data center 800 may include tools, services, software or other resources to train one or more machine learning models or predict or infer information using one or more machine learning models according to one or more embodiments described herein. For example, in at least one embodiment, a machine learning model may be trained by calculating weight parameters according to a neural network architecture using software and computing resources described above with respect to data center 800. In at least one embodiment, trained machine learning models corresponding to one or more neural networks may be used to infer or predict information using resources described above with respect to data center 800 by using weight parameters calculated through one or more training techniques described herein.

[0072] In at least one embodiment, data center may use CPUs, application-specific integrated circuits (ASICs), GPUs, FPGAs, or other hardware to perform training and/or inferencing using above-described resources. Moreover, one or more software and/or hardware resources described above may be configured as a service to allow users to train or performing inferencing of information, such as image recognition, speech recognition, or other artificial intelligence services.

[0073] Inference and/or training logic 715 are used to perform inferencing and/or training operations associated with one or more embodiments. Details regarding inference and/or training logic 715 are provided below in conjunction with FIGS. 7A and/or 7B. In at least one embodiment, inference and/or training logic 715 may be used in system FIG. 8 for inferencing or predicting operations based, at least in part, on weight parameters calculated using neural network training operations, neural network functions and/or architectures, or neural network use cases described herein.

[0074] Such components can be used to train one or more generation models to be able to generate scene graphs for unlabeled data, as well as to generate scenes based on those scene graphs that are indistinguishable from real scenes.

COMPUTER SYSTEMS

[0075] FIG. 9 is a block diagram illustrating an exemplary computer system, which may be a system with interconnected devices and components, a system-on-a-chip (SOC) or some combination thereof 900 formed with a processor that may include execution units to execute an instruction, according to at least one embodiment. In at least one embodiment, computer system 900 may include, without limitation, a component, such as a processor 902 to employ execution units including logic to perform algorithms for process data, in accordance with present disclosure, such as in embodiment described herein. In at least one embodiment, computer system 900 may include processors, such as PENTIUM® Processor family, Xeon™, Itanium®, XScale™ and/or StrongARM™, Intel® Core™, or Intel® Nervana™ microprocessors available from Intel Corporation of Santa Clara, California, although other systems (including PCs having other microprocessors, engineering workstations, set-top boxes and like) may also be used. In at least one embodiment, computer system 900 may execute a version of WINDOWS' operating system available from Microsoft Corporation of Redmond, Wash., although other operating systems (UNIX and Linux for example), embedded software, and/or graphical user interfaces, may also be used.

[0076] Embodiments may be used in other devices such as handheld devices and embedded applications. Some examples of handheld devices include cellular phones, Internet Protocol devices, digital cameras, personal digital assistants ("PDAs"), and handheld PCs. In at least one embodiment, embedded applications may include a microcontroller, a digital signal processor ("DSP"), system on a chip, network computers ("NetPCs"), set-top boxes, network hubs, wide area network ("WAN") switches, or any other system that may perform one or more instructions in accordance with at least one embodiment.

[0077] In at least one embodiment, computer system 900 may include, without limitation, processor 902 that may include, without limitation, one or more execution units 908 to perform machine learning model training and/or inferencing according to techniques described herein. In at least one embodiment, computer system 900 is a single processor desktop or server system, but in another embodiment computer system 900 may be a multiprocessor system. In at least

one embodiment, processor 902 may include, without limitation, a complex instruction set computer (“CISC”) microprocessor, a reduced instruction set computing (“RISC”) microprocessor, a very long instruction word (“VLIW”) microprocessor, a processor implementing a combination of instruction sets, or any other processor device, such as a digital signal processor, for example. In at least one embodiment, processor 902 may be coupled to a processor bus 910 that may transmit data signals between processor 902 and other components in computer system 900.

[0078] In at least one embodiment, processor 902 may include, without limitation, a Level 1 (“L1”) internal cache memory (“cache”) 904. In at least one embodiment, processor 902 may have a single internal cache or multiple levels of internal cache. In at least one embodiment, cache memory may reside external to processor 902. Other embodiments may also include a combination of both internal and external caches depending on particular implementation and needs. In at least one embodiment, register file 906 may store different types of data in various registers including, without limitation, integer registers, floating point registers, status registers, and instruction pointer register.

[0079] In at least one embodiment, execution unit 908, including, without limitation, logic to perform integer and floating point operations, also resides in processor 902. In at least one embodiment, processor 902 may also include a microcode (“ucode”) read only memory (“ROM”) that stores microcode for certain macro instructions. In at least one embodiment, execution unit 908 may include logic to handle a packed instruction set 909. In at least one embodiment, by including packed instruction set 909 in an instruction set of a general-purpose processor 902, along with associated circuitry to execute instructions, operations used by many multimedia applications may be performed using packed data in a general-purpose processor 902. In one or more embodiments, many multimedia applications may be accelerated and executed more efficiently by using full width of a processor’s data bus for performing operations on packed data, which may eliminate need to transfer smaller units of data across processor’s data bus to perform one or more operations one data element at a time.

[0080] In at least one embodiment, execution unit 908 may also be used in microcontrollers, embedded processors, graphics devices, DSPs, and other types of logic circuits. In at least one embodiment, computer system 900 may include, without limitation, a memory 920. In at least one embodiment, memory 920 may be implemented as a Dynamic Random Access Memory (“DRAM”) device, a Static Random Access Memory (“SRAM”) device, flash memory device, or other memory device. In at least one embodiment, memory 920 may store instruction(s) 919 and/or data 921 represented by data signals that may be executed by processor 902.

[0081] In at least one embodiment, system logic chip may be coupled to processor bus 910 and memory 920. In at least one embodiment, system logic chip may include, without limitation, a memory controller hub (“MCH”) 916, and processor 902 may communicate with MCH 916 via processor bus 910. In at least one embodiment, MCH 916 may provide a high bandwidth memory path 918 to memory 920 for instruction and data storage and for storage of graphics commands, data and textures. In at least one embodiment, MCH 916 may direct data signals between processor 902, memory 920, and other components in computer system 900 and to bridge data signals between processor bus 910, memory 920, and a system I/O 922. In at least one embodiment, system logic chip may provide a graphics port for coupling to a graphics controller. In at least one embodiment, MCH 916 may be coupled to memory 920 through a high bandwidth memory path 918 and graphics/video card 912 may be coupled to MCH 916 through an Accelerated Graphics Port (“AGP”) interconnect 914.

[0082] In at least one embodiment, computer system 900 may use system I/O 922 that is a proprietary hub interface bus to couple MCH 916 to I/O controller hub (“ICH”) 930. In at least one embodiment, ICH 930 may provide direct connections to some I/O devices via a local I/O bus. In at least one embodiment, local I/O bus may include, without limitation, a high-speed I/O bus for connecting peripherals to memory 920, chipset, and processor 902. Examples may include, without limitation, an audio controller 929, a firmware hub (“flash BIOS”) 928, a wireless transceiver 926, a data storage 924, a legacy I/O controller 923 containing user input and keyboard interfaces 925, a serial expansion port 927, such as Universal Serial Bus (“USB”),

and a network controller 934. Data storage 924 may comprise a hard disk drive, a floppy disk drive, a CD-ROM device, a flash memory device, or other mass storage device.

[0083] In at least one embodiment, FIG. 9 illustrates a system, which includes interconnected hardware devices or “chips”, whereas in other embodiments, FIG. 9 may illustrate an exemplary System on a Chip (“SoC”). In at least one embodiment, devices may be interconnected with proprietary interconnects, standardized interconnects (e.g., PCIe) or some combination thereof. In at least one embodiment, one or more components of computer system 900 are interconnected using compute express link (CXL) interconnects.

[0084] Inference and/or training logic 715 are used to perform inferencing and/or training operations associated with one or more embodiments. Details regarding inference and/or training logic 715 are provided below in conjunction with FIGS. 7A and/or 7B. In at least one embodiment, inference and/or training logic 715 may be used in system FIG. 9 for inferencing or predicting operations based, at least in part, on weight parameters calculated using neural network training operations, neural network functions and/or architectures, or neural network use cases described herein.

[0085] Such components can be used to train one or more generation models to be able to generate scene graphs for unlabeled data, as well as to generate scenes based on those scene graphs that are indistinguishable from real scenes.

[0086] FIG. 10 is a block diagram illustrating an electronic device 1000 for utilizing a processor 1010, according to at least one embodiment. In at least one embodiment, electronic device 1000 may be, for example and without limitation, a notebook, a tower server, a rack server, a blade server, a laptop, a desktop, a tablet, a mobile device, a phone, an embedded computer, or any other suitable electronic device.

[0087] In at least one embodiment, system 1000 may include, without limitation, processor 1010 communicatively coupled to any suitable number or kind of components, peripherals, modules, or devices. In at least one embodiment, processor 1010 coupled using a bus or interface, such as a I²C bus, a System Management Bus (“SMBus”), a Low Pin Count (LPC)

bus, a Serial Peripheral Interface (“SPI”), a High Definition Audio (“HDA”) bus, a Serial Advance Technology Attachment (“SATA”) bus, a Universal Serial Bus (“USB”) (versions 1, 2, 3), or a Universal Asynchronous Receiver/Transmitter (“UART”) bus. In at least one embodiment, FIG. 10 illustrates a system, which includes interconnected hardware devices or “chips”, whereas in other embodiments, FIG. 10 may illustrate an exemplary System on a Chip (“SoC”). In at least one embodiment, devices illustrated in FIG. 10 may be interconnected with proprietary interconnects, standardized interconnects (e.g., PCIe) or some combination thereof. In at least one embodiment, one or more components of FIG. 10 are interconnected using compute express link (CXL) interconnects.

[0088] In at least one embodiment, FIG 10 may include a display 1024, a touch screen 1025, a touch pad 1030, a Near Field Communications unit (“NFC”) 1045, a sensor hub 1040, a thermal sensor 1046, an Express Chipset (“EC”) 1035, a Trusted Platform Module (“TPM”) 1038, BIOS/firmware/flash memory (“BIOS, FW Flash”) 1022, a DSP 1060, a drive 1020 such as a Solid State Disk (“SSD”) or a Hard Disk Drive (“HDD”), a wireless local area network unit (“WLAN”) 1050, a Bluetooth unit 1052, a Wireless Wide Area Network unit (“WWAN”) 1056, a Global Positioning System (GPS) 1055, a camera (“USB 3.0 camera”) 1054 such as a USB 3.0 camera, and/or a Low Power Double Data Rate (“LPDDR”) memory unit (“LPDDR3”) 1015 implemented in, for example, LPDDR3 standard. These components may each be implemented in any suitable manner.

[0089] In at least one embodiment, other components may be communicatively coupled to processor 1010 through components discussed above. In at least one embodiment, an accelerometer 1041, Ambient Light Sensor (“ALS”) 1042, compass 1043, and a gyroscope 1044 may be communicatively coupled to sensor hub 1040. In at least one embodiment, thermal sensor 1039, a fan 1037, a keyboard 1046, and a touch pad 1030 may be communicatively coupled to EC 1035. In at least one embodiment, speaker 1063, headphones 1064, and microphone (“mic”) 1065 may be communicatively coupled to an audio unit (“audio codec and class d amp”) 1062, which may in turn be communicatively coupled to DSP 1060. In at least one embodiment, audio unit 1064 may include, for example and without limitation, an audio

coder/decoder (“codec”) and a class D amplifier. In at least one embodiment, SIM card (“SIM”) 1057 may be communicatively coupled to WWAN unit 1056. In at least one embodiment, components such as WLAN unit 1050 and Bluetooth unit 1052, as well as WWAN unit 1056 may be implemented in a Next Generation Form Factor (“NGFF”).

[0090] Inference and/or training logic 715 are used to perform inferencing and/or training operations associated with one or more embodiments. Details regarding inference and/or training logic 715 are provided below in conjunction with FIGs. 7a and/or 7b. In at least one embodiment, inference and/or training logic 715 may be used in system FIG. 10 for inferencing or predicting operations based, at least in part, on weight parameters calculated using neural network training operations, neural network functions and/or architectures, or neural network use cases described herein.

[0091] Such components can be used to train one or more generation models to be able to generate scene graphs for unlabeled data, as well as to generate scenes based on those scene graphs that are indistinguishable from real scenes.

[0092] FIG. 11 is a block diagram of a processing system, according to at least one embodiment. In at least one embodiment, system 1100 includes one or more processors 1102 and one or more graphics processors 1108, and may be a single processor desktop system, a multiprocessor workstation system, or a server system having a large number of processors 1102 or processor cores 1107. In at least one embodiment, system 1100 is a processing platform incorporated within a system-on-a-chip (SoC) integrated circuit for use in mobile, handheld, or embedded devices.

[0093] In at least one embodiment, system 1100 can include, or be incorporated within a server-based gaming platform, a game console, including a game and media console, a mobile gaming console, a handheld game console, or an online game console. In at least one embodiment, system 1100 is a mobile phone, smart phone, tablet computing device or mobile Internet device. In at least one embodiment, processing system 1100 can also include, couple with, or be integrated within a wearable device, such as a smart watch wearable device, smart

eyewear device, augmented reality device, or virtual reality device. In at least one embodiment, processing system 1100 is a television or set top box device having one or more processors 1102 and a graphical interface generated by one or more graphics processors 1108.

[0094] In at least one embodiment, one or more processors 1102 each include one or more processor cores 1107 to process instructions which, when executed, perform operations for system and user software. In at least one embodiment, each of one or more processor cores 1107 is configured to process a specific instruction set 1109. In at least one embodiment, instruction set 1109 may facilitate Complex Instruction Set Computing (CISC), Reduced Instruction Set Computing (RISC), or computing via a Very Long Instruction Word (VLIW). In at least one embodiment, processor cores 1107 may each process a different instruction set 1109, which may include instructions to facilitate emulation of other instruction sets. In at least one embodiment, processor core 1107 may also include other processing devices, such as a Digital Signal Processor (DSP).

[0095] In at least one embodiment, processor 1102 includes cache memory 1104. In at least one embodiment, processor 1102 can have a single internal cache or multiple levels of internal cache. In at least one embodiment, cache memory is shared among various components of processor 1102. In at least one embodiment, processor 1102 also uses an external cache (e.g., a Level-3 (L3) cache or Last Level Cache (LLC)) (not shown), which may be shared among processor cores 1107 using known cache coherency techniques. In at least one embodiment, register file 1106 is additionally included in processor 1102 which may include different types of registers for storing different types of data (e.g., integer registers, floating point registers, status registers, and an instruction pointer register). In at least one embodiment, register file 1106 may include general-purpose registers or other registers.

[0096] In at least one embodiment, one or more processor(s) 1102 are coupled with one or more interface bus(es) 1110 to transmit communication signals such as address, data, or control signals between processor 1102 and other components in system 1100. In at least one embodiment, interface bus 1110, in one embodiment, can be a processor bus, such as a version of a Direct Media Interface (DMI) bus. In at least one embodiment, interface 1110 is not limited to

a DMI bus, and may include one or more Peripheral Component Interconnect buses (e.g., PCI, PCI Express), memory busses, or other types of interface busses. In at least one embodiment processor(s) 1102 include an integrated memory controller 1116 and a platform controller hub 1130. In at least one embodiment, memory controller 1116 facilitates communication between a memory device and other components of system 1100, while platform controller hub (PCH) 1130 provides connections to I/O devices via a local I/O bus.

[0097] In at least one embodiment, memory device 1120 can be a dynamic random access memory (DRAM) device, a static random access memory (SRAM) device, flash memory device, phase-change memory device, or some other memory device having suitable performance to serve as process memory. In at least one embodiment memory device 1120 can operate as system memory for system 1100, to store data 1122 and instructions 1121 for use when one or more processors 1102 executes an application or process. In at least one embodiment, memory controller 1116 also couples with an optional external graphics processor 1112, which may communicate with one or more graphics processors 1108 in processors 1102 to perform graphics and media operations. In at least one embodiment, a display device 1111 can connect to processor(s) 1102. In at least one embodiment display device 1111 can include one or more of an internal display device, as in a mobile electronic device or a laptop device or an external display device attached via a display interface (e.g., DisplayPort, etc.). In at least one embodiment, display device 1111 can include a head mounted display (HMD) such as a stereoscopic display device for use in virtual reality (VR) applications or augmented reality (AR) applications.

[0098] In at least one embodiment, platform controller hub 1130 enables peripherals to connect to memory device 1120 and processor 1102 via a high-speed I/O bus. In at least one embodiment, I/O peripherals include, but are not limited to, an audio controller 1146, a network controller 1134, a firmware interface 1128, a wireless transceiver 1126, touch sensors 1125, a data storage device 1124 (e.g., hard disk drive, flash memory, etc.). In at least one embodiment, data storage device 1124 can connect via a storage interface (e.g., SATA) or via a peripheral bus, such as a Peripheral Component Interconnect bus (e.g., PCI, PCI Express). In at least one

embodiment, touch sensors 1125 can include touch screen sensors, pressure sensors, or fingerprint sensors. In at least one embodiment, wireless transceiver 1126 can be a Wi-Fi transceiver, a Bluetooth transceiver, or a mobile network transceiver such as a 3G, 4G, or Long Term Evolution (LTE) transceiver. In at least one embodiment, firmware interface 1128 enables communication with system firmware, and can be, for example, a unified extensible firmware interface (UEFI). In at least one embodiment, network controller 1134 can enable a network connection to a wired network. In at least one embodiment, a high-performance network controller (not shown) couples with interface bus 1110. In at least one embodiment, audio controller 1146 is a multi-channel high definition audio controller. In at least one embodiment, system 1100 includes an optional legacy I/O controller 1140 for coupling legacy (e.g., Personal System 2 (PS/2)) devices to system. In at least one embodiment, platform controller hub 1130 can also connect to one or more Universal Serial Bus (USB) controllers 1142 connect input devices, such as keyboard and mouse 1143 combinations, a camera 1144, or other USB input devices.

[0099] In at least one embodiment, an instance of memory controller 1116 and platform controller hub 1130 may be integrated into a discreet external graphics processor, such as external graphics processor 1112. In at least one embodiment, platform controller hub 1130 and/or memory controller 1116 may be external to one or more processor(s) 1102. For example, in at least one embodiment, system 1100 can include an external memory controller 1116 and platform controller hub 1130, which may be configured as a memory controller hub and peripheral controller hub within a system chipset that is in communication with processor(s) 1102.

[0100] Inference and/or training logic 715 are used to perform inferencing and/or training operations associated with one or more embodiments. Details regarding inference and/or training logic 715 are provided below in conjunction with FIGS. 7A and/or 7B. In at least one embodiment portions or all of inference and/or training logic 715 may be incorporated into graphics processor 1500. For example, in at least one embodiment, training and/or inferencing techniques described herein may use one or more of ALUs embodied in a graphics processor.

Moreover, in at least one embodiment, inferencing and/or training operations described herein may be done using logic other than logic illustrated in FIGS. 7A or 7B. In at least one embodiment, weight parameters may be stored in on-chip or off-chip memory and/or registers (shown or not shown) that configure ALUs of a graphics processor to perform one or more machine learning algorithms, neural network architectures, use cases, or training techniques described herein.

[0101] Such components can be used to train one or more generation models to be able to generate scene graphs for unlabeled data, as well as to generate scenes based on those scene graphs that are indistinguishable from real scenes.

[0102] FIG. 12 is a block diagram of a processor 1200 having one or more processor cores 1202A-1202N, an integrated memory controller 1214, and an integrated graphics processor 1208, according to at least one embodiment. In at least one embodiment, processor 1200 can include additional cores up to and including additional core 1202N represented by dashed lined boxes. In at least one embodiment, each of processor cores 1202A-1202N includes one or more internal cache units 1204A-1204N. In at least one embodiment, each processor core also has access to one or more shared cached units 1206.

[0103] In at least one embodiment, internal cache units 1204A-1204N and shared cache units 1206 represent a cache memory hierarchy within processor 1200. In at least one embodiment, cache memory units 1204A-1204N may include at least one level of instruction and data cache within each processor core and one or more levels of shared mid-level cache, such as a Level 2 (L2), Level 3 (L3), Level 4 (L4), or other levels of cache, where a highest level of cache before external memory is classified as an LLC. In at least one embodiment, cache coherency logic maintains coherency between various cache units 1206 and 1204A-1204N.

[0104] In at least one embodiment, processor 1200 may also include a set of one or more bus controller units 1216 and a system agent core 1210. In at least one embodiment, one or more bus controller units 1216 manage a set of peripheral buses, such as one or more PCI or PCI express busses. In at least one embodiment, system agent core 1210 provides management functionality

for various processor components. In at least one embodiment, system agent core 1210 includes one or more integrated memory controllers 1214 to manage access to various external memory devices (not shown).

[0105] In at least one embodiment, one or more of processor cores 1202A-1202N include support for simultaneous multi-threading. In at least one embodiment, system agent core 1210 includes components for coordinating and operating cores 1202A-1202N during multi-threaded processing. In at least one embodiment, system agent core 1210 may additionally include a power control unit (PCU), which includes logic and components to regulate one or more power states of processor cores 1202A-1202N and graphics processor 1208.

[0106] In at least one embodiment, processor 1200 additionally includes graphics processor 1208 to execute graphics processing operations. In at least one embodiment, graphics processor 1208 couples with shared cache units 1206, and system agent core 1210, including one or more integrated memory controllers 1214. In at least one embodiment, system agent core 1210 also includes a display controller 1211 to drive graphics processor output to one or more coupled displays. In at least one embodiment, display controller 1211 may also be a separate module coupled with graphics processor 1208 via at least one interconnect, or may be integrated within graphics processor 1208.

[0107] In at least one embodiment, a ring based interconnect unit 1212 is used to couple internal components of processor 1200. In at least one embodiment, an alternative interconnect unit may be used, such as a point-to-point interconnect, a switched interconnect, or other techniques. In at least one embodiment, graphics processor 1208 couples with ring interconnect 1212 via an I/O link 1213.

[0108] In at least one embodiment, I/O link 1213 represents at least one of multiple varieties of I/O interconnects, including an on package I/O interconnect which facilitates communication between various processor components and a high-performance embedded memory module 1218, such as an eDRAM module. In at least one embodiment, each of processor cores 1202A-

1202N and graphics processor 1208 use embedded memory modules 1218 as a shared Last Level Cache.

[0109] In at least one embodiment, processor cores 1202A-1202N are homogenous cores executing a common instruction set architecture. In at least one embodiment, processor cores 1202A-1202N are heterogeneous in terms of instruction set architecture (ISA), where one or more of processor cores 1202A-1202N execute a common instruction set, while one or more other cores of processor cores 1202A-1202N executes a subset of a common instruction set or a different instruction set. In at least one embodiment, processor cores 1202A-1202N are heterogeneous in terms of microarchitecture, where one or more cores having a relatively higher power consumption couple with one or more power cores having a lower power consumption. In at least one embodiment, processor 1200 can be implemented on one or more chips or as an SoC integrated circuit.

[0110] Inference and/or training logic 715 are used to perform inferencing and/or training operations associated with one or more embodiments. Details regarding inference and/or training logic 715 are provided below in conjunction with FIGs. 7a and/or 7b. In at least one embodiment portions or all of inference and/or training logic 715 may be incorporated into processor 1200. For example, in at least one embodiment, training and/or inferencing techniques described herein may use one or more of ALUs embodied in graphics processor 1512, graphics core(s) 1202A-1202N, or other components in FIG. 12. Moreover, in at least one embodiment, inferencing and/or training operations described herein may be done using logic other than logic illustrated in FIGS. 7A or 7B. In at least one embodiment, weight parameters may be stored in on-chip or off-chip memory and/or registers (shown or not shown) that configure ALUs of graphics processor 1200 to perform one or more machine learning algorithms, neural network architectures, use cases, or training techniques described herein.

[0111] Such components can be used to train one or more generation models to be able to generate scene graphs for unlabeled data, as well as to generate scenes based on those scene graphs that are indistinguishable from real scenes.

VIRTUALIZED COMPUTING PLATFORM

[0112] FIG. 13 is an example data flow diagram for a process 1300 of generating and deploying an image processing and inferencing pipeline, in accordance with at least one embodiment. In at least one embodiment, process 1300 may be deployed for use with imaging devices, processing devices, and/or other device types at one or more facilities 1302. Process 1300 may be executed within a training system 1304 and/or a deployment system 1306. In at least one embodiment, training system 1304 may be used to perform training, deployment, and implementation of machine learning models (e.g., neural networks, object detection algorithms, computer vision algorithms, etc.) for use in deployment system 1306. In at least one embodiment, deployment system 1306 may be configured to offload processing and compute resources among a distributed computing environment to reduce infrastructure requirements at facility 1302. In at least one embodiment, one or more applications in a pipeline may use or call upon services (e.g., inference, visualization, compute, AI, etc.) of deployment system 1306 during execution of applications.

[0113] In at least one embodiment, some of applications used in advanced processing and inferencing pipelines may use machine learning models or other AI to perform one or more processing steps. In at least one embodiment, machine learning models may be trained at facility 1302 using data 1308 (such as imaging data) generated at facility 1302 (and stored on one or more picture archiving and communication system (PACS) servers at facility 1302), may be trained using imaging or sequencing data 1308 from another facility(ies), or a combination thereof. In at least one embodiment, training system 1304 may be used to provide applications, services, and/or other resources for generating working, deployable machine learning models for deployment system 1306.

[0114] In at least one embodiment, model registry 1324 may be backed by object storage that may support versioning and object metadata. In at least one embodiment, object storage may be accessible through, for example, a cloud storage (e.g., cloud 1426 of FIG. 14) compatible application programming interface (API) from within a cloud platform. In at least one embodiment, machine learning models within model registry 1324 may uploaded, listed,

modified, or deleted by developers or partners of a system interacting with an API. In at least one embodiment, an API may provide access to methods that allow users with appropriate credentials to associate models with applications, such that models may be executed as part of execution of containerized instantiations of applications.

[0115] In at least one embodiment, training pipeline 1404 (FIG. 14) may include a scenario where facility 1302 is training their own machine learning model, or has an existing machine learning model that needs to be optimized or updated. In at least one embodiment, imaging data 1308 generated by imaging device(s), sequencing devices, and/or other device types may be received. In at least one embodiment, once imaging data 1308 is received, AI-assisted annotation 1310 may be used to aid in generating annotations corresponding to imaging data 1308 to be used as ground truth data for a machine learning model. In at least one embodiment, AI-assisted annotation 1310 may include one or more machine learning models (e.g., convolutional neural networks (CNNs)) that may be trained to generate annotations corresponding to certain types of imaging data 1308 (e.g., from certain devices). In at least one embodiment, AI-assisted annotations 1310 may then be used directly, or may be adjusted or fine-tuned using an annotation tool to generate ground truth data. In at least one embodiment, AI-assisted annotations 1310, labeled clinic data 1312, or a combination thereof may be used as ground truth data for training a machine learning model. In at least one embodiment, a trained machine learning model may be referred to as output model 1316, and may be used by deployment system 1306, as described herein.

[0116] In at least one embodiment, training pipeline 1404 (FIG. 14) may include a scenario where facility 1302 needs a machine learning model for use in performing one or more processing tasks for one or more applications in deployment system 1306, but facility 1302 may not currently have such a machine learning model (or may not have a model that is optimized, efficient, or effective for such purposes). In at least one embodiment, an existing machine learning model may be selected from a model registry 1324. In at least one embodiment, model registry 1324 may include machine learning models trained to perform a variety of different inference tasks on imaging data. In at least one embodiment, machine learning models in model

registry 1324 may have been trained on imaging data from different facilities than facility 1302 (e.g., facilities remotely located). In at least one embodiment, machine learning models may have been trained on imaging data from one location, two locations, or any number of locations. In at least one embodiment, when being trained on imaging data from a specific location, training may take place at that location, or at least in a manner that protects confidentiality of imaging data or restricts imaging data from being transferred off-premises. In at least one embodiment, once a model is trained – or partially trained – at one location, a machine learning model may be added to model registry 1324. In at least one embodiment, a machine learning model may then be retrained, or updated, at any number of other facilities, and a retrained or updated model may be made available in model registry 1324. In at least one embodiment, a machine learning model may then be selected from model registry 1324 – and referred to as output model 1316 – and may be used in deployment system 1306 to perform one or more processing tasks for one or more applications of a deployment system.

[0117] In at least one embodiment, training pipeline 1404 (FIG. 14), a scenario may include facility 1302 requiring a machine learning model for use in performing one or more processing tasks for one or more applications in deployment system 1306, but facility 1302 may not currently have such a machine learning model (or may not have a model that is optimized, efficient, or effective for such purposes). In at least one embodiment, a machine learning model selected from model registry 1324 may not be fine-tuned or optimized for imaging data 1308 generated at facility 1302 because of differences in populations, robustness of training data used to train a machine learning model, diversity in anomalies of training data, and/or other issues with training data. In at least one embodiment, AI-assisted annotation 1310 may be used to aid in generating annotations corresponding to imaging data 1308 to be used as ground truth data for retraining or updating a machine learning model. In at least one embodiment, labeled data 1312 may be used as ground truth data for training a machine learning model. In at least one embodiment, retraining or updating a machine learning model may be referred to as model training 1314. In at least one embodiment, model training 1314 – e.g., AI-assisted annotations 1310, labeled clinic data 1312, or a combination thereof – may be used as ground truth data for

retraining or updating a machine learning model. In at least one embodiment, a trained machine learning model may be referred to as output model 1316, and may be used by deployment system 1306, as described herein.

[0118] In at least one embodiment, deployment system 1306 may include software 1318, services 1320, hardware 1322, and/or other components, features, and functionality. In at least one embodiment, deployment system 1306 may include a software “stack,” such that software 1318 may be built on top of services 1320 and may use services 1320 to perform some or all of processing tasks, and services 1320 and software 1318 may be built on top of hardware 1322 and use hardware 1322 to execute processing, storage, and/or other compute tasks of deployment system 1306. In at least one embodiment, software 1318 may include any number of different containers, where each container may execute an instantiation of an application. In at least one embodiment, each application may perform one or more processing tasks in an advanced processing and inferencing pipeline (e.g., inferencing, object detection, feature detection, segmentation, image enhancement, calibration, etc.). In at least one embodiment, an advanced processing and inferencing pipeline may be defined based on selections of different containers that are desired or required for processing imaging data 1308, in addition to containers that receive and configure imaging data for use by each container and/or for use by facility 1302 after processing through a pipeline (e.g., to convert outputs back to a usable data type). In at least one embodiment, a combination of containers within software 1318 (e.g., that make up a pipeline) may be referred to as a virtual instrument (as described in more detail herein), and a virtual instrument may leverage services 1320 and hardware 1322 to execute some or all processing tasks of applications instantiated in containers.

[0119] In at least one embodiment, a data processing pipeline may receive input data (e.g., imaging data 1308) in a specific format in response to an inference request (e.g., a request from a user of deployment system 1306). In at least one embodiment, input data may be representative of one or more images, video, and/or other data representations generated by one or more imaging devices. In at least one embodiment, data may undergo pre-processing as part of data processing pipeline to prepare data for processing by one or more applications. In at least one

embodiment, post-processing may be performed on an output of one or more inferencing tasks or other processing tasks of a pipeline to prepare an output data for a next application and/or to prepare output data for transmission and/or use by a user (e.g., as a response to an inference request). In at least one embodiment, inferencing tasks may be performed by one or more machine learning models, such as trained or deployed neural networks, which may include output models 1316 of training system 1304.

[0120] In at least one embodiment, tasks of data processing pipeline may be encapsulated in a container(s) that each represents a discrete, fully functional instantiation of an application and virtualized computing environment that is able to reference machine learning models. In at least one embodiment, containers or applications may be published into a private (e.g., limited access) area of a container registry (described in more detail herein), and trained or deployed models may be stored in model registry 1324 and associated with one or more applications. In at least one embodiment, images of applications (e.g., container images) may be available in a container registry, and once selected by a user from a container registry for deployment in a pipeline, an image may be used to generate a container for an instantiation of an application for use by a user's system.

[0121] In at least one embodiment, developers (e.g., software developers, clinicians, doctors, etc.) may develop, publish, and store applications (e.g., as containers) for performing image processing and/or inferencing on supplied data. In at least one embodiment, development, publishing, and/or storing may be performed using a software development kit (SDK) associated with a system (e.g., to ensure that an application and/or container developed is compliant with or compatible with a system). In at least one embodiment, an application that is developed may be tested locally (e.g., at a first facility, on data from a first facility) with an SDK which may support at least some of services 1320 as a system (e.g., system 1400 of FIG. 14). In at least one embodiment, because DICOM objects may contain anywhere from one to hundreds of images or other data types, and due to a variation in data, a developer may be responsible for managing (e.g., setting constructs for, building pre-processing into an application, etc.) extraction and preparation of incoming data. In at least one embodiment, once validated by system 1400 (e.g.,

for accuracy), an application may be available in a container registry for selection and/or implementation by a user to perform one or more processing tasks with respect to data at a facility (e.g., a second facility) of a user.

[0122] In at least one embodiment, developers may then share applications or containers through a network for access and use by users of a system (e.g., system 1400 of FIG. 14). In at least one embodiment, completed and validated applications or containers may be stored in a container registry and associated machine learning models may be stored in model registry 1324. In at least one embodiment, a requesting entity – who provides an inference or image processing request – may browse a container registry and/or model registry 1324 for an application, container, dataset, machine learning model, etc., select a desired combination of elements for inclusion in data processing pipeline, and submit an imaging processing request. In at least one embodiment, a request may include input data (and associated patient data, in some examples) that is necessary to perform a request, and/or may include a selection of application(s) and/or machine learning models to be executed in processing a request. In at least one embodiment, a request may then be passed to one or more components of deployment system 1306 (e.g., a cloud) to perform processing of data processing pipeline. In at least one embodiment, processing by deployment system 1306 may include referencing selected elements (e.g., applications, containers, models, etc.) from a container registry and/or model registry 1324. In at least one embodiment, once results are generated by a pipeline, results may be returned to a user for reference (e.g., for viewing in a viewing application suite executing on a local, on-premises workstation or terminal).

[0123] In at least one embodiment, to aid in processing or execution of applications or containers in pipelines, services 1320 may be leveraged. In at least one embodiment, services 1320 may include compute services, artificial intelligence (AI) services, visualization services, and/or other service types. In at least one embodiment, services 1320 may provide functionality that is common to one or more applications in software 1318, so functionality may be abstracted to a service that may be called upon or leveraged by applications. In at least one embodiment, functionality provided by services 1320 may run dynamically and more efficiently, while also

scaling well by allowing applications to process data in parallel (e.g., using a parallel computing platform 1430 (FIG. 14)). In at least one embodiment, rather than each application that shares a same functionality offered by a service 1320 being required to have a respective instance of service 1320, service 1320 may be shared between and among various applications. In at least one embodiment, services may include an inference server or engine that may be used for executing detection or segmentation tasks, as non-limiting examples. In at least one embodiment, a model training service may be included that may provide machine learning model training and/or retraining capabilities. In at least one embodiment, a data augmentation service may further be included that may provide GPU accelerated data (e.g., DICOM, RIS, CIS, REST compliant, RPC, raw, etc.) extraction, resizing, scaling, and/or other augmentation. In at least one embodiment, a visualization service may be used that may add image rendering effects – such as ray-tracing, rasterization, denoising, sharpening, etc. – to add realism to two-dimensional (2D) and/or three-dimensional (3D) models. In at least one embodiment, virtual instrument services may be included that provide for beam-forming, segmentation, inferencing, imaging, and/or support for other applications within pipelines of virtual instruments.

[0124] In at least one embodiment, where a service 1320 includes an AI service (e.g., an inference service), one or more machine learning models may be executed by calling upon (e.g., as an API call) an inference service (e.g., an inference server) to execute machine learning model(s), or processing thereof, as part of application execution. In at least one embodiment, where another application includes one or more machine learning models for segmentation tasks, an application may call upon an inference service to execute machine learning models for performing one or more of processing operations associated with segmentation tasks. In at least one embodiment, software 1318 implementing advanced processing and inferencing pipeline that includes segmentation application and anomaly detection application may be streamlined because each application may call upon a same inference service to perform one or more inferencing tasks.

[0125] In at least one embodiment, hardware 1322 may include GPUs, CPUs, graphics cards, an AI/deep learning system (e.g., an AI supercomputer, such as NVIDIA's DGX), a cloud

platform, or a combination thereof. In at least one embodiment, different types of hardware 1322 may be used to provide efficient, purpose-built support for software 1318 and services 1320 in deployment system 1306. In at least one embodiment, use of GPU processing may be implemented for processing locally (e.g., at facility 1302), within an AI/deep learning system, in a cloud system, and/or in other processing components of deployment system 1306 to improve efficiency, accuracy, and efficacy of image processing and generation. In at least one embodiment, software 1318 and/or services 1320 may be optimized for GPU processing with respect to deep learning, machine learning, and/or high-performance computing, as non-limiting examples. In at least one embodiment, at least some of computing environment of deployment system 1306 and/or training system 1304 may be executed in a datacenter one or more supercomputers or high performance computing systems, with GPU optimized software (e.g., hardware and software combination of NVIDIA's DGX System). In at least one embodiment, hardware 1322 may include any number of GPUs that may be called upon to perform processing of data in parallel, as described herein. In at least one embodiment, cloud platform may further include GPU processing for GPU-optimized execution of deep learning tasks, machine learning tasks, or other computing tasks. In at least one embodiment, cloud platform (e.g., NVIDIA's NGC) may be executed using an AI/deep learning supercomputer(s) and/or GPU-optimized software (e.g., as provided on NVIDIA's DGX Systems) as a hardware abstraction and scaling platform. In at least one embodiment, cloud platform may integrate an application container clustering system or orchestration system (e.g., KUBERNETES) on multiple GPUs to enable seamless scaling and load balancing.

[0126] FIG. 14 is a system diagram for an example system 1400 for generating and deploying an imaging deployment pipeline, in accordance with at least one embodiment. In at least one embodiment, system 1400 may be used to implement process 1300 of FIG. 13 and/or other processes including advanced processing and inferencing pipelines. In at least one embodiment, system 1400 may include training system 1304 and deployment system 1306. In at least one embodiment, training system 1304 and deployment system 1306 may be implemented using software 1318, services 1320, and/or hardware 1322, as described herein.

[0127] In at least one embodiment, system 1400 (e.g., training system 1304 and/or deployment system 1306) may implemented in a cloud computing environment (e.g., using cloud 1426). In at least one embodiment, system 1400 may be implemented locally with respect to a healthcare services facility, or as a combination of both cloud and local computing resources. In at least one embodiment, access to APIs in cloud 1426 may be restricted to authorized users through enacted security measures or protocols. In at least one embodiment, a security protocol may include web tokens that may be signed by an authentication (e.g., AuthN, AuthZ, Gluecon, etc.) service and may carry appropriate authorization. In at least one embodiment, APIs of virtual instruments (described herein), or other instantiations of system 1400, may be restricted to a set of public IPs that have been vetted or authorized for interaction.

[0128] In at least one embodiment, various components of system 1400 may communicate between and among one another using any of a variety of different network types, including but not limited to local area networks (LANs) and/or wide area networks (WANs) via wired and/or wireless communication protocols. In at least one embodiment, communication between facilities and components of system 1400 (e.g., for transmitting inference requests, for receiving results of inference requests, etc.) may be communicated over data bus(es), wireless data protocols (Wi-Fi), wired data protocols (e.g., Ethernet), etc.

[0129] In at least one embodiment, training system 1304 may execute training pipelines 1404, similar to those described herein with respect to FIG. 13. In at least one embodiment, where one or more machine learning models are to be used in deployment pipelines 1410 by deployment system 1306, training pipelines 1404 may be used to train or retrain one or more (e.g. pre-trained) models, and/or implement one or more of pre-trained models 1406 (e.g., without a need for retraining or updating). In at least one embodiment, as a result of training pipelines 1404, output model(s) 1316 may be generated. In at least one embodiment, training pipelines 1404 may include any number of processing steps, such as but not limited to imaging data (or other input data) conversion or adaption. In at least one embodiment, for different machine learning models used by deployment system 1306, different training pipelines 1404 may be used. In at least one embodiment, training pipeline 1404 similar to a first example described with respect to

FIG. 13 may be used for a first machine learning model, training pipeline 1404 similar to a second example described with respect to FIG. 13 may be used for a second machine learning model, and training pipeline 1404 similar to a third example described with respect to FIG. 13 may be used for a third machine learning model. In at least one embodiment, any combination of tasks within training system 1304 may be used depending on what is required for each respective machine learning model. In at least one embodiment, one or more of machine learning models may already be trained and ready for deployment so machine learning models may not undergo any processing by training system 1304, and may be implemented by deployment system 1306.

[0130] In at least one embodiment, output model(s) 1316 and/or pre-trained model(s) 1406 may include any types of machine learning models depending on implementation or embodiment. In at least one embodiment, and without limitation, machine learning models used by system 1400 may include machine learning model(s) using linear regression, logistic regression, decision trees, support vector machines (SVM), Naïve Bayes, k-nearest neighbor (Knn), K means clustering, random forest, dimensionality reduction algorithms, gradient boosting algorithms, neural networks (e.g., auto-encoders, convolutional, recurrent, perceptrons, Long/Short Term Memory (LSTM), Hopfield, Boltzmann, deep belief, deconvolutional, generative adversarial, liquid state machine, etc.), and/or other types of machine learning models.

[0131] In at least one embodiment, training pipelines 1404 may include AI-assisted annotation, as described in more detail herein with respect to at least FIG. 15B. In at least one embodiment, labeled data 1312 (e.g., traditional annotation) may be generated by any number of techniques. In at least one embodiment, labels or other annotations may be generated within a drawing program (e.g., an annotation program), a computer aided design (CAD) program, a labeling program, another type of program suitable for generating annotations or labels for ground truth, and/or may be hand drawn, in some examples. In at least one embodiment, ground truth data may be synthetically produced (e.g., generated from computer models or renderings), real produced (e.g., designed and produced from real-world data), machine-automated (e.g., using feature analysis and learning to extract features from data and then generate labels), human

annotated (e.g., labeler, or annotation expert, defines location of labels), and/or a combination thereof. In at least one embodiment, for each instance of imaging data 1308 (or other data type used by machine learning models), there may be corresponding ground truth data generated by training system 1304. In at least one embodiment, AI-assisted annotation may be performed as part of deployment pipelines 1410; either in addition to, or in lieu of AI-assisted annotation included in training pipelines 1404. In at least one embodiment, system 1400 may include a multi-layer platform that may include a software layer (e.g., software 1318) of diagnostic applications (or other application types) that may perform one or more medical imaging and diagnostic functions. In at least one embodiment, system 1400 may be communicatively coupled to (e.g., via encrypted links) PACS server networks of one or more facilities. In at least one embodiment, system 1400 may be configured to access and referenced data from PACS servers to perform operations, such as training machine learning models, deploying machine learning models, image processing, inferencing, and/or other operations.

[0132] In at least one embodiment, a software layer may be implemented as a secure, encrypted, and/or authenticated API through which applications or containers may be invoked (e.g., called) from an external environment(s) (e.g., facility 1302). In at least one embodiment, applications may then call or execute one or more services 1320 for performing compute, AI, or visualization tasks associated with respective applications, and software 1318 and/or services 1320 may leverage hardware 1322 to perform processing tasks in an effective and efficient manner.

[0133] In at least one embodiment, deployment system 1306 may execute deployment pipelines 1410. In at least one embodiment, deployment pipelines 1410 may include any number of applications that may be sequentially, non-sequentially, or otherwise applied to imaging data (and/or other data types) generated by imaging devices, sequencing devices, genomics devices, etc. – including AI-assisted annotation, as described above. In at least one embodiment, as described herein, a deployment pipeline 1410 for an individual device may be referred to as a virtual instrument for a device (e.g., a virtual ultrasound instrument, a virtual CT scan instrument, a virtual sequencing instrument, etc.). In at least one embodiment, for a single

device, there may be more than one deployment pipeline 1410 depending on information desired from data generated by a device. In at least one embodiment, where detections of anomalies are desired from an MRI machine, there may be a first deployment pipeline 1410, and where image enhancement is desired from output of an MRI machine, there may be a second deployment pipeline 1410.

[0134] In at least one embodiment, an image generation application may include a processing task that includes use of a machine learning model. In at least one embodiment, a user may desire to use their own machine learning model, or to select a machine learning model from model registry 1324. In at least one embodiment, a user may implement their own machine learning model or select a machine learning model for inclusion in an application for performing a processing task. In at least one embodiment, applications may be selectable and customizable, and by defining constructs of applications, deployment and implementation of applications for a particular user are presented as a more seamless user experience. In at least one embodiment, by leveraging other features of system 1400 – such as services 1320 and hardware 1322 – deployment pipelines 1410 may be even more user friendly, provide for easier integration, and produce more accurate, efficient, and timely results.

[0135] In at least one embodiment, deployment system 1306 may include a user interface 1414 (e.g., a graphical user interface, a web interface, etc.) that may be used to select applications for inclusion in deployment pipeline(s) 1410, arrange applications, modify or change applications or parameters or constructs thereof, use and interact with deployment pipeline(s) 1410 during set-up and/or deployment, and/or to otherwise interact with deployment system 1306. In at least one embodiment, although not illustrated with respect to training system 1304, user interface 1414 (or a different user interface) may be used for selecting models for use in deployment system 1306, for selecting models for training, or retraining, in training system 1304, and/or for otherwise interacting with training system 1304.

[0136] In at least one embodiment, pipeline manager 1412 may be used, in addition to an application orchestration system 1428, to manage interaction between applications or containers of deployment pipeline(s) 1410 and services 1320 and/or hardware 1322. In at least one

embodiment, pipeline manager 1412 may be configured to facilitate interactions from application to application, from application to service 1320, and/or from application or service to hardware 1322. In at least one embodiment, although illustrated as included in software 1318, this is not intended to be limiting, and in some examples (e.g., as illustrated in FIG. 12cc) pipeline manager 1412 may be included in services 1320. In at least one embodiment, application orchestration system 1428 (e.g., Kubernetes, DOCKER, etc.) may include a container orchestration system that may group applications into containers as logical units for coordination, management, scaling, and deployment. In at least one embodiment, by associating applications from deployment pipeline(s) 1410 (e.g., a reconstruction application, a segmentation application, etc.) with individual containers, each application may execute in a self-contained environment (e.g., at a kernel level) to increase speed and efficiency.

[0137] In at least one embodiment, each application and/or container (or image thereof) may be individually developed, modified, and deployed (e.g., a first user or developer may develop, modify, and deploy a first application and a second user or developer may develop, modify, and deploy a second application separate from a first user or developer), which may allow for focus on, and attention to, a task of a single application and/or container(s) without being hindered by tasks of another application(s) or container(s). In at least one embodiment, communication, and cooperation between different containers or applications may be aided by pipeline manager 1412 and application orchestration system 1428. In at least one embodiment, so long as an expected input and/or output of each container or application is known by a system (e.g., based on constructs of applications or containers), application orchestration system 1428 and/or pipeline manager 1412 may facilitate communication among and between, and sharing of resources among and between, each of applications or containers. In at least one embodiment, because one or more of applications or containers in deployment pipeline(s) 1410 may share same services and resources, application orchestration system 1428 may orchestrate, load balance, and determine sharing of services or resources between and among various applications or containers. In at least one embodiment, a scheduler may be used to track resource requirements of applications or containers, current usage or planned usage of these resources, and resource

availability. In at least one embodiment, a scheduler may thus allocate resources to different applications and distribute resources between and among applications in view of requirements and availability of a system. In some examples, a scheduler (and/or other component of application orchestration system 1428) may determine resource availability and distribution based on constraints imposed on a system (e.g., user constraints), such as quality of service (QoS), urgency of need for data outputs (e.g., to determine whether to execute real-time processing or delayed processing), etc.

[0138] In at least one embodiment, services 1320 leveraged by and shared by applications or containers in deployment system 1306 may include compute services 1416, AI services 1418, visualization services 1420, and/or other service types. In at least one embodiment, applications may call (e.g., execute) one or more of services 1320 to perform processing operations for an application. In at least one embodiment, compute services 1416 may be leveraged by applications to perform super-computing or other high-performance computing (HPC) tasks. In at least one embodiment, compute service(s) 1416 may be leveraged to perform parallel processing (e.g., using a parallel computing platform 1430) for processing data through one or more of applications and/or one or more tasks of a single application, substantially simultaneously. In at least one embodiment, parallel computing platform 1430 (e.g., NVIDIA's CUDA) may enable general purpose computing on GPUs (GPGPU) (e.g., GPUs 1422). In at least one embodiment, a software layer of parallel computing platform 1430 may provide access to virtual instruction sets and parallel computational elements of GPUs, for execution of compute kernels. In at least one embodiment, parallel computing platform 1430 may include memory and, in some embodiments, a memory may be shared between and among multiple containers, and/or between and among different processing tasks within a single container. In at least one embodiment, inter-process communication (IPC) calls may be generated for multiple containers and/or for multiple processes within a container to use same data from a shared segment of memory of parallel computing platform 1430 (e.g., where multiple different stages of an application or multiple applications are processing same information). In at least one embodiment, rather than making a copy of data and moving data to different locations in

memory (e.g., a read/write operation), same data in same location of a memory may be used for any number of processing tasks (e.g., at a same time, at different times, etc.). In at least one embodiment, as data is used to generate new data as a result of processing, this information of a new location of data may be stored and shared between various applications. In at least one embodiment, location of data and a location of updated or modified data may be part of a definition of how a payload is understood within containers.

[0139] In at least one embodiment, AI services 1418 may be leveraged to perform inferencing services for executing machine learning model(s) associated with applications (e.g., tasked with performing one or more processing tasks of an application). In at least one embodiment, AI services 1418 may leverage AI system 1424 to execute machine learning model(s) (e.g., neural networks, such as CNNs) for segmentation, reconstruction, object detection, feature detection, classification, and/or other inferencing tasks. In at least one embodiment, applications of deployment pipeline(s) 1410 may use one or more of output models 1316 from training system 1304 and/or other models of applications to perform inference on imaging data. In at least one embodiment, two or more examples of inferencing using application orchestration system 1428 (e.g., a scheduler) may be available. In at least one embodiment, a first category may include a high priority/low latency path that may achieve higher service level agreements, such as for performing inference on urgent requests during an emergency, or for a radiologist during diagnosis. In at least one embodiment, a second category may include a standard priority path that may be used for requests that may be non-urgent or where analysis may be performed at a later time. In at least one embodiment, application orchestration system 1428 may distribute resources (e.g., services 1320 and/or hardware 1322) based on priority paths for different inferencing tasks of AI services 1418.

[0140] In at least one embodiment, shared storage may be mounted to AI services 1418 within system 1400. In at least one embodiment, shared storage may operate as a cache (or other storage device type) and may be used to process inference requests from applications. In at least one embodiment, when an inference request is submitted, a request may be received by a set of API instances of deployment system 1306, and one or more instances may be selected (e.g., for

best fit, for load balancing, etc.) to process a request. In at least one embodiment, to process a request, a request may be entered into a database, a machine learning model may be located from model registry 1324 if not already in a cache, a validation step may ensure appropriate machine learning model is loaded into a cache (e.g., shared storage), and/or a copy of a model may be saved to a cache. In at least one embodiment, a scheduler (e.g., of pipeline manager 1412) may be used to launch an application that is referenced in a request if an application is not already running or if there are not enough instances of an application. In at least one embodiment, if an inference server is not already launched to execute a model, an inference server may be launched. Any number of inference servers may be launched per model. In at least one embodiment, in a pull model, in which inference servers are clustered, models may be cached whenever load balancing is advantageous. In at least one embodiment, inference servers may be statically loaded in corresponding, distributed servers.

[0141] In at least one embodiment, inferencing may be performed using an inference server that runs in a container. In at least one embodiment, an instance of an inference server may be associated with a model (and optionally a plurality of versions of a model). In at least one embodiment, if an instance of an inference server does not exist when a request to perform inference on a model is received, a new instance may be loaded. In at least one embodiment, when starting an inference server, a model may be passed to an inference server such that a same container may be used to serve different models so long as inference server is running as a different instance.

[0142] In at least one embodiment, during application execution, an inference request for a given application may be received, and a container (e.g., hosting an instance of an inference server) may be loaded (if not already), and a start procedure may be called. In at least one embodiment, pre-processing logic in a container may load, decode, and/or perform any additional pre-processing on incoming data (e.g., using a CPU(s) and/or GPU(s)). In at least one embodiment, once data is prepared for inference, a container may perform inference as necessary on data. In at least one embodiment, this may include a single inference call on one image (e.g., a hand X-ray), or may require inference on hundreds of images (e.g., a chest CT). In at least one

embodiment, an application may summarize results before completing, which may include, without limitation, a single confidence score, pixel level-segmentation, voxel-level segmentation, generating a visualization, or generating text to summarize findings. In at least one embodiment, different models or applications may be assigned different priorities. For example, some models may have a real-time ($TAT < 1 \text{ min}$) priority while others may have lower priority (e.g., $TAT < 10 \text{ min}$). In at least one embodiment, model execution times may be measured from requesting institution or entity and may include partner network traversal time, as well as execution on an inference service.

[0143] In at least one embodiment, transfer of requests between services 1320 and inference applications may be hidden behind a software development kit (SDK), and robust transport may be provide through a queue. In at least one embodiment, a request will be placed in a queue via an API for an individual application/tenant ID combination and an SDK will pull a request from a queue and give a request to an application. In at least one embodiment, a name of a queue may be provided in an environment from where an SDK will pick it up. In at least one embodiment, asynchronous communication through a queue may be useful as it may allow any instance of an application to pick up work as it becomes available. Results may be transferred back through a queue, to ensure no data is lost. In at least one embodiment, queues may also provide an ability to segment work, as highest priority work may go to a queue with most instances of an application connected to it, while lowest priority work may go to a queue with a single instance connected to it that processes tasks in an order received. In at least one embodiment, an application may run on a GPU-accelerated instance generated in cloud 1426, and an inference service may perform inferencing on a GPU.

[0144] In at least one embodiment, visualization services 1420 may be leveraged to generate visualizations for viewing outputs of applications and/or deployment pipeline(s) 1410. In at least one embodiment, GPUs 1422 may be leveraged by visualization services 1420 to generate visualizations. In at least one embodiment, rendering effects, such as ray-tracing, may be implemented by visualization services 1420 to generate higher quality visualizations. In at least one embodiment, visualizations may include, without limitation, 2D image renderings, 3D

volume renderings, 3D volume reconstruction, 2D tomographic slices, virtual reality displays, augmented reality displays, etc. In at least one embodiment, virtualized environments may be used to generate a virtual interactive display or environment (e.g., a virtual environment) for interaction by users of a system (e.g., doctors, nurses, radiologists, etc.). In at least one embodiment, visualization services 1420 may include an internal visualizer, cinematics, and/or other rendering or image processing capabilities or functionality (e.g., ray tracing, rasterization, internal optics, etc.).

[0145] In at least one embodiment, hardware 1322 may include GPUs 1422, AI system 1424, cloud 1426, and/or any other hardware used for executing training system 1304 and/or deployment system 1306. In at least one embodiment, GPUs 1422 (e.g., NVIDIA's TESLA and/or QUADRO GPUs) may include any number of GPUs that may be used for executing processing tasks of compute services 1416, AI services 1418, visualization services 1420, other services, and/or any of features or functionality of software 1318. For example, with respect to AI services 1418, GPUs 1422 may be used to perform pre-processing on imaging data (or other data types used by machine learning models), post-processing on outputs of machine learning models, and/or to perform inferencing (e.g., to execute machine learning models). In at least one embodiment, cloud 1426, AI system 1424, and/or other components of system 1400 may use GPUs 1422. In at least one embodiment, cloud 1426 may include a GPU-optimized platform for deep learning tasks. In at least one embodiment, AI system 1424 may use GPUs, and cloud 1426 – or at least a portion tasked with deep learning or inferencing – may be executed using one or more AI systems 1424. As such, although hardware 1322 is illustrated as discrete components, this is not intended to be limiting, and any components of hardware 1322 may be combined with, or leveraged by, any other components of hardware 1322.

[0146] In at least one embodiment, AI system 1424 may include a purpose-built computing system (e.g., a super-computer or an HPC) configured for inferencing, deep learning, machine learning, and/or other artificial intelligence tasks. In at least one embodiment, AI system 1424 (e.g., NVIDIA's DGX) may include GPU-optimized software (e.g., a software stack) that may be executed using a plurality of GPUs 1422, in addition to CPUs, RAM, storage, and/or other

components, features, or functionality. In at least one embodiment, one or more AI systems 1424 may be implemented in cloud 1426 (e.g., in a data center) for performing some or all of AI-based processing tasks of system 1400.

[0147] In at least one embodiment, cloud 1426 may include a GPU-accelerated infrastructure (e.g., NVIDIA's NGC) that may provide a GPU-optimized platform for executing processing tasks of system 1400. In at least one embodiment, cloud 1426 may include an AI system(s) 1424 for performing one or more of AI-based tasks of system 1400 (e.g., as a hardware abstraction and scaling platform). In at least one embodiment, cloud 1426 may integrate with application orchestration system 1428 leveraging multiple GPUs to enable seamless scaling and load balancing between and among applications and services 1320. In at least one embodiment, cloud 1426 may tasked with executing at least some of services 1320 of system 1400, including compute services 1416, AI services 1418, and/or visualization services 1420, as described herein. In at least one embodiment, cloud 1426 may perform small and large batch inference (e.g., executing NVIDIA's TENSOR RT), provide an accelerated parallel computing API and platform 1430 (e.g., NVIDIA's CUDA), execute application orchestration system 1428 (e.g., KUBERNETES), provide a graphics rendering API and platform (e.g., for ray-tracing, 2D graphics, 3D graphics, and/or other rendering techniques to produce higher quality cinematics), and/or may provide other functionality for system 1400.

[0148] FIG. 15A illustrates a data flow diagram for a process 1500 to train, retrain, or update a machine learning model, in accordance with at least one embodiment. In at least one embodiment, process 1500 may be executed using, as a non-limiting example, system 1400 of FIG. 14. In at least one embodiment, process 1500 may leverage services 1320 and/or hardware 1322 of system 1400, as described herein. In at least one embodiment, refined models 1512 generated by process 1500 may be executed by deployment system 1306 for one or more containerized applications in deployment pipelines 1410.

[0149] In at least one embodiment, model training 1314 may include retraining or updating an initial model 1504 (e.g., a pre-trained model) using new training data (e.g., new input data, such as customer dataset 1506, and/or new ground truth data associated with input data). In at least

one embodiment, to retrain, or update, initial model 1504, output or loss layer(s) of initial model 1504 may be reset, or deleted, and/or replaced with an updated or new output or loss layer(s). In at least one embodiment, initial model 1504 may have previously fine-tuned parameters (e.g., weights and/or biases) that remain from prior training, so training or retraining 1314 may not take as long or require as much processing as training a model from scratch. In at least one embodiment, during model training 1314, by having reset or replaced output or loss layer(s) of initial model 1504, parameters may be updated and re-tuned for a new data set based on loss calculations associated with accuracy of output or loss layer(s) at generating predictions on new, customer dataset 1506 (e.g., image data 1308 of FIG. 13).

[0150] In at least one embodiment, pre-trained models 1406 may be stored in a data store, or registry (e.g., model registry 1324 of FIG. 13). In at least one embodiment, pre-trained models 1406 may have been trained, at least in part, at one or more facilities other than a facility executing process 1500. In at least one embodiment, to protect privacy and rights of patients, subjects, or clients of different facilities, pre-trained models 1406 may have been trained, on-premise, using customer or patient data generated on-premise. In at least one embodiment, pre-trained models 1406 may be trained using cloud 1426 and/or other hardware 1322, but confidential, privacy protected patient data may not be transferred to, used by, or accessible to any components of cloud 1426 (or other off premise hardware). In at least one embodiment, where a pre-trained model 1406 is trained at using patient data from more than one facility, pre-trained model 1406 may have been individually trained for each facility prior to being trained on patient or customer data from another facility. In at least one embodiment, such as where a customer or patient data has been released of privacy concerns (e.g., by waiver, for experimental use, etc.), or where a customer or patient data is included in a public data set, a customer or patient data from any number of facilities may be used to train pre-trained model 1406 on-premise and/or off premise, such as in a datacenter or other cloud computing infrastructure.

[0151] In at least one embodiment, when selecting applications for use in deployment pipelines 1410, a user may also select machine learning models to be used for specific applications. In at least one embodiment, a user may not have a model for use, so a user may select a pre-trained

model 1406 to use with an application. In at least one embodiment, pre-trained model 1406 may not be optimized for generating accurate results on customer dataset 1506 of a facility of a user (e.g., based on patient diversity, demographics, types of medical imaging devices used, etc.). In at least one embodiment, prior to deploying pre-trained model 1406 into deployment pipeline 1410 for use with an application(s), pre-trained model 1406 may be updated, retrained, and/or fine-tuned for use at a respective facility.

[0152] In at least one embodiment, a user may select pre-trained model 1406 that is to be updated, retrained, and/or fine-tuned, and pre-trained model 1406 may be referred to as initial model 1504 for training system 1304 within process 1500. In at least one embodiment, customer dataset 1506 (e.g., imaging data, genomics data, sequencing data, or other data types generated by devices at a facility) may be used to perform model training 1314 (which may include, without limitation, transfer learning) on initial model 1504 to generate refined model 1512. In at least one embodiment, ground truth data corresponding to customer dataset 1506 may be generated by training system 1304. In at least one embodiment, ground truth data may be generated, at least in part, by clinicians, scientists, doctors, practitioners, at a facility (e.g., as labeled clinic data 1312 of FIG. 13).

[0153] In at least one embodiment, AI-assisted annotation 1310 may be used in some examples to generate ground truth data. In at least one embodiment, AI-assisted annotation 1310 (e.g., implemented using an AI-assisted annotation SDK) may leverage machine learning models (e.g., neural networks) to generate suggested or predicted ground truth data for a customer dataset. In at least one embodiment, user 1510 may use annotation tools within a user interface (a graphical user interface (GUI)) on computing device 1508.

[0154] In at least one embodiment, user 1510 may interact with a GUI via computing device 1508 to edit or fine-tune (auto)annotations. In at least one embodiment, a polygon editing feature may be used to move vertices of a polygon to more accurate or fine-tuned locations.

[0155] In at least one embodiment, once customer dataset 1506 has associated ground truth data, ground truth data (e.g., from AI-assisted annotation, manual labeling, etc.) may be used by

during model training 1314 to generate refined model 1512. In at least one embodiment, customer dataset 1506 may be applied to initial model 1504 any number of times, and ground truth data may be used to update parameters of initial model 1504 until an acceptable level of accuracy is attained for refined model 1512. In at least one embodiment, once refined model 1512 is generated, refined model 1512 may be deployed within one or more deployment pipelines 1410 at a facility for performing one or more processing tasks with respect to medical imaging data.

[0156] In at least one embodiment, refined model 1512 may be uploaded to pre-trained models 1406 in model registry 1324 to be selected by another facility. In at least one embodiment, this process may be completed at any number of facilities such that refined model 1512 may be further refined on new datasets any number of times to generate a more universal model.

[0157] FIG. 15B is an example illustration of a client-server architecture 1532 to enhance annotation tools with pre-trained annotation models, in accordance with at least one embodiment. In at least one embodiment, AI-assisted annotation tools 1536 may be instantiated based on a client-server architecture 1532. In at least one embodiment, annotation tools 1536 in imaging applications may aid radiologists, for example, identify organs and abnormalities. In at least one embodiment, imaging applications may include software tools that help user 1510 to identify, as a non-limiting example, a few extreme points on a particular organ of interest in raw images 1534 (e.g., in a 3D MRI or CT scan) and receive auto-annotated results for all 2D slices of a particular organ. In at least one embodiment, results may be stored in a data store as training data 1538 and used as (for example and without limitation) ground truth data for training. In at least one embodiment, when computing device 1508 sends extreme points for AI-assisted annotation 1310, a deep learning model, for example, may receive this data as input and return inference results of a segmented organ or abnormality. In at least one embodiment, pre-instantiated annotation tools, such as AI-Assisted Annotation Tool 1536B in FIG. 15B, may be enhanced by making API calls (e.g., API Call 1544) to a server, such as an Annotation Assistant Server 1540 that may include a set of pre-trained models 1542 stored in an annotation model registry, for example. In at least one embodiment, an annotation model registry may store pre-

trained models 1542 (e.g., machine learning models, such as deep learning models) that are pre-trained to perform AI-assisted annotation on a particular organ or abnormality. These models may be further updated by using training pipelines 1404. In at least one embodiment, pre-installed annotation tools may be improved over time as new labeled clinic data 1312 is added.

[0158] Such components can be used to train one or more generation models to be able to generate scene graphs for unlabeled data, as well as to generate scenes based on those scene graphs that are indistinguishable from real scenes.

[0159] Other variations are within spirit of present disclosure. Thus, while disclosed techniques are susceptible to various modifications and alternative constructions, certain illustrated embodiments thereof are shown in drawings and have been described above in detail. It should be understood, however, that there is no intention to limit disclosure to specific form or forms disclosed, but on contrary, intention is to cover all modifications, alternative constructions, and equivalents falling within spirit and scope of disclosure, as defined in appended claims.

[0160] Use of terms “a” and “an” and “the” and similar referents in context of describing disclosed embodiments (especially in context of following claims) are to be construed to cover both singular and plural, unless otherwise indicated herein or clearly contradicted by context, and not as a definition of a term. Terms “comprising,” “having,” “including,” and “containing” are to be construed as open-ended terms (meaning “including, but not limited to,”) unless otherwise noted. Term “connected,” when unmodified and referring to physical connections, is to be construed as partly or wholly contained within, attached to, or joined together, even if there is something intervening. Recitation of ranges of values herein are merely intended to serve as a shorthand method of referring individually to each separate value falling within range, unless otherwise indicated herein and each separate value is incorporated into specification as if it were individually recited herein. Use of term “set” (e.g., “a set of items”) or “subset,” unless otherwise noted or contradicted by context, is to be construed as a nonempty collection comprising one or more members. Further, unless otherwise noted or contradicted by context, term “subset” of a corresponding set does not necessarily denote a proper subset of corresponding set, but subset and corresponding set may be equal.

[0161] Conjunctive language, such as phrases of form “at least one of A, B, and C,” or “at least one of A, B and C,” unless specifically stated otherwise or otherwise clearly contradicted by context, is otherwise understood with context as used in general to present that an item, term, etc., may be either A or B or C, or any nonempty subset of set of A and B and C. For instance, in illustrative example of a set having three members, conjunctive phrases “at least one of A, B, and C” and “at least one of A, B and C” refer to any of following sets: {A}, {B}, {C}, {A, B}, {A, C}, {B, C}, {A, B, C}. Thus, such conjunctive language is not generally intended to imply that certain embodiments require at least one of A, at least one of B, and at least one of C each to be present. In addition, unless otherwise noted or contradicted by context, term “plurality” indicates a state of being plural (e.g., “a plurality of items” indicates multiple items). A plurality is at least two items, but can be more when so indicated either explicitly or by context. Further, unless stated otherwise or otherwise clear from context, phrase “based on” means “based at least in part on” and not “based solely on.”

[0162] Operations of processes described herein can be performed in any suitable order unless otherwise indicated herein or otherwise clearly contradicted by context. In at least one embodiment, a process such as those processes described herein (or variations and/or combinations thereof) is performed under control of one or more computer systems configured with executable instructions and is implemented as code (e.g., executable instructions, one or more computer programs or one or more applications) executing collectively on one or more processors, by hardware or combinations thereof. In at least one embodiment, code is stored on a computer-readable storage medium, for example, in form of a computer program comprising a plurality of instructions executable by one or more processors. In at least one embodiment, a computer-readable storage medium is a non-transitory computer-readable storage medium that excludes transitory signals (e.g., a propagating transient electric or electromagnetic transmission) but includes non-transitory data storage circuitry (e.g., buffers, cache, and queues) within transceivers of transitory signals. In at least one embodiment, code (e.g., executable code or source code) is stored on a set of one or more non-transitory computer-readable storage media having stored thereon executable instructions (or other memory to store executable instructions)

that, when executed (i.e., as a result of being executed) by one or more processors of a computer system, cause computer system to perform operations described herein. A set of non-transitory computer-readable storage media, in at least one embodiment, comprises multiple non-transitory computer-readable storage media and one or more of individual non-transitory storage media of multiple non-transitory computer-readable storage media lack all of code while multiple non-transitory computer-readable storage media collectively store all of code. In at least one embodiment, executable instructions are executed such that different instructions are executed by different processors - for example, a non-transitory computer-readable storage medium store instructions and a main central processing unit (“CPU”) executes some of instructions while a graphics processing unit (“GPU”) executes other instructions. In at least one embodiment, different components of a computer system have separate processors and different processors execute different subsets of instructions.

[0163] Accordingly, in at least one embodiment, computer systems are configured to implement one or more services that singly or collectively perform operations of processes described herein and such computer systems are configured with applicable hardware and/or software that enable performance of operations. Further, a computer system that implements at least one embodiment of present disclosure is a single device and, in another embodiment, is a distributed computer system comprising multiple devices that operate differently such that distributed computer system performs operations described herein and such that a single device does not perform all operations.

[0164] Use of any and all examples, or exemplary language (e.g., “such as”) provided herein, is intended merely to better illuminate embodiments of disclosure and does not pose a limitation on scope of disclosure unless otherwise claimed. No language in specification should be construed as indicating any non-claimed element as essential to practice of disclosure.

[0165] All references, including publications, patent applications, and patents, cited herein are hereby incorporated by reference to same extent as if each reference were individually and specifically indicated to be incorporated by reference and were set forth in its entirety herein.

[0166] In description and claims, terms “coupled” and “connected,” along with their derivatives, may be used. It should be understood that these terms may be not intended as synonyms for each other. Rather, in particular examples, “connected” or “coupled” may be used to indicate that two or more elements are in direct or indirect physical or electrical contact with each other. “Coupled” may also mean that two or more elements are not in direct contact with each other, but yet still co-operate or interact with each other.

[0167] Unless specifically stated otherwise, it may be appreciated that throughout specification terms such as “processing,” “computing,” “calculating,” “determining,” or like, refer to action and/or processes of a computer or computing system, or similar electronic computing device, that manipulate and/or transform data represented as physical, such as electronic, quantities within computing system’s registers and/or memories into other data similarly represented as physical quantities within computing system’s memories, registers or other such information storage, transmission or display devices.

[0168] In a similar manner, term “processor” may refer to any device or portion of a device that processes electronic data from registers and/or memory and transform that electronic data into other electronic data that may be stored in registers and/or memory. As non-limiting examples, “processor” may be a CPU or a GPU. A “computing platform” may comprise one or more processors. As used herein, “software” processes may include, for example, software and/or hardware entities that perform work over time, such as tasks, threads, and intelligent agents. Also, each process may refer to multiple processes, for carrying out instructions in sequence or in parallel, continuously or intermittently. Terms “system” and “method” are used herein interchangeably insofar as system may embody one or more methods and methods may be considered a system.

[0169] In present document, references may be made to obtaining, acquiring, receiving, or inputting analog or digital data into a subsystem, computer system, or computer-implemented machine. Obtaining, acquiring, receiving, or inputting analog and digital data can be accomplished in a variety of ways such as by receiving data as a parameter of a function call or a call to an application programming interface. In some implementations, process of obtaining,

acquiring, receiving, or inputting analog or digital data can be accomplished by transferring data via a serial or parallel interface. In another implementation, process of obtaining, acquiring, receiving, or inputting analog or digital data can be accomplished by transferring data via a computer network from providing entity to acquiring entity. References may also be made to providing, outputting, transmitting, sending, or presenting analog or digital data. In various examples, process of providing, outputting, transmitting, sending, or presenting analog or digital data can be accomplished by transferring data as an input or output parameter of a function call, a parameter of an application programming interface or interprocess communication mechanism.

[0170] Although discussion above sets forth example implementations of described techniques, other architectures may be used to implement described functionality, and are intended to be within scope of this disclosure. Furthermore, although specific distributions of responsibilities are defined above for purposes of discussion, various functions and responsibilities might be distributed and divided in different ways, depending on circumstances.

[0171] Furthermore, although subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that subject matter claimed in appended claims is not necessarily limited to specific features or acts described. Rather, specific features and acts are disclosed as exemplary forms of implementing the claims.

CLAIMS

WHAT IS CLAIMED IS:

- 1 1. A computer-implemented method, comprising:
2 encoding, to a latent space, one or more features of a first set of labeled synthetic
3 data and one or more features of a second set of unlabeled real data;
4 providing the latent space as an input to train a scene graph prediction network;
5 aligning one or more features in the latent space and one or more features in an
6 output space of the scene prediction network;
7 aligning one or more labels of the synthetic data with the real data; and
8 training the scene graph generation model using the one or more aligned labels.
- 1 2. The computer-implemented method of claim 1, further comprising:
2 aligning the one or more features in the latent space and the one or more features
3 in the output space using at least one of: one or more gradient reversal layers (GRLs) or a
4 domain discriminator.
- 1 3. The computer-implemented method of claim 2, wherein aligning the one
2 or more features reduces one or more discrepancies in at least one of: an appearance or a content
3 between the set of labeled synthetic data and unlabeled real data.
- 1 4. The computer-implemented method of claim 1, further comprising:
2 aligning the one or more labels using pseudo statistics-based self-learning.
- 1 5. The computer-implemented method of claim 1, further comprising:
2 receiving an image; and
3 generating a scene graph for the image using the trained scene graph prediction
4 network.
- 1 6. The computer-implemented method of claim 1, wherein the training
2 comprises applying a network convergence criterion.

1 7. The computer-implemented method of claim 1, further comprising:
2 generating a scene graph using the trained scene graph prediction network; and
3 generating a synthesized image from the generated scene graph.

1 8. The computer-implemented method of claim 1, wherein the one or more
2 features are encoded to the latent space using one or more convolutional neural networks
3 (CNNs), the latent space including one or more features for both the first set of labeled synthetic
4 data and the second set of unlabeled real data.

1 9. The computer-implemented method of claim 1, wherein aligning the one
2 or more features reduces an appearance gap between the synthetic data and the real data, and
3 wherein aligning the one or more labels reduces a content gap between the synthetic data and the
4 real data.

1 10. A system, comprising:
2 at least one processor; and
3 memory including instructions that, when executed by the at least one processor,
4 cause the system to:
5 encode, to a latent space, one or more features of a first set of labeled
6 synthetic data and a second set of unlabeled real data;
7 provide the latent space as an input to train a scene graph prediction
8 network;
9 align the one or more features in the latent space and one or more features
10 of an output space of the scene prediction network;
11 align one or more labels between the synthetic data and the real data; and
12 train the scene graph generation model using the one or more aligned
13 labels.

1 11. The system of claim 10, wherein the instructions when executed further
2 cause the system to:

3 align the one or more features in the latent space and the one or more features in
4 the output space using at least one of: one or more gradient reversal layers (GRLs) or a domain
5 discriminator.

1 12. The system of claim 11, wherein aligning the one or more features reduces
2 one or more discrepancies in appearance and content between the set of labeled synthetic data
3 and unlabeled real data.

1 13. The system of claim 10, wherein the instructions when executed further
2 cause the system to:
3 align one or more labels using pseudo statistics-based self-learning.

1 14. The system of claim 10, wherein the instructions when executed further
2 cause the system to:
3 receive an unlabeled image; and
4 generate a scene graph for the image using the trained scene graph prediction
5 network.

1 15. The system of claim 10, wherein the system comprises at least one of:
2 a system for performing graphical rendering operations;
3 a system for performing simulation operations;
4 a system for performing simulation operations to test or validate autonomous
5 machine applications;
6 a system for performing deep learning operations;
7 a system implemented using an edge device;
8 a system incorporating one or more Virtual Machines (VMs);
9 a system implemented at least partially in a data center; or
10 a system implemented at least partially using cloud computing resources.

1 16. A non-transitory computer-readable storage medium including instructions
2 that, when executed, cause one or more processors to:

3 encode, to a latent space, one or more features of a first set of labeled synthetic
4 data and one or more features of a second set of unlabeled real data;
5 provide the latent space as an input to train a scene graph prediction network;
6 align the one or more features in the latent space and the one or more features in
7 the output space of the scene prediction network;
8 align one or more labels of the synthetic data with the real data; and
9 train the scene graph generation model using the one or more labels.

1 17. The non-transitory computer-readable storage medium of claim 16,
2 wherein the instructions when executed further cause one or more processors to:
3 align the one or more features in the latent space and the one or more features in
4 the output space using at least one of: gradient reversal layers (GRLs) or a domain discriminator.

1 18. The non-transitory computer-readable storage medium of claim 17,
2 wherein aligning the one or more features reduces one or more discrepancies in appearance and
3 content between the set of labeled synthetic data and unlabeled real data.

1 19. The non-transitory computer-readable storage medium of claim 16,
2 wherein the instructions when executed further cause one or more processors to:
3 align the one or more labels using pseudo statistics-based self-learning.

1 20. The non-transitory computer-readable storage medium of claim 16,
2 wherein the instructions when executed further cause one or more processors to:
3 receive an unlabeled image; and
4 generate a scene graph for the image using the trained scene graph prediction
5 network.




FIG. 1A




FIG. 1B




FIG. 2




FIG. 3A




FIG. 3B




FIG. 4A




FIG. 4B




FIG. 5




FIG. 6




FIG. 7A




FIG. 7B




FIG. 8





FIG. 9






FIG. 11




FIG. 12




FIG. 13




FIG. 14





FIG. 15A

**FIG. 15B**

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2021/034221

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06K9/62

ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06K

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

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

EP0-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X,P	AAYUSH PRAKASH ET AL: "Sim2SG: Sim-to-Real Scene Graph Generation for Transfer Learning", ARXIV.ORG, CORNELL UNIVERSITY LIBRARY, 201 OLIN LIBRARY CORNELL UNIVERSITY ITHACA, NY 14853, 30 November 2020 (2020-11-30), XP081824607, Sections 1 and 2; Appendix	1-20
A	SHANU KUMAR ET AL: "Adversarial Adaptation of Scene Graph Models for Understanding Civic Issues", ARXIV.ORG, CORNELL UNIVERSITY LIBRARY, 201 OLIN LIBRARY CORNELL UNIVERSITY ITHACA, NY 14853, 29 January 2019 (2019-01-29), XP081009440, Sections 1, 2, 5, 6 ----- -/--	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

8 October 2021

Date of mailing of the international search report

18/10/2021

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Miclea, Sorin

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2021/034221

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	JIUXIANG GU ET AL: "Scene Graph Generation with External Knowledge and Image Reconstruction", ARXIV.ORG, CORNELL UNIVERSITY LIBRARY, 201 OLIN LIBRARY CORNELL UNIVERSITY ITHACA, NY 14853, 1 April 2019 (2019-04-01), XP081163157, Section 3 -----	1-20
A	YU SUN ET AL: "Unsupervised Domain Adaptation through Self-Supervision", ARXIV.ORG, CORNELL UNIVERSITY LIBRARY, 201 OLIN LIBRARY CORNELL UNIVERSITY ITHACA, NY 14853, 26 September 2019 (2019-09-26), XP081483393, Sections 1, 3-5 -----	1-20
A	WEI-LUN CHANG ET AL: "All about Structure: Adapting Structural Information across Domains for Boosting Semantic Segmentation", ARXIV.ORG, CORNELL UNIVERSITY LIBRARY, 201 OLIN LIBRARY CORNELL UNIVERSITY ITHACA, NY 14853, 26 March 2019 (2019-03-26), XP081159494, Sections 1, 3 -----	1-20