



(51) International Patent Classification:

G06K 9/62 (2006.01)

G06T 11/00 (2006.01)

G06N 3/08 (2006.01)

(21) International Application Number:

PCT/US2020/015262

(22) International Filing Date:

27 January 2020 (27.01.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/261,395

29 January 2019 (29.01.2019)

US

(71) Applicant: **NVIDIA CORPORATION** [US/US]; 2788 San Tomas Expressway, Santa Clara, California 95051 (US).

(72) Inventors: **LIU, Ming-Yu**; 2788 San Tomas Expressway, Santa Clara, California 95051 (US). **HUANG, Xun**; 2788 San Tomas Expressway, Santa Clara, California 95051 (US). **KARRAS, Tero**; 2788 San Tomas Expressway, Santa Clara, California 95051 (US). **AILA, Timo**; 2788 San Tomas Expressway, Santa Clara, California 95051 (US). **LEHTINEN, Jaakko**; 2788 San Tomas Expressway, Santa Clara, California 95051 (US).

(74) Agent: **LOHR, Jason** et al.; 3 Embarcadero Center, Suite 1500, San Francisco, California 94111 (US).

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

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

(54) Title: METHOD FOR FEW-SHOT UNSUPERVISED IMAGE-TO-IMAGE TRANSLATION




FIG. 1

(57) Abstract: A few-shot, unsupervised image-to-image translation ("FUNIT") algorithm is disclosed that accepts as input images of previously-unseen target classes. These target classes are specified at inference time by only a few images, such as a single image or a pair of images, of an object of the target type. A FUNIT network can be trained using a data set containing images of many different object classes, in order to translate images from one class to another class by leveraging few input images of the target class. By learning to extract appearance patterns from the few input images for the translation task, the network learns a generalizable appearance pattern extractor that can be applied to images of unseen classes at translation time for a few-shot image-to-image translation task.

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))*

METHOD FOR FEW-SHOT UNSUPERVISED IMAGE-TO-IMAGE TRANSLATION

CROSS-REFERENCE TO RELATED APPLICATION

5 [0001] This application claims benefit of priority from U.S. Application No. 16/261,395, filed January 29, 2019, entitled “METHOD FOR FEW-SHOT UNSUPERVISED IMAGE-TO-IMAGE TRANSLATION” the full disclosure of which is incorporated herein by reference for all purposes.

BACKGROUND

10 [0002] Advances in processing power and image manipulation software have enabled an increasing variety of image creation and manipulation capabilities. For example, an image of a first type of object can be used to generate an image showing the first type of object having an aspect of a second type of object. In order to accomplish such generation, however, a user either has to manually generate or manipulate an image, or has to provide a large number of input
15 images of the first type of object that enable adequate generation of the target image. This may be complex and time consuming in the case of manual generation, and may not be practical in situations where a user might only have one or two images of the first type of object.

BRIEF DESCRIPTION OF THE DRAWINGS

20 [0003] Various embodiments in accordance with the present disclosure will be described with reference to the drawings, in which:

[0004] FIG. 1 illustrates an example image translation that can be performed in accordance with various embodiments.

[0005] FIG. 2 illustrates components of an example image translation network that can be
25 utilized in accordance with various embodiments.

[0006] FIG. 3 illustrates an example system that can be utilized to perform an image translation in accordance with various embodiments.

[0007] FIG. 4 illustrates an example process for performing an image translation in accordance with various embodiments.

5 [0008] FIG. 5 illustrates an example system for training an image synthesis network that can be utilized in accordance with various embodiments.

[0009] FIG. 6 illustrates layers of an example statistical model that can be utilized in accordance with various embodiments.

10 [0010] FIG. 7 illustrates example components of a computing device that can be used to implement aspects of the various embodiments.

DETAILED DESCRIPTION

[0011] In the following description, various embodiments will be described. For purposes of explanation, specific configurations and details are set forth in order to provide a thorough understanding of the embodiments. However, it will also be apparent to one skilled in the art
15 that the embodiments may be practiced without the specific details. Furthermore, well-known features may be omitted or simplified in order not to obscure the embodiment being described.

[0012] Approaches in accordance with various embodiments provide for the generation of images including representations of objects having one or more specific aspects. In particular, various embodiments provide a few-shot, unsupervised image-to-image translation (“FUNIT”) algorithm that works on previously-unseen target classes that are specified, at inference time, only
20 by a few example images, such as a single image or a pair of input images of an object of the target type. A FUNIT network can be trained using a data set containing images of many different object classes, in order to translate images from one class to another class by leveraging few input images of the target class. By learning to extract appearance patterns from the few input images for the
25 translation task, the network learns a generalizable appearance pattern extractor that can be applied to images of unseen classes at translation time for a few-shot image-to-image translation task.

[0013] Various other functions can be implemented within the various embodiments as well as discussed and suggested elsewhere herein.

[0014] FIG. 1 illustrates an example image translation 100 that can be performed in accordance with various embodiments. The translation can accept an input image 106 (or a digital representation of an image) including a representation of a type of object, in this case a breed of dog corresponding to a pug. It might be desired to generate an image of the pug in a different pose, such as to translate the representation of the pug from the pose of the target image 106 to a target pose, such as the pose in a selected or provided image 104. As illustrated, the image with the target pose includes a representation of a different type of object, in this example a lion. The image translation can take what it has learned from images of lions (and other such object classes) and use that learning to generate an image 108 of a pug exhibiting the target pose. Of interest in this example is the fact that the image translation network has been trained using images 102 of many classes of object types, including many images of the lion object type. The image translation network has not previously been trained on, or processed, any images including a representation of a pug.

[0015] Approaches in accordance with various embodiments can perform such an image generation, or translation as referred to herein, using one or more image-to-image translation methods. In various embodiments, unsupervised image-to-image translation techniques can be used that learn to map images in a given class to an analogous image in a different class, drawing on unstructured (i.e., non-registered) datasets of images. While some conventional approaches can provide images with translations in object pose, these approaches require access to many images in both the source class (e.g., exhibiting the target pose) and the destination class (including the previously unseen object to obtain the desired pose) either during training or at the time for translation. The need for a large number of input images of the desired object class greatly limits the usefulness of these conventional approaches. Drawing inspiration from the human capability of picking up the essence of a novel object from a small number of examples and generalizing from there, approaches in accordance with various embodiments utilize a few-shot, unsupervised image-to-image translation algorithm that works on previously unseen target classes that are specified, at translation time, only by a few example images, such as a single image or a pair of

input images of an object of the target type. In some embodiments the algorithm utilizes three or less, while in other embodiments such an algorithm can work with less than five input images of the target class. In some embodiments a single image can be used, but accuracy or photorealism may be increased using a second or third image of the target object class in different poses. An example model achieves this few-shot imagination capability by coupling an adversarial training scheme with a novel network design.

[0016] As mentioned, the human brain is remarkably good at generalization. When given a picture of a previously-unseen type of animal, for example, the human brain can form a vivid mental picture of the previously-unseen animal in various poses, particularly when that human has been exposed to images or views of similar, but different, animals in those poses before. For example, a person seeing a standing pug for the first time (as illustrated in the example input image 106 of FIG. 1) will have little to no trouble imagining what the pug would look like sitting on its hind legs (as in the example output image 108) or lying down, given a lifetime of experience of other animals in similar poses and how those animals appear in those poses.

[0017] While some conventional unsupervised image-to-image translation algorithms provide reasonable results in transferring complex appearance changes across image classes, the capability to generalize from a few samples of a new class based on prior knowledge is not possible using these algorithms. Concretely, these algorithms require large training sets over all classes of images on which translation is to be performed, and these algorithms do not support few-shot generalization.

[0018] Approaches in accordance with various embodiments attempt to bridge the gap between human and machine imagination capability using what is referred to herein as a Few-shot UN-supervised Image-to-image Translation (“FUNIT”) framework. A FUNIT framework can learn an image-to-image translation model for mapping an image of a source class to an analogous image of a target class by leveraging few images (i.e., two or less) of the target class given at the time for translation. The translation model in may never have encountered or processed images of the target class, during training or otherwise, but is asked to generate some of them at test time. The approach can be based on the notion that the few-shot generation capability of humans develops from their past visual experiences. For example, a person can better imagine views of a new object if the

person has seen many different object classes in the past. Using such an approach, a FUNIT model can be trained using a dataset containing images of many different object classes for simulating these past visual experiences. Specifically, the model can be trained to translate images from one class to another class by leveraging few input images of the target class. By learning to extract appearance patterns from the few input images for the translation task, the model learns a generalizable appearance pattern extractor that can be applied to images of unseen classes at translation time for a few-shot image-to-image translation task. In at least some embodiments, performance of such a few-shot translation approach improves as the number of classes in the training set increases. For discussion purposes herein, the source class will refer to an object class included in the training data that includes the desired pose, or other such aspect, of an object of the source class. A target class will refer to an object class that is not included in the training data, and is only seen at the time for testing, transformation, or inference.

[0019] An image translation framework in accordance with various embodiments can be based on a Generative Adversarial Networks (GAN). The use of an adversarial training scheme coupled with a specific type of network design can achieve the desired few-shot unsupervised image-to-image translation capability. Unsupervised and/or unpaired image-to-image translation can attempt to learn a conditional image generation function that can map an input image of a source class to an analogous image of a target class without pair supervision. Such an approach may not be sufficient as it attempts to recover the joint distribution using samples from marginal distributions. To deal with the problem, conventional approaches use additional constraints, such as to enforce the translation to preserve certain properties of the source data, such as pixel values, pixel gradients, semantic features, class labels, or pairwise sample distances. Other approaches are utilized as well.

[0020] Approaches in accordance with various embodiments utilize a partially-shared latent space assumption designed for a few-shot unsupervised image-to-image translation task. While capable of generating realistic translation outputs, conventional unsupervised image-to-image translation models are limited in at least two aspects. First, these models are sample inefficient, generating poor translation outputs if only few images are given at training time. Second, the learned models are limited to translating images between two classes. A trained model for one translation task cannot be directly reused for a new task despite similarity between the new task and the original

task. For example, a husky-to-cat translation model cannot be re-purposed for husky-to-tiger translation, even though a cat and a tiger may share a great similarity. While some conventional approaches utilize an unsupervised image-to-image translation framework for partially addressing the first aspect, approaches in accordance with various embodiments assume many source class images but few target class images. Moreover, it can be assumed that the few target class images are only available at test or translation time, and can be from many different object classes.

[0021] Multi-class unsupervised image-to-image translation extends the unsupervised image-to-image translation methods to multiple classes. Approaches in accordance with various embodiments utilize a training dataset consisting of images of multiple classes, but instead of translating images among the observed classes, the focus is placed on translating images of these observed classes to analogous images of previously unseen classes.

[0022] A FUNIT framework in accordance with various embodiments can map an image of a source class to an analogous image of an unseen target class by leveraging a few target class images that are made available at the time for translation. The training data can consist of unpaired images from a set of object classes (*e.g.*, images of various breeds of dogs and cats), which are referred to herein as the source classes. These images can be used to train an unsupervised image-to-image translation model, translating images between any two source classes. During testing, a few images from a novel object class (*e.g.*, a mountain lion class), referred to as the target class, can be provided as input to the model. The model can leverage the few target images to translate any source class image to an analogous image of the target class.

[0023] A FUNIT framework in accordance with various embodiments can include a conditional image generator, G , and a multi-task adversarial discriminator, D . Unlike conditional image generators in conventional unsupervised image-to-image translation frameworks which take one image as input, a generator in accordance with various embodiments can simultaneously accept as input a content image, x , belonging to a first class, and a set of images ($y_1, \dots, y_K$) belonging to a second class. An output image can be produced via the following equation:

$$\bar{x} = G(x, \{y_1, \dots, y_K\})$$

It can be assumed that the content image belongs to destination object class while each of the K class images belongs to the source object class. The generator can map an input content image to an output image, such that the output image has the appearance of an image belonging to destination object class but with a pose of the source object class, where the source and destination objects share at least some structural similarity. During training, the generator can learn to translate images between two randomly sampled source classes. At translation time, the generator can accept as input one or a few images from an unseen target class as the destination class images, and map an image sampled from any of the source classes to an analogous image of the target or destination class.

[0024] FIG. 2 illustrates components of an example FUNIT framework 200 that can be utilized in accordance with various embodiments. As mentioned, a training set can consist of images of various object classes, or source classes. A model, such as a GAN, can be trained to translate images between any or all of these source object classes. In deployment, the trained model can receive as input very few images of the target class, which is sufficient to translate images of source classes to one or more analogous images of the target class. An example FUNIT generator can accept as input a content image and one or more target class images. The generator can generate a translation of the input content image that resembles, or has an appearance of, the image(s) of the target class. As illustrated in FIG. 2, the FUNIT generator can accept as input an image of a first object class and a set of images of a second object class, and can generate a translation output. The translation output should resemble images of the second object class and share the same class-invariant appearance with the input image.

[0025] An example generator G consists of four primary components: a content encoder, a class encoder, an adaptive instance-norm (AdaIN) decoder, and an image decoder. Instance-norm and rectified linear units (ReLUs) can be applied to each convolutional and fully-connected layer of the network. The content encoder can contain several convolutional layers followed by several residual blocks. The content encoder can map an input content image, x , to a spatially distributed feature map z_x , referred to herein as the content latent code. The class encoder can comprise several convolutional layers followed by an average pooling layer. The class encoder can accept the class images as input and outputs a class latent code. The average pooling layer

can average activations first across spatial dimensions (*e.g.*, height and width) and then across the set of images. The image decoder can comprise several AdaIN residual blocks followed by a couple of upscale convolutional layers. The AdaIN residual block can be a residual block using the AdaIN layer as the normalization layer. For each sample, the AdaIN layer (also referred to as a normalization layer) can first normalize the activations in each channel to a zero mean and unit variance distribution. The normalization layer can then transform the distribution, through a de-normalization process, to have specific mean and variance values. A primary goal of the image decoder is to decode the content latent code and the class latent code to generate a translation of the input content image. In some embodiments the AdaIN decoder is a multilayer perceptron. It decodes the class latent code to a set of mean and variance vectors that are used as the new means and variances for the respective channels in the respective AdaIN residual block in the image decoder. Using such a generator design, a class-invariant latent representation (*e.g.*, an object pose) can be extracted using the content encoder, and a class-specific latent representation (*e.g.*, an object appearance) can be extracted using the class encoder. By feeding the class latent code to the image decoder via the AdaIN layers, the class images are enabled to control the spatially invariant means and variances, while the content image determines the remaining information. At training time, the class encoder can learn to extract a class-specific latent representation from the images of the source classes. At testing or translation time, this generalizes to images of previously unseen class.

[0026] An example multi-task adversarial discriminator D is trained by solving multiple adversarial classification tasks simultaneously. The discriminator in some embodiments is a patch GAN discriminator that can render an output spatial map for an input image, where each entry in the map indicates the score for the corresponding patch in the input image. Each of the tasks to be solved can be a binary classification task in some embodiments, determining whether an input image to D is a real image of a source class or a translation output coming from the generator. As there are a number of source classes, the discriminator can be designed to produce a corresponding number of outputs. When updating D for a real image of source class, D can be penalized if a certain output is negative. For a translation output yielding a fake image of the source class, D can be penalized if a corresponding output is positive. D may not be penalized for not predicting negatives for images of other classes. When updating the generator G , G may only be penalized if

the specified output of D is negative. The discriminator D can be designed in some embodiments based on a class-conditional discriminator that consists of several residual blocks followed by a global average pooling layer. The feature produced by the global average pooling layer is called the discriminator feature, from which classification scores can be produced using linear mappings.

- 5 [0027] An example FUNIT framework can be trained by solving a minimax optimization problem given by:

$$\min_D \max_G L_{GAN}(D, G) + \lambda_R L_R(G) + \lambda_{FM} L_{FM} + \lambda_{GP} L_{GP}(D)$$

- 10 where L_{GAN} , L_R , L_{FM} , and L_{GP} are the GAN loss, content reconstruction loss, feature matching loss, and gradient penalty loss, respectively. The scalars, λ_R , λ_{FM} , and λ_{GP} control the relative importance of the terms, and their values are set to 0.1, 1, and 10, respectively. The GAN loss can be a conditional loss. The discriminator can be trained to predict a value such as “true” for a real image and “false” for a translation output, while the generator can be trained to make the discriminator predict true for an appropriate translation output.

- 15 [0028] In some embodiments FUNIT model training can be regularized using the image reconstruction loss. Specifically, when using the same image as the input content image and input class image, the model should generate an output image identical to the input. Feature matching loss L_{FM} can be used to further improve the few-shot translation performance. Specifically, a feature extractor, referred to as D_f , can be constructed by removing the last (i.e., prediction) layer from D . D_f can be used to extract features from the translation output and the class images and minimize the L1 loss given. A zero-centered gradient penalty loss can be applied to real images for a more stable GAN training in some embodiments. This can involve penalizing the discriminator for having a non-zero gradient field.
- 20

- [0029] FIG. 3 illustrates an example environment 300 that can be utilized to implement aspects of the various embodiments. In some embodiments, a user may utilize a client device 602 to provide an input image, which may be an image including a representation of an object that has not previously been encountered by an image generator network 328 as discussed herein. The user may also utilize the client device to select an image indicating a target pose for which a translation is to be performed for the object in the input image. The client device can be any
- 25

appropriate computing device capable of enabling a user to select and/or provide images for processing, such as may include a desktop computer, notebook computer, smart phone, tablet computer, computer workstation, gaming console, and the like. A user can select, provide, or otherwise specify the transformation input a user interface (UI) of an image editor application 5 306 (or other image manipulation or generation software package) running on the client device, although at least some functionality may also operate on a remote device, networked device, or in “the cloud” in some embodiments. The user can provide input to the UI, such as through a touch-sensitive display 304 or by moving a mouse cursor displayed on a display screen, among other such options. As mentioned, the user may be able to provide an input image of a target 10 class, and may select an image indicating a desired pose or other such aspect, which may be exhibited by an object in an image of a source class. The client device can include at least one processor 308 (*e.g.*, a CPU or GPU) to execute the application and/or perform tasks on behalf of the application, and memory 310 for including the non-transitory computer-readable instructions for execution by the processor. Images provided to, or generated via, the application can be 15 stored locally to local storage 312, such as a hard drive or flash memory, among other such options.

[0030] In some embodiments, input images received or selected on the client device 302 can be processed on the client device in order to generate an image with the desired translation, such as to apply the appearance of a target image to a pose extracted from a set of source images. In 20 other embodiments, the client device 302 may send the input images, data extracted from the images, or data specifying the images over at least one network 314 to be received by a remote computing system, as may be part of a resource provider environment 316. The at least one network 314 can include any appropriate network, including an intranet, the Internet, a cellular network, a local area network (LAN), or any other such network or combination, and 25 communication over the network can be enabled via wired and/or wireless connections. The provider environment 316 can include any appropriate components for receiving requests and returning information or performing actions in response to those requests. As an example, the provider environment might include Web servers and/or application servers for receiving and processing requests, then returning data or other content or information in response to the 30 request.

[0031] Communications received to the provider environment 316 can be received to an interface layer 318. The interface layer 318 can include application programming interfaces (APIs) or other exposed interfaces enabling a user to submit requests to the provider environment. The interface layer 318 in this example can include other components as well, such as at least one Web server, routing components, load balancers, and the like. Components of the interface layer 318 can determine a type of request or communication, and can direct the request to the appropriate system or service. For example, if a communication is to train an image translation network for classes of image content, such as birds, animals, or people, the communication can be directed to an image manager 320, which can be a system or service provided using various resources of the provider environment 316. The communication, or information from the communication, can be directed to a training manager 324, which can select an appropriate model or network and then train the model using relevant training images and/or data 324. Once a network is trained and successfully evaluated, the network can be stored to a model repository 326, for example, that may store different models or networks for different types of image translation or processing. If a request is received to the interface layer 318 that includes input to be used for an image translation, information for the request can be directed to an image generator 328, also referred to herein as part of an image translation network or service, that can obtain the corresponding trained network, such as a trained generative adversarial network (GAN) as discussed herein, from the model repository 326 if not already stored locally to the generator 328. The image generator 328 can take as input the target image (or few images) and data indicating the target pose, as may be exhibited by a selected source image as discussed herein. The image generator 328 can then cause the input to be processed to generate an image representing the target transformation. As mentioned, this can involve the input being processed by the content encoder 330 to extract a class-invariant latent representation, such as may correspond to the selected pose. The class encoder 332 can extract the class-specific latent representation (e.g., an object appearance) from the input target image. The class latent code can be fed to AdaIN decoder 334, which decodes the class latent code to a set of mean and variance vectors that are used as the new means and variances for the respective channels in the respective AdaIN residual block in the image decoder 336. The generated image can then be transmitted to the client device 302 for display on the display element 304, or for other such usage. If the user

wants to modify any aspects of the image, the user can provide additional input to the application 306, which can cause a new or updated image to be generated using the same process for the new or updated input, such as an additional image of the target class or specification of a different pose, among other such options. In some embodiments, an image generation network can utilize a deep generative model that can learn to sample images given a training dataset. The models used can include, for example, generative adversarial networks (GANs) and variational auto-encoder (VAE) networks while aiming for an image translation task. An image translation network, or translator 340, can comprise a GAN in various embodiments that consists of a generator 328 and a discriminator 338. The generator 328 can be used to produce translated images so that the discriminator cannot differentiate between real and generated.

[0032] In various embodiments the processor 308 (or a processor of the training manager 322 or image translator 340) will be a central processing unit (CPU). As mentioned, however, resources in such environments can utilize GPUs to process data for at least certain types of requests. With thousands of cores, GPUs are designed to handle substantial parallel workloads and, therefore, have become popular in deep learning for training neural networks and generating predictions. While the use of GPUs for offline builds has enabled faster training of larger and more complex models, generating predictions offline implies that either request-time input features cannot be used or predictions must be generated for all permutations of features and stored in a lookup table to serve real-time requests. If the deep learning framework supports a CPU-mode and the model is small and simple enough to perform a feed-forward on the CPU with a reasonable latency, then a service on a CPU instance could host the model. In this case, training can be done offline on the GPU and inference done in real-time on the CPU. If the CPU approach is not a viable option, then the service can run on a GPU instance. Because GPUs have different performance and cost characteristics than CPUs, however, running a service that offloads the runtime algorithm to the GPU can require it to be designed differently from a CPU based service.

[0033] FIG. 4 illustrates an example process 400 for generating an image representing a translation of a representation of a previously unseen object to have a pose learned from a known source object class that can be utilized in accordance with various embodiments. It should be

understood for this and other processes discussed herein that there can be additional, alternative, or fewer steps performed in similar or alternative orders, or in parallel, within the scope of the various embodiments unless otherwise stated. Further, while pose is used as a primary example, there can be other aspects of the source image class that are utilized to generate a transformation of an input object as discussed and suggested herein. In this example, an input digital image is received 402, or otherwise obtained or specified, that includes a representation of an object of interest. In this example, the object is of a source class that has not yet been encountered by a neural network to be used to process the image, or at least is of a type that was not used to train the network. A selection of a source image can also be obtained 404, where the source image includes a representation of a source object of an object class that was used to train the neural network. In some embodiments a user might select the source object, while in others the system or network may select a source image or pose, among other options discussed and suggested herein.

[0034] Once obtained, the input image and source image (or indication of the source image) can be provided 406 as input to the trained neural network. As discussed herein, the network can have been trained using many images of multiple object classes, including the source class, learning to translate between objects of the various classes. One or more encoders of a generator of the network can extract 408 a class-invariant latent representation corresponding to the target pose from the source image or class. A class-specific latent representation can also be extracted 410 that corresponds to an appearance of the target object from the input image. The neural network can then be used to infer 412, based at least in part upon these latent representations, a translated image including a representation of an object of the target class having the target pose. As mentioned, such a process can produce a quality translated image without requiring a large number of input images of the target class.

[0035] As mentioned, various embodiments take advantage of machine learning. As an example, deep neural networks (DNNs) developed on processors have been used for diverse use cases, from self-driving cars to faster drug development, from automatic image captioning in online image databases to smart real-time language translation in video chat applications. Deep learning is a technique that models the neural learning process of the human brain, continually

learning, continually getting smarter, and delivering more accurate results more quickly over time. A child is initially taught by an adult to correctly identify and classify various shapes, eventually being able to identify shapes without any coaching. Similarly, a deep learning or neural learning system needs to be trained in object recognition and classification for it get
5 smarter and more efficient at identifying basic objects, occluded objects, etc., while also assigning context to objects.

[0036] At the simplest level, neurons in the human brain look at various inputs that are received, importance levels are assigned to each of these inputs, and output is passed on to other neurons to act upon. An artificial neuron or perceptron is the most basic model of a neural
10 network. In one example, a perceptron may receive one or more inputs that represent various features of an object that the perceptron is being trained to recognize and classify, and each of these features is assigned a certain weight based on the importance of that feature in defining the shape of an object.

[0037] A deep neural network (DNN) model includes multiple layers of many connected
15 perceptrons (e.g., nodes) that can be trained with enormous amounts of input data to quickly solve complex problems with high accuracy. In one example, a first layer of the DNN model breaks down an input image of an automobile into various sections and looks for basic patterns such as lines and angles. The second layer assembles the lines to look for higher level patterns such as wheels, windshields, and mirrors. The next layer identifies the type of vehicle, and the
20 final few layers generate a label for the input image, identifying the model of a specific automobile brand. Once the DNN is trained, the DNN can be deployed and used to identify and classify objects or patterns in a process known as inference. Examples of inference (the process through which a DNN extracts useful information from a given input) include identifying handwritten numbers on checks deposited into ATM machines, identifying images of friends in
25 photos, delivering movie recommendations to over fifty million users, identifying and classifying different types of automobiles, pedestrians, and road hazards in driverless cars, or translating human speech in real-time.

[0038] During training, data flows through the DNN in a forward propagation phase until a prediction is produced that indicates a label corresponding to the input. If the neural network

does not correctly label the input, then errors between the correct label and the predicted label are analyzed, and the weights are adjusted for each feature during a backward propagation phase until the DNN correctly labels the input and other inputs in a training dataset. Training complex neural networks requires massive amounts of parallel computing performance, including floating-point multiplications and additions that are supported. Inferencing is less compute-intensive than training, being a latency-sensitive process where a trained neural network is applied to new inputs it has not seen before to classify images, translate speech, and generally infer new information.

[0039] Neural networks rely heavily on matrix math operations, and complex multi-layered networks require tremendous amounts of floating-point performance and bandwidth for both efficiency and speed. With thousands of processing cores, optimized for matrix math operations, and delivering tens to hundreds of TFLOPS of performance, a computing platform can deliver performance required for deep neural network-based artificial intelligence and machine learning applications.

[0040] FIG. 5 illustrates an example system 500 that can be used to classify data, or generate inferences, in accordance with various embodiments. Various predictions, labels, or other outputs can be generated for input data as well, as should be apparent in light of the teachings and suggestions contained herein. Further, both supervised and unsupervised training can be used in various embodiments discussed herein. In this example, a set of classified data 502 is provided as input to function as training data. The classified data can include instances of at least one type of object for which a statistical model is to be trained, as well as information that identifies that type of object. For example, the classified data might include a set of images that each includes a representation of a type of object, where each image also includes, or is associated with, a label, metadata, classification, or other piece of information identifying the type of object represented in the respective image. Various other types of data may be used as training data as well, as may include text data, audio data, video data, and the like. The classified data 502 in this example is provided as training input to a training manager 504. The training manager 504 can be a system or service that includes hardware and software, such as one or more computing devices executing a training application, for training the statistical

model. In this example, the training manager 504 will receive an instruction or request indicating a type of model to be used for the training. The model can be any appropriate statistical model, network, or algorithm useful for such purposes, as may include an artificial neural network, deep learning algorithm, learning classifier, Bayesian network, and the like. The training manager 504 can select a base model, or other untrained model, from an appropriate repository 506 and utilize the classified data 502 to train the model, generating a trained model 508 that can be used to classify similar types of data. In some embodiments where classified data is not used, the appropriate based model can still be selected for training on the input data per the training manager.

[0041] The model can be trained in a number of different ways, as may depend in part upon the type of model selected. For example, in one embodiment a machine learning algorithm can be provided with a set of training data, where the model is a model artifact created by the training process. Each instance of training data contains the correct answer (e.g., classification), which can be referred to as a target or target attribute. The learning algorithm finds patterns in the training data that map the input data attributes to the target, the answer to be predicted, and a machine learning model is output that captures these patterns. The machine learning model can then be used to obtain predictions on new data for which the target is not specified.

[0042] In one example, a training manager can select from a set of machine learning models including binary classification, multiclass classification, and regression models. The type of model to be used can depend at least in part upon the type of target to be predicted. Machine learning models for binary classification problems predict a binary outcome, such as one of two possible classes. A learning algorithm such as logistic regression can be used to train binary classification models. Machine learning models for multiclass classification problems allow predictions to be generated for multiple classes, such as to predict one of more than two outcomes. Multinomial logistic regression can be useful for training multiclass models. Machine learning models for regression problems predict a numeric value. Linear regression can be useful for training regression models.

[0043] In order to train a machine learning model in accordance with one embodiment, the training manager must determine the input training data source, as well as other information such

as the name of the data attribute that contains the target to be predicted, required data transformation instructions, and training parameters to control the learning algorithm. During the training process, a training manager in some embodiments may automatically select the appropriate learning algorithm based on the type of target specified in the training data source.

- 5 Machine learning algorithms can accept parameters used to control certain properties of the training process and of the resulting machine learning model. These are referred to herein as training parameters. If no training parameters are specified, the training manager can utilize default values that are known to work well for a large range of machine learning tasks. Examples of training parameters for which values can be specified include the maximum model
- 10 size, maximum number of passes over training data, shuffle type, regularization type, learning rate, and regularization amount. Default settings may be specified, with options to adjust the values to fine-tune performance.

- [0044] The maximum model size is the total size, in units of bytes, of patterns that are created during the training of model. A model may be created of a specified size by default, such as a
- 15 model of 100 MB. If the training manager is unable to determine enough patterns to fill the model size, a smaller model may be created. If the training manager finds more patterns than will fit into the specified size, a maximum cut-off may be enforced by trimming the patterns that least affect the quality of the learned model. Choosing the model size provides for control of the trade-off between the predictive quality of a model and the cost of use. Smaller models can cause
- 20 the training manager to remove many patterns to fit within the maximum size limit, affecting the quality of predictions. Larger models, on the other hand, may cost more to query for real-time predictions. Larger input data sets do not necessarily result in larger models because models store patterns, not input data; if the patterns are few and simple, the resulting model will be small. Input data that has a large number of raw attributes (input columns) or derived features
- 25 (outputs of the data transformations) will likely have more patterns found and stored during the training process.

[0045] In some embodiments, the training manager can make multiple passes or iterations over the training data to discover patterns. There may be a default number of passes, such as ten passes, while in some embodiments up to a maximum number of passes may be set, such as up to

one hundred passes. In some embodiments there may be no maximum set, or there may be a convergence or other criterion set which will trigger an end to the training process. In some embodiments the training manager can monitor the quality of patterns (i.e., the model convergence) during training, and can automatically stop the training when there are no more data points or patterns to discover. Data sets with only a few observations may require more passes over the data to obtain higher model quality. Larger data sets may contain many similar data points, which can reduce the need for a large number of passes. The potential impact of choosing more data passes over the data is that the model training can takes longer and cost more in terms of resources and system utilization.

[0046] In some embodiments the training data is shuffled before training, or between passes of the training. The shuffling in many embodiments is a random or pseudo-random shuffling to generate a truly random ordering, although there may be some constraints in place to ensure that there is no grouping of certain types of data, or the shuffled data may be reshuffled if such grouping exists, etc. Shuffling changes the order or arrangement in which the data is utilized for training so that the training algorithm does not encounter groupings of similar types of data, or a single type of data for too many observations in succession. For example, a model might be trained to predict a product type, where the training data includes movie, toy, and video game product types. The data might be sorted by product type before uploading. The algorithm can then process the data alphabetically by product type, seeing only data for a type such as movies first. The model will begin to learn patterns for movies. The model will then encounter only data for a different product type, such as toys, and will try to adjust the model to fit the toy product type, which can degrade the patterns that fit movies. This sudden switch from movie to toy type can produce a model that does not learn how to predict product types accurately. Shuffling can be performed in some embodiments before the training data set is split into training and evaluation subsets, such that a relatively even distribution of data types is utilized for both stages. In some embodiments the training manager can automatically shuffle the data using, for example, a pseudo-random shuffling technique.

[0047] When creating a machine learning model, the training manager in some embodiments can enable a user to specify settings or apply custom options. For example, a user may specify

one or more evaluation settings, indicating a portion of the input data to be reserved for evaluating the predictive quality of the machine learning model. The user may specify a recipe that indicates which attributes and attribute transformations are available for model training. The user may also specify various training parameters that control certain properties of the training process and of the resulting model.

[0048] Once the training manager has determined that training of the model is complete, such as by using at least one end criterion discussed herein, the trained model 508 can be provided for use by a classifier 514 in classifying unclassified data 512. In many embodiments, however, the trained model 508 will first be passed to an evaluator 510, which may include an application or process executing on at least one computing resource for evaluating the quality (or another such aspect) of the trained model. The model is evaluated to determine whether the model will provide at least a minimum acceptable or threshold level of performance in predicting the target on new and future data. Since future data instances will often have unknown target values, it can be desirable to check an accuracy metric of the machine learning on data for which the target answer is known, and use this assessment as a proxy for predictive accuracy on future data.

[0049] In some embodiments, a model is evaluated using a subset of the classified data 502 that was provided for training. The subset can be determined using a shuffle and split approach as discussed above. This evaluation data subset will be labeled with the target, and thus can act as a source of ground truth for evaluation. Evaluating the predictive accuracy of a machine learning model with the same data that was used for training is not useful, as positive evaluations might be generated for models that remember the training data instead of generalizing from it. Once training has completed, the evaluation data subset is processed using the trained model 508 and the evaluator 510 can determine the accuracy of the model by comparing the ground truth data against the corresponding output (or predictions/observations) of the model. The evaluator 510 in some embodiments can provide a summary or performance metric indicating how well the predicted and true values match. If the trained model does not satisfy at least a minimum performance criterion, or other such accuracy threshold, then the training manager 504 can be instructed to perform further training, or in some instances try training a new or different model,

among other such options. If the trained model 508 satisfies the relevant criteria, then the trained model can be provided for use by the classifier 514.

[0050] When creating and training a machine learning model, it can be desirable in at least some embodiments to specify model settings or training parameters that will result in a model
5 capable of making the most accurate predictions. Example parameters include the number of passes to be performed (forward and/or backward), regularization, model size, and shuffle type. As mentioned, however, selecting model parameter settings that produce the best predictive performance on the evaluation data might result in an overfitting of the model. Overfitting occurs when a model has memorized patterns that occur in the training and evaluation data sources, but
10 has failed to generalize the patterns in the data. Overfitting often occurs when the training data includes all of the data used in the evaluation. A model that has been over fit may perform well during evaluation, but may fail to make accurate predictions on new or otherwise unclassified data. To avoid selecting an over fitted model as the best model, the training manager can reserve additional data to validate the performance of the model. For example, the training data set might
15 be divided into 60 percent for training, and 40 percent for evaluation or validation, which may be divided into two or more stages. After selecting the model parameters that work well for the evaluation data, leading to convergence on a subset of the validation data, such as half the validation data, a second validation may be executed with a remainder of the validation data to ensure the performance of the model. If the model meets expectations on the validation data,
20 then the model is not overfitting the data. Alternatively, a test set or held-out set may be used for testing the parameters. Using a second validation or testing step helps to select appropriate model parameters to prevent overfitting. However, holding out more data from the training process for validation makes less data available for training. This may be problematic with smaller data sets as there may not be sufficient data available for training. One approach in such
25 a situation is to perform cross-validation as discussed elsewhere herein.

[0051] There are many metrics or insights that can be used to review and evaluate the predictive accuracy of a given model. One example evaluation outcome contains a prediction accuracy metric to report on the overall success of the model, as well as visualizations to help explore the accuracy of the model beyond the prediction accuracy metric. The outcome can also

provide an ability to review the impact of setting a score threshold, such as for binary classification, and can generate alerts on criteria to check the validity of the evaluation. The choice of the metric and visualization can depend at least in part upon the type of model being evaluated.

5 [0052] Once trained and evaluated satisfactorily, the trained machine learning model can be used to build or support a machine learning application. In one embodiment building a machine learning application is an iterative process that involves a sequence of steps. The core machine learning problem(s) can be framed in terms of what is observed and what answer the model is to predict. Data can then be collected, cleaned, and prepared to make the data suitable for
10 consumption by machine learning model training algorithms. The data can be visualized and analyzed to run sanity checks to validate the quality of the data and to understand the data. It might be the case that the raw data (e.g., input variables) and answer (e.g., the target) are not represented in a way that can be used to train a highly predictive model. Therefore, it may be desirable to construct more predictive input representations or features from the raw variables.
15 The resulting features can be fed to the learning algorithm to build models and evaluate the quality of the models on data that was held out from model building. The model can then be used to generate predictions of the target answer for new data instances.

[0053] In the example system 500 of FIG. 5, the trained model 510 after evaluation is provided, or made available, to a classifier 514 that is able to use the trained model to process
20 unclassified data. This may include, for example, data received from users or third parties that are not classified, such as query images that are looking for information about what is represented in those images. The unclassified data can be processed by the classifier using the trained model, and the results 516 (i.e., the classifications or predictions) that are produced can be sent back to the respective sources or otherwise processed or stored. In some embodiments,
25 and where such usage is permitted, the now classified data instances can be stored to the classified data repository, which can be used for further training of the trained model 508 by the training manager. In some embodiments the model will be continually trained as new data is available, but in other embodiments the models will be retrained periodically, such as once a day or week, depending upon factors such as the size of the data set or complexity of the model.

[0054] The classifier can include appropriate hardware and software for processing the unclassified data using the trained model. In some instances the classifier will include one or more computer servers each having one or more graphics processing units (GPUs) that are able to process the data. The configuration and design of GPUs can make them more desirable to use in processing machine learning data than CPUs or other such components. The trained model in some embodiments can be loaded into GPU memory and a received data instance provided to the GPU for processing. GPUs can have a much larger number of cores than CPUs, and the GPU cores can also be much less complex. Accordingly, a given GPU may be able to process thousands of data instances concurrently via different hardware threads. A GPU can also be configured to maximize floating point throughput, which can provide significant additional processing advantages for a large data set.

[0055] Even when using GPUs, accelerators, and other such hardware to accelerate tasks such as the training of a model or classification of data using such a model, such tasks can still require significant time, resource allocation, and cost. For example, if the machine learning model is to be trained using 100 passes, and the data set includes 1,000,000 data instances to be used for training, then all million instances would need to be processed for each pass. Different portions of the architecture can also be supported by different types of devices. For example, training may be performed using a set of servers at a logically centralized location, as may be offered as a service, while classification of raw data may be performed by such a service or on a client device, among other such options. These devices may also be owned, operated, or controlled by the same entity or multiple entities in various embodiments.

[0056] FIG. 6 illustrates an example neural network 600, or other statistical model, that can be utilized in accordance with various embodiments. In this example the statistical model is an artificial neural network (ANN) that includes a multiple layers of nodes, including an input layer 602, an output layer 606, and multiple layers 604 of intermediate nodes, often referred to as “hidden” layers, as the internal layers and nodes are typically not visible or accessible in conventional neural networks. As discussed elsewhere herein, there can be additional types of statistical models used as well, as well as other types of neural networks including other numbers of selections of nodes and layers, among other such options. In this network, all nodes of a given

layer are interconnected to all nodes of an adjacent layer. As illustrated, the nodes of an intermediate layer will then each be connected to nodes of two adjacent layers. The nodes are also referred to as neurons or connected units in some models, and connections between nodes are referred to as edges. Each node can perform a function for the inputs received, such as by using a specified function. Nodes and edges can obtain different weightings during training, and individual layers of nodes can perform specific types of transformations on the received input, where those transformations can also be learned or adjusted during training. The learning can be supervised or unsupervised learning, as may depend at least in part upon the type of information contained in the training data set. Various types of neural networks can be utilized, as may include a convolutional neural network (CNN) that includes a number of convolutional layers and a set of pooling layers, and have proven to be beneficial for applications such as image recognition. CNNs can also be easier to train than other networks due to a relatively small number of parameters to be determined.

[0057] In some embodiments, such a complex machine learning model can be trained using various tuning parameters. Choosing the parameters, fitting the model, and evaluating the model are parts of the model tuning process, often referred to as hyperparameter optimization. Such tuning can involve introspecting the underlying model or data in at least some embodiments. In a training or production setting, a robust workflow can be important to avoid overfitting of the hyperparameters as discussed elsewhere herein. Cross-validation and adding Gaussian noise to the training dataset are techniques that can be useful for avoiding overfitting to any one dataset. For hyperparameter optimization it may be desirable in some embodiments to keep the training and validation sets fixed. In some embodiments, hyperparameters can be tuned in certain categories, as may include data preprocessing (in other words, translating words to vectors), CNN architecture definition (for example, filter sizes, number of filters), stochastic gradient descent parameters (for example, learning rate), and regularization (for example, dropout probability), among other such options.

[0058] In an example pre-processing step, instances of a dataset can be embedded into a lower dimensional space of a certain size. The size of this space is a parameter to be tuned. The architecture of the CNN contains many tunable parameters. A parameter for filter sizes can

represent an interpretation of the information that corresponds to the size of a instance that will be analyzed. In computational linguistics, this is known as the n-gram size. An example CNN uses three different filter sizes, which represent potentially different n-gram sizes. The number of filters per filter size can correspond to the depth of the filter. Each filter attempts to learn something different from the structure of the instance, such as the sentence structure for textual data. In the convolutional layer, the activation function can be a rectified linear unit and the pooling type set as max pooling. The results can then be concatenated into a single dimensional vector, and the last layer is fully connected onto a two-dimensional output. This corresponds to the binary classification to which an optimization function can be applied. One such function is an implementation of a Root Mean Square (RMS) propagation method of gradient descent, where example hyperparameters can include learning rate, batch size, maximum gradient normal, and epochs. With neural networks, regularization can be an extremely important consideration. As mentioned, in some embodiments the input data may be relatively sparse. A main hyperparameter in such a situation can be the dropout at the penultimate layer, which represents a proportion of the nodes that will not “fire” at each training cycle. An example training process can suggest different hyperparameter configurations based on feedback for the performance of previous configurations. The model can be trained with a proposed configuration, evaluated on a designated validation set, and the performance reporting. This process can be repeated to, for example, trade off exploration (learning more about different configurations) and exploitation (leveraging previous knowledge to achieve better results).

[0059] As training CNNs can be parallelized and GPU-enabled computing resources can be utilized, multiple optimization strategies can be attempted for different scenarios. A complex scenario allows tuning the model architecture and the preprocessing and stochastic gradient descent parameters. This expands the model configuration space. In a basic scenario, only the preprocessing and stochastic gradient descent parameters are tuned. There can be a greater number of configuration parameters in the complex scenario than in the basic scenario. The tuning in a joint space can be performed using a linear or exponential number of steps, iteration through the optimization loop for the models. The cost for such a tuning process can be significantly less than for tuning processes such as random search and grid search, without any significant performance loss.

[0060] Some embodiments can utilize backpropagation to calculate a gradient used for determining the weights for the neural network. Backpropagation is a form of differentiation, and can be used by a gradient descent optimization algorithm to adjust the weights applied to the various nodes or neurons as discussed above. The weights can be determined in some
5 embodiments using the gradient of the relevant loss function. Backpropagation can utilize the derivative of the loss function with respect to the output generated by the statistical model. As mentioned, the various nodes can have associated activation functions that define the output of the respective nodes. Various activation functions can be used as appropriate, as may include radial basis functions (RBFs) and sigmoids, which can be utilized by various support vector
10 machines (SVMs) for transformation of the data. The activation function of an intermediate layer of nodes is referred to herein as the inner product kernel. These functions can include, for example, identity functions, step functions, sigmoidal functions, ramp functions, and the like. Activation functions can also be linear or non-linear, among other such options.

[0061] FIG. 7 illustrates a set of basic components of a computing device 700 that can be
15 utilized to implement aspects of the various embodiments. In this example, the device includes at least one processor 702 for executing instructions that can be stored in a memory device or element 704. As would be apparent to one of ordinary skill in the art, the device can include many types of memory, data storage or computer-readable media, such as a first data storage for program instructions for execution by the processor 702, the same or separate storage can be
20 used for images or data, a removable memory can be available for sharing information with other devices, and any number of communication approaches can be available for sharing with other devices. The device typically will include some type of display element 706, such as a touch screen, organic light emitting diode (OLED) or liquid crystal display (LCD), although devices such as portable media players might convey information via other means, such as through audio
25 speakers. As discussed, the device in many embodiments will include at least communication component 708 and/or networking components 710, such as may support wired or wireless communications over at least one network, such as the Internet, a local area network (LAN), Bluetooth®, or a cellular network, among other such options. The components can enable the device to communicate with remote systems or services. The device can also include at least one
30 additional input device 712 able to receive conventional input from a user. This conventional

input can include, for example, a push button, touch pad, touch screen, wheel, joystick, keyboard, mouse, trackball, keypad or any other such device or element whereby a user can input a command to the device. These I/O devices could even be connected by a wireless infrared or Bluetooth or other link as well in some embodiments. In some embodiments, however, such a device might not include any buttons at all and might be controlled only through a combination of visual and audio commands such that a user can control the device without having to be in contact with the device.

[0062] The various embodiments can be implemented in a wide variety of operating environments, which in some cases can include one or more user computers or computing devices which can be used to operate any of a number of applications. User or client devices can include any of a number of general purpose personal computers, such as desktop or laptop computers running a standard operating system, as well as cellular, wireless and handheld devices running mobile software and capable of supporting a number of networking and messaging protocols. Such a system can also include a number of workstations running any of a variety of commercially-available operating systems and other known applications for purposes such as development and database management. These devices can also include other electronic devices, such as dummy terminals, thin-clients, gaming systems and other devices capable of communicating via a network.

[0063] Most embodiments utilize at least one network that would be familiar to those skilled in the art for supporting communications using any of a variety of commercially-available protocols, such as TCP/IP or FTP. The network can be, for example, a local area network, a wide-area network, a virtual private network, the Internet, an intranet, an extranet, a public switched telephone network, an infrared network, a wireless network and any combination thereof. In embodiments utilizing a Web server, the Web server can run any of a variety of server or mid-tier applications, including HTTP servers, FTP servers, CGI servers, data servers, Java servers and business application servers. The server(s) may also be capable of executing programs or scripts in response requests from user devices, such as by executing one or more Web applications that may be implemented as one or more scripts or programs written in any programming language, such as Java[®], C, C# or C++ or any scripting language, such as Python,

as well as combinations thereof. The server(s) may also include database servers, including without limitation those commercially available from Oracle[®], Microsoft[®], Sybase[®] and IBM[®].

[0064] The environment can include a variety of data stores and other memory and storage media as discussed above. These can reside in a variety of locations, such as on a storage medium local to (and/or resident in) one or more of the computers or remote from any or all of the computers across the network. In a particular set of embodiments, the information may reside in a storage-area network (SAN) familiar to those skilled in the art. Similarly, any necessary files for performing the functions attributed to the computers, servers or other network devices may be stored locally and/or remotely, as appropriate. Where a system includes computerized devices, each such device can include hardware elements that may be electrically coupled via a bus, the elements including, for example, at least one central processing unit (CPU), at least one input device (e.g., a mouse, keyboard, controller, touch-sensitive display element or keypad) and at least one output device (e.g., a display device, printer or speaker). Such a system may also include one or more storage devices, such as disk drives, optical storage devices and solid-state storage devices such as random access memory (RAM) or read-only memory (ROM), as well as removable media devices, memory cards, flash cards, etc.

[0065] Such devices can also include a computer-readable storage media reader, a communications device (e.g., a modem, a network card (wireless or wired), an infrared communication device) and working memory as described above. The computer-readable storage media reader can be connected with, or configured to receive, a computer-readable storage medium representing remote, local, fixed and/or removable storage devices as well as storage media for temporarily and/or more permanently containing, storing, transmitting and retrieving computer-readable information. The system and various devices also typically will include a number of software applications, modules, services or other elements located within at least one working memory device, including an operating system and application programs such as a client application or Web browser. It should be appreciated that alternate embodiments may have numerous variations from that described above. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, software (including

portable software, such as applets) or both. Further, connection to other computing devices such as network input/output devices may be employed.

[0066] Storage media and other non-transitory computer readable media for containing code, or portions of code, can include any appropriate media known or used in the art, such as but not
5 limited to volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data, including RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices or
10 any other medium which can be used to store the desired information and which can be accessed by a system device. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments

[0067] The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may
15 be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims.

CLAIMS

WHAT IS CLAIMED IS:

- 1 1. A computer-implemented method, comprising:
2 receiving an input image including a representation of a target object of an
3 unknown type and having a first pose;
4 receiving indication of a source image including a representation of a source
5 object of a source class, the source object of the source image having a second pose;
6 providing the input image and the source image to a neural network, the neural
7 network trained using images of a plurality of object classes, including the source class, to
8 extract a class-invariant representation of the second pose and a class-specific latent
9 representation of an appearance of the target object; and
10 inferring an output image representing the target object in the second pose.
- 1 2. The computer-implemented method of claim 1, wherein the neural
2 network is a generative adversarial network (GAN) including a conditional image generator and
3 a multi-task adversarial discriminator.
- 1 3. The computer-implemented method of claim 2, further comprising:
2 extracting the class-invariant representation of the second pose using a content
3 encoder of the conditional image generator; and
4 extracting the class-specific latent representation of the appearance of the target
5 object using a class encoder of the conditional image generator.
- 1 4. The computer-implemented method of claim 2, further comprising:
2 normalizing, by a normalization layer of the multi-task adversarial discriminator,
3 layer activations to zero mean and unit variance distribution; and
4 de-normalizing the normalized layer activations using an affine transformation.
- 1 5. The computer-implemented method of claim 1, further comprising:

2 performing adversarial training of the neural network to translate poses between
3 the plurality of object classes.

1 6. A computer-implemented method, comprising:
2 receiving a digital representation of an image including an object in a first pose;
3 and
4 inferring, using a neural network, an output image representing the object in a
5 second pose, wherein the neural network includes a conditional generator and an adversarial
6 discriminator for extracting a class-invariant representation of the second pose and a class-
7 specific representation of an appearance of the object.

1 7. The computer-implemented method of claim 6, wherein the neural
2 network was not trained using images of an object class of the object.

1 8. The computer-implemented method of claim 6, further comprising:
2 performing unsupervised training of the neural network using a training data set
3 containing a plurality of images of multiple object classes, wherein the neural network learns to
4 translate poses between the object classes.

1 9. The computer-implemented method of claim 6, wherein the neural
2 network is a generative adversarial network (GAN).

1 10. The computer-implemented method of claim 6, further comprising:
2 extracting a class-invariant representation of the second pose using a content
3 encoder of the conditional generator; and
4 extracting the class-specific latent representation of an appearance of the target
5 object using a class encoder of the conditional generator.

1 11. The computer-implemented method of claim 6, further comprising:
2 normalizing, by a normalization layer of the adversarial discriminator, layer
3 activations to zero mean and unit variance distribution; and
4 de-normalizing the normalized layer activations using an affine transformation.

1 12. The computer-implemented method of claim 11, further comprising:
2 feeding class latent code to an adaptive instance-norm (AdaIN) decoder
3 configured to decode the class latent code to a set of mean and variance vectors functioning as new
4 means and variances for the respective channels in a respective AdaIN residual block of an image
5 decoder of the neural network.

1 13. The computer-implemented method of claim 6, wherein the conditional
2 generator includes a content encoder comprising a plurality of convolutional layers followed by a
3 plurality of residual blocks, and
4 wherein the conditional generator includes a class encoder comprising a plurality
5 of convolutional layers followed by an average pooling layer.

1 14. The computer-implemented method of claim 13, further comprising:
2 causing the average pooling layer to average activations across spatial dimensions
3 before averaging activations across a set of images.

4 15. A system, comprising:
5 at least one processor; and
6 memory including instructions that, when executed by the at least one processor,
7 cause the system to:
8 receive a digital representation of an image including an object in a first
9 pose; and
10 infer, using a neural network, an output image representing the object in a
11 second pose, wherein the neural network includes a conditional generator and an
12 adversarial discriminator for extracting a class-invariant representation of the second pose
13 and a class-specific representation of an appearance of the object.

1 16. The system of claim 15, wherein the neural network was not trained using
2 images of an object class of the object.

1 17. The system of claim 15, wherein the instructions when executed further
2 cause the system to:

3 perform unsupervised training of the neural network using a training data set
4 containing a plurality of images of different object classes, wherein the neural network learns to
5 translate poses between the different object classes.

1 18. The system of claim 15, wherein the neural network is a generative
2 adversarial network (GAN).

1 19. The system of claim 15, wherein the instructions when executed further
2 cause the system to:

3 extract a class-invariant representation of the second pose using a content encoder
4 of the conditional generator; and

5 extract the class-specific latent representation of an appearance of the target
6 object using a class encoder of the conditional generator.

1 20. The system of claim 15, wherein the instructions when executed further
2 cause the system to:

3 normalize, by a normalization layer of the adversarial discriminator, layer
4 activations to zero mean and unit variance distribution; and

5 de-normalize the normalized layer activations using an affine transformation.




FIG. 1

200




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2020/015262

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06K9/62 G06N3/08 G06T11/00
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 G06T G06N

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	TRAN LUAN ET AL: "Disentangled Representation Learning GAN for Pose-Invariant Face Recognition", IEEE COMPUTER SOCIETY CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION. PROCEEDINGS, IEEE COMPUTER SOCIETY, US, 21 July 2017 (2017-07-21), pages 1283-1292, XP033249467, ISSN: 1063-6919, DOI: 10.1109/CVPR.2017.141 [retrieved on 2017-11-06]	6-11, 13-20
Y	Abstract section 1, par. 3, first sentence section 1, par. 3, lines 9-11 section 1, par. 4, lines 2-6 and 10-12 section 1, par. 5, lines 12-13 section 1, par. 6, lines 4-6 section 3.3, lines 2-3 after equation (6)] section 3.4, last sentence -/--	1-5,12



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

6 May 2020

Date of mailing of the international search report

14/05/2020

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

Mukasa, Oliver

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2020/015262

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
	section 4.1, par. 1; page 1; figure 1 ----- X US 2018/314716 A1 (KIM TAEK SOO [KR] ET AL) 1 November 2018 (2018-11-01) Y paragraphs [0026], [0036] - [0040], [0059] - [0062]; figures 2A, 2B, 3, 4, 5A, 5B, 5C ----- X CN 107 292 813 A (UNIV ZHEJIANG) 24 October 2017 (2017-10-24) the whole document ----- Y Xun Huang ET AL: "Workshop track -ICLR 2017 ARBITRARY STYLE TRANSFER IN REAL-TIME WITH ADAPTIVE INSTANCE NORMALIZATION", 31 December 2017 (2017-12-31), XP055691247, Retrieved from the Internet: URL: https://openreview.net/pdf?id=B1fUVMzKg [retrieved on 2020-05-04] section 1, par. 3, lines 6-10 section 3, par. 2, lines 1-6 section 5.1, par. 1, lines 4-7 -----	6,15 1-5 6,15 12

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2020/015262

Patent document cited in search report	Publication date	Patent family member(s)	Publication date	
US 2018314716	A1	01-11-2018	KR 20180120478 A	06-11-2018
			US 2018314716 A1	01-11-2018
			US 2019205334 A1	04-07-2019

CN 107292813	A	24-10-2017	NONE	
