

Method and System for Training Generative Models

Technical Field

[0001] The present disclosure generally relates to a method and system for training generative models.

Background Art

[0002] Generative models (also known as generators) refer to a class of neural networks that can generate new data instances. In contrast, conditional models (also known as discriminative models or discriminators) refer to another class of neural networks that can differentiate between different data instances. For example, a generative model can generate new photos of fruits, while a conditional model can distinguish between an orange and an apple.

[0003] A generative adversarial network (GAN) is a type of generative model that can create new data instances resembling source data instances. A GAN typically includes a generative model paired with a conditional model. The generative model is trained to produce a target output while the conditional model is trained to distinguish the target output from real data. In other words, the generative model is trained to generate plausible target outputs that can trick the conditional model into misclassifying the target outputs as real data, while the conditional model is trained to accurately differentiate between target outputs and real data. For example, a generative model in a GAN can be trained with photographs of human faces to create synthetic images that look like photographs of real human faces even though the faces do not belong to an actual person, while a conditional model in the GAN can be trained to distinguish the synthetic images from photographs of real persons.

[0004] Generative models such as GANs can also be trained to perform image-to-image translations. An image-to-image translation generally includes generating a synthetic version of a source image with a specific modification. For example, a generative model can be trained to translate a black-and-white image to a colour image. However, large quantities of labelled training data sets (i.e. training examples) are typically required to train the generative model for image-to-image translation, to ensure quality of results, accuracy of the model, as well as to remove and/or pre-empt biases in the model, so that the trained neural network model can work sufficiently well in most situations. This is particularly important if the generative model (or data instances generated by the generative model) are used in applications with widespread uses (e.g. in face recognition systems). Nonetheless, it can be

time-consuming and expensive to collate labelled training data sets typically required for training the neural network model (e.g. a generative model or a face recognition model). In some cases, labelled data sets may be practically impossible to collate, as training data may not exist, or is insufficient in quantity to adequately train the neural network model. For example, while there may be ample labelled training data for training the neural network model in one domain, labelled training data in another domain may be absent or insufficient in quantity to adequately train the neural network model.

Summary

[0005] In an embodiment, there is provided a method for training generative models. The method includes receiving, by a processing device, a training facial image in a first complexion, generating, using the processing device, a synthetic facial image in a second complexion using a first generative model, based on the training facial image in the first complexion, and generating, using the processing device, a synthetic facial image in the first complexion using a second generative model, based on the synthetic facial image in the second complexion. The method also includes calculating, using the processing device, an identity classification loss value using an identity loss minimisation model, based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in a face recognition model to establish an identity of a subject and updating, using the processing device, the first generative model and the second generative model based on the identity classification loss value.

[0006] In an embodiment, there is provided a system for training generative models. The system includes a processing device configured to receive a training facial image in a first complexion, generate a synthetic facial image in a second complexion using a first generative model, based on the training facial image in the first complexion and generate a synthetic facial image in the first complexion using a second generative model, based on the synthetic facial image in the second complexion. The processor is also configured to calculate an identity classification loss value using an identity loss minimisation model, based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in a face recognition model to establish an identity of a subject and update the first generative model and the second generative model based on the identity classification loss value.

[0007] In an embodiment, there is provided a method for training a face recognition model. The method includes receiving, by a processing device, an identified facial image in a source complexion, generating, using the processing device, an identified facial image in a target complexion based on the received facial image using a generative model, the generative

model trained using a cycle generative adversarial network, and training, using the processing device, the face recognition model with the identified facial images in the source and the target complexions. The cycle generative adversarial network includes a first generative model configured to generate, based on a training facial image in a first complexion, a synthetic facial image in a second complexion, a second generative model configured to generate, based on the synthetic facial image in the second complexion, a synthetic facial image in the first complexion and an identity loss minimisation model configured to calculate an identity classification loss value based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in the face recognition model to establish an identity of a subject. The first and second generative models are trained using the cycle generative adversarial network to minimise the identity classification loss, and the generative model is the first generative model.

[0008] In an embodiment, there is provided a system comprising a processing device. The processing device is configured to receive an identified facial image in a source complexion, generate an identified facial image in a target complexion based on the received facial image using a generative model, the generative model trained using a cycle generative adversarial network, and train a face recognition model with the identified facial images in the source and the target complexions. The cycle generative adversarial network includes a first generative model configured to generate, based on a training facial image in a first complexion, a synthetic facial image in a second complexion, a second generative model configured to generate, based on the synthetic facial image in the second complexion, a synthetic facial image in the first complexion and an identity loss minimisation model configured to calculate an identity classification loss value based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in the face recognition model to establish an identity of a subject. The first and second generative models are trained using the cycle generative adversarial network to minimise the identity classification loss, and the generative model is the first generative model.

[0009] Details of one or more embodiments of the subject matter of this specification are set forth in the accompanying drawings and the description below. Other desirable features and characteristics will become apparent from the subsequent description and the appended claims, taken in conjunction with the accompanying drawings and the background of the disclosure.

Brief Description of Drawings

[0010] Embodiments will be better understood and readily apparent to one of ordinary skill in the art from the following written description, by way of example only, and in conjunction with the drawings, in which:

Fig.1

[0011] Fig. 1 shows a schematic diagram illustrating an example of a system for training generative models, and exchange of information in the system for training generative models, in accordance with embodiments of the disclosure.

Fig.2

[0012] Fig. 2 shows a flowchart illustrating an example of a method for training generative models, in accordance with embodiments of the disclosure.

Fig.3

[0013] Fig. 3 shows a flowchart illustrating an example of a method for training a face recognition model, in accordance with embodiments of the disclosure.

Fig.4

[0014] Fig. 4 shows a schematic diagram illustrating an example of a system for training generative models and for training a face recognition model, in accordance with embodiments of the disclosure.

Fig.5

[0015] Fig. 5 shows a schematic diagram of an example of a computing device used to realise the systems of Figs. 4 and 5.

[0016] Skilled artisans will appreciate that elements in the figures are illustrated for simplicity and clarity and have not necessarily been depicted to scale. For example, the dimensions of some of the elements in the illustrations, block diagrams or flowcharts may be exaggerated in respect to other elements to help to improve understanding of the present embodiments.

Description of Embodiments

[0017] Embodiments of the present disclosure will be described, by way of example only, with reference to the drawings. Like reference numerals and characters in the drawings refer to like elements or equivalents.

[0018] Some portions of the description which follows are explicitly or implicitly presented in terms of algorithms and functional or symbolic representations of operations on data within a computer memory. These algorithmic descriptions and functional or symbolic representations are the means used by those skilled in the data processing arts to convey most effectively the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities, such as electrical, magnetic or optical signals capable of being stored, transferred, combined, compared, and otherwise manipulated.

[0019] Unless specifically stated otherwise, and as apparent from the following, it will be appreciated that throughout the present specification, discussions utilizing terms such as “associating”, “calculating”, “comparing”, “determining”, “extracting”, “forwarding”, “generating”, “identifying”, “including”, “inserting”, “modifying”, “receiving”, “recording”, “replacing”, “scanning”, “transmitting”, “updating” or the like, refer to the action and processes of a computer system, or similar electronic device, that manipulates and transforms data represented as physical quantities within the computer system into other data similarly represented as physical quantities within the computer system or other information storage, transmission or display devices.

[0020] The present specification also discloses apparatus for performing the operations of the methods. Such apparatus may be specially constructed for the required purposes or may include a computer or other computing device selectively activated or reconfigured by a computer program stored therein. The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various machines may be used with programs in accordance with the teachings herein. Alternatively, the construction of more specialized apparatus to perform the required method steps may be appropriate. The structure of a computer will appear from the description below.

[0021] In addition, the present specification also implicitly discloses a computer program, in that it would be apparent to the person skilled in the art that the individual steps of the method described herein may be put into effect by computer code. The computer program is not intended to be limited to any particular programming language and implementation

thereof. It will be appreciated that a variety of programming languages and coding thereof may be used to implement the teachings of the disclosure contained herein. Moreover, the computer program is not intended to be limited to any particular control flow. There are many other variants of the computer program, which can use different control flows without departing from the spirit or scope of the disclosure.

[0022] Furthermore, one or more of the steps of the computer program may be performed in parallel rather than sequentially. Such a computer program may be stored on any computer readable medium. The computer readable medium may include storage devices such as magnetic or optical disks, memory chips, or other storage devices suitable for interfacing with a computer. The computer readable medium may also include a hard-wired medium such as exemplified in the Internet system, or wireless medium such as exemplified in the mobile telephone system. The computer program when loaded and executed on a computer effectively results in an apparatus that implements the steps of the preferred method.

[0023] In embodiments of the present disclosure, use of the term 'server' may mean a single computing device or at least a computer network of interconnected computing devices which operate together to perform a particular function. In other words, the server may be contained within a single hardware unit or be distributed among several or many different hardware units.

[0024] The term "configured to" is used in the specification in connection with systems, apparatus, and computer program components. For a system of one or more computers to be configured to perform particular operations or actions means that the system has installed on it software, firmware, hardware, or a combination of them that in operation cause the system to perform the operations or actions. For one or more computer programs to be configured to perform particular operations or actions means that the one or more programs include instructions that, when executed by data processing apparatus, cause the apparatus to perform the operations or actions. For special-purpose logic circuitry to be configured to perform particular operations or actions means that the circuitry has electronic logic that performs the operations or actions.

[0025] Embodiments of the disclosure seek to provide methods and systems that can generate labelled data in a target domain, using labelled data in a source domain and unlabelled data in the target domain. For example, embodiments of the disclosure can generate identified facial images in a target complexion, using identified facial images in another complexion and unlabelled data in a target complexion. In other words, embodiments of the disclosure can translate the complexion of a facial image from one to

another, while preserving identifying features of the facial image (i.e. preserving identity of the person in the image).

[0026] The present disclosure includes various techniques that can improve and enhance generation of synthetic labelled data (e.g. identified synthetic facial images), by generating synthetic labelled data in a manner that provides technical advantages over conventional and manual (e.g., human-directed, -assisted or -guided) approaches. Technical advantages can include reduced processing times required to generate synthetic labelled data, as well as improvements in processing, memory, and network resources usage required to generate and/or process such data. Moreover, the synthetic labelled data (e.g. identified synthetic facial images) can be used in training data sets to improve accuracy of the neural network models, and to remove biases in the models (e.g. improve accuracy and to reduce biases in facial recognition models). Synthetic labelled data can also reduce computational resources and time typically required to train neural network models.

[0027] Fig. 1 shows a schematic diagram illustrating an example of a system 100 for training generative models, in accordance with embodiments of the disclosure. Fig. 1 also shows exchange of information in the system 100. The system 100 can be implemented with processing device 402 as shown in Fig. 4. The system 100 can include a first generative model G12, a second generative model G21, a first conditional model D1, a second conditional model D2 and an identity loss minimisation model IP1. The system 100 can also include a first cycle consistency loss minimisation model CCL1 and a second cycle consistency loss minimisation model CCL2. In some embodiments, the first and second generative models G12, G21 can be deep artificial neural networks with a plurality of hidden layers. The first generative model G12 can receive a training facial image F1 in a first complexion, and generate a synthetic facial image F12 in a second complexion based on the training facial image F1. The training facial image F1 can be a facial image of a person, and the identity of the person can be known. That is, the training facial image F1 can be an identified facial image of a person in a first complexion. The second generative model G21 can use the synthetic facial image F12 in the second complexion to generate a synthetic facial image F121 in the first complexion. That is, the first generative model G12 can take facial images F1 in a first domain (e.g. first complexion) as inputs and can output synthetic images F12 in a second domain (e.g. second complexion). The second generative model G21 can take synthetic facial images F12 in the second domain (e.g. second complexion), and output synthetic facial images F121 in the first domain (e.g. first complexion).

[0028] The identity loss minimisation model IP1 can be a classification model, and can use the training facial image F1 and the synthetic facial image F121 in the first complexion as

inputs to calculate an identity classification loss value, based on differences in facial features of the training facial image F1 and the synthetic facial image F121. The facial features can be used in a face recognition model to establish an identity of a person. In an embodiment, the identity loss minimisation model IP1 can be a multi-class classification model, and the identity classification loss value can be a multi-class classification loss value. For example, training facial images in the source domain (e.g. in the first complexion) can come with labels for identities (say, image 0 with label 0, image 1 with label 1, image 2 with label 2, ..., image N with label N) and each class in the identity loss minimisation model IP1 can be associated with respective labels of the facial images in the source domain. The identity loss minimisation model IP1 can distinguish each of the training facial images F1 (e.g. identify each training facial image separately), and calculate a separate identity classification loss value for each respective training facial images F1 and respective synthetic facial images F121 in the first domain. That is, a first identity classification loss value can be calculated for training image 0 with label 0 and associated output synthetic facial in the first domain, and a nth identity classification loss value can be calculated for training image n with label n and associated output synthetic facial in the first domain.

[0029] The first and second generative models G12, G21 can be updated using the processing device 402 based on the identity classification loss value. In some embodiments, the system 100 can update weights applied to nodes of hidden layers of the first and second generative models G12, G21 based on the identity classification loss value, so that the identity classification loss value is minimised. The process described above (e.g. generation of synthetic facial images in the first and second complexions, and calculation of the identity classification loss values) can be repeated, and the weights applied to nodes of the hidden layers of the first and second generative models G12, G21 can be iteratively adjusted so that the training facial images F1 and respective synthetic facial images F121 in the first domain (i.e. in the first complexion) match.

[0030] In some embodiments, the first conditional model D1 can be a deep artificial neural network model. The first conditional model D1 can classify if a facial image (e.g. the training facial image F1 and the synthetic facial image F121) in the first complexion is a synthetic facial image, and is trained to minimise an error rate associated with misclassification of facial images in the first complexion. In some embodiments, the first conditional model D1 can use binary cross-entropy loss as a loss function to determine a first conditional loss associated with the classification accuracy of facial images in the first complexion. The first conditional model D1 is trained to minimise the first conditional loss while the first and the second generative models G12, G21 are trained to maximise the first conditional loss. For example, the first conditional model D1 can be trained with the training facial images F1 that

are assigned with label 1 and synthetic facial images in the first complexion (e.g. facial images F121) assigned with label 0. The first conditional model D1 can take a facial image in the first complexion as input, and output a probability score ranging from 0 to 1, representative of a probability of whether the facial image in the first complexion is real (i.e. not synthetic). That is, the synthetic facial images F121 in the first complexion can be used as inputs to the first conditional model D1, and the first conditional model D1 can determine if the synthetic facial images F121 in the first complexion are synthetic (i.e. generated by a neural network). In some embodiments, the first conditional model D1 is trained to minimise an error rate associated with misclassification of facial images in the first complexion, while the first and the second generative models G12, G21 are trained to maximise the error rate associated with misclassification of the facial image in the first complexion. In other words, the first conditional model D1 is set up to compete against the first and the second generative models G12, G21. The first and the second generative models G12, G21 are iteratively trained to trick the first conditional model D1, while the first conditional model D1 is iteratively trained to better detect synthetic facial images in the first complexion.

[0031] The first cycle consistency loss minimisation model CCL1 can determine a first cycle consistency loss value based on differences in pixel values of the training facial image F1 and the synthetic facial image F121 in the first complexion. In an embodiment, the difference between the training facial image F1 and the synthetic facial image F121 in the first complexion can be calculated as L1 norm, which is the summed absolute difference in pixel values between the image F1 and the image F121. The first and second generative models G12, G21 are trained to minimise the first cycle consistency loss value determined by the first cycle consistency loss minimisation model CCL1. In some embodiments, the system 100 can update weights applied to nodes of the hidden layers of the first generative model G12 and the second generative model G21 based on the first cycle consistency loss value, so that the first cycle consistency loss value is minimised. The process described above (e.g. generation of synthetic facial images in the first and second complexions, and calculation of the first cycle consistency loss value) can be repeated and the weights in the first and second generative models G12, G21 can be iteratively adjusted to minimise the first cycle consistency loss value.

[0032] In some embodiments, the second generative model G21 can receive a training facial image F2 in a second complexion, and generate a synthetic facial image F21 in the first complexion based on the training facial image F2. The training facial image F2 can be a facial image of an unidentified person in the second complexion. The training facial image F2 can be obtained from public sources. The first generative model G12 can use the synthetic facial image F21 in the first complexion to generate a synthetic facial image F212 in the

second complexion. That is, the second generative model G21 can take facial images F2 in the second domain (e.g. second complexion) as inputs and can output synthetic images F21 in the first domain (e.g. first complexion). The first generative model G12 can take synthetic facial images F21 in the first domain (e.g. first complexion), and output synthetic facial images F212 in the second domain (e.g. second complexion).

[0033] In some embodiments, the second conditional model D2 can be a deep artificial neural network model. The second conditional model D2 can classify if a facial image (e.g. the training facial image F2 and the synthetic facial image F212) in the second complexion is a synthetic facial image, and is trained to minimise an error rate associated with misclassification of facial images in the second complexion. In some embodiments, the second conditional model D2 can use binary cross-entropy loss as a loss function to determine a second conditional loss associated with the classification accuracy of facial images in the second complexion. The second conditional model D2 is trained to minimise the second conditional loss while the first and the second generative models G12, G21 are trained to maximise the second conditional loss. For example, the second conditional model D2 can be trained with the training facial images F2 that are assigned with label 1 and synthetic facial images in the second complexion (e.g. facial images F212) assigned with label 0. The second conditional model D2 can take a facial image in the second complexion as input, and output a probability score ranging from 0 to 1, representative of a probability of whether the facial image in the second complexion is real (i.e. not synthetic). That is, the synthetic facial images F212 in the second complexion can be used as inputs to the second conditional model D2, and the second conditional model D2 can determine if the synthetic facial images F212 in the second complexion are synthetic (i.e. generated by a neural network). In some embodiments, the second conditional model D2 is trained to minimise an error rate associated with misclassification of facial images in the second complexion, while the first and the second generative models G12, G21 are trained to maximise the error rate associated with misclassification of the facial image in the second complexion. In other words, the second conditional model D2 is set up to compete against the first and the second generative models G12, G21. The first and the second generative models G12, G21 are iteratively trained to trick the second conditional model D2, while the second conditional model D2 is iteratively trained to better detect synthetic facial images in the second complexion.

[0034] The second cycle consistency loss minimisation model CCL2 can determine a second cycle consistency loss value based on differences in pixel values of the training facial image F2 and the synthetic facial image F212 in the second complexion. In an embodiment, the difference between the training facial image F2 and the synthetic facial

image F212 in the second complexion can be calculated as L1 norm, which is the summed absolute difference in pixel values between the image F2 and image F212. The first and second generative models G12, G21 are trained to minimise the second cycle consistency loss value determined by the second cycle consistency loss minimisation model CCL2. In some embodiments, the system 100 can update weights applied to nodes of the hidden layers of the first generative model G12 and the second generative model G21 based on the second cycle consistency loss value, so that the second cycle consistency loss value is minimised. The process described above (e.g. generation of synthetic facial images in the first and second complexions, and calculation of the second cycle consistency loss value) can be repeated and the weights in the first and second generative models G12, G21 iteratively adjusted to minimise the second cycle consistency loss value.

[0035] Thus, the system 100 can involve simultaneous training of the first generative model G12, the second generative model G21, the first conditional model D1, the second conditional model D2 and the identity loss minimisation model IP1. The first generative model G12 can use facial images F1 in the first domain (the first complexion) as input and generate images F12 in the second domain (the second complexion), while the second generative model G21 can use facial images F2 in the second domain as input and generate images F21 in the first domain. In embodiments, the output facial images F12 from the first generative model G12 can be used as input to the second generative model G21 to generate the output facial images F121. Conversely, the output facial images F21 from the second generative model G21 can be fed as input to the first generative model G12 to generate the output facial images F212. The first and second cycle consistency loss minimisation models CCL1, CCL2 can compare the output facial images F121, F212 with the respective source (original) facial images F1, F2, and the generative models G12, G21 are trained to minimise the first and second cycle consistency loss values determined by the cycle consistency loss minimisation models CCL1, CCL2. Thus, in some embodiments, the cycle consistency loss minimisation models CCL1, CCL2 can guide the generative models G12, G21 toward image translation. The first and second conditional models D1 and D2 can classify if the generated images F121, F212 are synthetic images. The first and second conditional models D1, D2 are trained to increase classification accuracy, while the first and second generative models G12, G21 are trained to compete with the first and second conditional models D1, D2 to reduce the classification accuracy of the first and second conditional models D1, D2.

[0036] In an implementation, Stochastic Gradient Decent (SGD) can be used to train the system 100, to minimise the first and second conditional losses, as well as to minimise the first and second cycle consistency losses. Training can be done iteratively with the first and

second generative models G12, G21, the conditional models D1, D2 and identity loss minimisation model IP1. The trained first and second generative models G12, G21 can be used to generate labelled data in a target domain, using labelled data in a source domain and unlabelled data in the target domain. For example, the trained first generative model G12 can generate identified facial images in the second complexion using identified facial images in the first complexion, while the trained second generative model G21 can generate identified facial images in the first complexion using identified facial images in the second complexion. In other words, the trained first and second generative models G12, G21 can preserve identity of a person in a facial image while translating the complexion of the person from one to another.

[0037] Fig. 2 shows a flowchart illustrating an example of a method 200 for training generative models, in accordance with embodiments of the disclosure. The method 200 can be implemented with the system 400 as shown in Fig. 4, which includes the processing device 402. The method 200 is described in conjunction with Fig. 1, which illustrates the exchange of information in the method 200 for training generative models. The method 200 broadly includes step 202 of receiving, by the processing device 402, a training facial image F1 in a first complexion. The method 200 also includes step 204 of generating, using the processing device 402, a synthetic facial image F12 in a second complexion using the first generative model G12, based on the identified training facial image F1 in the first complexion, and step 206 of generating, using the processing device 402, a synthetic facial image F121 in the first complexion using the second generative model G21, based on the synthetic facial image F12 in the second complexion. The method 100 also includes step 208 of calculating, using the processing device 402, an identity classification loss value using an identity loss minimisation model IP1, based on differences in facial features of the training facial image F1 and the synthetic facial image F121 in the first complexion, and step 1110 of updating, using the processing device 402, the first and second generative models G12, G21 based on the identity classification loss value. The facial features can be used in a face recognition model to establish an identity of a subject. The first and second generative models G12, G21 can be deep artificial neural networks, and the step of updating the first and second generative models G12, G21 can include updating weights applied to nodes of hidden layers of the first and second generative models G12, G21 based on the identity classification loss value, so that the identity classification loss value is minimised.

[0038] In some embodiments, the method 200 can further include determining, using the processing device 402, if a facial image in the first complexion (e.g. the training facial image F1 and the synthetic facial image F121) is a synthetic facial image using the first conditional model D1, and updating using the processing device 402, the first and the second

generative models G12, G21 based on the determination, to maximise an error rate associated with misclassification of the facial image in the first complexion by the first conditional model D1. The first conditional model D1 is configured to minimise the error rate associated with misclassification of the facial image in the first complexion. The method 200 can further include determining, using the processing device 402, a first cycle consistency loss value based on differences in pixel values of the training facial image F1 and the synthetic facial image F121 in the first complexion, and updating, using the processing device 402, the first and the second generative models G12, G21 based on the first cycle consistency loss value, to minimise the first cycle consistency loss value.

[0039] In some embodiments, the method 200 can also include receiving, by the processing device 402, a training facial image F2 in the second complexion, generating, using the processing device 402, a synthetic facial image F21 in the first complexion using the second generative model G21, based on the training facial image F2 in the second complexion, and generating, using the processing device 402, a synthetic facial image F212 in the second complexion using the first generative model G12, based on the synthetic facial image F21 in the first complexion. The method 200 can further include determining, using the processing device 402, if a facial image (e.g. the training facial image F2 and the synthetic facial image F212) in the second complexion is a synthetic facial image using the second conditional model D2, the second conditional model D2 configured to minimise an error rate associated with misclassification of the facial image in the second complexion, and updating, using the processing device, the first and the second generative models G12, G21 based on the determination, to maximise the error rate associated with misclassification of the facial image in the second complexion.

[0040] In some embodiments, the method 200 can further include determining, using the processing device 402, a second cycle consistency loss value based on differences in pixel values of the training facial image F2 and the synthetic facial image F212 in the second complexion, and updating, using the processing device 402, the first and the second generative models G12, G21 based on the second cycle consistency loss value, to minimise the second cycle consistency loss value.

[0041] Fig. 3 shows a flowchart illustrating an example of a method 300 for training a face recognition model, in accordance with embodiments of the disclosure. The method 300 can be implemented with the system 400 shown in Fig. 4, which includes the processing device 402. The method 300 includes step 302 of receiving, by the processing device 402, an identified facial image in a source complexion, generating, using a processing device 402, an identified facial image in a target complexion based on the received facial image using a

generative model, and training the face recognition model with the identified facial images in the source and the target complexions. The generative model can be the first generative model G12 shown in Fig. 1. In other words, the method 300 can be implemented with the first generative model G12 shown in Fig. 1. Thus, large quantities of labelled training data sets (e.g. identified facial images in the first and second complexion) can be made available to the face recognition model. Availability of labelled training data sets can improve accuracy of the model, as well as to remove and/or pre-empt biases in the model.

[0042] Fig. 5 depicts an example of a computing device 500, hereinafter interchangeably referred to as a computer system 500, where one or more such computing devices 500 may be used to execute the methods 200 and 300 of Figs. 2 and 3. One or more components of the exemplary computing device 500 can also be used to implement the systems 100, 400 as well as the processing device 402. The following description of the computing device 500 is provided by way of example only and is not intended to be limiting.

[0043] As shown in Fig. 5, the example computing device 500 includes a processor 507 for executing software routines. Although a single processor is shown for the sake of clarity, the computing device 500 may also include a multi-processor system. The processor 507 is connected to a communication infrastructure 506 for communication with other components of the computing device 500. The communication infrastructure 506 may include, for example, a communications bus, cross-bar, or network.

[0044] The computing device 500 further includes a main memory 508, such as a random access memory (RAM), and a secondary memory 510. The secondary memory 510 may include, for example, a storage drive 512, which may be a hard disk drive, a solid state drive or a hybrid drive and/or a removable storage drive 517, which may include a magnetic tape drive, an optical disk drive, a solid state storage drive (such as a USB flash drive, a flash memory device, a solid state drive or a memory card), or the like. The removable storage drive 517 reads from and/or writes to a removable storage medium 577 in a well-known manner. The removable storage medium 577 may include magnetic tape, optical disk, non-volatile memory storage medium, or the like, which is read by and written to by removable storage drive 517. As will be appreciated by persons skilled in the relevant art(s), the removable storage medium 577 includes a computer readable storage medium having stored therein computer executable program code instructions and/or data.

[0045] In an alternative implementation, the secondary memory 510 may additionally or alternatively include other similar means for allowing computer programs or other instructions to be loaded into the computing device 500. Such means can include, for example, a removable storage unit 522 and an interface 550. Examples of a removable

storage unit 522 and interface 550 include a program cartridge and cartridge interface (such as that found in video game console devices), a removable memory chip (such as an EPROM or PROM) and associated socket, a removable solid state storage drive (such as a USB flash drive, a flash memory device, a solid state drive or a memory card), and other removable storage units 522 and interfaces 550 which allow software and data to be transferred from the removable storage unit 522 to the computer system 500.

[0046] The computing device 500 also includes at least one communication interface 527. The communication interface 527 allows software and data to be transferred between computing device 500 and external devices via a communication path 526. In embodiments of the disclosure, the communication interface 527 permits data to be transferred between the computing device 500 and a data communication network, such as a public data or private data communication network. The communication interface 527 may be used to exchange data between different computing devices 500 which such computing devices 500 form part an interconnected computer network. Examples of a communication interface 527 can include a modem, a network interface (such as an Ethernet card), a communication port (such as a serial, parallel, printer, GPIB, IEEE 1394, RJ45, USB), an antenna with associated circuitry and the like. The communication interface 527 may be wired or may be wireless. Software and data transferred via the communication interface 527 are in the form of signals which can be electronic, electromagnetic, optical or other signals capable of being received by communication interface 527. These signals are provided to the communication interface via the communication path 526.

[0047] As shown in Fig. 5, the computing device 500 further includes a display interface 502 which performs operations for rendering images to an associated display 555 and an audio interface 552 for performing operations for playing audio content via associated speaker(s) 557.

[0048] As used herein, the term "computer program product" may refer, in part, to removable storage medium 577, removable storage unit 522, a hard disk installed in storage drive 512, or a carrier wave carrying software over communication path 526 (wireless link or cable) to communication interface 527. Computer readable storage media refers to any non-transitory, non-volatile tangible storage medium that provides recorded instructions and/or data to the computing device 500 for execution and/or processing. Examples of such storage media include magnetic tape, CD-ROM, DVD, Blu-ray™ Disc, a hard disk drive, a ROM or integrated circuit, a solid state storage drive (such as a USB flash drive, a flash memory device, a solid state drive or a memory card), a hybrid drive, a magneto-optical disk, or a computer readable card such as a PCMCIA card and the like, whether or not such devices

are internal or external of the computing device 500. Examples of transitory or non-tangible computer readable transmission media that may also participate in the provision of software, application programs, instructions and/or data to the computing device 500 include radio or infra-red transmission channels as well as a network connection to another computer or networked device, and the Internet or Intranets including e-mail transmissions and information recorded on Websites and the like.

[0049] The computer programs (also called computer program code) are stored in main memory 508 and/or secondary memory 510. Computer programs can also be received via the communication interface 527. Such computer programs, when executed, enable the computing device 500 to perform one or more features of embodiments discussed herein. In embodiments, the computer programs, when executed, enable the processor 507 to perform features of the above-described embodiments. Accordingly, such computer programs represent controllers of the computer system 500.

[0050] Software may be stored in a computer program product and loaded into the computing device 500 using the removable storage drive 517, the storage drive 512, or the interface 550. The computer program product may be a non-transitory computer readable medium. Alternatively, the computer program product may be downloaded to the computer system 500 over the communication path 526. The software, when executed by the processor 507, causes the computing device 500 to perform the necessary operations to execute the methods 100, 200 as shown in Figs. 1 and 2.

[0051] It is to be understood that the embodiment of Fig. 5 is presented merely by way of example to explain the operation and structure of the system 500. Therefore, in some embodiments one or more features of the computing device 500 may be omitted. Also, in some embodiments, one or more features of the computing device 500 may be combined together. Additionally, in some embodiments, one or more features of the computing device 500 may be split into one or more component parts.

[0052] It will be appreciated that the elements illustrated in Fig. 5 function to provide means for performing the various functions and operations of the system as described in the above embodiments.

[0053] When the computing device 500 is configured to realise the system 100 to train a generative model, the system 100 will have a non-transitory computer readable medium 512 having stored thereon an application which when executed causes the system 100 to perform steps comprising: receiving, by the processing device, a training facial image in a first complexion, generating, using the processing device, a synthetic facial image in a

second complexion using a first generative model, based on the training facial image in the first complexion, generating, using the processing device, a synthetic facial image in the first complexion using a second generative model, based on the synthetic facial image in the second complexion, calculating, using the processing device, an identity classification loss value using an identity loss minimisation model, based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in a face recognition model to establish an identity of a subject, and updating, using the processing device, the first generative model and the second generative model based on the identity classification loss value.

[0054] In embodiments, the computing device 500 can include at least one processor 507 and a non-transitory computer-readable storage medium 512 coupled to the at least one processor 507 and storing programming instructions for execution by the at least one processor 507. The programming instructions can instruct the at least one processor 507 to, receive a training facial image in a first complexion, generate a synthetic facial image in a second complexion using a first generative model, based on the training facial image in the first complexion and generate a synthetic facial image in the first complexion using a second generative model, based on the synthetic facial image in the second complexion. The programming instructions can also instruct the at least one processor 507 to, calculate an identity classification loss value using an identity loss minimisation model, based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in a face recognition model to establish an identity of a subject, and update the first generative model and the second generative model based on the identity classification loss value.

[0055] In some embodiments, the programming instructions can instruct the at least one processor 507 to determine if a facial image in the first complexion is a synthetic facial image using a first conditional model. The first conditional model can minimise an error rate associated with misclassification of the facial image in the first complexion. The programming instructions can instruct the at least one processor 507 to update the first and the second generative models based on the first cycle consistency loss value, to minimise the first cycle consistency loss value.

[0056] In some embodiments, the programming instructions can instruct the at least one processor 507 to determine a first cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic facial image in the first complexion, and update the first and the second generative models based on the first cycle consistency loss value, to minimise the first cycle consistency loss value.

[0057] In some embodiments, the programming instructions can instruct the at least one processor 507 to receive a training facial image in the second complexion, generate a synthetic facial image in the first complexion using the second generative model, based on the training facial image in the second complexion and generate a synthetic facial image in the second complexion using the first generative model, based on the synthetic facial image in the first complexion. The programming instructions can also instruct the at least one processor 507 to determine if a facial image in the second complexion is a synthetic facial image using a second conditional model. The second conditional model can minimise an error rate associated with misclassification of the facial image in the second complexion and update the first and the second generative models based on the determination, to maximise the error rate associated with misclassification of the facial image in the second complexion.

[0058] In some embodiments, the programming instructions can instruct the at least one processor 507 to determine a second cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic facial image in the second complexion, and update the first and the second generative models based on the second cycle consistency loss value, to minimise the second cycle consistency loss value.

[0059] In some embodiments, the computing device can include at least one processor 507 and a non-transitory computer-readable storage medium 512 coupled to the at least one processor 507 and storing programming instructions for execution by the at least one processor 507. The programming instructions can instruct the at least one processor 507 to, receive an identified facial image in a source complexion, generate an identified facial image in a target complexion based on the received facial image using a generative model, and train the face recognition model with the identified facial images in the source and the target complexions. The generative model can be trained using a cycle generative adversarial network. The cycle generative adversarial network can include a first generative model, a second generative model and an identity loss minimisation model. The first generative model can generate, based on a training facial image in a first complexion, a synthetic facial image in a second complexion. The second generative model can generate, based on the synthetic facial image in the second complexion, a synthetic facial image in the first complexion. The identity loss minimisation model can calculate an identity classification loss value based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in a face recognition model to establish an identity of a subject. The first and second generative models can be trained using the cycle generative adversarial network to minimise the identity classification loss. The generative model can be the first generative model.

[0060] In some embodiments, the cycle generative adversarial network can further include a first conditional model. The first conditional model can classify if a facial image in the first complexion is a synthetic facial image, the first conditional model trained to minimise an error rate associated with misclassification of the facial image in the first complexion. The first and the second generative models are trained using the cycle generative adversarial network to maximise the error rate associated with misclassification of the facial image in the first complexion.

[0061] In some embodiments, the cycle generative adversarial network can further include a first cycle consistency loss minimisation model. The first cycle consistency loss minimisation model can determine a first cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic facial image in the first complexion. The first and second generative models are trained using the cycle generative adversarial network to minimise the first cycle consistency loss value determined by the first cycle consistency loss minimisation model.

[0062] In some embodiments, the second generative model can generate a synthetic facial image in the first complexion, based on a training facial image in the second complexion. The first generative model can generate a synthetic facial image in the second complexion, based on the synthetic facial image in the first complexion. The cycle generative adversarial network can include a second conditional model. The second conditional model can classify if a facial image in the second complexion is a synthetic facial image, the second conditional model trained to minimise an error rate associated with misclassification of the facial image in the second complexion. The first and the second generative models are trained using the cycle generative adversarial network to maximise the error rate associated with misclassification of the facial image in the second complexion.

[0063] In some embodiments, the cycle generative adversarial network can include a second cycle consistency loss minimisation model. The second cycle consistency loss minimisation model can determine a second cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic image of the subject in the second complexion. The first and second generative models are trained using the cycle generative adversarial network to minimise the second cycle consistency loss value determined by the second cycle consistency loss minimisation model.

[0064] It will be appreciated by a person skilled in the art that numerous variations and/or modifications may be made to the present disclosure as shown in the specific embodiments without departing from the spirit or scope of the disclosure as broadly described. The present

embodiments are, therefore, to be considered in all respects to be illustrative and not restrictive.

Claims

1. A method for training generative models, the method comprising:
 - receiving, by a processing device, a training facial image in a first complexion;
 - generating, using the processing device, a synthetic facial image in a second complexion using a first generative model, based on the training facial image in the first complexion;
 - generating, using the processing device, a synthetic facial image in the first complexion using a second generative model, based on the synthetic facial image in the second complexion;
 - calculating, using the processing device, an identity classification loss value using an identity loss minimisation model, based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in a face recognition model to establish an identity of a subject;
 - updating, using the processing device, the first generative model and the second generative model based on the identity classification loss value;
 - generating, using a first conditional model, a probability score indicating whether a facial image in the first complexion is synthetic; and
 - determining, based on the probability score, that the facial image in the first complexion is synthetic.
2. The method as claimed in claim 1, wherein the first conditional model is configured to minimise an error rate associated with misclassification of the facial image in the first complexion, and wherein the method further comprises:
 - updating, using the processing device, the first and the second generative models based on the determination, to maximise the error rate associated with misclassification of the facial image in the first complexion.
3. The method as claimed in claim 1 or 2, further comprising:
 - determining, using the processing device, a first cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic facial image in the first complexion; and
 - updating, using the processing device, the first and the second generative models based on the first cycle consistency loss value, to minimise the first cycle consistency loss value.

4. The method as claimed in claim 1, further comprising:
 - receiving, by the processing device, a training facial image in the second complexion;
 - generating, using the processing device, a synthetic facial image in the first complexion using the second generative model, based on the training facial image in the second complexion;
 - generating, using the processing device, a synthetic facial image in the second complexion using the first generative model, based on the synthetic facial image in the first complexion;
 - determining, using the processing device, if a facial image in the second complexion is a synthetic facial image using a second conditional model, the second conditional model configured to minimise an error rate associated with misclassification of the facial image in the second complexion; and
 - updating, using the processing device, the first and the second generative models based on the determination, to maximise the error rate associated with misclassification of the facial image in the second complexion.
5. The method as claimed in claim 4, further comprising:
 - determining, using the processing device, a second cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic facial image in the second complexion; and
 - updating, using the processing device, the first and the second generative models based on the second cycle consistency loss value, to minimise the second cycle consistency loss value.
6. A system for training generative models, the system comprising:
 - a processing device configured to:
 - receive a training facial image in a first complexion;
 - generate a synthetic facial image in a second complexion using a first generative model, based on the training facial image in the first complexion;
 - generate a synthetic facial image in the first complexion using a second generative model, based on the synthetic facial image in the second complexion;
 - calculate an identity classification loss value using an identity loss minimisation model, based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in a face recognition model to establish an identity of a subject;

update the first generative model and the second generative model based on the identity classification loss value;

generate, using a first conditional model, a probability score indicating whether a facial image in the first complexion is synthetic; and

determine, based on the probability score, that the facial image in the first complexion is synthetic.

7. The system as claimed in claim 6, wherein the first conditional model is configured to minimise an error rate associated with misclassification of the facial image in the first complexion, and wherein the processing device is further configured to:
 - update the first and the second generative models based on the determination, to maximise the error rate associated with misclassification of the facial image in the first complexion.
8. The system as claimed in claim 6 or 7, the processing device further configured to:
 - determine a first cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic facial image in the first complexion; and
 - update the first and the second generative models based on the first cycle consistency loss value, to minimise the first cycle consistency loss value.
9. The system as claimed in claim 6, the processing device further configured to:
 - receive a training facial image in the second complexion;
 - generate a synthetic facial image in the first complexion, using the second generative model, based on the training facial image in the second complexion;
 - generate a synthetic facial image in the second complexion, using the first generative model, based on the synthetic facial image in the first complexion;
 - determine if a facial image in the second complexion is a synthetic facial image using a second conditional model, the second conditional model configured to minimise an error rate associated with misclassification of the facial image in the second complexion; and
 - update first and the second generative models based on the determination, to maximise the error rate associated with misclassification of the facial image in the second complexion.
10. The system as claimed in claim 9, wherein the processing device is further configured to:

determine a second cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic facial image in the second complexion; and

update the first and the second generative models based on the second cycle consistency loss value, to minimise the second cycle consistency loss value.

11. A method for training a face recognition model, the method comprising:

receiving, by a processing device, an identified facial image in a source complexion;

generating, using the processing device, an identified facial image in a target complexion based on the received facial image using a generative model, the generative model trained using a cycle generative adversarial network; and

training, using the processing device, the face recognition model with the identified facial images in the source and the target complexions;

wherein the cycle generative adversarial network comprises:

a first generative model configured to generate, based on a training facial image in a first complexion, a synthetic facial image in a second complexion;

a second generative model configured to generate, based on the synthetic facial image in the second complexion, a synthetic facial image in the first complexion by converting the second complexion of the synthetic facial image to the first complexion of the training facial image;

an identity loss minimisation model configured to calculate an identity classification loss value based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in the face recognition model to establish an identity of a subject; and

a first conditional model configured to:

generate a probability score indicating whether a facial image in the first complexion is synthetic; and

determine, based on the probability score, that the facial image in the first complexion is synthetic;

wherein the first and second generative models are trained using the cycle generative adversarial network to minimise the identity classification loss; and

wherein the generative model is the first generative model.

12. The method as claimed in claim 11, wherein the first conditional model is trained to minimise an error rate associated with misclassification of the facial image in the first complexion; and

wherein the first and the second generative models are trained using the cycle generative adversarial network to maximise the error rate associated with misclassification of the facial image in the first complexion.

13. The method as claimed in claim 11 or 12, wherein the cycle generative adversarial network further comprises:

a first cycle consistency loss minimisation model configured to determine a first cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic facial image in the first complexion;

wherein the first and second generative models are trained using the cycle generative adversarial network to minimise the first cycle consistency loss value determined by the first cycle consistency loss minimisation model.

14. The method as claimed in claim 11, wherein the second generative model is configured to generate a synthetic facial image in the first complexion, based on a training facial image in the second complexion;

wherein the first generative model configured to generate a synthetic facial image in the second complexion, based on the synthetic facial image in the first complexion; and

wherein the cycle generative adversarial network further comprises:

a second conditional model configured to classify if a facial image in the second complexion is a synthetic facial image, the second conditional model trained to minimise an error rate associated with misclassification of the facial image in the second complexion;

wherein the first and the second generative models are trained using the cycle generative adversarial network to maximise the error rate associated with misclassification of the facial image in the second complexion.

15. The method as claimed in claim 14, wherein the cycle generative adversarial network further comprises:

a second cycle consistency loss minimisation model configured to determine a second cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic image of the subject in the second complexion; and

wherein the first and second generative models are trained using the cycle generative adversarial network to minimise the second cycle consistency loss value determined by the second cycle consistency loss minimisation model.

16. A system, comprising:
 - a processing device configured to:
 - receive an identified facial image in a source complexion;
 - generate an identified facial image in a target complexion based on the received facial image using a generative model, the generative model trained using a cycle generative adversarial network; and
 - train a face recognition model with the identified facial images in the source and the target complexions;
 - wherein the cycle generative adversarial network comprises:
 - a first generative model configured to generate, based on a training facial image in a first complexion, a synthetic facial image in a second complexion;
 - a second generative model configured to generate, based on the synthetic facial image in the second complexion, a synthetic facial image in the first complexion by converting the second complexion of the synthetic facial image to the first complexion of the training facial image;
 - an identity loss minimisation model configured to calculate an identity classification loss value based on differences in facial features of the training facial image and the synthetic facial image in the first complexion, wherein the facial features are used in the face recognition model to establish an identity of a subject; and
 - a first conditional model configured to:
 - generate a probability score indicating whether a facial image in the first complexion is synthetic; and
 - determine, based on the probability score, that the facial image in the first complexion is synthetic;
 - wherein the first and second generative models are trained using the cycle generative adversarial network to minimise the identity classification loss; and
 - wherein the generative model is the first generative model.

17. The system as claimed in claim 16, wherein the first conditional model is trained to minimise an error rate associated with misclassification of the facial image in the first complexion, and
 - wherein the first and the second generative models are trained using the cycle generative adversarial network to maximise the error rate associated with misclassification of the facial image in the first complexion.

18. The system as claimed in claim 16 or 17, wherein the cycle generative adversarial network further comprises:
 - a first cycle consistency loss minimisation model configured to determine a first cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic facial image in the first complexion;
 - wherein the first and second generative models are trained using the cycle generative adversarial network to minimise the first cycle consistency loss value determined by the first cycle consistency loss minimisation model.

19. The system as claimed in claim 18, wherein the second generative model is configured to generate a synthetic facial image in the first complexion, based on a training facial image in the second complexion;
 - wherein the first generative model configured to generate a synthetic facial image in the second complexion, based on the synthetic facial image in the first complexion; and
 - wherein the cycle generative adversarial network further comprises:
 - a second conditional model configured to classify if a facial image in the second complexion is a synthetic facial image, the second conditional model trained to minimise an error rate associated with misclassification of the facial image in the second complexion;
 - wherein the first and the second generative models are trained using the cycle generative adversarial network to maximise the error rate associated with misclassification of the facial image in the second complexion.

20. The system as claimed in claim 19, wherein the cycle generative adversarial network further comprises:
 - a second cycle consistency loss minimisation model configured to determine a second cycle consistency loss value based on differences in pixel values of the training facial image and the synthetic image of the subject in the second complexion;
 - and
 - wherein the first and second generative models are trained using the cycle generative adversarial network to minimise the second cycle consistency loss value determined by the second cycle consistency loss minimisation model.

21. The method as claimed in claim 1, wherein the first conditional model is trained using a plurality of training facial images in the first complexion having a first label and a

plurality of synthetic facial images in the first complexion having a second label, and wherein the first label and the second label are different.