



(51) International Patent Classification:

G06K 9/62 (2006.01) ***G06F 21/62*** (2013.01)
G06K 9/00 (2006.01)

(21) International Application Number:

PCT/IL2020/051291

(22) International Filing Date:

15 December 2020 (15.12.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/948,364 16 December 2019 (16.12.2019) US

(71) Applicant: DE-IDENTIFICATION LTD. [IL/IL]; 3/2
Yael St., 6437216 Tel Aviv (IL).

(72) **Inventors:** KUTA, Eliran; 18 Stephan Zweig St., 6964208 Tel Aviv (IL). BLONDHEIM, Sella; 68 Shlomo Hamelech Street, 6451111 Tel Aviv (IL). PERRY, Gil; 14 Nehemia St., Tel Aviv (IL). BEN YOSEF, Matan; KORAZIM 106, KORAZIM (IL). HACHOHEN, Yoav; 56/1 Koreh HaDorot

St., 93393 Jerusalem (IL). **GORODISSKY, Or**; 13/9 Israelis Street, 6438213 Tel Aviv (IL).

(74) **Agent: FRYDMAN, Idan** et al.; PEARL COHEN ZEDEK LATZER BARATZ, P.O. Box 7198, 6107121 Tel Aviv (IL).

(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, IT, 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,

(54) Title: SYSTEM AND METHOD FOR RECONSTRUCTION OF FACES FROM ANONYMIZED MEDIA USING NEURAL NETWORK BASED STEGANOGRAPHY




Fig. 5

(57) Abstract: A system and method for concealing and revealing a human face in media objects may include obtaining a first media object capturing an image of the face; employing a first unit to: extract, from the first media object, a set of features representing the face, and generate a second media object, by embedding the extracted features in an anonymized media object; and employing a second unit to recognize the face based on the second media object.

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).

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

SYSTEM AND METHOD FOR RECONSTRUCTION OF FACES FROM ANONYMIZED MEDIA USING NEURAL NETWORK BASED STEGANOGRAPHY

FIELD OF THE INVENTION

[001] The present invention relates to face anonymization. More particularly, the present invention
5 relates to systems and methods for anonymizing faces in images and for reconstruction of faces from
anonymized media using neural network based steganography.

BACKGROUND OF THE INVENTION

[002] Recent improvements in commercial facial recognition systems has brought these systems to
achieve an unprecedented level of accuracy in many face recognition tasks. As these systems became
10 abundant and widely used, major concerns have been raised regarding the privacy threats these systems
may impose. As such, different vendors have begun offering systems that can automatically replace
faces of people who appear in different media, such as photos and videos, with anonymized versions of
them, and thus preventing the ability to recognize these people by third-parties who may later consume
this media. However, such practice can limit certain companies who wish to comply with recent privacy
15 regulations, but who still wish to preserve the ability to identify the individuals that originally appeared
in the media (before it has undergone any anonymization procedure).

SUMMARY OF THE INVENTION

[003] In some embodiments method of concealing and revealing a human face in media objects may
include obtaining a first media object capturing an image of the face; employing a first unit to: extract,
20 from the first media object, a set of features representing the face, and generate a second media object,
by embedding the extracted features in an anonymized media object, wherein perceptual differences
between the second media object and the anonymized media object are minimized by the first unit; and
employing a second unit to recognize the face based on the second media object.

[004] Some embodiments may employ the second unit to: extract, from the second media object, a at
25 least some of the embedded features, and generate a reconstructed media object by reconstruct the image
of the face based on the extracted features. Reconstructing the image of the face may be such that the
face is recognizable by a human. Reconstructing the image of the face may be such that the face is
recognizable by a computerized face recognition unit.

[005] At least one of the first and second units may include a model generated by adversarial training
30 of a neural network using a discriminator adapted to classify a media object as: either the first media
object or the second media object, or either the second media object or the reconstructed media object.
Embedding the extracted features in the anonymized media object may include at least one of:
modifying, according to the extracted features, a predefined set of pixels in the anonymized media
object; and modifying, according to the extracted features, a set of pixels in the anonymized media

object, wherein the set of pixels may be selected based on a function of at least one of: a scene, a location, an intensity, a color and a frequency calculated based on a set of pixels.

[006] The first and second units may include respective first and second models generated by training respective first and second neural networks. The first and second neural networks may be trained simultaneously. At least one of the first and second neural networks may be trained using a training set of pairs of a media object capturing an image of a face and corresponding anonymized media object, wherein in each of the anonymized media objects, the face from the corresponding media object may be replaced by an anonymized face.

[007] The first media object may capture a set of faces and, for each of at least some of the faces in the set, an embodiment may extract, from the first media object, a set of features, and embed the extracted features in an anonymized media object; and extract from the second media object, a at least some of the embedded features, and reconstruct images of at least some of the faces in the set based on the extracted features. The first media object may be one of: an image and a multimedia object. Other aspects and/or advantages of the present invention are described herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[008] The subject matter regarded as the invention is particularly pointed out and distinctly claimed in the concluding portion of the specification. The invention, however, both as to organization and method of operation, together with objects, features, and advantages thereof, may best be understood by reference to the following detailed description when read with the accompanying drawings in which:

[009] Fig. 1 shows a block diagram of an exemplary computing device, according to some embodiments of the invention;

[010] Fig. 2 shows a block diagram of a system for face reconstruction from anonymized media, according to some embodiments of the invention;

[011] Fig. 3 schematically illustrates a flowchart of information flow in the system shown in Fig. 2, according to some embodiments of the invention;

[012] Fig. 4 shows a flowchart of method of face reconstruction from anonymized media, according to some embodiments of the invention; and

[013] Fig. 5 shows a block diagram of a system according to illustrative embodiments of the invention.

[014] It will be appreciated that, for simplicity and clarity of illustration, elements shown in the figures have not necessarily been drawn to scale. For example, the dimensions of some of the elements may be exaggerated relative to other elements for clarity. Further, where considered appropriate, reference numerals may be repeated among the figures to indicate corresponding or analogous elements.

DETAILED DESCRIPTION OF THE INVENTION

[015] In the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of the invention. However, it will be understood by those skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known methods, procedures, and components, modules, units and/or circuits have not been described in detail so as not to obscure the invention. Some features or elements described with respect to one embodiment may be combined with features or elements described with respect to other embodiments. For the sake of clarity, discussion of same or similar features or elements may not be repeated.

[016] Although embodiments of the invention are not limited in this regard, discussions utilizing terms such as, for example, “processing”, “computing”, “calculating”, “determining”, “establishing”, “analyzing”, “checking”, or the like, may refer to operation(s) and/or process(es) of a computer, a computing platform, a computing system, or other electronic computing device, that manipulates and/or transforms data represented as physical (e.g., electronic) quantities within the computer’s registers and/or memories into other data similarly represented as physical quantities within the computer’s registers and/or memories or other information non-transitory storage medium that may store instructions to perform operations and/or processes. Although embodiments of the invention are not limited in this regard, the terms “plurality” and “a plurality” as used herein may include, for example, “multiple” or “two or more”. The terms “plurality” or “a plurality” may be used throughout the specification to describe two or more components, devices, elements, units, parameters, or the like. The term set when used herein may include one or more items. Unless explicitly stated, the method embodiments described herein are not constrained to a particular order or sequence. Additionally, some of the described method embodiments or elements thereof can occur or be performed simultaneously, at the same point in time, or concurrently.

[017] Reference is made to Fig. 1, which is a schematic block diagram of an example computing device, according to some embodiments of the invention. Computing device 100 may include a controller or processor 105 (e.g., a central processing unit processor (CPU), a chip or any suitable computing or computational device), an operating system 115, memory 120, executable code 125, storage 130, input devices 135 (e.g. a keyboard or touchscreen), and output devices 140 (e.g., a display), a communication unit 145 (e.g., a cellular transmitter or modem, a Wi-Fi communication unit, or the like) for communicating with remote devices via a communication network, such as, for example, the Internet. Controller 105 may be configured to execute program code to perform operations described herein. The system described herein may include one or more computing device 100, for example, to act as the various devices or the components shown in Fig. 2. For example, system 200 may be, or may include computing device 100 or components thereof.

[018] Operating system 115 may be or may include any code segment (e.g., one similar to executable code 125 described herein) designed and/or configured to perform tasks involving coordinating, scheduling, arbitrating, supervising, controlling or otherwise managing operation of computing device 100, for example, scheduling execution of software programs or enabling software programs or other modules or units to communicate.

[019] Memory 120 may be or may include, for example, a Random Access Memory (RAM), a read only memory (ROM), a Dynamic RAM (DRAM), a Synchronous DRAM (SD-RAM), a double data rate (DDR) memory chip, a Flash memory, a volatile memory, a non-volatile memory, a cache memory, a buffer, a short term memory unit, a long term memory unit, or other suitable memory units or storage units. Memory 120 may be or may include a plurality of possibly different memory units. Memory 120 may be a computer or processor non-transitory readable medium, or a computer non-transitory storage medium, e.g., a RAM.

[020] Executable code 125 may be any executable code, e.g., an application, a program, a process, task or script. Executable code 125 may be executed by controller 105 possibly under control of operating system 115. For example, executable code 125 may be a software application that performs methods as further described herein. Although, for the sake of clarity, a single item of executable code 125 is shown in Fig. 1, a system according to some embodiments of the invention may include a plurality of executable code segments similar to executable code 125 that may be stored into memory 120 and cause controller 105 to carry out methods described herein.

[021] Storage 130 may be or may include, for example, a hard disk drive, a universal serial bus (USB) device or other suitable removable and/or fixed storage unit. In some embodiments, some of the components shown in Fig. 1 may be omitted. For example, memory 120 may be a non-volatile memory having the storage capacity of storage 130. Accordingly, although shown as a separate component, storage 130 may be embedded or included in memory 120.

[022] Input devices 135 may be or may include a keyboard, a touch screen or pad, one or more sensors or any other or additional suitable input device. Any suitable number of input devices 135 may be operatively connected to computing device 100. Output devices 140 may include one or more displays or monitors and/or any other suitable output devices. Any suitable number of output devices 140 may be operatively connected to computing device 100. Any applicable input/output (I/O) devices may be connected to computing device 100 as shown by blocks 135 and 140. For example, a wired or wireless network interface card (NIC), a universal serial bus (USB) device or external hard drive may be included in input devices 135 and/or output devices 140.

[023] Some embodiments of the invention may include an article such as a computer or processor non-transitory readable medium, or a computer or processor non-transitory storage medium, such as for

example a memory, a disk drive, or a USB flash memory, encoding, including or storing instructions, e.g., computer-executable instructions, which, when executed by a processor or controller, carry out methods disclosed herein. For example, an article may include a storage medium such as memory 120, computer-executable instructions such as executable code 125 and a controller such as controller 105.

5 Such a non-transitory computer readable medium may be, for example, a memory, a disk drive, or a USB flash memory, encoding, including or storing instructions, e.g., computer-executable instructions, which, when executed by a processor or controller, carry out methods disclosed herein. The storage medium may include, but is not limited to, any type of disk including, semiconductor devices such as read-only memories (ROMs) and/or random-access memories (RAMs), flash memories, electrically
10 erasable programmable read-only memories (EEPROMs) or any type of media suitable for storing electronic instructions, including programmable storage devices. For example, in some embodiments, memory 120 is a non-transitory machine-readable medium.

[024] A system according to some embodiments of the invention may include components such as, but not limited to, a plurality of central processing units (CPU) or any other suitable multi-purpose or
15 specific processors or controllers (e.g., controllers similar to controller 105), a plurality of input units, a plurality of output units, a plurality of memory units, and a plurality of storage units. A system may additionally include other suitable hardware components and/or software components. In some embodiments, a system may include or may be, for example, a personal computer, a desktop computer, a laptop computer, a workstation, a server computer, a network device, or any other suitable computing
20 device. For example, a system as described herein may include one or more facility computing device 100 and one or more remote server computers in active communication with one or more facility computing device 100 such as computing device 100, and in active communication with one or more portable or mobile devices such as smartphones, tablets and the like.

[025] Reference is now made to Fig. 2, which shows a block diagram of a system 200 for face
25 reconstruction from anonymized media, according to some embodiments. In Fig. 2, hardware elements are indicated with a solid line and the direction of arrows may indicate the direction of information flow.

[026] The system 200 may include at least one processor 201 (e.g., such as controller 105 shown in Fig. 1), configured to apply at least one neural network for face reconstruction. In some embodiments, the at least one processor 201 may employ Steganography, by hiding secret messages (or images) within
30 ordinary public messages. Thus, Steganography may be used to hide a face of a subject in anonymized media.

[027] In some embodiments, original (non-anonymized) facial data may be compressed and/or embedded into anonymized media, while keeping the differences between newly generated media and the input anonymized media, visually not perceptible (e.g., by a human observer). In some embodiments,

the at least one processor 201 may employ Steganography to reconstruct the original faces from the re-generated anonymized media, even if this media travels through compression mechanisms after-which the reconstructed media may slightly differ from the uncompressed original media, as further described hereinafter.

5 [028] According to some embodiments, the at least one processor 201 may be coupled to a database 202 comprising pairs of source media files 203 and corresponding anonymized media files 204. Given a source media file 203 (e.g., a video) and an anonymized media file 204 (e.g., a video), a modified version of the anonymized media file 204 may be generated by a dedicated neural network to closely resembles the original anonymized media file 204, but from which the original faces from the source
10 media file 203 may be reconstructed.

[029] The processor 201 may be configured to apply a first neural network 205 for training a machine learning algorithm on data received from the database 202. In some embodiments, the first neural network 205 may extract face representations 206 from the received one or more source media files 203. The first neural network 205 may accordingly generate a new media file 207 based on the received one
15 or more anonymized media files 204, where the new media file 207 may be generated for face reconstruction of a face of the received one or more source media file 203.

[030] In some embodiments, at the end of the training process, the first neural network 205 may embed the extracted face representations 206 into the new media file 207, where perceptual differences between the new media file 207 and the received anonymized media file 204 may be minimized by the first neural
20 network 205 (e.g., with a loss function).

[031] After training, the at least one processor 201 may apply a second neural network 208 configured to extract the face representations 206 embedded into the new media file 207, and reconstruct at least one face image from the extracted face representations 206. In some embodiments, the reconstructed at least one face image may be reconstructed to be recognized by a human observer.

25 [032] Reference is now made to Fig. 3, which schematically illustrates a flowchart of information flow in the system 200 shown in Fig. 2, according to some embodiments. The data from the pairs of source media files 203 and corresponding anonymized media files 204 may be fed to the first neural network 205 for training to generate a (modified) new media file 207. After training, the new media file may be accordingly fed to the second neural network 208 in order to reconstruct the source face image.

30 [033] Referring back to Fig. 2. The training may include jointly training the first neural network 205 and the second neural network 208 on a collection of source media 203 and anonymized media 204 pairs. Specifically, a set of K source videos $V = \{v^{(k)}\}_{k \in [K]}$, and a corresponding set of anonymized videos $V_A = \{v_A^{(k)}\}_{k \in [K]}$ may be used, where $v_A^{(k)}$ is a modified version of the source video $v^{(k)}$, in which all faces have been replaced by an anonymized version by some external anonymization method.

Given a frame x from a source video $v \in V$, and frame x_A from an anonymized video $v_A \in V_A$, the first neural network 205 may be trained to generate a new frame $\widehat{x}_A = E(x, x_A)$ which is perceptually similar to the anonymized version x_A , but in which dense representations of the original faces. While the second neural network 208 may be trained to reconstruct the original faces from the new frame $\widehat{x}_A$ generated by the first neural network 205. The two neural networks may be trained using the Adam optimizer on the following loss function formulas :

$$L_{sim}(x_A, \widehat{x}_A) = ||x_A - \widehat{x}_A||_1$$

Formula 1

$$L_{recon}(x, \widehat{x}_A) = \sum_{f \in F(x)} ||f - \hat{f}||_1$$

Formula 2

where $F(x)$ is the set of all faces found in the video frame x , and $\hat{f}$ is the reconstruction of a source face $f \in F(x)$, as obtained from the second neural network 208.

[034] It will be understood that the Adam optimizer describe herein is just one example and that any other methods or techniques, e.g., stochastic gradients may be used. It will also be understood that, for the sake of clarity and simplicity, only one face in an image is described, however, the method described herein may be applied to images with any number of faces, that is, any number of faces in an image may be concealed and then revealed as described with respect to a single face in an image.

[035] In order to increase the fidelity of the frames generated by the first neural network 205, adversarial training may be employed and another neural network 'D' may be used that may be trained to discriminate between original video frames belonging to the anonymized videos, and those video frames generated by the first neural network 205. In some embodiments, the first neural network 205 may be trained to fool the new neural network 'D' by generating more convincing video frames. This adversarial loss may be written as:

$$L_D(x_A, \widehat{x}_A) = \log(D(\widehat{x}_A)) - \log(D(x_A))$$

Formula 3

$$L_G(\widehat{x}_A) = -\log(D(\widehat{x}_A))$$

Formula 4

$$L_{adv}(x_A, \widehat{x}_A) = L_D(x_A, \widehat{x}_A) + L_G(\widehat{x}_A)$$

Formula 5

[036] The final loss term may accordingly become:

$$L(V, V_A) = \sum_{k \in [K]} \sum_{x \in v^{(k)}} L_{sim}(x_A, \widehat{x}_A) + \lambda * L_{recon}(x, \widehat{x}_A) + L_{adv}(x_A, \widehat{x}_A)$$

Formula 6

where λ is a tuneable hyper-parameter that controls the balance between the two loss terms.

[037] Reference is now made to Fig. 4, which shows a flowchart of method of face reconstruction from anonymized media, according to some embodiments. In Step 401, a source media file and a corresponding anonymized media file may be received.

[038] In Step 402, a first neural network may be applied, the first neural network may be configured to
5 extract face representations from the received source media file, generate a new media file based on the received anonymized media file, and embed the extracted face representations into the new media file.

[039] In some embodiments, the new media file may be generated for face reconstruction of a face of the received source media file. In some embodiments, perceptual differences between the new media file and the received anonymized media file may be minimized by the first neural network.

10 [040] In Step 403, a second neural network may be applied, the second neural network may be configured to apply a second neural network configured to extract the face representations embedded into the new media file, and reconstruct at least one face image from the extracted face representations, wherein the reconstructed at least one face image is reconstructed to be recognized by a human observer.

[041] In some embodiments, a method of face reconstruction from anonymized media may include the
15 steps of receiving a source media file and a corresponding anonymized media file; applying a first neural network configured to: extract face representations from the received source media file; generate a new media file based on the received anonymized media file, wherein the new media file is generated for face reconstruction of a face of the received source media file; embed the extracted face representations into the new media file, wherein perceptual differences between the new media file and the received
20 anonymized media file are minimized by the first neural network; and applying a second neural network configured to: extract the face representations embedded into the new media file; and reconstruct at least one face image from the extracted face representations, wherein the reconstructed at least one face image is reconstructed to be recognized by a human.

[042] In some embodiments, at least one of the first neural network and the second neural network are
25 trained with a machine learning algorithm using a training set of media file pairs of source media and corresponding anonymized media, and wherein in each anonymized media file, the face from the corresponding source media file is replaced by an anonymized face.

[043] In order to train the at least one of the first and second neural networks, some embodiments may apply adversarial training with a third neural network trained to discriminate between media frames of
30 the anonymized media and media frames generated by the first neural network. A third or fourth neural network may be trained to discriminate between images reconstructed by the second network and original images.

[044] In some embodiments, a system for face reconstruction from anonymized media may include at least one processor, configured to apply at least one neural network for face reconstruction; and a

database, coupled to the processor, and comprising pairs of source media files and corresponding anonymized media files; wherein the processor is configured to: apply a first neural network configured to: extract face representations from the received source media file; generate a new media file based on the received anonymized media file, wherein the new media file is generated for face reconstruction of a face of the received source media file; and embed the extracted face representations into the new media file, wherein perceptual differences between the new media file and the received anonymized media file are minimized by the first neural network; and apply a second neural network configured to: extract the face representations embedded into the new media file; and reconstruct at least one face image from the extracted face representations, wherein the reconstructed at least one face image is reconstructed to be recognized by a human.

[045] In some embodiments, at least one of the first neural network and the second neural network in the system are trained with a machine learning algorithm using a training set from the database, and wherein in each anonymized media file the face from the corresponding source media file is replaced by an anonymized face.

[046] In some embodiments, the processor is further configured to apply adversarial training with a third neural network trained to discriminate between media frames of the anonymized media and media frames generated by the first neural network.

[047] As described, available (e.g., commercial) facial recognition systems are able to identify or recognize faces in images or videos with high accuracy. Accordingly, storing or otherwise keeping images or videos of people reflects on their privacy since stored images can be stolen or otherwise obtained by a third party (e.g., an entity other than the owner of an image and/or other than a person appearing in the image). Merely replacing, anonymizing or obscuring faces in stored images (as done by some known systems and methods) cannot adequately solve the challenge of safely keeping images of people since the actual (original) faces cannot be reconstructed from the stored images.

[048] There is therefore a need in the art for a system and method that can anonymize an image, photo or video and, in addition, provide the ability to see, recognize or identify an individual (or face) that originally appeared in the image, photo or video (before it has undergone an anonymization procedure).

[049] As described herein, some embodiments extract (and possibly compress) original, non-anonymized, facial data and embed the extracted data into an input anonymized media to thus create a new anonymized media, while keeping the differences between the newly generated media (the new anonymized media) and the input anonymized media, visually imperceptible. As further described, some embodiments may restore an original, non-anonymized media, given the new anonymized media. Some embodiments may reconstruct an original face from an anonymized media, even if the anonymized media traveled through, or was subject to manipulations such as compression.

[050] Accordingly, some embodiments of the invention may enable an owner of images of sensitive material (e.g., an individual, organization or any other entity who owns images of faces of employees or private individuals) to store and/or keep the sensitive images such that if the images are obtained by a third party, the third party cannot see or identify the sensitive material, e.g., even if a third party obtains
5 images produced as described herein, the third party cannot (or is highly unlikely to) see and/or identify the individuals' or faces captured by the images.

[051] In addition to concealing an identity of a person in an image, e.g., by producing an anonymized image as described, some embodiments of the invention enable producing (or reproducing) an image of the person (or other entity or object) from an anonymized image.

10 [052] Accordingly, embodiments of the invention enable an owner of images to safely store or keep images of individuals without worrying about the privacy of the individuals. For example, using embodiments of the invention, rather than storing in its database images of its employees, an organization can store in its database anonymized images as described. If or when the need to actually see the face of an employee in an anonymized image, the organization can use an embodiment of the
15 invention that produces (or reproduces) an image of an individual (or of any other object in an anonymized image) from an anonymized image as described herein. Where applicable, the terms "digital media" and/or simply "media" as used herein may relate to a digital object that captures a face of a person and/or a scene, for example, a media may be a digital image or photo and/or a video clip.

[053] Some embodiments of the invention receive as input, or operate on, pairs of source and
20 anonymized objects or media, e.g., a pair of images or videos. For example, a face of a person may appear in, or be captured by, a source media. Given a source media and an anonymized media, some embodiments modify the input anonymized media to thus generate a modified anonymized media closely resembles the input, original anonymized media, but from which a face appearing in, or captured by, an input, source media can be reconstructed.

25 [054] Some embodiments may include, train and use artificial neural networks (NN). For example, as described, in some embodiments, a first NN (referred to herein as "embedder") may be trained to extract face representations of faces from an input, original media, and create a modified anonymized media by embedding the extracted data into an anonymized media. The first NN may be trained to keep perceptual differences between an input anonymized media and a modified anonymized media to a minimum.
30 Perceptual differences as referred to herein may be, or may include any difference that can be perceived, e.g., by a human or a computerized face recognition unit.

[055] In some embodiments, a second NN (referred to herein as "reconstructor") may be trained to extract data in a modified anonymized media (e.g., one created by the embedder NN) and use extracted data to reconstruct a viewable, recognizable or identifiable image of a face. It is noted that embodiments

of the invention may use, support and/or operate with any anonymization techniques, systems or methods. In other words, embodiments of the invention may be combined with, or may include any face anonymization system or method. Accordingly, the scope of the invention is not limited by the technique, system and/or method used for anonymizing media.

5 [056] Reference is made to Fig. 5 which shows a block diagram of a system 500 according to illustrative embodiments of the invention. System 500 may include one or more instances of computing device 100. For example, in some embodiments, NNs 530, 535, 550 and 555 may each be, or may each include a computing device 100, in other embodiments, one or more computing devices 100 may be used to implement or realize NNs 530, 535, 550 and 555. As shown, system 500 may include, use and/or or
10 train, four NN units 530, 535, 550 and 555. As shown, an anonymized media 510 and a source face image 520 may be provided as input to NN 530. For the sake of clarity and simplicity, anonymized media 510 may be referred to herein as media 510 and source face image 520 may be referred to herein as image 520. Generally, an anonymized media 510 may be any image which does not include any, or enough, information enabling determining an identity of a person. Conceptually, an anonymized media
15 510 may not include an image of a person, e.g., it may be a randomly or artificially generated image (e.g., by a computer that generates an image by randomly setting pixel values), or an anonymized media 510 may be an image of a nature scene or it may be an image of a person that was processed or modified such that the identity of the person cannot be determined.

[057] For the sake of simplicity and clarity, a single image (520), single anonymized media (510), single
20 modified anonymized media (540) and single reconstructed face image (560) are shown in Fig. 5 and mainly described herein. It will however be understood that embodiments of the invention may process videos or sequences of images or multimedia content as described herein for a single image 520. For example, an embodiment may iteratively process frames in a video or other multimedia content as described herein with respect to image 520, that is, frames in an input video may each be treated and
25 processed by embodiments of the invention as described with reference to image 520, either using the same media 510 for all frames or selecting a different media 510 for different frames. It is noted that media objects 510, 520, 540 and 560 may be, or may include, images or frames of a video. Some embodiments may create a media 540 for each frame in a video, element of a multimedia content object image in a sequence of images, and may further reconstruct each frame or image by creating, for each
30 frame or image, an image 560 as further described herein.

[058] For example, anonymized media 510 may be created by anonymizing source face image 520. As described, any method of anonymizing, blurring or obscuring a face seen or included in, or captured by, source face image 520 may be used in order to create anonymized media 510. In some embodiments, a model may be trained based on, or using, anonymized images created using a specific method, e.g.,

blurring, pixelization, face swapping and the like. In some embodiments anonymized media 510 may be created without any reference or respect to source face image 520, e.g., anonymized media 510 may be a template anonymized media created from scratch or based on a randomly selected image of a face.

[059] As shown, based on source face image 520 and anonymized media 510, the embedder NN 530 may create modified anonymized media 540 (also referred to herein as media 540). For example, to create modified anonymized media 540, NN 530 may extract face representations from source face image 520 and modify anonymized media 510 by embedding or including in anonymized media 510 the face representations extracted from source face image 520. Face representations as referred to herein may be any information (e.g., pixel values, colors and the like) that may be extracted from an image and/or may be calculated, determined or derived based on an image.

[060] Embedder NN 530 may be trained to create modified anonymized media 540 such that a visual or other difference between modified anonymized media 540 and anonymized media 510 is kept to a minimum. Otherwise described, the goal of NN 530 is to create a modified anonymized media 540 which is indistinguishable from anonymized media 510, either by man or by machine.

[061] In some embodiments, NN 535 may act or function as a discriminator or classifier. For example, NN 535 may be trained to identify or classify anonymized media objects 510 as such, and to further identify or classify modified media objects 540 as such. Accordingly, given an input media object, NN 535 may classify the input media as either belonging to a class of anonymized media objects (the class of media object 510) or belonging to a class of modified media objects (the class of modified media object 540). In this context, media objects 510 and 540 are referred to in the art as real and fake objects respectively, that is, an original, input (media 510 in this case) is considered, or referred to in the art, as the “real” thing while a modified original object (media 540 in this case) is considered, or referred to as “fake”

[062] Staying with this (real/fake) terminology, NN 535 may be trained to classify a media object as either real or fake. In some embodiments, NN 535 and NN 530 may be trained together, simultaneously or in an interleaved fashion or manner. For example, after NN 530 creates a modified media object 540 the object is provided as input to NN 535 and output from NN 535 may be provided back to NN 530 informing NN 530 whether NN 535 classified the modified media object 540 as anonymized media 510 or modified media object 540 (real or fake).

[063] Accordingly, a loop may be formed where as NN 535 gets better in classifying media objects it drives NN 530 to generate modified media objects 540 that are closer and closer in similarity to input media objects 510. For example, in simultaneous training, in each iteration of a training, a modified media object 540 created by NN 530 is provided to NN 535 for classification and NN 535 informs NN 530 how well it did, that is, NN 530 did well if NN 535 classified the (fake) modified media object 540

as (real) anonymized media object 510 and NN 530 did poorly or insufficiently well if NN 535 classified the (fake) modified media object 540 as fake. Training NNs 530 and 535 together may be interleaved or carried out in an alternative manner. For example, NN 530 may be trained over a set of images or iterations and then, e.g., after 100 iterations, a media object 540 may be provided to NN 535 for classification as described. In other embodiments, a set of modified media objects 540 previously created by NN 530 may be provided to NN 535 and used for training NN 535, in yet other embodiments, success or failure indications may be provided to NN 530, from NN 535, based on an average or ratio of a plurality of modified media objects 540 classified as fake or real as described.

[064] In some embodiments, adversarial training may be used to train NN 530. For example, adversary NN 535 may be trained to discriminate between modified anonymized media 540 and anonymized media 510. An adversary NN as referred to herein may be an NN that challenges another NN. For example, while NN 530 attempts (is trained to) produce a media 540 that is as close to a media 510, adversary NN 535 challenges NN 530 by trying to identify or determine that the output media 540 of NN 530 is not similar to media 510. Training an NN using an adversary NN is referred to in the art and herein as adversarial training.

[065] For example and as illustrated by the arrows connecting anonymized media 510 and modified anonymized media 540 with NN 535, NN 535 may be provided with anonymized media 510 and modified anonymized media 540 and may be trained to determine whether or not they are different or same. Accordingly, NN 535 may challenge NN 530, that is, NN 530's goal is to fool NN 535 into determining ("thinking") that anonymized media 510 and modified anonymized media 540 are the same media (since NN 535 could not find any (or enough) differences between the two media objects). As illustrated, NN 530 may receive input from NN 535, e.g., NN 530 may "know" whether or not it managed to fool NN 535 into thinking that anonymized media 510 and modified anonymized media 540 are the same.

[066] Accordingly, in a training according to embodiments of the invention, adversary NN 535 keeps learning to identify differences between anonymized media 510 and modified anonymized media 540 while NN 530 keeps learning how to better create modified anonymized media 540 such that NN 535 will fail in identifying differences between anonymized media 510 and modified anonymized media 540, this loop or race can be kept running until NN 530 creates modified anonymized media 540 that is up to a requirement. For example, training of NN's 530 and 535 may be conducted until some third party unit is unable to distinguish between anonymized media 510 and modified anonymized media 540.

[067] In some embodiments, NN 555 may act or function as a discriminator or classifier. For example, NN 555 may be trained to identify reconstructed face images 560 and further identify source face images

520 and thus, given an input media object, NN 555 may classify the input media as either belonging to a class of reconstructed face images (the class of reconstructed face image 560) or belonging to a class of source face images (the class of source face images 520). In this context, media objects or images 520 and 560 are referred to in the art as real and fake objects respectively, staying with this terminology, NN 555 may be trained to classify a media object as real or fake. In some embodiments, NN 555 and NN 550 may be trained together, simultaneously or in an interleaved fashion. For example, after NN 550 creates a reconstructed face image 560 the image is provided as input to NN 555 and output from NN 555 may be provided back to NN 550 informing NN 550 whether NN 555 classified the reconstructed face image 560 as reconstructed face image 560 or as a source face image 520 (fake or real).

[068] Accordingly, a loop is formed where as NN 555 gets better in classifying media objects it drives NN 550 to generate reconstructed face images 560 that are closer and closer in similarity to source face images 520. For example, in simultaneous training, in each iteration of a training, a reconstructed face images 560 created by NN 550 is provided to NN 555 for classification and NN 555 informs NN 550 how well it did, that is, NN 550 did well if NN 555 classified the (fake) reconstructed face image 560 as a (real) face image 520 and NN 550 did poorly or insufficiently well if NN 555 classified the (fake) reconstructed face image 560 as fake. Training NNs 550 and 555 together may be interleaved, or carried out in an alternative manner. For example, NN 550 may be trained over a set of 540 images or iterations and then, e.g., after 100 iterations, a reconstructed face images 560 may be provided to NN 555 for classification as described. In other embodiments, a set of modified reconstructed face images 560 previously created by NN 550 may be provided to NN 555 and used for training NN 555, in yet other embodiments, success or failure indications may be provided to NN 550, from NN 555, based on an average or ratio of reconstructed face images 560 classified as fake or real as described.

[069] In some embodiments, NN 535 (also referred to herein as a discriminator) may be adapted or trained to distinguish between anonymized and anonymized images into which information representing a face was added. More specifically, NN 535 may be adapted or trained to determine how different is an anonymized media 510 from a modified anonymized media 540.

[070] Otherwise described, NN 535 may be trained and/or adapted to classify similar but non-identical pairs of 510/540 media objects as “not-different” if the difference between them is marginal, less or below some level and NN 535 may trained and/or adapted to classify less similar pairs of 510/540 media objects as “too different”. For example, in some embodiments, differences between 510 and 530 images may be minimized using a loss function which transforms source images 520 and respective modified images 540 to a selected representation (e.g., a feature vector or some function representing the images), measure the distance between the representations of the images, and attempts to minimize the distance.

[071] Output from NN 535 may be provided as input to NN 530 and used in training NN 530 to create a 540 media object such that (a) the identity of the human who's face is shown in image 520 is concealed and (b) modified image 540 highly resembles, or is as similar or close as possible to, anonymized image 510. Otherwise described, provided with a 510 anonymized image and a 540 modified anonymized image, NN 535 may be adapted to determine or identify which of the two input images is the input 510 image and which is the modified 540 image. Output from NN 535 may be provided back to NN 530 thus indicating, to NN 530, how close or similar images 510 and 540 are, and NN 530 may use input from NN 535 in order to improve its performance, that is, in order to create modified images 540 which conceal identities and, at the same time, are as close as possible to respective 510 images.

[072] As illustrated in Fig. 5, modified anonymized media 540 may be provided to reconstructor NN 550 which may be trained to recreate or reconstruct, from modified anonymized media 540 the face in source face image 520 in the form of reconstructed face image 560 (also referred to herein as image 560). In some embodiments, adversarial training may be used to train NN 550. For example, adversary NN 555 may be trained to identify differences between the face in image 560 and the face in image 520. For example and as illustrated in Fig. 5, NN 555 may receive image 520 and image 560 and may be trained to determine whether or not the two images are the same. Accordingly, in a training according some embodiments, NN 550 tries, learns or is trained to, fool NN 555 into thinking image 520 and image 560 are the same while NN 555 tries, learns or is trained to identify differences between image 520 and image 560.

[073] As illustrated, NN 550 may receive input from NN 555, e.g., NN 550 may "know" whether or not it managed to fool NN 555 into thinking that image 560 and image 520 are the same image.

[074] In some embodiments, NN 555 may be adapted, or trained to identify or determine, given two images, which of the two is an original image 520 and which is a reconstructed image 560, and, as shown, output from NN 555 may be provided back to NN 550 thus indicating to NN 550 how close or similar image 560 is to image 520. Accordingly, during training, NN 555 keeps informing NN 550 how well it performs, that is, how close an image 560 is to its respective image 520 thus, during training, NN 550 keeps getting better at generating images 560 which are as close as possible to images 520. In some embodiments, differences between 520 images and respective 560 images may be minimized using a loss function which transforms input images 520 and reconstructed images 560 to a selected representation (e.g., a feature vector or some function representing the images), measure the distance between the representations of the images, and attempts to minimize the distance.

[075] Accordingly, in some embodiments, loop or race as described with reference to NN 530 and NN 535 is created for NN 550 and NN 555 in a training, this loop or race can be kept running until NN 550

creates learns or is trained to adequately reconstruct image 520 from image 560, e.g., such that a third party unit cannot distinguish, or find differences, between image 520 and image 560.

[076] In some embodiments, a method of concealing and revealing a human face in media objects may include: obtaining a first media object capturing an image of the face; employing a first unit to: extract, from the first media object, a set of features representing the face, and generate a second media object, by embedding the extracted features in an anonymized media object, wherein perceptual differences between the second media object and the anonymized media object are minimized by the first unit. The method may further include employing a second unit to recognize the face based on the second media object. For example, recognizing or identifying a face in the second (anonymized) media object may include reconstructing the first (original image) media based on the second (anonymized) media and associating the reconstructed media or image with the first media object (e.g., first original image), e.g., determining that both first and reconstructed media objects include a representation of the same face. In some embodiments, recognizing a face in an anonymized image may not include reconstructing an image as described, rather, an embodiment (e.g., NN 550) may extract data (features) from an anonymized image (e.g., from media 540) and use the extracted data to identify a person or to determine that the face of person whose features are embedded or included in media 540 is same as the face in media 520.

[077] Otherwise described, an embodiment may determine whether or not the face shown or captured in an input, first image (e.g., a face captured by an image 520) and the face represented in a second image (e.g., image 540) are of the same person, thus, recognizing a face in the second image may be, or may include, determining the face in the second image is of the person whose face is captured or depicted in the first image. The terms or phrases “features” and “features representing a face” used herein may relate or refer to any information extracted from an image or derived based on an image. For example, a feature may be the color of eyes, the distance between eyes, a point of view of a camera taking a picture, a color density, an orientation of a face in an image and the like. Generally, where NNs are involved, features may be any aspect, information, parameter or value in an image that an NN learns, or is trained to identify. Accordingly, features may be extracted from an image and/or embedded into an image. For example, embedding a feature into an image may include modifying the color of eyes in an image of a face, changing skin color, changing hair style, changing an orientation of a face in an image and so on.

[078] The method may further include employing the second unit to: extract, from the second media object, at least some of the embedded features, and reconstruct the image of the face based on the extracted features.

[079] For example, acting as the first unit, NN 530 may be trained to conceal a human face by obtaining image 520 (first media object), extract from image 520 a set of features representing a face captured by image 520 and generate or create media 540 (a second media object) by embedding the extracted features in media 510 such that perceptual differences between media 510 and media 540 are minimized.

5 As further described, training an NN may include developing a model that may subsequently be used as part of units in any (possibly large number of) number systems. In some embodiments, reconstructing an image of a face is such that the face is recognizable by a human. For example, NN 550 may be trained such that a human is easily able to determine that the face in media 560 is the same face in media 520. For example, NN 555 may be trained and/or configured to mimic human perception of images and may
10 thus cause NN 550 to produce images 560 with accuracy that would satisfy humans.

[080] In some embodiments, reconstructing an image of a face is such that the face is recognizable by a computerized face recognition unit. For example, the standard set for NN 550 is that a face recognition unit or software is able to determine that the face in image 560 is the same face in image 520. For example, NN 555 may mimic or include face recognition software and may thus cause NN 550 to keep
15 improving until the face recognition software or unit is able to determine that the faces in images 560 and 520 are of the same person.

[081] In some embodiments, embedding extracted features in an anonymized media object includes at least one of: modifying, according to the extracted features, a predefined set of pixels in the anonymized media object; and modifying, according to the extracted features, a set of pixels in the anonymized media
20 object, wherein the set of pixels is selected based on a function of at least one of: a scene, a location, an intensity, a color and a frequency calculated based on a set of pixels.

[082] For example, to create media 540, NN 540 may embed, insert or include pixels (or pixel values) in media 510 such that the inserted or modified pixels represent features, or information extracted from image 520. The specific pixels in media 510 which NN 530 modifies or replaces may be selected based
25 on any combination of logic and pixels' attributes. For example, NN 530 may replace or modify only pixels in a scene (but not in an anonymized face), or NN 530 may replace pixels in specific locations (e.g., pixels in specific offsets from a reference pixel where the offsets are calculated according to some function). In other cases, NN 530 may replace, insert or modify pixels in media 510 according to a frequency of pixel values in media 510, a color and so on. It will be noted that the logic for selecting
30 which pixels to replace may be predefined or it may be autonomously developed by NN 530, e.g., during a training in which NN 530 attempts to fool NN 535 into thinking that medias 510 and 540 are the same media NN 530 may autonomously devise ways of embedding information in media 510 when creating media 540. Of course, training as described may include repeating the process of creating media 540 using many pairs of media 510 and image 520.

[083] In some embodiments, by training an NN as described, a model may be created and such model may subsequently be used by many units, e.g., by a plurality of systems adapted to receive and be configured based on a model. Accordingly, it will be understood that units described herein as concealing and revealing a human face in images may be units that include a model created by training an NN as described.

[084] In some embodiments, some units may include a model which was developed using adversarial training. For example, training of NN 530 using adversarial NN 535 is an adversarial training that develops a model in NN 530. In some embodiments, training NNs, e.g., training NNs 530, 535, 550 and 555, may be done using a training set of pairs, each pair including a media object capturing an image of a face and corresponding anonymized media object, e.g., pairs of media 510 and image 520.

[085] In some embodiments, an anonymized media may be created by replacing a face in an image with an anonymized face. For example, in some embodiments, a pair of image and anonymized media object (e.g., a 520/510 pair) may be created by anonymizing the image to create the anonymized media object. For example, starting with a clear input image in which John is clearly seen (image 520), an embodiment may modify the input image of John such that John can no longer be seen or identified in the modified image and the embodiment may then pair the modified image with the input image, e.g., use the input and modified images as image 520 and media 510 in a training cycle. Training an NN may include any (possibly very large) number of cycles involving a large number of 520/510 image/media pairs.

[086] Although for the sake of clarity and simplicity, a single face in an images and media objects is described, it will be understood that any number of faces may be included or shown in images and medias 510, 520, 540 and 560. In some embodiments, a first media object (e.g., 510) may capture, include or show, a set of faces and, for some, or even all of the faces in the set, an embodiment may: extract, from the first media object (e.g., from 510), a set of features, and embed the extracted features in an anonymized media object, e.g., in media 510 thus creating media 540. Accordingly, any number of faces in an input image may be concealed in a media 540.

[087] In some embodiments, a plurality or set of faces concealed in a media (e.g., in media 540 as described) may be revealed, e.g., made visible, recognizable or identifiable in image 560. For example, NN 550 may be trained to extract, from media 540, embedded features of a plurality of faces and use the extracted features to reconstruct (e.g., in image 560) the set of faces. For example, in an exemplary case of using an embodiment of the invention, image 510 clearly shows a number of family members, the family members cannot be seen or identified in media 540 (which may be safely stored without worrying that unauthorized people are able to see who is in the picture) and image 560 may be an image that, similar to image 520, clearly shows the family members.

[088] In some embodiments, a set of NNs may be trained together, simultaneously, concurrently and/or at the same time. For example, a training session of NN 530 may be viewed as a sequence of steps where, in each step, NN 530 creates a media 540 based on a (possibly different) input image 520 and gets input from NN 535 indicating how well it performed. In order to train NNs 530 and 550 together or simultaneously, the output of NN 530 in each step may be provided as input to the next step of NN 550, thus the system shown in Fig. 5 may be trained as a unit and the training may improve the unit's capability of concealing, anonymizing, obfuscating or obscuring a face in an image as well improve the unit's capability of reconstructing or revealing a face concealed in an anonymized object. Of course, parts of the system shown in Fig. 5 may be used separately, for example, at a first time or stage, NN 530 (or a model developed by training NN 530) may be use, e.g., by an organization in order to conceal or hide faces of employees in images stored in a database of the organization (e.g., as media 540 objects). At some later time or stage, NN 550 (or a model developed by training NN 550) may be used, by the organization, in order to reveal or see the face of a specific employee, e.g., by creating an image 560 from the media 540 of the specific employee.

[089] While certain features of the invention have been illustrated and described herein, many modifications, substitutions, changes, and equivalents may occur to those skilled in the art. It is, therefore, to be understood that the appended claims are intended to cover all such modifications and changes as fall within the invention.

[090] Various embodiments have been presented. Each of these embodiments may, of course, include features from other embodiments presented, and embodiments not specifically described may include various features described herein.

CLAIMS

1. A method of concealing and revealing a human face in media objects, the method comprising:
obtaining a first media object capturing an image of the face;
employing a first unit to:
5 extract, from the first media object, a set of features representing the face, and
 generate a second media object, by embedding the extracted features in an anonymized
media object, wherein perceptual differences between the second media object and the
anonymized media object are minimized by the first unit; and
employing a second unit to recognize the face based on the second media object.
10
2. The method of claim 1, comprising:
employing the second unit to:
 extract, from the second media object, at least some of the embedded features, and
 generate a reconstructed media object by reconstruct the image of the face based on
15 the extracted features.
3. The method of claim 2, wherein reconstructing the image of the face is such that the face is
recognizable by a human.
- 20 4. The method of claim 2, wherein reconstructing the image of the face is such that the face is
recognizable by a computerized face recognition unit.
5. The method of claim 2, wherein at least one of the first and second units includes a model
generated by adversarial training of a neural network using a discriminator adapted to classify a
25 media object as:
 either the first media object or the second media object, or
 either the second media object or the reconstructed media object.
6. The methods of claims 1-5, wherein embedding the extracted features in the anonymized
30 media object includes at least one of:
 modifying, according to the extracted features, a predefined set of pixels in the
anonymized media object; and

modifying, according to the extracted features, a set of pixels in the anonymized media object, wherein the set of pixels is selected based on a function of at least one of: a scene, a location, an intensity, a color and a frequency calculated based on a set of pixels.

5 7. The method of claim 1-6, wherein the first and second units include respective first and second models generated by training respective first and second neural networks.

8. The method of claim 1-7, wherein the first and second neural networks are trained simultaneously.

10

9. The method of claim 1-8, wherein at least one of the first and second neural networks is trained using a training set of pairs of a media object capturing an image of a face and corresponding anonymized media object, and wherein in each of the anonymized media objects, the face from the corresponding media object is replaced by an anonymized face.

15

10. The method of claim 1-9, wherein the first media object captures a set of faces and wherein, for each of at least some of the faces in the set, the method includes:

extracting, from the first media object, a set of features, and
embedding the extracted features in an anonymized media object;

20

and wherein the method further includes:

extracting from the second media object, at least some of the embedded features, and
reconstructing images of at least some of the faces in the set based on the extracted features.

11. The method of claim 1-10, wherein the first media object is one of: an image and a
25 multimedia object.

12. A computer readable medium comprising instructions which, when executed by a processor, cause the processor to implement any of the methods of claims 1-11.

1/5




Fig. 1

2/5




Fig. 2

3/5




Fig. 3

4/5





FIG. 4

**Fig. 5**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IL2020/051291

A. CLASSIFICATION OF SUBJECT MATTER

IPC (20210101) G06K 9/62, G06K 9/00, G06F 21/62

CPC (20130101) G06K 9/6217, G06K 9/00248, G06F 21/6254

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)

IPC (20210101) G06K 9/62, G06K 9/00, G06F 21/62

CPC (20130101) G06K 9/6217, G06K 9/00248, G06F 21/6254

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)

Databases consulted: Esp@cenet, Google Patents, Derwent Innovation, SIMILARI

Search terms used: image, face, extract, features, embedding, anonymized, recognize, restore, reconstruct

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2018225061 A1 De-Identification ltd. 13 Dec 2018 (2018/12/13) The whole document	1
Y	The whole document	2-12
Y	WO 2015039084 A1 EYEVERIFY 19 Mar 2015 (2015/03/19) The whole document	1-12
Y	US 6278791 B1 Honsinger et al. 21 Aug 2001 (2001/08/21) The whole document	1-12
A	US 2014328547 A1 Whitehill et al. 06 Nov 2014 (2014/11/06) The whole document	1-12

☒ 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

“D” document cited by the applicant in the international application

“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

24 Jan 2021

Date of mailing of the international search report

25 Jan 2021

Name and mailing address of the ISA:

Israel Patent Office

Technology Park, Bldg.5, Malcha, Jerusalem, 9695101, Israel

Email address: pctoffice@justice.gov.il

Authorized officer

GORBUNOVA Yelena

Telephone No. 972-73-3927220

INTERNATIONAL SEARCH REPORT

International application No.
PCT/IL2020/051291

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	MX 2014015326 A INSTITUTO POLIT?CNICO NACIONAL 14 Jun 2016 (2016/06/14) The whole document	1-12

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IL2020/051291

Patent document cited search report			Publication date	Patent family member(s)			Publication Date
WO	2018225061	A1	13 Dec 2018	WO	2018225061	A1	13 Dec 2018
				CN	111344712	A	26 Jun 2020
				EP	3635621	A1	15 Apr 2020
				IL	252657	D0	31 Aug 2017
				IL	271165	D0	30 Jan 2020
				JP	2020522829	A	30 Jul 2020
				KR	20200014889	A	11 Feb 2020
				US	2020097767	A1	26 Mar 2020
US	2014328547	A1	06 Nov 2014	US	2014328547	A1	06 Nov 2014
				US	9105119	B2	11 Aug 2015
				US	2015324633	A1	12 Nov 2015
				US	9639743	B2	02 May 2017
				US	2017301121	A1	19 Oct 2017
				US	10319130	B2	11 Jun 2019
WO	2015039084	A1	19 Mar 2015	WO	2015039084	A1	19 Mar 2015
				BR	112016005604	A8	18 Feb 2020
				CN	104823203	A	05 Aug 2015
				CN	104823203	B	19 Mar 2019
				CN	105917353	A	31 Aug 2016
				CN	105917353	B	21 Sep 2018
				CN	108846390	A	20 Nov 2018
				CN	108846390	B	14 Feb 2020
				CN	109165593	A	08 Jan 2019
				CN	109165593	B	10 Mar 2020
				CN	109190531	A	11 Jan 2019
				CN	109190531	B	05 May 2020
				CN	109284683	A	29 Jan 2019
				CN	110048832	A	23 Jul 2019

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IL2020/051291

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
		CN 110048832 B	18 Sep 2020
		CN 110084013 A	02 Aug 2019
		CN 110084013 B	11 Aug 2020
		EP 3047425 A1	27 Jul 2016
		EP 3047426 A1	27 Jul 2016
		HK 1213670 A1	09 Sep 2016
		HK 1257828 A1	01 Nov 2019
		JP 2016532945 A	20 Oct 2016
		JP 6039863 B2	07 Dec 2016
		JP 2017054532 A	16 Mar 2017
		JP 6255463 B2	27 Dec 2017
		JP 2016538661 A	08 Dec 2016
		JP 6353910 B2	04 Jul 2018
		JP 2018077862 A	17 May 2018
		JP 6599421 B2	30 Oct 2019
		JP 2018092684 A	14 Jun 2018
		JP 6651565 B2	19 Feb 2020
		JP 2020074183 A	14 May 2020
		JP 6774580 B2	28 Oct 2020
		KR 20160075509 A	29 Jun 2016
		KR 101808467 B1	12 Dec 2017
		KR 20160064138 A	07 Jun 2016
		KR 101853218 B1	27 Apr 2018
		KR 20180045054 A	03 May 2018
		KR 101967124 B1	08 Apr 2019
		KR 20170138587 A	15 Dec 2017
		KR 102113911 B1	21 May 2020
		US 8965066 B1	24 Feb 2015
		US 2015078630 A1	19 Mar 2015
		US 2015078629 A1	19 Mar 2015
		US 9053365 B2	09 Jun 2015

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IL2020/051291

Patent document cited search report			Publication date		Patent family member(s)		Publication Date			
					US	2015286864	A1	08 Oct 2015		
					US	9390327	B2	12 Jul 2016		
					US	2016086030	A1	24 Mar 2016		
					US	9454700	B2	27 Sep 2016		
					US	2015186721	A1	02 Jul 2015		
					US	9495588	B2	15 Nov 2016		
					US	2017154202	A1	01 Jun 2017		
					US	9940517	B2	10 Apr 2018		
					US	2017140204	A1	18 May 2017		
					US	2018211092	A9	26 Jul 2018		
					US	10210388	B2	19 Feb 2019		
					US	2018189543	A1	05 Jul 2018		
					US	10664700	B2	26 May 2020		
					US	2020175250	A1	04 Jun 2020		
					US	10726260	B2	28 Jul 2020		
					US	2020184187	A1	11 Jun 2020		
					US	10789465	B2	29 Sep 2020		
					WO	2015039102	A1	19 Mar 2015		
			US	6278791	B1	21 Aug 2001	US	6278791	B1	21 Aug 2001
			-----			-----		-----		-----
MX	2014015326	A	14 Jun 2016	MX	2014015326	A	14 Jun 2016			
-----			-----		-----		-----			
					MX	354612	B	13 Mar 2018		
								</		