



(51) International Patent Classification:

G06V 10/30 (2022.01) G06V 10/24 (2022.01)
G06V 10/72 (2022.01) G06V 10/40 (2022.01)
G06V 10/75 (2022.01) G06V 10/42 (2022.01)
G06V 10/77 (2022.01) G06V 40/16 (2022.01)
G06V 10/80 (2022.01) G06V 40/19 (2022.01)
G06T 7/80 (2017.01)

(21) International Application Number:

PCT/US2022/073014

(22) International Filing Date:

17 June 2022 (17.06.2022)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

63/212,528 18 June 2021 (18.06.2021) US

(71) Applicants: **THE BOARD OF TRUSTEES OF THE LE-
LAND STANFORD JUNIOR UNIVERSITY** [US/US];
Building 170, 3rd Floor, Main Quad, P.O. Box 20386,
Stanford, CA 94305-2038 (US). **UNIVERSIDAD IN-
DUSTRIAL DE SANTANDER** [CO/CO]; Carrera 27,
Calle 9, Bucaramanga, Santander (CO).

(72) Inventors: **NIEBLES, Juan, Carlos**; Building 170, 3rd
Floor, Main Quad, P.O. Box 20386, Stanford, CA
94305-2038 (US). **HINOJOSA, Carlos**; Universidad In-
dustrial de Santander, Carrera 27, Calle 9, Salon Lp 334, Bu-
caramanga, Santander (CO). **ARGUELLO, Henry**; Uni-
versidad Industrial de Santander, Carrera 27, Calle 9, Salon
Lp 333, Bucaramanga, Santander (CO).

(74) Agent: **KAVEH, David**; KPPB LLP, 2190 S. Towne Cen-
tre Place, Suite 300, Anaheim, CA 92806 (US).

(54) Title: SYSTEMS AND METHODS FOR PRIVACY-PRESERVING OPTICS




FIG. 3A

(57) **Abstract:** Systems and methods for privacy-preserving optics are described. An embodiments includes a method of preserving-privacy on captured images while performing a computer vision task that includes generating an optimal set of parameters to parameterize an encoding optical element to produce optical distortions such that acquired images by the camera are distorted, where the optimal set of parameters is learned via end-to-end learning that jointly optimizes from a camera optics model to a computational process that performs a computer vision task on the distorted images acquired by the camera, where the distorted images visually obscure a privacy attribute of people to protect their privacy but still preserve features to perform the computer vision task, acquiring several distorted images, and performing a computer vision task directly on the distorted images where the distortions generated by the camera are optimal and allow obtaining high performance on the computer vision task.

(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, IQ, IR, IS, IT, JM, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

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

Published:

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

SYSTEMS AND METHODS FOR PRIVACY-PRESERVING OPTICS

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application claims benefit of and priority under 35 U.S.C. 119(e) to U.S. Provisional Application Serial Number 63/212,528, entitled “Systems and Methods for Privacy Preserving Optical Systems” by Niebles et al., filed June 18, 2021, the disclosure of which is herein incorporated by reference in its entirety.

FIELD OF THE INVENTION

[0002] The present invention generally relates to the design and implementation of optimized optics for performing a machine vision task while preserving privacy on captured images..

BACKGROUND

[0003] Cameras are optical devices that capture visual images. The majority of cameras use one or more lenses to focus light on a light-sensitive surface. Digital cameras use optical sensors as the light-sensitive surface in contrast to photosensitive film. Digital cameras capture images as image data. While there are many different ways to encode image data, a common way is as a grid of pixels, where each pixel records intensity values. A time-series of images can be viewed (and encoded) as video data.

[0004] Computer vision is a scientific field concerned with understanding how computers can gain understanding from image and video data. Computer vision is a broad field, and includes (non-exhaustively) scene reconstruction, object detection, event detection, video tracking, object recognition, pose estimation, motion estimation, visual serving, and many other applications.

SUMMARY OF THE INVENTION

[0005] Systems and method in accordance with embodiments of the invention capture images while preserving-privacy attributes and perform a computer vision task on the captured images. One embodiment includes a system, that includes: a camera including an encoding optical element, at least one lens and a focal plane array (FPA) detector; at

least one processor; memory that includes an image processing pipeline application; where the image processing pipeline application directs the at least one processor to: generate an optimal set of parameters to parametrize the encoding optical element in the camera to produce optical distortions such that acquired images by the camera are distorted, where the optimal set of parameters is learned via end-to-end learning that jointly optimizes from a camera optics model to a computational process that performs a computer vision task on the distorted images acquired by the camera, where the distorted images visually obscure a privacy attribute of people to protect their privacy but still preserve features to perform the computer vision task; acquire several distorted images using the camera; perform the computer vision task directly on the distorted images where the distortions generated by the camera are optimal and allow obtaining high performance on the computer vision task.

[0006] In a further embodiment, the image processing pipeline application includes a deep neural network (DNN), where the DNN is trained to detect several features directly on the distorted images, and where the DNN includes several layers for detecting the several features based on several nodes that perform a calculation on degraded input image data.

[0007] In a further embodiment, the encoding optical element includes an ensemble of optical elements that generate image distortion via light modulation.

[0008] In a further embodiment, the distortion produced by the encoding optical element is optimal for a specific computer vision task and allows obtaining high performance.

[0009] In a further embodiment, the functionality of the camera, including the encoding optical element, is emulated in software to generate synthetic degraded images.

[0010] In a further embodiment, the camera model is coupled with the image processing pipeline software to emulate the system.

[0011] In a further embodiment, the set of parameters for the encoding optical element add optical aberrations to the camera.

[0012] In a further embodiment the distorted images obscure at least one privacy-preserving human attribute selected from a group consisting of, for example, facial features, gender, race, and age.

[0013] In a further embodiment, the distorted images are robust to blind and non-blind deconvolution attacks.

[0014] In a further embodiment, performing a computer vision task includes performing human pose estimation (HPE) on the distorted images.

[0015] In a further embodiment, the image processing pipeline application includes a computer vision decoder that has been jointly trained to generate the set of parameters of the camera encoding optical element α^* and a convolutional neural network for human pose estimation h^* such that:

$$\alpha^*, h^* = \arg \min_{\alpha, h} L_T(h) + L_P(\alpha)$$

where L_T is a loss function for the human pose estimation task, and L_P is a loss function that encourages privacy-preservation.

[0016] In a further embodiment, the set of parameters of the camera encoding optical element includes a lens surface profile ϕ , in terms of Zernike coefficients α , and a corresponding point spread function (PSF) $\mathbf{H}$ for the camera lens.

[0017] In a further embodiment, the set of parameters parametrize a lens surface profile ϕ with the Zernike basis, such that :

$$\phi = \sum_{j=1}^q \alpha_j \mathbf{Z}_j,$$

where $\mathbf{Z}_j$ is the j -th Zernike polynomial in Noll notation, and α_j is the corresponding coefficient, wherein each Zernike polynomial describes a wavefront aberration such that the surface profile ϕ is formed by a linear combination of all aberrations.

[0018] In a further embodiment, the computer vision decoder is a convolutional neural network (CNN) decoder that is trained using a backbone network and several branches of convolutional layers, where the backbone network extracts features from an image of size $w \times h$, which are then fed into the several branches convolutional layers, where a first branch predicts a set of confidence maps, where each map represents a specific body part location, and a second branch predicts a set of Part Affinity Fields (PAFs), where each field represents the degree of association between parts.

[0019] In a further embodiment, the optical encoder and the computer vision decoder share a loss function; and the optical encoder and computer vision decoder are

configured to jointly optimize the privacy-preserving optical system and a computer vision model.

[0020] Another embodiment includes a method of preserving-privacy on captured images while performing a computer vision task, including generating an optimal set of parameters to parameterize an encoding optical element of a camera to produce optical distortions such that acquired images by the camera are distorted, where the optimal set of parameters is learned via end-to-end learning that jointly optimizes from a camera optics model to a computational process that performs a computer vision task on the distorted images acquired by the camera, where the distorted images visually obscure a privacy attribute of people to protect their privacy but still preserve features to perform the computer vision task; acquiring several distorted images using the camera; performing a computer vision task directly on the distorted images where the distortions generated by the camera are optimal and allow obtaining high performance on the computer vision task.

[0021] In a further embodiment, the method further includes using a deep neural network (DNN), where the DNN is trained to detect several features directly on the distorted images, and where the DNN includes several layers for detecting the several features based on several nodes that perform a calculation on degraded input image data.

[0022] In a further embodiment, encoding optical element includes an ensemble of optical elements that generate image distortion via light modulation.

[0023] In a further embodiment, the distortion produced by the encoding optical element is optimal for a specific computer vision task and allows obtaining high performance.

[0024] In a further embodiment, the computer vision task is a human pose estimation (HPE) task on the distorted images.

BRIEF DESCRIPTION OF THE DRAWINGS

[0025] The description and claims will be more fully understood with reference to the following figures and data graphs, which are presented as exemplary embodiments of the invention and should not be construed as a complete recitation of the scope of the invention.

[0026] FIG. 1 conceptually illustrates a system for generating a privacy preserving optic and computer vision model in accordance with an embodiment of the invention.

[0027] FIG. 2 is a block diagram illustrating a camera using a privacy generating optic and a computer vision model in accordance with an embodiment of the invention.

[0028] FIG. 3A and Fig. 3B conceptually illustrates an end-end framework for privacy-preserving optics for machine vision.

[0029] FIG. 4 illustrates a block diagram of a privacy-preserving system for machine vision in accordance with an embodiment of the invention.

DETAILED DESCRIPTION

[0030] Privacy is a significant concern for many consumers, however more consumer products are including cameras in order to provide functionality that is in fact desired by consumers. By way of example, a video gaming system may include one or more cameras (e.g. the Microsoft Kinect) that is functionally an internet connected camera always pointed at a consumer's living room. Further, augmented reality glasses, security cameras, home automation cameras, medical patient monitoring cameras, and many more devices produce similar concerns. As can be readily appreciated given the ubiquity of cameras, there are many instances where privacy should be incentivized. While these cameras might not have nefarious intent, there is no guarantee that they could not be coopted by bad actors. Therefore, it is desirable to in fact degrade the capabilities of these cameras so they cannot capture quality images of private spaces, but also to maintain functionality which has traditionally required relatively high fidelity images. Systems and methods described herein are capable of generating privacy-preserving optics that produce images that are difficult for humans to interpret visually, and also difficult for computers to deblur or otherwise enhance.

[0031] Deep optics is the concept of building an end-to-end optimization of a computational camera such that the optics and the computer vision algorithms are optimized for each other to achieve higher resolution and higher fidelity in the image while simultaneously improving the performance of the algorithm. Systems and methods described herein utilize a “reverse” deep optics approach, where the goal is not to improve image quality, but in fact to obscure information in the image. In many embodiments, a specialized lens or lens stack can be generated in concert with a computer vision model which achieves the goal of a conventional computer vision model in concert with a conventional lens, but produces and operates on images that preserve the privacy of individuals in the images. For example, in many embodiments, a pose estimation computer vision model can be generated to be used with a specifically generated optic, where the optic produces blurry images that cannot be processed to show the faces of individuals in the image, but still produces valid, high accuracy pose estimation when processed. As can be readily appreciated, this is not limited to pose estimation, but the computer vision model can be swapped out for any computer vision application as appropriate to the requirements of specific applications of embodiments of the invention, with a similarly appropriate custom optic.

[0032] Turning now to FIG 1, a system for generating a privacy preserving optic and computer vision model in accordance with an embodiment of the invention is illustrated. Generation system 100 includes an optical encoder 110, a computer vision decoder 120, and a loss function 130. In many embodiments, the optical encoder is a machine learning system which generates simulated images from labeled training data as passed through a model of a lens. In various embodiments, the optical encoder can include a convex thin lens and a refractive/diffractive optical element (freeform lens) add-on. In contrast to the traditional lens design approach, the parametrization of the freeform lens can be learned during training such that aberrations are added instead of removing them. Then, the optical encoder can generate privacy-preserving images from labeled training data as passed through the lens model. In many embodiments, the computer vision decoder is a machine learning system which attempts to perform a specific computer vision task on images provided by the optical encoder. A loss function can be designed to be privacy preserving and rewards privacy-related attributes such as, but not limited to, face clarity,

race, gender, age, and/or any other privacy-related attribute as appropriate to the requirements of specific applications of embodiments of the invention. In many embodiments, accuracy of the computer vision model is also rewarded. In various embodiments, the optical encoder and computer vision decoder attempt to obtain the highest accuracy in the computer vision model while incrementing the image distortion as much as possible. Both the optical encoder and the computer vision decoder can be updated via back-propagation. Specific embodiments of system 100 are discussed in further detail below.

[0033] In many embodiments, output of optical encoder 110 is an optic design 112. Similarly, the output of the computer vision decoder 120 is a computer vision model 122. The optic design and computer vision model can be designed to operate with each other in a complete system. Turning now to FIG. 2, a camera utilizing an optic design and computer vision model produced as discussed above is illustrated. Camera 200 includes a privacy-preserving optic 210. Privacy preserving optic 210 may function “poorly” as a lens. That is, it may not preserve high fidelity images as commonly understood, and may be highly irregular. It should be understood that the standard convex lens illustrated in the instant figures is illustrative only, and is not representative of any specific optic design. Light passes through the privacy-preserving optic and strikes image sensor 220. The image sensor passes image data to the computer vision model 230 which has been generated in concert with the privacy preserving optic. The computer vision model produces an output which is passed to an input/output (I/O) interface for use in other parts of the system and/or in other systems.

[0034] As can be readily appreciated, while one lens is shown in the above, any number of lens can be generated for use in a privacy-preserving lens stack. Further, it is important to note that conventional cameras with conventional optics can be made into privacy-preserving camera systems by replacing the optical encoder with a firmware encoder which, when embedded into the image sensor, can manipulate the image sensor’s output to directly produce a privacy-preserving image. However, given that the privacy preservation occurs at the sensor level rather than the optic level, there is a possibility of an additional vulnerability if the device were able to be accessed by a bad actor.

[0035] Many embodiments provide for a camera that can directly obviate sensitive data while still obtaining useful information for a given task. In many embodiments, the entire system, including the camera's optical elements and image processing algorithms parameters can be optimized in an end-to-end fashion, enabling the design of domain-specific computational cameras. The end-to-end optimization of domain-specific computational cameras has been known as Deep Optics, which aims to improve the optical elements to acquire high-resolution/high-fidelity images and simultaneously improve the performance of computer vision algorithms. Many embodiments of the system can extend this philosophy to design privacy-preserving optical systems.

[0036] Many embodiments provide for a privacy-preserving computational camera via end-to-end optimization to capture useful information to perceive humans in the scene while hiding privacy-sensitive information. Since many computer vision applications need to analyze humans as the first step in their frameworks, the system can jointly optimize a freeform lens (e.g., the spatially varying surface height of a lens, among various other attributes) together with a human pose estimation (HPE) network to develop a privacy-preserving HPE system. Accordingly, the system can provide a privacy-preserving end-to-end optimization framework to extract useful information from a scene yet preventing the imaging system from obtaining detailed and privacy-sensitive visual data. The system can use an end-to-end optimization framework to optimize an optical encoder (e.g., Hardware-level protection) with a software decoder (e.g., convolutional neural net) to add a visual privacy protection layer to HPE. The optical elements of the camera lens can be jointly optimized and the backbone of a HPE network can be fine-tuned. In many embodiments, it may not be necessary to retrain the HPE network layers to achieve privacy preservation. Many embodiments can perform extensive simulations on a dataset (e.g., the COCO dataset) to validate the privacy-preserving deep optics approach for HPE.

[0037] In many embodiments, the optical system lens can be designed to degrade the image quality and obscure sensitive private information, which can be opposite from the traditional approach of improving the imaging quality. The system can add a visual privacy protection layer to an already trained HPE network using the designed optics and fine-

tune the backbone layers. There can be a trade-off between the attained scene degradation and the HPE precision.

Privacy-preserving Pose Estimation

[0038] Many embodiments of the system provide for privacy-preserving human pose estimation and/or other machine vision applications. In certain embodiments, the system can optimize the camera optics and the human pose estimation network jointly to achieve privacy protection via image degradation.

[0039] To preserve privacy, many embodiments of the system modify a camera lens to degrade an image in such a way that the identity of the subjects is obscured while preserving important features for pose estimation. An end-to-end framework of a system for privacy-preserving pose estimation in accordance with an embodiment is conceptually illustrated in Fig. 3A and Fig. 3B. In particular, the system can include two components, an Optical Encoder as illustrated in Fig. 3A and a Convolutional Neural Network (CNN) Decoder as illustrated in Fig. 3B in accordance with an embodiment of the invention. The Optical Encoder module can be parametrized appropriately to allow for learning of the camera lens. The CNN Decoder can perform the task of human pose estimation and/or other machine vision operations on an optically degraded image.

[0040] In many embodiments, during training, the Optical Encoder module and the CNN Decoder can be jointly optimized to provide for the privacy-preserving pose estimation system. In many embodiments, the result of the training process is two-fold: the camera lens parameters α^* and the convolutional network for pose estimation h^* . To achieve this, a loss function can be formulated for learning that combines the two goals as provided by equation (1) below:

[0041]
$$\alpha^*, h^* = \arg \min_{\alpha, h} L_T(h) + L_P(\alpha). \quad (1)$$

[0042] where L_T is the loss function for the pose estimation task, and L_P is a loss function that encourages privacy-preservation. During inference, the system can be deployed in hardware by constructing a camera lens using the optimal parameters α^* that acquires degraded images on which the network h^* can perform pose estimation. Certain embodiments of the system can be deployed as a software-only solution and can implement the image degradation post-acquisition.

OPTICAL ENCODER

[0043] An Optical Encoder module as illustrated in Fig. 3A in accordance with an embodiment of the invention can be responsible for the image acquisition process in a privacy-preserving human pose estimation (HPE) system. In many embodiments of the system can modify the optical system of the camera during training to provide for privacy-preservation. Accordingly, the system can produce images that visually obscure the identity of the person but still preserve important features for pose estimation. Many embodiments of the system achieve this by adopting deep optics, and use end-to-end training to jointly optimize the camera optics and the HPE network.

[0044] In many embodiments, the camera optics can be optimized by adding optical aberrations directly on the surface of the lens (e.g., a freeform lens among various other types of lenses) instead of removing them. Many embodiments may not perform image reconstruction and instead work directly with acquired low-quality images.

[0045] Many embodiments of the system can appropriately parametrize the camera lens to enable end-to-end learning and to perform back-propagation. A training signal to optimize the camera optics can back-propagate from the privacy-preserving loss $L_p(\alpha)$.

[0046] In several embodiments of the system, there are several parts to the parametrization: the lens surface profile ϕ , which can be in terms of Zernike coefficients α , and the corresponding point spread function (PSF) $\mathbf{H}$ for the camera lens. Described below are further details regarding the relationship between ϕ and $\mathbf{H}$ by the image formation model in accordance with various embodiments of the invention. Also described below are the parametrization of the lens surface profile ϕ in terms of coefficients α for the Zernike polynomials in accordance with various embodiments of the invention.

Image Formation Model

[0047] In many embodiments of the system, a wave-based image formation model for natural scenes can be derived to write the PSF $\mathbf{H}$ in terms of ϕ , assuming spatially incoherent light. The light transport in the camera can be modeled using a differentiable Fourier optics model.

[0048] Fig. 3A illustrates an optical system in accordance with an embodiment of the invention. The optical system can include a convex thin lens with a custom refractive optical element add-on with surface profile ϕ . Similar to a photographic filter, such an optical element can be mounted directly in front of the lens. The response of the camera system to a point light source can be described by the point spread function (PSF) created by the lens. The sensing process can be modeled as a 2D convolution operation between the scene and PSF as provided by equation (2) below:

[0049]
$$\mathbf{y} = g(\mathbf{H} * \mathbf{x}) + \boldsymbol{\eta}, \quad (2)$$

[0050] where $\mathbf{x} \in \mathbb{R}_+^{w \times h}$ is the scene and it is represented as a discrete color image with $w \times h$ pixels, and each pixel has value in $[0,1]$; $\boldsymbol{\eta}$ represents the Gaussian noise in the sensor, and $g(\cdot)$ is the camera response function, which can be assumed to be linear. In many embodiments, the PSF can be assumed to be shift-invariant, but the model could be generalized.

[0051] Assuming that the thin lens has a focal length f at a distance d_2 from the sensor, the relationship between the in-focus distance and the sensor distance in the paraxial ray approximation is given by the thin-lens equation: $1/f = 1/d_1 + 1/d_2$.

[0052] Therefore, an object at a distance d_1 in front of the lens appears in focus at a distance d_2 behind the lens. Assuming that the scene is at optical infinity, many embodiments can first propagate the light emitted by the point, represented as a spherical wave, to the lens. The complex-valued wave field immediately before the lens is given by equation (3) below:

[0053]
$$U(x, y) = \exp(ik\sqrt{x^2 + y^2 + z^2}), \quad (3)$$

[0054] where $k = 2\pi/\lambda$ is the wavenumber. The refractive optical element can first delay the phase of this incident wavefront by an amount proportional to the surface profile ϕ of the optical element at each point (x, y) . Equivalently, the optical element may be represented by a multiplicative phase transformation of the form described by equation (4) below:

[0055]
$$t_\phi(x, y) = \exp(ik(n(\lambda) - 1)\phi(x, y)), \quad (4)$$

[0056] where $n(\lambda)$ is the wavelength-dependent refractive index of the optical element material.

[0057] The light wave continues to propagate to the camera lens, which induces the following phase transformation described by equation (5) below:

$$\mathbf{[0058]} \quad t_l(x, y) = \exp\left(-i \frac{k}{2f}(x^2 + y^2)\right) \quad (5)$$

[0059] Considering that a lens has a finite aperture size, many embodiments of the system can use a binary circular mask $A(x, y)$ with diameter D to model the aperture and block light in regions outside the open aperture. To find the electric field immediately after the lens, many embodiments of the system can multiply the amplitude and phase modulations of the refractive optical element and lens with the input electric field using equation (6) below:

$$\mathbf{[0060]} \quad \tilde{U}(x, y) = A(x, y)t_\phi t_l(x, y)U(x, y). \quad (6)$$

[0061] Finally, the field propagates a distance d_2 to the sensor with the exact transfer function:

$$\mathbf{[0062]} \quad T_{d_2}(f_x, f_y) = \exp\left[ikd_2\sqrt{1 - (\lambda f_x)^2 - (\lambda f_y)^2}\right], \quad (7)$$

[0063] where (f_x, f_y) are spatial frequencies. This transfer function can be applied in the Fourier domain as:

$$\mathbf{[0064]} \quad \bar{U}(x', y') = \mathcal{F}^{-1}\{\mathcal{F}\{\tilde{U}(x, y)\} \cdot T_{d_2}(f_x, f_y)\}, \quad (8)$$

[0065] where $\mathcal{F}$ denotes the 2D Fourier transform. Since the sensor can measure light intensity, many embodiments of the system take the magnitude-squared to find the values of the PSF $\mathbf{H}$ at each position (x, y) as:

$$\mathbf{[0066]} \quad H(x', y') = |\bar{U}(x', y')|^2. \quad (9)$$

Lens Parametrization

[0067] Many embodiments of the system can parametrize a lens surface profile ϕ with the Zernike basis, which can lead to smoother surfaces, as provided by equation 10 below:

$$\mathbf{[0068]} \quad \phi = \sum_{j=1}^q \alpha_j \mathbf{Z}_j, \quad (10)$$

[0069] where $\mathbf{Z}_j$ is the j -th Zernike polynomial in Noll notation, and α_j is the corresponding coefficient. Each Zernike polynomial can describe a wavefront aberration; hence the surface profile ϕ can be formed by the linear combination of all aberrations. In

this regard, the optical element parameterized by ϕ can be seen as an optical encoder, where the coefficients α_j determine the data transformation. Therefore, the end-to-end training in accordance with many embodiments of the system can find a set of coefficients $\alpha^* = \{\alpha_j\}_{j=1}^q$ that provides the maximum visual distortion of the scene but allows to extract relevant features to perform HPE. In many embodiments, different types of parameters and techniques can be used to determine wavefront aberrations as appropriate to the requirements of specific applications in accordance with embodiments of the invention.

CNN DECODER

[0070] Many embodiments of the system can perform HPE using an available network architecture (e.g., OpenPose network architecture). The network (e.g., OpenPose network) in accordance with many embodiments of the system can be composed of a backbone (e.g., VGG-19 backbone), and have one or more branches of convolutional layers (e.g., two branches of convolutional layers). The backbone network can extract features from an image of size $w \times h$, which can then be fed into the one or more branches. In certain embodiments, one branch can predict a set of confidence maps, where each map represents a specific body part location; a second branch can predict a set of Part Affinity Fields (PAFs), where each field represents the degree of association between parts. Successive stages can be performed to refine the predictions made by each branch. In many embodiments of the system, the confidence maps and the PAFs can be parsed by greedy inference to produce the 2D locations of body keypoints for each person in the image.

HPE Loss Function L_T

[0071] In many embodiments, the network (e.g., OpenPose) loss can account for both the body and face keypoints to improve human pose estimation in an image. Let $\mathcal{S} = \{\mathbf{S}_1, \mathbf{S}_2, \dots, \mathbf{S}_E\}$ be the set of confidence maps, where each map $\mathbf{S}_e \in \mathbb{R}^{w \times h}$ represents a specific keypoint location, $e \in \{1, \dots, E\}$. Similarly, let $\mathcal{V} = \{\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_C\}$ be the set of PAFs, where each affinity field $\mathbf{V}_c \in \mathbb{R}^{w \times h \times 2}$ represents the degree of association between the keypoints, and $c \in \{1, \dots, C\}$. Certain embodiments of the system can split the

confidence maps as $\mathcal{S} = \{\mathcal{S}_B, \mathcal{S}_F\}$, where $\mathcal{S}_B \subseteq \mathcal{S}$ includes the maps $\mathbf{S}_e$ for the body keypoints, and $\mathcal{S}_F \subseteq \mathcal{S}$ includes the maps $\mathbf{S}_e$ for the face keypoint locations. Similarly, certain embodiments of the system can split the PAFs as $\mathcal{V} = \{\mathcal{V}_B, \mathcal{V}_F\}$, where $\mathcal{V}_B \subseteq \mathcal{V}$ includes the affinity fields $\mathbf{V}_c$ of the body limbs, and $\mathcal{V}_F \subseteq \mathcal{V}$ includes the affinity fields $\mathbf{V}_c$ that represent the degree of association between two face parts. Certain embodiments can define the loss function for a subset of keypoints χ at stage τ is

$$\text{[0072]} \quad F_\tau(\chi) = \sum_{\delta=1}^{|\chi|} \sum_{\mathbf{p}} \mathbf{B}(\mathbf{p}) \cdot \|\chi_\delta^\tau(\mathbf{p}) - \chi_\delta^*(\mathbf{p})\|_2^2 \quad (11)$$

[0073] where $|\chi|$ is the number of keypoints in the subset. For instance, if $\chi = \mathcal{S}_B$ then $|\chi|$ will be the total number of body-related confidence maps. $\mathbf{B}$ is a binary mask with $\mathbf{B}(p) = 0$ when the annotation is missing at the pixel p , and χ_δ^* denotes the groundtruth. Then, the overall OpenPose L_T is

$$\text{[0074]} \quad L_T = \sum_{\tau=1}^{\Gamma_1} F_\tau(\mathcal{V}_B) + F_\tau(\mathcal{V}_F) + \sum_{\tau=\Gamma_1+1}^{\Gamma_1+\Gamma_2} F_\tau(\mathcal{S}_B) + F_\tau(\mathcal{S}_F), \quad (12)$$

[0075] where Γ_1 and Γ_2 denote the total of PAF and confidence map stages, respectively. Although the above description provides a particular set of computations to determine an HPE Loss Function, any of a variety of computational techniques can be utilized as appropriate to the requirements of specific applications in accordance with embodiments of the invention.

PRIVACY-PRESERVING LOSS FUNCTION L_P

[0076] In many embodiments, defining a privacy-preserving loss function can depend on concrete application contexts. There can be various privacy-related attributes, such as the face, race, gender, or age among various other attributes. In many embodiments of the system, the face can be the main attribute to obscure in a privacy-preserving vision task. Accordingly, certain embodiments of the system can define the privacy-preserving loss taking into account face keypoints detection in the images. In many embodiments, a user may not be interested in obtaining an accurate localization of face keypoints, and the system can obscure such face regions from the image. Accordingly, the system may preserve the body keypoints and let the end-to-end training degrade most if not all of the image's spatial details (e.g., including the faces). To further enforce image degradation,

many embodiments of the system can maximize the ℓ_2 norm error between the original image $\mathbf{x}$ and the acquired image $\mathbf{y}$, defined as

[0077]
$$L_E = \sum_b \frac{1}{2} \|\mathbf{y}_b - \mathbf{x}_b\|_2^2, \quad (13)$$

[0078] where the subscript b denotes the color bands of the images (e.g., RGB images, among various other types of images in different color coordinates). In many embodiments, the privacy-preserving loss function can be defined as provided in equation (14) below:

[0079]
$$L_P = -\sum_{\tau=1}^{\Gamma_1} F_{\tau}(\mathcal{V}_F) - \sum_{\tau=\Gamma_1+1}^{\Gamma_1+\Gamma_2} F_{\tau}(\mathcal{S}_F) - \beta_2 L_E. \quad (14)$$

[0080] Considering Eq. 1, many embodiments of the system can compute the total loss at the end of the framework as follows in equation (15) below:

[0081]
$$L = L_T + L_P = \beta_1 (\sum_{\tau=1}^{\Gamma_1} F_{\tau}(\mathcal{V}_B) + \sum_{\tau=\Gamma_1+1}^{\Gamma_1+\Gamma_2} F_{\tau}(\mathcal{S}_B)) - \beta_2 L_E. \quad (15)$$

[0082] Although a particular set of computations are described for a privacy preserving loss function, any of a set of computations can be utilized for preserving privacy for different attributes, including race, gender, age among others as appropriate to the requirements of specific applications in accordance with embodiments of the invention. Training details for training a network in accordance with embodiments of the invention are described in detail below.

TRAINING DETAILS

Optics Layer Simulation

[0083] In certain embodiments, the system can include simulating a sensor with a particular pixel size (e.g., pixel size of $3.40\mu\text{m}$) and resolution (e.g., resolution of 864×864 pixels). The system can use the first Zernike coefficients in Noll notation (e.g., first Zernike coefficients $q = 350$ Zernike coefficients in Noll notation) to shape the surface profile ϕ . The fourth Zernike coefficient (the defocus term) can be initialized, such that the lens has a particular focal length (e.g., focal length of $f = 25\text{mm}$). The optical element can be discretized with a particular feature size (e.g., $3.40\mu\text{m}$ feature size) on a particular grid (e.g., 864×864 grid). Although a specific set of pixel size, resolution and Zernike coefficients are set forth above, any of a variety of pixel size, resolution and coefficients

can be specified as appropriate to the requirements of specific applications in accordance with embodiments of the invention.

Fine-tuning

[0084] Many embodiments of the system can add a privacy protection layer to a pre-trained network (e.g., OpenPose network). In certain embodiments, to perform training, the system can assume an aberration-free freeform lens and use pretrained weights of a particular network implementation (e.g., Tensorflow implementation of OpenPose) as a starting point. After initialization with the pre-trained weights, the system can freeze one or more branches (e.g., the two branches of OpenPose) and only fine-tune some layers of the backbone (e.g., VGG-19 backbone) with a lower learning rate to learn extracting human body features from the private image y . Figure 3B illustrates the frozen and fine-tuned layers in accordance with an embodiment of the invention.

Training

[0085] In many embodiments, during training, the system can first perform one forward pass through the network by convolving the images from the training set with the PSF H to obtain the optically-encoded sensor image y , as described by Eq. 2. Next, the backbone (e.g., VGG-19 backbone) can extract features from y , and then the features are fed into the two branches of the architecture. Certain embodiments of the system can split the confidence maps $\mathcal{S}$ and PAFs $\mathcal{V}$ into body-related and face-related features as described, and compute the loss described in Eq. 15. In many embodiments, after computing L , the system can use the automatic differentiation capabilities to back-propagate the error and update the parameters of the backbone (e.g., VGG-19 backbone) and the coefficients α_j that model surface profile ϕ of the lens using Eq. 10.

[0086] In many embodiments, the end-to-end model can be trained using an optimizer, batch size, and learning rate (e.g., the Adam optimizer with a batch size of 22 and an initial learning rate of 2×10^{-5}). In certain embodiments, an exponential learning rate decay can be applied with a particular decay factor (e.g., decay factor of 0.666) that is triggered after a certain number of training steps (e.g., 15K, 20K, 25K, 28K, and 35K training steps). In certain embodiments, the network can be trained for a number of steps (e.g., trained

the network for 50K steps (gradient updates) as appropriate to the requirements of specific applications in accordance with embodiments of the invention.

Different Formulations for Privacy-Preserving Loss L_P .

[0087] Certain embodiments of the system can define a privacy-preserving loss L_P as:

[0088]
$$L_{P_1} = -\beta_3 (\sum_{\tau=1}^{\Gamma_1} F_{\tau}(\mathcal{V}_F) - \sum_{\tau=\Gamma_1+1}^{\Gamma_1+\Gamma_2} F_{\tau}(\mathcal{S}_F)) - \beta_2 L_E \quad (16)$$

[0089]
$$L_{P_2} = -\sum_{\tau=1}^{\Gamma_1} F_{\tau}(\mathcal{V}_F) - \sum_{\tau=\Gamma_1+1}^{\Gamma_1+\Gamma_2} F_{\tau}(\mathcal{S}_F) - \beta_4 L_F, \quad (17)$$

[0090] where $\beta_3 > 1$,

[0091]
$$2pt2pt L_F = Sim_{cos}(a_f(\mathbf{x}), a_f(\mathbf{y})), \quad (18)$$

[0092] im_{cos} denotes the cosine similarity, and $a_f(\cdot)$ stand for the model (e.g., ArcFace model). To compute L_F , certain embodiments can use a pretrained model (e.g., ArcFace model) on faces extracted from an input image $\mathbf{x}$ and distorted image $\mathbf{y}$.

Compute L_F in Eq. 18

[0093] Certain embodiments can use one or more different loss functions. Two additional loss function approaches in accordance with certain embodiments of the invention include: L_{P_1} , and L_{P_2} , shown in Eq. 16 and Eq. 17. In particular, for L_{P_2} certain embodiments can compute L_F loss in Eq. 18, which can measure face similarity/dissimilarity between original or input image $\mathbf{x}$ and “private” image $\mathbf{y}$. To extract the face regions from the images, certain embodiments can use an available detector (e.g., the RetinaFace detector) with a particular backbone (e.g., ResNet50 backbone). Certain embodiments of the system can generate multiple files including face labels for a particular dataset (e.g., COCO 2017 dataset) for the training, validation, and testing sets, respectively. A file can include lines with the image filename and coordinates $(x_1, y_1), (x_2, y_2)$ specifying the upper-left and lower-right corners of the face rectangle. In certain embodiments, coordinates can be presented as floating-point numbers in the range $[0, 1]$ relative to the specific image’s width and height. Then, with these face region annotations, faces can be cropped from both $\mathbf{x}$ and $\mathbf{y}$ and extract the embeddings using a model (e.g., ArcFace model) with pretrained weights loaded. The L_F can be obtained

by comparing both embedding vectors using the cosine similarity. Although particular loss functions are described that can measure face similarity/dissimilarity, any of a variety of loss functions can be utilized including for other privacy-preserving attributes as appropriate to the requirements of specific applications in accordance with embodiments of the invention.

Lightweight OpenPose

[0094] In many embodiments of the system, different human pose estimation (HPE) networks can be adopted as the CNN-decoder. Certain embodiments of the system can use a single branch for PAF and keypoints predictions instead of the two branches of OPPS; and replacing expensive convolutions (e.g., 7×7 convolutions) with less expensive processing (e.g., 3×3 , 1×1 , and 3×3 with dilation of 2 convolutions blocks). In certain embodiments, The backbone (e.g., VGG-19 backbone) network of OPPS can be replaced with other networks (e.g., MobileNet family networks). The other networks (e.g., MobileNets) can be built primarily from depthwise separable convolutions to reduce the computation in the first few layers. Certain embodiments can use the well-known ResNet-50 network as a replacement for the VGG-19 backbone.

Training Details

[0095] Certain embodiments of the system may assume an aberration-free freeform lens and use pretrained weights of a particular implementation (e.g., Tensorflow implementation of LOPPS) as a starting point. In certain embodiments, once initializing the network with the previously learned weights, the single branch of LOPPS can be frozen and the first 68 trainable layers of the Resnet backbone fine-tuned to learn to extract human body features from the privacy image.

[0096] Certain embodiments of the system can simulate a sensor with a particular pixel size (e.g., pixel size of $3.40\mu\text{m}$) and a particular resolution (e.g., resolution of 864×864 pixels). Certain embodiments of the system can consider various coefficients (e.g., the first $q = 350$ Zernike coefficients in Noll notation) to shape the surface profile ϕ . In certain embodiments, a defocus coefficient (e.g., fourth Zernike coefficient) can be initialized, such that the lens has a particular focal length (e.g., focal length of $f = 25\text{mm}$).

In many embodiments, the optical element can be discretized with a particular feature size (e.g., $3.40\mu\text{m}$ feature size) on a particular grid (e.g., 864×864 grid). In many embodiments, the system can be trained using a particular optimizer with a particular batch size (e.g., batch size of 24) and a particular initial learning rate (e.g., initial learning rate of 4×10^{-5}). Certain embodiments can apply an exponential learning rate decay with a particular decay factor (e.g., decay factor of 0.666) that is triggered after certain training steps (e.g., 15K, 20K, 25K, and 28K training steps). Certain embodiments of the system can train the network for a particular number of steps (e.g., 50K steps (gradient updates)).

[0097] A privacy-preserving system for machine vision in accordance with an embodiment of the invention is illustrated in Fig. 4. The system 400 can include one or more cameras 405 that capture images of a surrounding environment, a network interface 410 that communicates with an external device, a processor, and memory storing one or more different applications. The applications can include an optical encoder 420, one or more machine vision applications 425 such as human pose estimation (HPE) application, and a software decoder 425 such as a Convolutional Neural Network. In many embodiments, the optical encoder 420 is parametrized for a particular camera lens of the camera 405 in order to generate visually degraded images that protect privacy with respect to one or more attributes (e.g., face, age, race, gender, among others). For example, the optical encoder 420 can acquire images and produce images that obscure the identity of a person but still preserve important features for pose estimation.

[0098] The software decoder 425 can perform one or more different machine vision tasks, such as human pose estimation on optically degraded images provided by the optical encoder 420. Although Fig. 4 illustrates a particular privacy-preserving system for generating visually degraded images and performing machine vision on them, any of a variety of architectures can be utilized as appropriate to the requirements of specific applications in accordance with embodiments of the invention.

[0099] Although specific systems and methods for privacy-preserving optics are discussed above with respect to Figs. 1-4, many different systems and methods can be implemented for a variety of different machine vision tasks in accordance with many different embodiments of the invention. It is therefore to be understood that the present invention may be practiced in ways other than specifically described, without departing

from the scope and spirit of the present invention. Thus, embodiments of the present invention should be considered in all respects as illustrative and not restrictive. Accordingly, the scope of the invention should be determined not by the embodiments illustrated, but by the appended claims and their equivalents. Additional disclosure can be found in the appendices below.

WHAT IS CLAIMED IS:

1. A system, comprising:
 - a camera comprising an encoding optical element and at least one lens;
 - at least one processor; and
 - memory comprising an image processing pipeline application;wherein the image processing pipeline application directs the at least one processor to:
 - generate an optimal set of parameters to parametrize the encoding optical element in the camera to produce optical distortions such that acquired images by the camera are distorted, wherein the optimal set of parameters is learned via end-to-end learning that jointly optimizes from a camera optics model to a computational process that performs a computer vision task on distorted images acquired by the camera, wherein the distorted images visually obscure a privacy attribute of people to protect their privacy but still preserve features to perform the computer vision task;
 - acquire a plurality of distorted images using the camera; and
 - perform the computer vision task directly on the plurality of distorted images, wherein the distortions generated by the camera are optimal and allow obtaining high performance on the computer vision task.
2. The system of claim 1, wherein the image processing pipeline application comprises a deep neural network (DNN), wherein the DNN is trained to detect a plurality of features directly on the distorted images, and wherein the DNN comprises a plurality of layers for detecting the plurality of features based on a plurality of nodes that perform a calculation on degraded input image data.
3. The system of claim 1, wherein the encoding optical element comprises an ensemble of optical elements that generate image distortion via light modulation.

4. The system of claim 3, wherein the distortion produced by the encoding optical element is optimal for a specific computer vision task and allows obtaining high performance.

5. The system of claim 1, where functionality of the camera, including the encoding optical element, is emulated in software to generate synthetic degraded images.

6. The system of claim 5, where the camera is coupled with the image processing pipeline application to emulate the system.

7. The system of claim 1, where the set of parameters for the encoding optical element add optical aberrations to the camera lens.

8. The system of claim 1, wherein the distorted images obscure at least one privacy-preserving human attribute selected from a group consisting of, for example, facial features, gender, race, and age.

9. The system of claim 1, wherein the distorted images are robust to blind and non-blind deconvolution attacks.

10. The system of claim 1, wherein performing a computer vision task comprises performing human pose estimation (HPE) on the distorted images.

11. The system of claim 1, wherein the image processing pipeline application comprises a computer vision decoder that has been jointly trained to generate the set of parameters of the camera encoding optical element α^* and a convolutional neural network for human pose estimation h^* such that:

$$\alpha^*, h^* = \arg \min_{\alpha, h} L_T(h) + L_P(\alpha).$$

where L_T is a loss function for the human pose estimation task, and L_P is a loss function that encourages privacy-preservation.

12. The system of claim 11, wherein the set of parameters of the camera encoding optical element comprises a lens surface profile ϕ , in terms of Zernike coefficients α , and a corresponding point spread function (PSF) $\mathbf{H}$ for the camera lens.

13. The system of claim 12, wherein the set of parameters parametrize a lens surface profile ϕ with the Zernike basis, such that:

$$\phi = \sum_{j=1}^q \alpha_j \mathbf{Z}_j,$$

where $\mathbf{Z}_j$ is the j -th Zernike polynomial in Noll notation, and α_j is the corresponding coefficient, wherein each Zernike polynomial describes a wavefront aberration such that the surface profile ϕ is formed by a linear combination of all aberrations.

14. The system of claim 11, wherein the computer vision decoder is a convolutional neural network (CNN) decoder that is trained using a backbone network and a plurality of branches of convolutional layers, wherein the backbone network extracts features from an image of size $w \times h$, which are then fed into the plurality of branches convolutional layers, wherein a first branch predicts a set of confidence maps, where each map represents a specific body part location, and a second branch predicts a set of Part Affinity Fields (PAFs), where each field represents the degree of association between parts.

15. The system of claim 11, wherein:
the optical encoder and the computer vision decoder share a loss function; and
the optical encoder and computer vision decoder are configured to jointly optimize the privacy-preserving optical system and a computer vision model.

16. A method of preserving-privacy on captured images while performing a computer vision task, comprising:

generating an optimal set of parameters to parameterize an encoding optical element of a camera to produce optical distortions such that acquired images by the camera are distorted, wherein the optimal set of parameters is learned via end-to-end learning that jointly optimizes from a camera optics model to a computational process that

performs a computer vision task on distorted images acquired by the camera, wherein the distorted images visually obscure a privacy attribute of people to protect their privacy but still preserve features to perform the computer vision task;

acquiring a plurality of distorted images using the camera; and

performing a computer vision task directly on the plurality of distorted images wherein the distortions generated by the camera are optimal and allow obtaining high performance on the computer vision task.

17. The method of claim 16, further comprising using a deep neural network (DNN), wherein the DNN is trained to detect a plurality of features directly on the distorted images, and wherein the DNN comprises a plurality of layers for detecting the plurality of features based on a plurality of nodes that perform a calculation on degraded input image data.

18. The method of claim 16, wherein the encoding optical element comprises an ensemble of optical elements that generate image distortion via light modulation.

19. The method of claim 18, wherein the distortion produced by the encoding optical element is optimal for a specific computer vision task and allows obtaining high performance.

20. The method of claim 16, wherein the computer vision task is a human pose estimation (HPE) task on the distorted images.




FIG. 1




FIG. 2




FIG. 3A




FIG. 3B




FIG. 4

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US22/73014

A. CLASSIFICATION OF SUBJECT MATTER

IPC - INV. G06V 10/30; G06V 10/72; G06V 10/75; G06V 10/77; G06V 10/80 (2022.01)

ADD. G06T 7/80; G06V 10/24; G06V 10/40; G06V 10/42; G06V 40/16; G06V 40/19 (2022.01)

CPC - INV. G06V 10/30; G06V 10/72; G06V 10/754; G06V 10/772; G06V 10/80

ADD. G06K 7/10732; G06T 7/75; G06T 7/80; G06T 7/97; G06V 10/243; G06V 10/40; G06V 10/42; G06V 40/16; G06V 40/193

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)

See Search History document

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

See Search History document

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

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 10,535,120 B2 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 14 January 2020; Entire document	1-20
A	US 2021/0076966 A1 (SURGICAL SAFETY TECHNOLOGIES INC.) 18 March 2021; Entire document	1-20
A	WO 2017/141102 A1 (6115187 CANADA, D/B/A IMMERSION, INC.) 24 August 2017; Entire document	1-20

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:

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

"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

23 August 2022 (23.08.2022)

Date of mailing of the international search report

SEP 06 2022

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450
Facsimile No. 571-273-8300

Authorized officer

Shane Thomas

Telephone No. PCT Helpdesk: 571-272-4300