



(12) **DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION**

(13) **A1**

(86) Date de dépôt PCT/PCT Filing Date: 2021/03/05
 (87) Date publication PCT/PCT Publication Date: 2021/09/10
 (85) Entrée phase nationale/National Entry: 2022/09/06
 (86) N° demande PCT/PCT Application No.: US 2021/021227
 (87) N° publication PCT/PCT Publication No.: 2021/178911
 (30) Priorité/Priority: 2020/03/06 (US62/986,552)

(51) Cl.Int./Int.Cl. *G06N 20/00* (2019.01)
 (71) Demandeur/Applicant:
 THE REGENTS OF THE UNIVERSITY OF CALIFORNIA,
 US
 (72) Inventeurs/Inventors:
 MIRESHGHALLAH, FATEMEHSADAT, US;
 ESMAEILZADEH, HADI, US
 (74) Agent: SMART & BIGGAR LLP

(54) Titre : PROCÉDES DE FOURNITURE DE CONFIDENTIALITE DE DONNEES POUR INFERENCE BASEE SUR UN RESEAU DE NEURONES ARTIFICIELS
 (54) Title: METHODS OF PROVIDING DATA PRIVACY FOR NEURAL NETWORK BASED INFERENCE

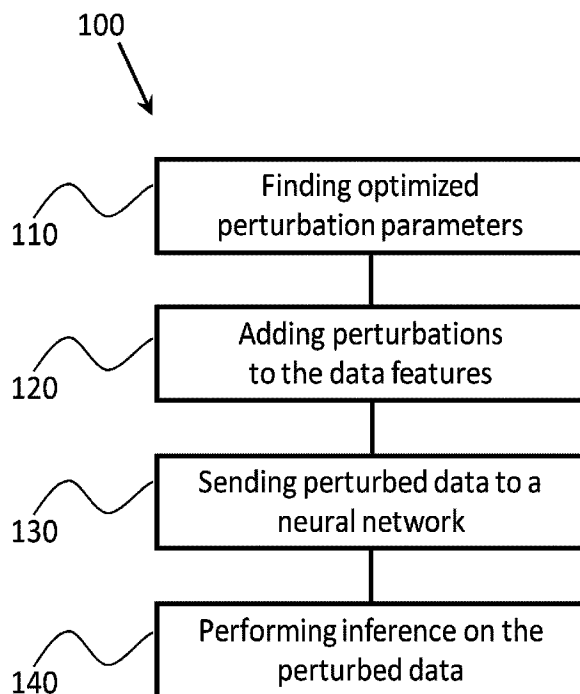


FIG. 1

(57) **Abrégé/Abstract:**

Methods and systems that provide data privacy for implementing a neural network-based inference are described. A method includes injecting stochasticity into the data to produce perturbed data, wherein the injected stochasticity satisfies an ϵ -differential privacy criterion and transmitting the perturbed data to a neural network or to a partition of the neural network for inference.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau



(10) International Publication Number
WO 2021/178911 A1

(43) International Publication Date
10 September 2021 (10.09.2021)

- (51) International Patent Classification:
G06N 20/00 (2019.01)
- (21) International Application Number:
PCT/US2021/021227
- (22) International Filing Date:
05 March 2021 (05.03.2021)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/986,552 06 March 2020 (06.03.2020) US
- (71) Applicant: THE REGENTS OF THE UNIVERSITY OF CALIFORNIA [US/US]; 1111 Franklin Street, 5th Floor, Oakland, California 94607-5200 (US).
- (72) Inventors: MIRESHGHALLAH, Fatemehsadat; 9500 Gilman Drive, La Jolla, California 92093 (US). ES-MAEILZADEH, Hadi; 9500 Gilman Drive, La Jolla, California 92093 (US).
- (74) Agent: SATHE, Vinay et al.; Perkins Coie LLP, P.O. Box 1247, Seattle, Washington 98111-1247 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, 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, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

(54) Title: METHODS OF PROVIDING DATA PRIVACY FOR NEURAL NETWORK BASED INFERENCE

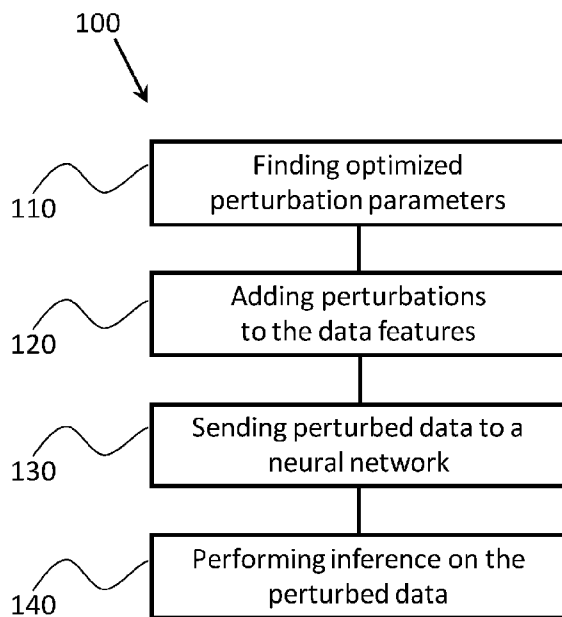


FIG. 1

(57) Abstract: Methods and systems that provide data privacy for implementing a neural network-based inference are described. A method includes injecting stochasticity into the data to produce perturbed data, wherein the injected stochasticity satisfies an ϵ -differential privacy criterion and transmitting the perturbed data to a neural network or to a partition of the neural network for inference.

WO 2021/178911 A1

WO 2021/178911 A1 

Published:

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

METHODS OF PROVIDING DATA PRIVACY FOR NEURAL NETWORK BASED INFERENCE

[0001] This patent document claims priority to and benefits of U.S. Provisional Application
5 No. 62/986,552 entitled “METHODS OF PROVIDING DATA PRIVACY FOR NEURAL NETWORK BASED INFERENCE” filed on March 6, 2020. The entire content of the before-mentioned patent applications is incorporated by reference as part of the disclosure of this document.

10 STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR DEVELOPMENT

[0002] This invention was made with government support under CNS-1703812 and ECCS-1609823 awarded by the National Science Foundation. The government has certain rights in the invention.

15 TECHNICAL FIELD

[0003] This patent document relates to computer technologies including machine learning techniques.

BACKGROUND

[0004] Artificial neural networks (ANN) are computing systems which learn to perform
20 tasks using examples and generally without being pre-programmed with task-specific rules. A deep neural network (DNN) is an artificial neural network with multiple layers of artificial neurons between an input layer and an output layer. DNNs can learn linear and non-linear relationships between their inputs and outputs.

SUMMARY

25 [0005] The techniques disclosed herein can be implemented in various embodiments to achieve, among other features and benefits, finding optimal stochastic perturbations to obfuscate

features of the private data before it is sent to a neural network for inference.

[0006] Methods and systems which allow finding optimal stochastic perturbations to obfuscate features of the private data before the data is sent to a neural network which performs an inference task on the data are described. To this end, the methods according to the disclosed technology may be used by embodiments to reduce the information content of the transmitted perturbed data relative to the unperturbed data while conserving essential pieces of the data that enable an inference request to be serviced accurately by the neural network

[0007] One aspect of the disclosed embodiments relates to a method of providing data privacy that includes specifying a value of a parameter λ related to trade-off between an accuracy of inference done by a neural network (NN) on a perturbed input of the NN and a degree of mutual information degradation between a raw input of the NN and the perturbed input of the NN. The method further includes, for each feature of a data instance, providing a Laplace distribution corresponding to the feature, wherein the Laplace distribution has parameters of location M and scale B . The method also includes forming a tensor of locations M_T and a tensor of scales B_T using the parameters of the Laplace distributions provided for the features of the data instance. The method further includes providing a loss function L having a term proportional to $\log(B)$, wherein tensor B is related to the tensor of scales B_T , and a term proportional to a product of the value of the parameter λ and a utility loss function L_{nn} of the neural network, wherein the utility loss function L_{nn} can be used to train the neural network to perform an inference task T . The method also includes finding optimized elements of B_T and M_T by optimizing, using the loss function L , accuracy of the inference task T performed by the neural network. The method further includes, for each feature of the data instance, determining values of optimized parameters of location M and scale B of the Laplace distribution corresponding to the feature using the optimized elements of B_T and M_T .

[0008] Another aspect of the disclosed embodiments relates to a method of providing data privacy for neural network computations that includes, for a given privacy budget ϵ and for a target inference task to be performed by a neural network, finding optimized parameters of a set of statistical distributions which optimize performance of the neural network on the inference task and with respect to the privacy budget ϵ using a loss function L , wherein the loss function L has a term related to a tensor, wherein the tensor is related to a parameter of at least one

distribution from the set of statistical distributions and another term related to a utility loss function L_{nn} of the neural network for the target inference task. The method further includes selecting, for each feature of an input, a perturbation value drawn from a distribution in the set of statistical distributions which has the optimized parameters and adding the perturbation value to a value associated with the feature to obtain a perturbed input. The method also includes sending the perturbed input to the neural network. The method further includes performing the target inference task by the neural network on the perturbed input.

[0009] Yet another aspect of the disclosed embodiments relates to a data processing method that includes determining parameters of a distortion operation by which a source data is converted into a distorted data. The method also includes performing a data processing task on the distorted data.

[0010] An aspect of the disclosed embodiments relates to a method of providing privacy for data that includes injecting stochasticity into the data to produce perturbed data, wherein the injected stochasticity satisfies an ϵ -differential privacy criterion. The method further includes transmitting the perturbed data to a neural network or to a partition of the neural network for inference.

[0011] Another aspect of the disclosed embodiments relates to a method of providing privacy for data that includes determining, for a pre-trained deep neural network (DNN), an amount of stochastic perturbation, applying the amount of the stochastic perturbation to source data to obtain perturbed source data, and transmitting the perturbed source data to the DNN or to a partition of the DNN.

[0012] Yet another aspect of the disclosed embodiments relates to a method of providing privacy for data that includes determining an amount of stochastic perturbation in a source data without accessing sensitive information or labels associated with the source data. The method further includes transmitting, to a neural network or to a partition of the neural network, a perturbed data obtained by perturbing the source data using the amount of stochastic perturbation.

[0013] An aspect of the disclosed embodiments relates to a communication apparatus that includes a memory and a processor, wherein the processor is configured to read code from the memory and implement a method according to the technology disclosed in this patent document.

[0014] Another aspect of the disclosed embodiments relates to a non-transitory computer-readable medium storing instructions that, when executed by a computer, cause the computer to perform a method according to the technology disclosed in this patent document.

BRIEF DESCRIPTION OF THE DRAWINGS

5 [0015] FIG. 1 shows a flow diagram of an example embodiment of a method according to the technology disclosed in this patent document.

[0016] FIG. 2 shows a flow diagram of an example embodiment of another method according to the technology disclosed in this patent document.

10 [0017] FIG. 3 shows an example system which can be used to implement one or more methods according to the technology disclosed in this patent document.

[0018] FIG. 4 shows another example system which can be used to implement one or more methods according to the technology disclosed in this patent document.

[0019] FIG. 5 shows a flow diagram of an example embodiment of a data processing method according to the disclosed technology.

15 [0020] FIG. 6 shows a flow diagram of an example embodiment of a method of providing privacy for data according to the technology disclosed in this patent document.

[0021] FIG. 7 shows a flow diagram of another example embodiment of a method of providing privacy for data according to the technology disclosed herein.

20 [0022] FIG. 8 shows a flow diagram of yet another example embodiment of a method of providing privacy for data according to the technology disclosed in this patent document.

[0023] FIG. 9 illustrates effects of an example method according to the disclosed technology on input images as the average scale of perturbations is increased.

[0024] FIGS. 10A-10D illustrate accuracy loss vs. the privacy budget (ϵ) for four benchmark datasets, measured on the test set of each of the four datasets.

25 [0025] FIGS. 11A-11C illustrate accuracy loss vs. the remaining mutual information for a fixed differential-privacy budget of $\epsilon = 2.5$ for three benchmark datasets.

[0026] FIG. 12 illustrates that methods according to the technology disclosed in this patent

document are protecting the user against malicious use of their private information.

[0027] FIGS. 13A-13D illustrate accuracy loss vs. the privacy budget (ϵ) for a privacy budget of up to 10 for the four benchmark datasets, measured on the test set of each of the four datasets.

5

DETAILED DESCRIPTION

[0028] INFERENCE-as-a-Service (INFaaS) in the cloud has enabled the prevalent use of Deep Neural Networks (DNNs) in home automation, targeted advertising, machine vision, etc. The cloud receives an inference request as an input which can contain a rich set of private information that can be misused or leaked, possibly inadvertently. This prevalent setting can compromise the privacy of users during the inference phase. Therefore, there is a need in data processing and data communication methods that can ensure privacy of the user data without any significant sacrifice of the neural network's performance on inference tasks.

[0029] The technology disclosed in this patent document provides methods and systems which allow finding optimal stochastic perturbations to obfuscate features of the private data before it is sent to the cloud. To this end, the methods according to the disclosed technology reduce the information content of the transmitted perturbed data relative to the unperturbed data while conserving essential pieces of the data that enable an inference request to be serviced accurately by a DNN. Methods and systems according to the technology disclosed in this patent document improve computer technology by improving capabilities of computers to ensure privacy of data they process while preserving accuracy of the computations that the computers perform on the data.

[0030] Methods according to the disclosed technology can use gradient-based optimization of a pre-trained DNN (the DNN having known weights optimized for an inference task) to obtain optimal parameters of distributions of the stochastic perturbations used to obfuscate the data features. The obfuscated data is subsequently sent to the neural network which performs the inference task on the data.

[0031] For example, some embodiments of the methods disclosed in this patent document use Laplace distribution as a parametric model for the stochastic perturbations and can obtain optimal parameters for a set of perturbation Laplace distributions using Monte Carlo sampling of

those distributions during a gradient descent-based optimization using a pre-trained neural network. The obtained set of Laplace distributions is used to perturb private data while retaining the neural network's ability to perform the inference task the neural network is pre-trained to perform.

5 [0032] Methods according to the technology disclosed in this patent application also allow guaranteeing that the stochasticity injected into the input data using the methods satisfies the ϵ -differential privacy criterion.

[0033] In some embodiments of the technology disclosed herein, the desired privacy budget ϵ is incorporated as a constraint in a loss function L of a pre-trained neural network which also
10 incorporates parameters of stochastic perturbations (e.g., parameters of a set of Laplace distributions). Optimal values of the parameters of stochastic perturbations are obtained by using L to maximize the accuracy of an inference task performed by the pre-trained neural network with respect to the privacy budget constraint.

[0034] In certain embodiments, methods according to the disclosed technology provide a
15 way to control a tradeoff between the mutual information (MI) between an input of a pre-trained DNN and a perturbed representation of the input and accuracy of the inference task performed by the DNN. In some implementations, this tradeoff is controlled using a Lagrange multiplier incorporated into a loss function of the DNN.

[0035] According to some embodiments, parameters of the distributions of the stochastic
20 perturbations are determined before providing a sensitive input to the pre-trained DNN followed by sampling a random perturbation tensor from the determined distributions, adding the tensor to the sensitive input, and sending thus perturbed input to the DNN for inference.

[0036] In some implementations of the technology disclosed in this patent document, a data
25 instance is considered as a dataset and features of the data instance are considered as records of the dataset. For instance, for images, the data instance can be an image and the features can be the image pixels. Protecting sensitive information in an image could be achieved by adding noise to the image features (e.g., pixels) through an ϵ -differentially private mechanism according to an implementation of the technology disclosed in this patent document, which makes the features of the image less distinguishable.

[0037] FIG. 1 illustrates steps of a method 100 of providing data privacy for neural network based computations according to an embodiment of the technology disclosed in this patent document.

[0038] Step 110 of the method 100 includes, for a given privacy budget ϵ and for a target inference task to be performed by a neural network, finding parameters of a set of statistical distributions which optimize performance of the neural network on the inference task and with respect to the privacy budget ϵ using a loss function L , wherein the loss function L has a term related to a tensor (e.g., a term proportional to a logarithm of the tensor), wherein the tensor is related to a parameter (e.g., a scale) of at least one distribution (e.g., a Laplace distribution) from the set of statistical distributions (for example, elements of the tensor are scales of the Laplace distributions) and another term related to (e.g., proportional to) a utility loss function L_{nn} of the neural network for the target inference task.

[0039] Step 120 of the method 100 includes selecting, for each feature of an input, a perturbation value drawn from a distribution in the set of statistical distributions which has parameters optimized in Step 110 of the method and adding the perturbation value to a value associated with the feature to obtain a perturbed input.

[0040] In Step 130 of the method 100, the perturbed input is sent to the neural network which performs the target inference task on the perturbed input in Step 140 of the method 100.

[0041] FIG. 2 illustrates steps of a method 200 of providing data privacy for neural network based computations according to an embodiment of the technology disclosed in this patent document.

[0042] Step 210 of the method 200 includes specifying a value of a parameter λ related to trade-off between an accuracy of inference done by a neural network on a perturbed input and a degree of mutual information degradation between a raw input and the perturbed input of the neural network.

[0043] The less mutual information exists between the raw and perturbed inputs, the higher the level of data privacy provided through perturbing the raw inputs. However, increasing the privacy level can negatively affect the accuracy of an inference task performed by the neural network on the perturbed inputs.

[0044] The methods disclosed in this patent document, such as the method 200, for example, allow controlling a ratio between the degree of loss of private information in the data achieved through an obfuscation mechanism implemented according to the disclosed technology and the degree of accuracy loss for a target inference task a user wants a neural network to perform on the obfuscated data.

[0045] Step 220 of the method 200 includes providing, for each feature of a data instance, a Laplace distribution corresponding to the feature, wherein the Laplace distribution has parameters of location M and scale B . For example, the data instance can be an image (e.g. a photo or a frame of a video) and pixels of the image can be its features. Methods according to the disclosed technology apply differential privacy to obfuscate personal/sensitive features within a data instance.

[0046] Step 230 of the method 200 includes forming a tensor of locations M_T and a tensor of scales B_T using the parameters of the Laplace distributions provided for the features of the data instance. Tensors M_T and B_T can have the same shape as the data tensor D of the data instance.

[0047] Step 240 of the method 200 includes providing a loss function L having a term proportional to $\log(B)$, wherein tensor B is obtained using the tensor B_T , and a term proportional to the product of the value of the parameter λ and a utility loss function L_{nn} of the neural network for a target inference task T .

[0048] In some embodiments of the methods disclosed in this patent document, the neural network can be pre-trained to perform the inference task T . In some embodiments, the neural network is pre-trained to perform the inference task T using the utility loss function L_{nn} .

[0049] Step 250 of the method 200 includes finding optimized elements of tensors B_T and M_T by optimizing accuracy of the inference task T performed by the neural network using the loss function L .

[0050] In some implementations, Step 250 can include a sub-step 250A of drawing multiple samples of a noise tensor (a perturbation tensor) using the Laplace distributions provided for the data instance features, averaging over the losses associated with the samples of the perturbation tensor, and applying an update using the obtained average. In other implementations, only a single sample of the perturbation tensor is used. Generally, the perturbation tensor can have the

same shape as the data tensor D of the data instance.

[0051] Step 260 of the method 200 includes determining, for each feature of the data instance, values of optimized parameters of location M and scale B of the Laplace distribution corresponding to the feature using the optimized elements of B_T and M_T .

5 **[0052]** Methods according to the disclosed technology can further include a step of selecting, for each feature of a raw input data instance D_R of the neural network, a perturbation value x from the Laplace distribution corresponding to the feature and adding the perturbation value x to a value of the feature to produce a perturbed input data instance D_P .

10 **[0053]** Methods according to the disclosed technology can also include a step of sending the perturbed input data instance D_P to the neural network for inference. In some embodiments, the neural network can reside in a cloud while the methods according to the disclosed technology can be implemented on a user device (e.g., a smartphone). The perturbed data can be sent to the neural network via one or more communication networks (e.g., Internet and/or a cellular network).

15 **[0054]** In some embodiments, the neural network can reside on the same device (e.g., a smartphone, a personal computer, or a computer server) which implements the methods disclosed in this patent application. In such a case, applying the disclosed methods to obfuscate raw data obtained using the device (e.g., using a smartphone camera) or transmitted to the device from an external data source (e.g., sent to a server over a network) can be beneficial to avoid
20 potential breaches of privacy of the data when, for example, storage of the data on the device (in any form suitable for data storage) is performed in connection with processing of the data by the neural network on the device.

[0055] FIG. 3 shows an example embodiment of a system 300 which can be used to implement one or more methods according to the technology disclosed in this patent application.
25 System 300 can include, for example, an edge device 310, a communication network 320, and a cloud-based device 330.

[0056] The edge device 310 can include a memory 311, a processor 312, and a data input/output (I/O) interface 313. The memory 311 of the device 310 stores a code which, when executed by the processor 312 makes the device 310 implement one or more methods according

to the technology disclosed herein. In some embodiments, device 310 can be a smartphone, a laptop, a personal computer, or an embedded computer.

[0057] By implementing the one or more methods according to the technology disclosed in this patent document, the device 310 transforms raw data into perturbed data which is sent by the device 310 via its I/O 313 and over a communication network 320 (e.g., the Internet and/or a wireless network) to a cloud-based device 330 (e.g., a server). The device 330 hosts a neural network 331 which is configured to perform an inference task T. The device 330 receives the perturbed data sent to it by the edge device 310 over the communication network 320 and provides the perturbed data as an input data to the neural network 331 which performs the inference task T on the data. Results of the inference task T performed by the neural network 331 can be sent to the device 310 over the communication network 320 and, for example, displayed to a user of the device 310 using, for example, a display unit 314 of the device 310.

[0058] FIG. 4 shows an example embodiment of a system 400 which can be used to implement one or more methods according to the technology disclosed in this patent application. System 400 can include, for example, an edge device 410, a communication network 420, and a cloud-based device 430.

[0059] The edge device 410 can include a memory 411, a processor 412, and a data input/output (I/O) interface 413. The memory 411 of the device 410 stores a code which, when executed by the processor 412 makes the device 410 implement one or more methods according to the technology disclosed herein. In some embodiments, device 410 can be a smartphone, a laptop, a personal computer, or an embedded computer.

[0060] By implementing the one or more methods according to the technology disclosed in this patent document, the edge device 410 transforms raw data into perturbed data. The perturbed data or a part of the perturbed data can be provided to Partition 1 of a neural network which resides in the processor 412 of the device 410 through a data bus and/or an application programming interface (API) call. The perturbed data or a part of the perturbed data can also be provided to Partition 2 of the neural network which resides at a remote location (e.g., a remote cloud-based server) through I/O 413 of the device 410 and over the communication network 420.

[0061] According to the embodiments illustrated in FIG. 3 and FIG. 4, the disclosed technology may be implemented in systems in which the neural network may be implemented

entirely on a processor or a processor cluster that determines the perturbation, or the neural network may be partitioned into a first partition that is executed on the processor / processor cluster and a second partition that is remotely located using remote computational resources such as a cloud based computing service, or a combination of the two.

5 **[0062]** FIG. 5 shows a flow diagram of an example embodiment of a data processing method 500 according to the disclosed technology. Step 510 of the method 500 includes determining parameters of a distortion operation by which a source data is converted into a distorted data. Step 520 of the method 500 includes performing a data processing task on the distorted data.

10 **[0063]** FIG. 6 shows a flow diagram of an example embodiment of a method 600 of providing privacy for data according to the technology disclosed in this patent document. Step 610 of the method 600 includes injecting stochasticity into the data to produce perturbed data, wherein the injected stochasticity satisfies an ϵ -differential privacy criterion. Step 620 of the method 600 includes transmitting the perturbed data to a neural network or to a partition of the neural network for inference.

15 **[0064]** FIG. 7 shows a flow diagram of an example embodiment of a method 700 of providing privacy for data according to the technology disclosed herein. Step 710 of the method 700 includes determining, for a pre-trained deep neural network (DNN), an amount of stochastic perturbation. Step 720 of the method 700 includes applying the amount of the stochastic perturbation to source data to obtain perturbed source data. Step 730 of the method 700 includes
20 transmitting the perturbed source data to the DNN or to a partition of the DNN.

[0065] FIG. 8 shows a flow diagram of an example embodiment of a method 800 of providing privacy for data according to the technology disclosed in this patent document. Step 810 of the method 800 includes determining, an amount of stochastic perturbation in a source data without accessing sensitive information or labels associated with the source data. Step 820
25 of the method 800 includes transmitting, to a neural network or to a partition of the neural network, a perturbed data obtained by perturbing the source data using the amount of stochastic perturbation.

[0066] As mentined above, the success of deep learning in many areas including vision, recommendation systems, natural language processing, etc., has heralded the adoption of Deep
30 Neural Networks (DNNs) in production systems. However, the computational complexity of

DNNs has pushed their execution to mostly cloud infrastructure, where an edge device on the user side captures and sends the raw inputs (requests) to the cloud for inference. This execution model, called INFERENCE-as-a-Service (INFaaS), has become de-facto (e.g., mobile phones), yet it poses serious privacy concerns.

5 [0067] The threat is that as soon as the raw data is sent to the cloud, it can be misused or leaked through security vulnerabilities even if the cloud provider and/or the communication link is trusted. Such a risk is present for every single request (input) and is exacerbated by the fact that the raw inputs contains a rich set of information that is not directly relevant to the inference task. The technology disclosed in this patent document (referred to as “Cloak” below) provides a systematic approach towards providing a differentially private inference mechanism by adding
10 perturbations to the primary input (request). An unprincipled addition of perturbations will lead to significant loss in the inference accuracy, putting the service in the position of questionable utility.

[0068] To address these challenges, the technology disclosed in this patent document which
15 we termed “Cloak” formulates the discovery of the perturbation as an *offline gradient-based optimization problem* that reformulates a pre-trained DNN (with optimized known weights) as an analytical function of the stochastic perturbations. Using Laplace distribution as a parametric model for the stochastic perturbations, Cloak learns the optimal parameters using *gradient descent* and Monte Carlo sampling. With this setup, these *learned* distributions can obfuscate
20 the private data while retaining the cloud's ability to perform the inference task. This set of optimized Laplace distributions further guarantee that the learned perturbations meet the ϵ -differential privacy criterion. To meet ϵ (the desired privacy budget), we incorporate it as a hard constraint in Cloak's loss function to maximize the accuracy with respect to this constraint. Cloak takes a further step after achieving the highest accuracy for a given ϵ and explores an
25 alternative space by reducing Mutual Information (MI) between the original input and its perturbed representation. This approach opens a trade-off between accuracy and information content of the perturbed input that is controlled using a Lagrange multiplier that acts as a knob.

[0069] It worth emphasizing that Cloak's learning process is offline and is not invoked during inference. For each inference request, a distinct random perturbation tensor is sampled
30 from the learned distributions and is added to the raw input.

[0070] Cloak offers these capabilities with a stark contrast with recently emerging line of work on using noise for inference privacy that require to know what they are protecting against (i.e., a private label). Cloak, however, does not need this extra labeling and it attempts at removing any excessive information that is not conducive to the main inference task/label. Additionally, unlike these techniques, Cloak does not need to change parameters or the architecture of the pre-trained networks. Instead, it produces a significantly perturbed representation of the input, aiming to only provide just enough information to serve the inference task (FIG. 1). Furthermore, the aforementioned works do not directly learn the perturbations. Finally, Cloak does not impose the prohibitive computational cost of techniques that use homomorphic encryption which can increase the execution time of a neural network by three orders of magnitude.

[0071] FIG. 9 illustrates an effect of an example method according to the disclosed technology on the input images as we increase the average scale of perturbations \bar{B} that is proportional to the standard deviation. The non-sensitive classification task is smile detection. The first row of the image in FIG. 9 shows the heat map of perturbation scales for each pixel. At higher scales, the example method according to the disclosed technology obfuscates the features that are non-conducive to the smile detection task e.g. the background, the hair, etc., while only the lips and some minor facial attributes remain recognizable.

[0072] We provide an analytical formulation of the stochasticity learning problem as a constrained convex optimization program that maximizes the privacy by minimizing the mutual information between the raw input and the data sent to the cloud subject to a restriction on the degradation of the DNN utility (accuracy). Cloak's objective function and the constraint both depend on the prior distribution of the perturbation. However, for a fixed parameterized family of distributions and a fixed value of the Lagrange multiplier knob, the problem has a unique global optima. The convexity of the formulation guarantees that the gradient descent will converge to that global optima. This proves that Cloak maximizes the privacy under the utility preserving constraint. Further, we prove Cloak satisfies differential privacy guarantees with respect to the features of the input. The per-feature guarantee ensures that an adversary cannot use the spatio-temporal correlations to reverse engineer the process and gain access to the sensitive information.

[0073] Experimental evaluation with real-world datasets of UTK-Face, CIFAR-100, and MNIST shows that Cloak can reduce the mutual information content of input images by 80.07% with accuracy loss of 7.12% for an ϵ of 2.5. The large dimensions of images in the CelebA dataset prevent the accurate estimation of mutual information. As such, we take a practical
 5 approach towards evaluation of this dataset. First, we visualize the effects of Cloak pictorially on FIG. 9. Then, we select the inference service as "smile detection" on this dataset. The results show that Cloak reduces the accuracy of a malevolent DNN classifier that aims to infer "gender," from 96.7% to 65.7%—reduction of 31.0%. This significant reduction is achieved while Cloak does not assume or incorporate anything about the label, feature, or sensitivity (privacy) of
 10 "gender." Additionally, Cloak preserves the accuracy of "smile detection" service at 87.2%—only 4.9% accuracy degradation.

[0074] We define function f , to be the target function the output of which is to be perturbed. This could be any pre-processing (e.g., normalization) carried out on the raw input data, before being handed to the untrusted party. In particular, this function can be the identity
 15 function if there is no pre-processing steps. We also define function g to be the neural network computation which is to be executed by the untrusted party.

[0075] Definition 1. ϵ -Differential Privacy (ϵ -DP). For $\epsilon \geq 0$, an algorithm A satisfies ϵ -DP if and only if for any pair of datasets D and D' that differ in only one element:

$$P[A(D) = t] \leq e^\epsilon P[A(D') = t] \forall t \quad (1)$$

20 where, $P[A(D) = t]$ denotes the probability that the algorithm A outputs t . DP tries to approximate the effect of an individual opting out of contributing to the dataset, by ensuring that any effect due to the inclusion of one's data is small.

[0076] Definition 2. Laplace Mechanism. Given a target function f and a fixed $\epsilon \geq 0$, the randomizing algorithm $A_f(D) = f(D) + x$ where x is a perturbation random variable drawn
 25 from a Laplace distribution $Lap\left(\mu, \frac{\Delta f}{\epsilon}\right)$, is called the Laplace Mechanism and is ϵ -DP. Here, Δf is the global sensitivity of function f , and is defined as $\Delta f = \sup |f(D) - f(D')|$ over the all dataset pairs (D, D') that differ in only one element.

[0077] The amount of mutual information between the raw data and the representation that is to be publicized is another measure of privacy that is widely used in literature. Cloak aims to minimize the mutual information between the raw input (D) and its perturbed representation, $A_f(D)$, that was attained via the randomization algorithm A_f . We can bound this mutual information by

$$\begin{aligned}
I(D; A_f(D)) &\leq I(f(D); A_f(D)) \\
&= H(A_f(D)) - H(A_f(D)|f(D)) \\
&= H(A_f(D)) - H(x) \\
&\leq H(f(D)) - H(x) \\
&= H(f(D)) - \log(2be)
\end{aligned} \tag{2}$$

where, $H(\cdot)$ represents the Shannon entropy. Here, we assume x is distributed as $Lap(\mu, b)$, i.e., b is the scale of the Laplace distribution. The first inequality follows from the data processing inequality. The first term in the last line of Equation (2) depends on the target function f and is a constant with respect to the optimization variable x in our case. Hence, we do not consider it for our optimization. The second term can be written as $-\log(2be) = -\log(2e) - \log(b)$, where $-\log(2e)$ is constant. Thus, if we minimize $-\log(b)$, i.e., maximize the scale of the distribution of perturbations, we have minimized an upper bound on the mutual information of the leaked data.

[0078] During the inference phase, for each input data tensor (D), Cloak adds a distinct randomly generated perturbation tensor (X). This addition yields a perturbed representation of the input that is sent out to the cloud for classification with the DNN denoted by g . During the inference phase no optimization is performed. X is generated by sampling each element, x , independently, from a Laplace distribution $Lap(\mu, b)$, where μ and b are the corresponding elements of the locations (M) and scales tensors (B).

[0079] In a separate *offline* process, Cloak finds these M and B tensors by solving an optimization problem. Cloak trains the locations and scales tensors in such a way that: (1)

provides the required differential privacy guarantee, as defined shortly; (2) incurs minimum degradation to the neural network's objective, and, (3) decreases the amount of mutual information between the original tensor D and its noisy, perturbed version, $A_f(D)$. When the desired tensors M and B are found, they are used for perturbing input data for inference, in
 5 deployed neural networks as mentioned above.

[0080] The goal of differential privacy is to guarantee that one cannot infer whether an individual data instance was in the dataset by observing an algorithm's output. In Cloak's problem setting, we apply differential privacy to obfuscate personal/sensitive features, within each data instance. This problem can be looked upon as trying to make the changes in one
 10 feature within the data instance indistinguishable. As such, we define an instantiation of differential privacy, where a data instance is considered as a dataset and features are considered records of this dataset, as *feature differential privacy*. For instance, in images, the features can be the pixels, and one attempt at protecting sensitive information in the image could be adding noise through a differentially private mechanism, so as to make pixels less distinguishable. In
 15 this case, the datasets D and D' can be two identical images that only differ in one pixel.

[0081] Cloak aims at casting the noise distribution parameters as a trainable tensor, and using conventional gradient based methods to train them. To be able to define gradients over the locations and scales, we rewrite the noise sampling to be $X = B \circ E + M$, instead of $X \sim Lap(M, B)$, where E is a tensor with the same shape as the data tensor (D) and is
 20 sampled i.i.d from $Lap(0,1)$ and \circ is the element-wise multiplication. This redefinition enables us to formulate the problem as an analytical function for which we can calculate the gradients.

[0082] To achieve the differential privacy guarantee, the elements of B must be larger than $\frac{\Delta f}{\epsilon}$. We define a to-be-tuned hyper-parameter M_x as an upper bound and reparameterize the scales tensor B as P :

$$B = \frac{1.0 + \tanh(P)}{2} \left(M_x - \frac{\Delta f}{\epsilon} \right) + \frac{\Delta f}{\epsilon} \quad (3)$$

[0083] As discussed experimentally later, M_x is usually between 1.5 to 2.5. We put this extra constraint on the distributions, so that none of the noise elements gets a very high scale and becomes an outlier compared to the scale of other noise elements.

[0084] Differential privacy offers a worst-case privacy guarantee. Cloak, however, decreases the mutual information between the input data and its perturbed representation, while maintaining the desired differential privacy guarantee. The reduction in mutual information can be considered as an average-case privacy measure.

5 [0085] As discussed, the mutual information between tensors D and $A_f(D)$ where the latter is acquired by injecting noise to $f(D)$, is proportional to the log of the standard deviation of the noise that is added to each feature (element) in $f(D)$. Therefore, Cloak needs to maximize the scales of the noise distributions to minimize the leaked mutual information.

10 [0086] Increasing the scale parameters of noise elements through training must be done in consideration of the neural network's objective. For instance, in a setting where the task is smile detection, some facial features like the mouth are conducive to this detection. Other features, however, such as that person's hair, the background, their eye color, etc. are irrelevant. Thus, adding more noise to these features can achieve higher privacy in terms of information loss, while incurring minimum degradation to the utility.

15 [0087] An example of noise injection can be seen in FIG. 9, where the effect of increase in the average of the noise scales is illustrated for two sample images from the CelebA dataset. In this illustration, Cloak causes non-conducive features to get blurry by increasing the scales. For the rightmost images, only the smile and some minor facial attributes are present. For instance, notice the second example image where the other person in the background fades away as the
20 noise increases. The mean accuracy drop over the test portion of the dataset for the four perturbation distribution sets (four sets of tensors B and M) in this Figure are 3.4%, 3.7%, 4.5% and 12.7% respectively, with standard deviations in the 0.12%–0.68% range. The privacy budget (ϵ) is 2.5.

25 [0088] Such a rather targeted (non-uniform) obfuscation behavior that leads to low degradation in accuracy for a given ϵ is the result of our formulation that incorporates mutual information in Cloak's loss function. If only the differential privacy constraint was considered, the noise would have just been increased uniformly across the image and accuracy loss would have been much higher.

[0089] Cloak's objective function is to minimize the mutual information (or maximize the

noise variance as discussed before) subject to a bounded degradation in the neural network's utility. Using Lagrange multiplier λ and given our reparameterization of \mathbf{B} , we can express the loss function of Cloak as

$$L(\mathbf{M}, \mathbf{B}) = -\log \mathbf{B} + \lambda L_{nn}(\mathbf{M}, \mathbf{B}) \quad (4)$$

5 where, L_{nn} represents the utility loss function of the neural network. For instance, in the case of C -class classification, this utility can be represented as $L_{nn}(\mathbf{M}, \mathbf{B}) = -\frac{1}{n} \sum_{i=1}^n \sum_{c=1}^C y_c^i \log(p_c^i)$, which is the cross entropy loss, where p_c^i is the i^{th} Monte Carlo sample of the probability the neural network has assigned to observation i belonging to class c , and, y_c^i is the indicator function of whether the i^{th} sample belongs to the class c . Converting the averaging over n 10 samples into the population expectation, we can rewrite our loss function as

$$L(\mathbf{M}, \mathbf{B}) = -\log \mathbf{B} + \lambda E_{\mathbf{X} \sim L(\mathbf{M}, \mathbf{B})} \left[\sum_{c=1}^C y_c \log \left(g \left(f(\mathbf{D}) + \mathbf{X} \right)_c \right) \right] \quad (5)$$

[0090] We take sufficiently large n noise samples drawn from the current locations and scales tensors to approximate the second term. This means that to apply a single update to the trainable parameters, Cloak runs multiple forward passes on the entire neural network, at each 15 pass draws new samples for the noise (perturbation) tensor, and averages over the losses and applies the update using the average. However, in practice, if mini-batch training is used, we observed that using only a single noise sample for each update can yield desirable results, since a new noise tensor is sampled for each mini-batch.

Algorithm 1 Cloak's Perturbation Training Workflow

Input: $\mathbf{D}, y, f, \epsilon, \lambda$

Initialize $\mathbf{M} = \mathbf{0}$, $\mathbf{P} = -\infty$ and $M_x > \frac{\Delta_f}{\epsilon}$

repeat

 Sample $\mathbf{E} \sim \text{Lap}(0, 1)$

 Let $\mathbf{B} = \frac{1.0 + \tanh(\mathbf{P})}{2} (M_x - \frac{\Delta_f}{\epsilon}) + \frac{\Delta_f}{\epsilon}$

 Let $\mathbf{X} = \mathbf{B} \circ \mathbf{E} + \mathbf{M}$

 Take gradient step on \mathbf{M} and \mathbf{P} from Equation (4)

until Algorithm converges

Return: \mathbf{M}, \mathbf{B}

[0091] The parameter λ is the knob, which could be tuned to reflect whether we want to trade-off accuracy for information degradation (more privacy), or vice versa. We name the $-\log B$, the mutual information decay term. One important feature of Cloak is that it does not need access to sensitive information/labels. In other words, it does not need to be given specific information about what it is protecting against, it only needs access to the non-sensitive labels of the main task, and it tries to remove any extra information with the trained perturbations.

[0092] Cloak's optimization process can be seen in Algorithm 1. Before beginning the process, the locations tensor (M) is initialized to 0, and tensor P is initialized in such a way that the scales take the $\left(\frac{\Delta f}{\epsilon}\right)$ value. The rest of the parameters of the neural network are set to their pre-trained values. The functions f and g are the target function and neural network function, as defined above.

[0093] Once the training is finished, the optimized locations and scales are saved. During deployment (inference), for each input that is supposed to be given to the inference service, a set of perturbations is sampled from the optimized distributions. Each element in the perturbations tensor X is sampled independently and added to input before it is uploaded for inference.

[0094] Proposition 1. Cloak is ϵ -DP w.r.t. a single feature.

[0095] *Proof.* For a single feature D_i in the dataset D , we have $A_f(D_i) = f(D_i) + X_i$, where X_i is drawn from $Lap(M_i, B_i)$. This is a Laplace Mechanism and hence it is ϵ -DP.

[0096] Proposition 2. Cloak is ϵ -feature differentially private with respect to all features.

[0097] *Proof.* By Proposition 3.1, each single feature i is ϵ_i -DP for some $\epsilon_i \geq 0$. Assume D and D' are two datasets that are different only on j^{th} feature. Since Cloak samples each noise perturbation independently, we have

$$\begin{aligned}
P[A_f(D)=t] &= \prod_i P[A_i(D_i)=t_i] \\
&= P[A_j(D'_j)=t_j] \prod_{i \neq j} P[A_i(D'_i)=t_i] \\
&\leq e^{\epsilon_j} P[A_j(D'_j)=t_j] \prod_{i \neq j} P[A_i(D'_i)=t_i] \\
&\leq e^{\max(\epsilon_i)} P[A(D')=t].
\end{aligned}$$

[0098] Hence, Cloak is ϵ -feature DP with respect to all features.

[0099] To evaluate Cloak, we use four real-world datasets on four DNNs. Namely, we use VGG-16 on CelebA, AlexNet on CIFAR-100, a modified version of VGG-16 model on UTKFace, and LeNet-5 on MNIST.

[00100] We define a set of non-sensitive tasks as inference services over these datasets. Specifically, we use smile detection on CelebA, the 20 super-class classification on CIFAR-100, and gender detection on UTKFace. For MNIST, we use a classifier that detects if the input is greater than five. The pre-trained accuracy of the networks for smile detection, super-class classification, gender detection and greater than five detection are 91.8%, 55.7%, 87.87% and 99.29%. The accuracy numbers reported herein are all on a held-out test set, which has not been seen during training by the neural networks. For Cloak results, since the output is not deterministic, we repeatedly run the inference ten times on the test set with the batch size of one and report the *mean accuracy*. Since the *standard deviation* of the accuracy numbers is small (consistently less than 1.5%) the confidence bars are not visible on the graphs.

[00101] The input image sizes for CelebA, CIFAR-100, UTKFace and MNIST are $224 \times 224 \times 3$, $32 \times 32 \times 3$, $32 \times 32 \times 3$, and 32×32 , respectively. In addition, in our experiments, the inputs are all normalized to 1. In all of the experiments, we add the perturbations directly to the input image and create a noisy representation (similar to those in Figure 1) which we then be given to the neural network. Therefore, the function f defined above is the identity function and the sensitivity Δ_f is 1. The experiments are all carried out using Python 3.6 and PyTorch 1.3.1. We use Adam optimizer for perturbation training. The mutual information numbers reported herein are estimated over the test set using the Shannon Mutual Information estimator provided by the Python ITE toolbox.

[00102] FIGS. 10A-10D shows the mean accuracy loss for a given privacy budget ϵ . Two methods are compared here. First, the *Basic - No Training* method that we consider as the baseline in which perturbation elements (the noise for each pixel) are independently generated from a Laplace distribution with location of zero and scale of $\frac{\Delta f}{\epsilon}$. $\frac{\Delta f}{\epsilon}$ is the minimum scale that the perturbation distributions can have, in order to achieve a certain privacy budget of ϵ for the feature-differential privacy criterion described above.

[00103] For the second method, Cloak is used to observe the effect of perturbation distribution training on the accuracy. To do so we removed the mutual information term from Equation (4). The trainable parameters are only the locations and scales (tensors M and B) of perturbation distributions. This means that the number of trainable parameters are $2 \times$ *input_feature_dimensions*, and the rest of the neural network parameters are frozen.

[00104] We observe that the scales (B) of the perturbation elements do not change during training, since Cloak is only optimizing for accuracy and is not trying to decrease the mutual information. So, the scales remain at their $\frac{\Delta f}{\epsilon}$ value that delivers the ϵ -feature differential privacy guarantee. This implies that the improvement over the *No Training* method is merely caused by the training of the elements of the locations tensor. The vertical arrows in the graphs indicate the amount of accuracy improvement achieved through the perturbation training. For a budgets ϵ of less than 0.25, the amount of noise is so high that the perturbation training can only modestly mitigate the situation and all the methods yield similar results. FIGS. 10A-10D are depicted for ϵ below 5, to zoom into the region of interest and show the details. FIGS. 13A-D show graphs with ϵ of up to 10.

[00105] FIGS. 10A-10D illustrate accuracy loss vs. the privacy budget (ϵ) for the four benchmark datasets, measured on the test set of each of the four datasets. The datasets are: CelebA (FIG. 10A), CIFAR-100 (FIG. 10B), MNIST (FIG. 10C), and UTKFace (FIG. 10D). Solid lines (labeled “Basic - No Training” in FIGS. 10A-D) represent example results obtained after adding random differentially-private perturbation. Lighter dashed lines (labeled “Training Perturbations” in FIGS. 10A-10D) show example effects of training the perturbation distribution parameters. The neural network parameters remain frozen at all times. The gaps between the solid and dashed lines shown with vertical arrows in FIGS. 10A-10D indicate the improvement

in accuracy brought by perturbation training using an example embodiment of a method according to the technology disclosed in this patent document for a given ϵ .

[00106] FIGS. 11A-11C shows the accuracy loss vs. the remnant mutual information in the original image. Remnant mutual information denotes the mutual information between the original image and its noisy representation, divided by the amount of information in bits in the original image. For the *With MI Term* method, the λ in Equation (4) is tuned, so that as the accuracy grows, the mutual information in the original images would degrade. The privacy budget is $\epsilon = 2.5$ for all the points depicted for this method, since the mutual information degradation does not violate the differential privacy guarantee by pushing any scale below $\frac{\Delta f}{\epsilon}$.

[00107] The *Budget Tightening* method shows how much accuracy loss would be incurred, if the same level of mutual information is to be achieved through merely tightening the differential privacy guarantee, using Cloak without the MI term ($-\log(B)$), similar to the training method of described above. The graphs show that using the MI term helps achieve a certain level of mutual information with less damage to model utility. The vertical arrows show the difference in utility (accuracy) for these two methods. The worse utility brought by naively tightening the differential privacy budget is because DP tries to make all the features in the image similar, and adds perturbations with the same scale to all the pixels.

[00108] For both cases, the graphs show a trade-off between accuracy and the loss in mutual information. For CIFAR-100, since the task is classifying the 20 superclasses, there is inherently more need for information for classification, compared to the other benchmarks. Therefore, for a given level of accuracy loss, the information content cannot be decreased as much as it can in the other benchmarks. We do not present any numerical results for the CelebA dataset here, since the input images have an extremely large number of features and the mutual information estimator tool is not capable of estimating the mutual information.

[00109] FIGS. 11A-11C illustrate accuracy loss vs. the remaining mutual information for a fixed differential-privacy budget of $\epsilon = 2.5$ for three datasets: CIFAR-100 (FIG. 11A), MNIST (FIG. 11B), and UTKFace (FIG. 11C). Data points labeled by downward-directed triangles (labeled “The Budget Tightening” in FIGS. 11A-11C) show example results obtained using the “Budget Tightening” method. Corresponding values of ϵ are indicated next to each “Budget

Tightening” data point in FIGS. 11A-11C. As FIGS. 11A-11C illustrate, the “Budget Tightening” method degrades the accuracy more for a given remnant mutual information level compared to the “With MI Term” method. Example data obtained using the “With MI Term” method are shown in FIGS. 11A-11C by circles connected by a solid line (labeled “With MI Term” in FIGS. 11A-11C). The “With MI Term” method exploits the mutual information term in the loss to reach the desired mutual information level, without tightening the differential privacy guarantee by reducing ϵ . The gaps between the “Budget Tightening” and “With MI Term” data points are shown by vertical arrows in FIGS. 11A-11C and indicate the differences in utility of these approaches for a given level of mutual information reduction.

10 **[00110]** To evaluate the obfuscation that Cloak provides, we devise an experiment in which a malevolent party tries to infer a sensitive private label from an input that is sent to an inference service. To this end, we use CelebA dataset and use gender as the sensitive label while the inference service is set to detect smiles from the perturbed inputs. Cloak does not access to the private labels while learning the perturbations, and it only optimizes for the accuracy of the smile
15 detection task, and decreasing the mutual information.

[00111] FIG. 12 shows the results of this experiment. The malevolent party is assumed to have a pre-trained VGG-16 neural network for gender classification, that has an accuracy of 96.7% on the test set, shown by the *Gender Classification Baseline* line. Cloak trains perturbation parameters with $\epsilon = 2.5$, but with different levels of inter-element intensity, i.e., it
20 uses mutual information term, similar to the ones depicted in Figure 1. The *Noisy Gender Classification* depicts the accuracy achieved by the malevolent party’s neural network, when it receives the perturbed representations (of the test set) and runs inference on them. The *Noisy Gender Classification with Re-training* depicts the same accuracy, but when the malevolent neural network is re-trained using the perturbed representations of the training dataset.

25 **[00112]** We first assume a malevolent party that would re-train its entire neural network (mark all parameters as trainable), but it would not learn anything. We then assume a malevolent party that would limit the training to the last fully connected layer, which yielded the results seen under the *Noisy Gender Classification with Re-training*. We believe the reason that limiting the number of parameters is helpful is that it limits the fluctuations in the network, by fixing a huge
30 portion of the model parameters. The last line in the Figure is the accuracy that a random

predictor would provide, which is 50%. For smile detection with accuracy of 86.9%, the gender classifier suffers an accuracy loss of 30.7% and does 15.8% better than a random classifier. The drop in gender classification accuracy, and the disability in learning and gaining higher accuracy is caused by the perturbation obfuscating unrelated features. However, since the nonsensitive
 5 task is smile detection, there could be some features that both gender detection and smile detection rely on (like the mouth area). This means it might not be possible to completely remove gender information while still maintaining smile detection accuracy.

[00113] FIG. 12 illustrates that an example method according to the disclosed technology learns perturbations for the inference service of smile detection on CelebA dataset with 4.9%
 10 degradation to the accuracy of this service. If a malevolent party attempts to use the same perturbed image to infer sensitive (private) information, e.g. gender, the performance of this new inference service will be degraded by at least 31% while originally the accuracy was 96.7%. This shows that methods according to the technology disclosed in this patent document are protecting the user against malicious use of their private information.

[00114] Consider a deterministic pre-processing function $f(\mathbf{D})$ that operates on the input
 15 tensor data \mathbf{D} . We can model a DNN by a deterministic function $g(f(\mathbf{D});\theta^*)$, where, $f(\mathbf{D})$ is the tensor of pre-processed input data and θ^* is the fixed optimized parameter that fully determines the function g (and maximizes its utility over the distribution of the input tensor). We assume the DNN is trained and the parameter θ^* remains the same throughout the process.
 20 As such, for the ease of the notation, we might drop the parameter θ^* .

[00115] Suppose $L_{nn}(D, g(f(D);\theta))$ is a loss function such that

$$\theta^* \in \arg \min_{\theta} L_{nn}(D, g(f(D);\theta))$$

[00116] If L_{nn} is the loss function corresponding to the DNN, then we are assuming that the network is trained and θ^* is the optimal parameter.

[00117] We define the privacy P as the negative of the mutual information between the input
 25 of the function g and the raw data \mathbf{D} , i.e.,

$$P = -I(\mathbf{D}; f(\mathbf{D}))$$

where $I(\cdot; \cdot)$ represents mutual information. The lower mutual information implies higher privacy. Since f is a deterministic function, we have

$$P = -H(f(\mathbf{D})) + H(f(\mathbf{D})|\mathbf{D}) = -H(f(\mathbf{D})) \quad (6)$$

5 where $H(\cdot)$ denotes the entropy function.

[00118] Denote a perturbation tensor by \mathbf{X} which is independent of \mathbf{D} . We intend to perturb the input to function g by changing $f(\mathbf{D})$ to $f(\mathbf{D}) + \mathbf{X}$, i.e., the output of function g changes from $g(f(\mathbf{D}); \theta^*)$ to $g(f(\mathbf{D}) + \mathbf{X}; \theta^*)$. It is worth reemphasizing that the parameter θ^* remains unchanged. Now, we need to change the privacy measure P to \hat{P} where

10 $\hat{P} = -I(\mathbf{D}; g(f(\mathbf{D}) + \mathbf{X}))$. We can provide the following lower bound

$$\begin{aligned} \hat{P} &= -I(\mathbf{D}; f(\mathbf{D}) + \mathbf{X}) \\ &\geq -I(f(\mathbf{D}); f(\mathbf{D}) + \mathbf{X}) \\ &= -H(f(\mathbf{D})) + H(f(\mathbf{D})|f(\mathbf{D}) + \mathbf{X}) \\ &= P + H(f(\mathbf{D})|f(\mathbf{D}) + \mathbf{X}) \end{aligned} \quad (7)$$

where the last equality is derived from (6). This equation implies that by noise injection process, we can improve the privacy at least by $H(f(\mathbf{D})|f(\mathbf{D}) + \mathbf{X})$.

[00119] After noise injection, we would like to find the optimal perturbation \mathbf{X}^* that

15 maximizes the privacy while keeping the performance at an admissible error range. In other words, we are looking for the solution to the following optimization problem

$$\mathbf{X}^* = \arg \max_{\mathbf{X}} \hat{P} \text{ s.t. } L_m(\mathbf{D}, g(f(\mathbf{D}) + \mathbf{X}; \theta^*)) \leq \gamma$$

[00120] Given (7), we use

$$I(f(\mathbf{D}); f(\mathbf{D}) + \mathbf{X}) = H(f(\mathbf{D}) + \mathbf{X}) - H(f(\mathbf{D}) + \mathbf{X}|f(\mathbf{D})) = H(f(\mathbf{D}) + \mathbf{X}) - H(\mathbf{X}) \text{ as a}$$

surrogate objective and reformulate the problem in terms of a Lagrange multiplier λ as

$$\begin{aligned} \mathbf{X}^* = \arg \min_{\mathbf{X}} & H(f(\mathbf{D}) + \mathbf{X}) - H(\mathbf{X}) \\ & + \lambda L_{nn}(\mathbf{D}, g(f(\mathbf{D}) + \mathbf{X}; \theta^*)) \end{aligned} \quad (8)$$

where the dual parameter λ is to be determined via cross validation on the generalization error.

[00121] Optimization problem (8) has three terms. The first term controls the amount of information that leaks to the remote part of the DNN and we want to minimize this information. The second term controls the amount of uncertainty that we are injecting in the form of iid noise. This term is typically proportional to the variance of the noise and we want this term to be maximized (and hence the negative sign). The last term controls the amount of degradation in the performance of the DNN and we want to minimize that with respect to the knob λ .

[00122] The loss function $L_{nn}(\cdot)$ is determined at the training time of the DNN to obtain the optimal parameter θ^* . The loss function for a regression problem can be a Lebesgue L_p -norm and for a q -class classification problem with $Z = g(f(\mathbf{D}) + \mathbf{X}; \theta^*)$, it can be $L_{nn}(\mathbf{D}, Z) = -\sum_{j=1}^q \mathbf{1}_{\mathbf{X} \in C_j} \log(Z_j)$ where $\mathbf{1}_{\cdot}$ is the indicator function, C_j represents the j^{th} class and Z_j is the j^{th} logit representing the probability that $\mathbf{D} \in C_j$. Suppose $\mathbf{1}_{\mathbf{D}}$ is a one-hot-encoded q -vector with j^{th} element being 1 if $\mathbf{D} \in C_j$ and the rest are zero. We then can write the classification loss in vector form as $L_{nn}(\mathbf{D}, Z) = -\mathbf{1}_{\mathbf{X}}^T \log(Z)$. For the remainder of this paper, we target a q -class classification and rewrite (8) as

$$\begin{aligned} \mathbf{X}^* = \arg \min_{\mathbf{X}} & H(f(\mathbf{D}) + \mathbf{X}) - H(\mathbf{X}) \\ & + \lambda \mathbf{1}_{\mathbf{X}}^T \log(g(f(\mathbf{D}) + \mathbf{X}; \theta^*)) \end{aligned} \quad (9)$$

[00123] Assuming \mathbf{X} is element-wise distributed as $Lap(\mathbf{M}, \mathbf{B})$, we can establish that the first and second term in (9) are proportional to $\log(\mathbf{B})$ and rewrite the optimization problem as

$$\begin{aligned} (M^*, B^*) = \arg \min_{M, B} & -\log(B) \\ & + \lambda E_{X \sim Lap(M, B)} \left[\mathbf{1}_X^T \log \left(g \left(f(D) + X; \theta^* \right) \right) \right] \end{aligned} \quad (10)$$

where $E[\cdot]$ represents the expectation that can be approximated via Monte Carlo sampling of X . Solving this optimization problem, we get the optimal noise distribution parameters.

[00124] In order to do the inference, we generate a noise tensor X drawn from the
5 distribution we learned in (10), and add it to the output of function f before sending it to the server.

[00125] The optimization problem (10) guarantees that for a given value of the knob λ , one can find the distribution parameters that minimizes the information sent to the cloud, while preserving the quality of the service. Since this optimization problem is convex, we are
10 guaranteed to converge to the optimal point via gradient descent algorithm. That is why Cloak uses this optimization technique.

Table 1. Neural networks used in experiments described above.

Neural Network	Conv Layers	FC Layers	Dataset
VGG-16	13	3	CelebA
AlexNet	5	3	CIFAR-100
LeNet-5	13	3	MNIST
VGG-16 (modified)	3	2	UTKFace

[00126] We have run the experiments for CelebA and CIFAR-100 datasets on an Nvidia RTX
15 2080 Ti GPU, with 11GB VRAM, paired with 10 Intel Core i9-9820X processors with 64GBs of memory. The experiments for MNIST and UTKFace datasets were run on the CPU. The systems runs an Ubuntu 18.04 OS, with CUDA version V 10.2.89.

[00127] Table 1 shows the DNNs used for the experiments. The modified VGG-16 is
20 different from the conventional one in the size of the last 3 fully connected layers. They are (512,256), (256,256) and (256,2).

[00128] As mentioned above, the mutual information between the input images and their noisy representations are estimated over the test set images using ITE toolbox's Shannon mutual information estimator. For MNIST images, our dataset had inputs of size 32×32 pixels, which we flattened to 1024 element vectors, for estimating the mutual information. For other datasets, since the images were larger ($32 \times 32 \times 3$), there were more dimensions and mutual information estimation was not accurate. So, what we did here was calculate mutual information channel by channel (i.e., we estimated the mutual information between the red channel of the image and its noisy representation, then the green channel and then blue), and we averaged over all channels.

[00129] Tables 2, 4 and 3 show the hyperparameters used for training in the experiments of described above. For the last two, the *Point#* indicates the experiment that produced the given point in the graph, if the points were numbered from left to right. In our implementation, for ease of use and without loss of generality, we have introduced a variable γ to the loss function in Equation (4), in a way that $\gamma = \frac{1}{\lambda}$, and it is the coefficient of the mutual information term, not the cross-entropy loss. With this introduction, we do not directly assign a λ (as if λ were removed and replaced by γ as a coefficient of the other term) and we need not set λ very high when we want to remove the mutual information term. We just set γ to zero. In the tables, we have used lambda to be consistent, and in the cells where the number for λ is not given, it means that the mutual information term was not used. But in the Code, the coefficient is set on the other term and is $1/\lambda$ s reported here. The batch sizes used for training are 128 for CIFAR-100, MNIST and UTKFace and 40 for CelebA. For testing, as mentioned in the text above, the batch size is 1, so as to sample a new noise tensor for each image and capture the stochasticity. Also, the number of samples taken for each update in optimization is 1, as mentioned above, since we do mini-batch training and for each mini-batch we take a new sample. Finally, M_x is set to 2.0 for all benchmarks, except for MNIST where it is set to be 1.5.

[00130] FIGS. 13A-13D re-illustrate FIGS. 10A-10D, respectively, with a wider privacy budget of up to 10. It can be seen in FIGS. 13A-13D that for most benchmarks, for privacy budgets larger than 5, since the amount of injected noise is less, the accuracy loss is less and the training the perturbation parameters makes a smaller improvement.

Table 2. Hyperparameters for experiments illustrated in FIGS. 10A-10D.

Dataset	Epsilon	Details
CelebA	10.00	0.3 epoch w/ lr=0.0001
	5.00	0.3 epoch w/ lr=0.01, 0.3 epoch w/ lr=0.0001
	2.50	0.3 epoch w/ lr=0.01, 0.3 epoch w/ lr=0.0001
	2.00	1 epoch w/ lr=0.01, 1 epoch w/ lr=0.00001
	1.52	1 epoch w/ lr=0.01, 1 epoch w/ lr=0.00001
	1.00	1 epoch w/ lr=0.01, 0.5 epoch w/ lr=0.00001
	0.50	1 epoch w/ lr=0.01, 0.5 epoch w/ lr=0.00001
	0.25	1 epoch w/ lr=0.01, 0.5 epoch w/ lr=0.00001
CIFAR-100	10.00	1 epoch w/ lr= 0.00001
	5.00	1 epoch w/ lr= 0.01, 2 epoch w/ lr=0.001, 3 epoch w/ lr=0.00001
	2.50	9 epoch w/ lr= 0.01, 6 epoch w/ lr= 0.001, 3 epoch w/ lr =0.0001
	2.00	20 epoch w/ lr=0.000001
	1.52	20 epoch w/ lr=0.000001
	1.00	20 epoch w/ lr=0.000001
	0.50	20 epoch w/ lr=0.000001
	0.25	20 epoch w/ lr=0.000001
MNIST	10.00	20 epoch w/ lr=0.001
	5.00	20 epoch w/ lr=0.001
	2.50	30 epoch w/ lr=0.001, 25 w/ lr=0.0001
	2.00	20 epoch w/ lr=0.001, 32 w/ lr=0.0001
	1.52	30 epoch w/ lr=0.001, 25 w/ lr=0.0001
	1.00	30 epoch w/ lr=0.001, 25 w/ lr=0.0001
	0.50	10 epoch w/ lr=0.00001
	0.25	10 epoch w/ lr=0.00001
UTKFace	10.00	1 epoch w/ lr= 0.01, 5 epoch w/ lr=0.0001
	5.00	1 epoch w/ lr= 0.01, 5 epoch w/ lr=0.0001
	2.50	1 epoch w/ lr= 0.01, 8 epoch w/ lr=0.0001

Dataset	Epsilon	Details
	2.00	6 epoch w/ lr= 0.01, 6 epoch w/ lr=0.0001
	1.52	6 epoch w/ lr= 0.01, 6 epoch w/ lr=0.0001
	1.00	6 epoch w/ lr= 0.01, 6 epoch w/ lr=0.0001
	0.50	12 epoch w/ lr= 0.01, 6 epoch w/ lr=0.0001
	0.25	12 epoch w/ lr= 0.01, 6 epoch w/ lr=0.0001

Table 3. Hyperparameters for experiments illustrated in FIG. 12.

Task	Point#	Details
Gender Classification	1	1 epoch w/ lr=0.0001
	2	1 epoch w/ lr=0.01, 2 epoch w/ lr=0.0001
	3	1 epoch w/ lr=0.01, 2 epoch w/ lr=0.0001, 1 epoch w/ lr=0.00001
	4	1 epoch w/ lr=0.01, 2 epoch w/ lr=0.0001, 1 epoch w/ lr=0.00001
	5	1 epoch w/ lr=0.01, 2 epoch w/ lr=0.0001, 1 epoch w/ lr=0.00001
	6	1 epoch w/ lr=0.01, 2 epoch w/ lr=0.0001, 3 epoch w/ lr=0.00001
	7	1 epoch w/ lr=0.01, 2 epoch w/ lr=0.0001, 2 epoch w/ lr=0.00001
	8	1 epoch w/ lr=0.01, 2 epoch w/ lr=0.0001, 3 epoch w/ lr=0.00001
Smile Detection	1	1 epoch w/ lr=0.01, 0.1 epoch w/ lr=0.001
	2	0.5 epoch w/ lr=0.01 lambda=1, 0.2 epoch w/ lr=0.001 lambda=1
	3	0.5 epoch w/ lr=0.01 lambda=1, 0.4 epoch w/ lr=0.001 lambda=1
	4	0.5 epoch w/ lr=0.01 lambda=1, 0.5 epoch w/ lr=0.001 lambda=1

Task	Point#	Details
	5	0.5 epoch w/ lr=0.01 lambda=1, 0.7 epoch w/ lr=0.001 lambda=1
	6	0.5 epoch w/ lr=0.01 lambda=1, 0.8 epoch w/ lr=0.001 lambda=1
	7	0.5 epoch w/ lr=0.01 lambda=1, 0.8 epoch w/ lr=0.001 lambda=1, 0.2 epoch w/ lr=0.001 lambda=5
	8	0.5 epoch w/ lr=0.01 lambda=1, 0.8 epoch w/ lr=0.001 lambda=1, 0.2 epoch w/ lr=0.001 lambda=100

Table 4. Hyperparameters for experiments illustrated in FIGS. 11A-11C.

Dataset	Point#	Details
CIFAR-100	1	6 epoch w/ lr=0.001 6, 14 epoch w/ lr=0.00001
	2	10 epoch w/ lr= 0.001 lambda = 1, 2 epoch w/ lr=0.001 lambda =10
	3	10 epoch w/ lr= 0.001 lambda = 1, 2 epoch w/ lr=0.001 lambda = 10
	4	12 epoch w/ lr= 0.001 lambda = 1, 2 epoch w/ lr=0.001 lambda = 10
	5	17 epoch w/ lr= 0.001 lambda = 1, 3 epoch w/ lr=0.001 lambda = 10
	6	24 epoch w/ lr= 0.001 lambda = 1, 2 epoch w/ lr=0.001 lambda = 10
	7	30 epoch w/ lr= 0.001 lambda = 1, 2 epoch w/ lr=0.001 lambda = 10
	8	40 epoch w/ lr= 0.001 lambda = 0.2, 2 epoch w/ lr=0.001 lambda = 10

	9	140 epoch w/ lr= 0.001 lambda = 0.2, 2 epoch w/ lr=0.001 lambda = 10
MNIST	10	300 epoch w/ lr= 0.001 lambda = 0.01, 20 epoch w/ lr=0.0001 lambda = 10
	1	25 epoch w/ lr =0.001, 30 epoch w/ lr=0.0001
	2	50 epoch w/ lr= 0.01 lambda = 100, 40 epoch w/ lr=0.001 lambda = 1000
	3	50 epoch w/ lr= 0.01 lambda = 100, 50 epoch w/ lr=0.001 lambda = 1000
	4	50 epoch w/ lr= 0.01 lambda = 100, 60 epoch w/ lr=0.001 lambda = 200
	5	50 epoch w/ lr= 0.01 lambda = 100, 90 epoch w/ lr=0.001 lambda = 200
	6	50 epoch w/ lr= 0.01 lambda = 100, 160 epoch w/ lr=0.001 lambda = 200
	7	50 epoch w/ lr= 0.01 lambda = 100, 180 epoch w/ lr=0.001 lambda = 200
	8	50 epoch w/ lr= 0.01 lambda = 100, 260 epoch w/ lr=0.001 lambda = 100
	9	50 epoch w/ lr= 0.01 lambda = 100, 290 epoch w/ lr=0.001 lambda = 100
	10	50 epoch w/ lr= 0.01 lambda = 100, 300 epoch w/ lr=0.001 lambda = 100
UTKFace	1	6 epoch w/ lr = 0.01, 6 epoch w/ lr=0.0001
	2	4 epoch w/ lr=0.01 lambda=0.1, 2 epoch w/ lr=0.0001 lambda = 100
	3	8 epoch w/ lr=0.01 lambda=0.1, 2 epoch w/ lr=0.0001 lambda = 100
	4	10 epoch w/ lr=0.01 lambda=0.1, 2 epoch w/ lr=0.0001 lambda = 100

	5	12 epoch w/ lr=0.01 lambda=0.1, 2 epoch w/ lr=0.0001 lambda = 100
--	---	--

[00131] One aspect of the disclosed embodiments relates to a method of providing data privacy that includes specifying a value of a parameter λ related to trade-off between an accuracy of inference done by a neural network (NN) on a perturbed input of the NN and a degree of mutual information degradation between a raw input of the NN and the perturbed input of the NN; for each feature of a data instance, providing a Laplace distribution corresponding to the feature, wherein the Laplace distribution has parameters of location M and scale B ; forming a tensor of locations M_T and a tensor of scales B_T using the parameters of the Laplace distributions provided for the features of the data instance; providing a loss function L having a term proportional to $\log(B)$, wherein tensor B is related to the tensor of scales B_T , and a term proportional to a product of the value of the parameter λ and a utility loss function L_{nn} of the neural network, wherein the utility loss function L_{nn} can be used to train the neural network to perform an inference task T ; finding optimized elements of B_T and M_T by optimizing, using the loss function L , accuracy of the inference task T performed by the neural network; for each feature of the data instance, determining values of optimized parameters of location M and scale B of the Laplace distribution corresponding to the feature using the optimized elements of B_T and M_T .

[00132] In some example embodiments of the method of providing data privacy, the method further includes selecting, for each feature of an input data instance D of the neural network, a perturbation value x from the Laplace distribution corresponding to the feature and adding the perturbation value x to a value of the feature to produce a perturbed input data instance D_P ; and sending the D_P to the NN for inference. In other example embodiments of the method of providing data privacy, the neural network is pre-trained to perform an inference task. According to some example embodiments, the neural network is pre-trained to perform the inference task using the utility loss function L_{nn} . In certain example embodiments, the input data instance D is obtained by pre-processing a raw input D_R . In some example embodiments, the pre-processing is a normalization. In some example embodiments, the method of providing data privacy further includes specifying a value of a differential privacy budget ϵ as a constraint. According to some example embodiments, the data instance is an image. In certain example embodiments, the

feature of the data instance is a pixel of the image. In some example embodiments, sending the D_P to the NN for inference comprising sending the D_P over a network.

[00133] Another aspect of the disclosed embodiments relates to a method of providing data privacy for neural network computations that includes: for a given privacy budget ϵ and for a target inference task to be performed by a neural network, finding optimized parameters of a set of statistical distributions which optimize performance of the neural network on the inference task and with respect to the privacy budget ϵ using a loss function L , wherein the loss function L has a term related to a tensor, wherein the tensor is related to a parameter of at least one distribution from the set of statistical distributions and another term related to a utility loss function L_{nn} of the neural network for the target inference task; selecting, for each feature of an input, a perturbation value drawn from a distribution in the set of statistical distributions which has the optimized parameters and adding the perturbation value to a value associated with the feature to obtain a perturbed input; sending the perturbed input to the neural network; and performing the target inference task by the neural network on the perturbed input.

[00134] In some example embodiments of the method of providing data privacy for neural network computations, the term related to the tensor includes a logarithm of the tensor. In other example embodiments, the statistical distributions in the set of statistical distributions are of a same type. According to certain example embodiments, the statistical distributions in the set of statistical distributions are Laplace distributions. In some example embodiments, the parameter of at least one distribution is a scale of the at least one distribution. According to certain example embodiments of the method of providing data privacy for neural network computations, elements of the tensor are scales of the statistical distributions. In some example embodiments, sending the perturbed input to the neural network comprises sending the perturbed input over a communication network.

[00135] Yet another aspect of the disclosed embodiments relates to a data processing method that includes determining parameters of a distortion operation by which a source data is converted into a distorted data; and performing a data processing task on the distorted data.

[00136] In some example embodiments of the data processing method, the data processing task satisfies an ϵ -differential privacy criterion when applied to the distorted data. In other example embodiments, the parameters are determined using a pre-trained neural network as an

analytical function of the parameters. According to certain example embodiments, the parameters are determined using a gradient-based optimization of the analytical function of the parameters. According to some example embodiments, weights of the neural network do not change their values. In some example embodiments, the data processing method further includes refraining from specifying privacy of which data features needs to be protected. In some example
5 embodiments, the distortion operation reduces mutual information between the source data and the distorted data. According to certain example embodiments, the distortion operation incurs a limited degradation to an accuracy of the data processing task due to the distortion operation. In some example embodiments, the data processing task includes an inference task performed by a
10 neural network. In other example embodiments, the parameters are parameters of a set of statistical distributions. In some example embodiments, the data processing method further includes transferring the distorted data over an external network to a remote server. According to some example embodiments, the external network is a communications network.

[00137] An aspect of the disclosed embodiments relates to a method of providing privacy for data that includes injecting stochasticity into the data to produce perturbed data, wherein the
15 injected stochasticity satisfies an ϵ -differential privacy criterion, and transmitting the perturbed data to a neural network or to a partition of the neural network for inference.

[00138] In some example embodiments, the neural network is a deep neural network. According to certain example embodiments, the neural network is a pre-trained neural network.
20 According to some example embodiments, the amount of stochasticity is such that information content of the perturbed data retains essential pieces that enable the inference to be serviced accurately by the neural network. In some example embodiments, the amount of stochasticity is determined by discovering the stochasticity via an offline gradient-based optimization that reformulates the neural network as an analytical function of the stochasticity. In certain example
25 embodiments, weights of the neural network do not change their values.

[00139] Another aspect of the disclosed embodiments relates to a method of providing privacy for data that includes determining, for a pre-trained deep neural network (DNN), an amount of stochastic perturbation, applying the amount of the stochastic perturbation to source data to obtain perturbed source data, and transmitting the perturbed source data to the DNN or to a
30 partition of the DNN.

[00140] In some example embodiments, the amount of stochastic perturbation is such that information content of the perturbed source data retains essential pieces that enable an inference request to be serviced accurately by the DNN. In certain example embodiments, the amount of stochastic perturbation is determined by discovering the stochasticity via an offline gradient-based optimization problem that reformulates the pre-trained DNN as an analytical function of the stochastic perturbation. According to some example embodiments, weights of the pre-trained DNN do not change their values.

[00141] Yet another aspect of the disclosed embodiments relates to a method of providing privacy for data that includes determining an amount of stochastic perturbation in a source data without accessing sensitive information or labels associated with the source data, and transmitting, to a neural network or to a partition of the neural network, a perturbed data obtained by perturbing the source data using the amount of stochastic perturbation.

[00142] In some example embodiments, the neural network is a deep neural network. In certain example embodiments, the neural network is a pre-trained neural network. According to some example embodiments, the amount of stochastic perturbation is such that information content of the perturbed data retains essential pieces that enable an inference request to be serviced accurately by the neural network. In some example embodiments, the amount of stochastic perturbation is determined by discovering the stochasticity via an offline gradient-based optimization problem that reformulates the neural network as an analytical function of the stochastic perturbation. In certain example embodiments, weights of the neural network do not change their values.

[00143] In some example embodiments of a method of providing privacy for data according to the disclosed technology, the neural network resides on a device which performs the method. In other example embodiments of a method of providing privacy for data according to the disclosed technology, the neural network resides on a device other than the device which performs the method. According to some example embodiments of a method of providing privacy for data according to the disclosed technology, the neural network is a distributed neural network. According to certain example embodiments of a method of providing privacy for data according to the disclosed technology, a part of the neural network resides on a device which performs the method and another part of the neural network resides on a device other than the

device which performs the method.

[00144] An aspect of the disclosed embodiments relates to a communication apparatus that includes a memory and a processor, wherein the processor is configured to read code from the memory and implement a method according to the technology disclosed in this patent document.

5 **[00145]** Another aspect of the disclosed embodiments relates to a non-transitory computer-readable medium storing instructions that, when executed by a computer, cause the computer to perform a method according to the technology disclosed in this patent document.

[00146] Implementations of the subject matter and the functional operations described in this patent document can be implemented in various systems, digital electronic circuitry, or in
10 computer software, firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Implementations of the subject matter described in this specification can be implemented as one or more computer program products, i.e., one or more modules of computer program instructions encoded on a tangible and non-transitory computer readable medium for execution by, or to control the
15 operation of, data processing apparatus. The computer readable medium can be a machine-readable storage device, a machine-readable storage substrate, a memory device, a composition of matter effecting a machine-readable propagated signal, or a combination of one or more of them. The term “data processing unit” or “data processing apparatus” encompasses all apparatus, devices, and machines for processing data, including by way of example a
20 programmable processor, a computer, or multiple processors or computers. The apparatus can include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

[00147] A computer program (also known as a program, software, software application, script, or code) can be written in any form of programming language, including compiled or
25 interpreted languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program does not necessarily correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data (e.g., one or more scripts stored in a
30 markup language document), in a single file dedicated to the program in question, or in multiple

coordinated files (e.g., files that store one or more modules, sub programs, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

5 **[00148]** The processes and logic flows described in this specification can be performed by one or more programmable processors executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated
10 circuit).

[00149] Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are
15 a processor for performing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. However, a computer need not have such devices. Computer readable media suitable for storing computer program instructions and data include all
20 forms of nonvolatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[00150] It is intended that the specification, together with the drawings, be considered
25 exemplary only, where exemplary means an example. As used herein, the singular forms “a”, “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. Additionally, the use of “or” is intended to include “and/or”, unless the context clearly indicates otherwise.

[00151] While this patent document contains many specifics, these should not be construed as
30 limitations on the scope of any invention or of what may be claimed, but rather as descriptions of

features that may be specific to particular embodiments of particular inventions. Certain features that are described in this patent document in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple
5 embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

[00152] Similarly, while operations are depicted in the drawings in a particular order, this
10 should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. Moreover, the separation of various system components in the embodiments described in this patent document should not be understood as requiring such separation in all
embodiments.

[00153] Only a few implementations and examples are described and other implementations, enhancements and variations can be made based on what is described and illustrated in this
15 patent document.

CLAIMS

What is claimed is:

1. A method of providing data privacy, comprising:

5 specifying a value of a parameter λ related to trade-off between an accuracy of inference done by a neural network (NN) on a perturbed input of the NN and a degree of mutual information degradation between a raw input of the NN and the perturbed input of the NN;

for each feature of a data instance, providing a Laplace distribution corresponding to the feature, wherein the Laplace distribution has parameters of location M and scale B ;

10 forming a tensor of locations M_T and a tensor of scales B_T using the parameters of the Laplace distributions provided for the features of the data instance;

providing a loss function L having a term proportional to $\log(B)$, wherein tensor B is related to the tensor of scales B_T , and a term proportional to a product of the value of the parameter λ and a utility loss function L_{nn} of the neural network, wherein the utility loss function L_{nn} can be used to train the neural network to perform an inference task T ;

15 finding optimized elements of B_T and M_T by optimizing, using the loss function L , accuracy of the inference task T performed by the neural network;

for each feature of the data instance, determining values of optimized parameters of location M and scale B of the Laplace distribution corresponding to the feature using the optimized elements of B_T and M_T .

2. The method as in claim 1, comprising:

25 selecting, for each feature of an input data instance D of the neural network, a perturbation value x from the Laplace distribution corresponding to the feature and adding the perturbation value x to a value of the feature to produce a perturbed input data instance D_p ; and

sending the D_p to the NN for inference.

3. The method as in any of claims 1-2, wherein the neural network is pre-trained to perform an inference task.

4. The method as in claim 3, wherein the neural network is pre-trained to perform the inference task using the utility loss function L_{nn} .
5. The method as in any of claims 1-4, wherein the input data instance D is obtained by pre-processing a raw input D_R .
- 5 6. The method as in any of claims 1-5, wherein the pre-processing is a normalization.
7. The method as in any of claims 1-6, comprising:
specifying a value of a differential privacy budget ϵ as a constraint.
8. The method as in any of claims 1-7, wherein the data instance is an image.
9. The method as in claim 8, wherein the feature of the data instance is a pixel of the image.
- 10 10. The method as in any of claims 2-9, wherein sending the D_P to the NN for inference comprising sending the D_P over a network.
11. A method of providing data privacy for neural network computations, comprising:
for a given privacy budget ϵ and for a target inference task to be performed by a neural network, finding optimized parameters of a set of statistical distributions which
15 optimize performance of the neural network on the inference task and with respect to the privacy budget ϵ using a loss function L , wherein the loss function L has
a term related to a tensor, wherein the tensor is related to a parameter of at least one distribution from the set of statistical distributions
and another term related to a utility loss function L_{nn} of the neural network for
20 the target inference task;
selecting, for each feature of an input, a perturbation value drawn from a distribution in the set of statistical distributions which has the optimized parameters and adding the perturbation value to a value associated with the feature to obtain a perturbed input;
sending the perturbed input to the neural network; and
25 performing the target inference task by the neural network on the perturbed input.

12. The method as in claim 11, wherein the term related to the tensor includes a logarithm of the tensor.
13. The method as in any of claims 11-12, wherein the statistical distributions in the set of statistical distributions are of a same type.
- 5 14. The method as in any of claims 11-13, wherein the statistical distributions in the set of statistical distributions are Laplace distributions.
15. The method as in any of claims 11-14, wherein the parameter of at least one distribution is a scale of the at least one distribution.
16. The method as in any of claims 11-15, wherein elements of the tensor are scales of the
10 statistical distributions.
17. The method as in any of claims 11-16, wherein sending the perturbed input to the neural network comprises sending the perturbed input over a communication network.
18. A data processing method, comprising:
determining parameters of a distortion operation by which a source data is
15 converted into a distorted data; and
performing a data processing task on the distorted data.
19. The data processing method as in claim 18, wherein the data processing task satisfies an ϵ -differential privacy criterion when applied to the distorted data.
20. The data processing method as in any of claims 18-19, wherein the parameters are
20 determined using a pre-trained neural network as an analytical function of the parameters.
21. The data processing method as in claim 20, wherein the parameters are determined using a gradient-based optimization of the analytical function of the parameters.
22. The data processing method as in any of claims 20-21, wherein weights of the neural network do not change their values.

23. The data processing method as in any of claims 18-22, comprising refraining from specifying privacy of which data features needs to be protected.
24. The data processing method as in any of claims 18-23, wherein the distortion operation reduces mutual information between the source data and the distorted data.
- 5 25. The data processing method as in any of claims 18-24, wherein the distortion operation incurs a limited degradation to an accuracy of the data processing task due to the distortion operation.
26. The data processing method as in any of claims 18-25, wherein the data processing task includes an inference task performed by a neural network.
- 10 27. The data processing method as in any of claims 18-26, wherein the parameters are parameters of a set of statistical distributions.
28. The data processing method as in any of claims 18-27, comprising transferring the distorted data over an external network to a remote server.
29. The data processing method as in claim 28, wherein the external network is a
15 communications network.
30. A method of providing privacy for data, comprising:
injecting stochasticity into the data to produce perturbed data, wherein the injected stochasticity satisfies an ϵ -differential privacy criterion, and
transmitting the perturbed data to a neural network or to a partition of the neural
20 network for inference.
31. The method as in claim 30, wherein the neural network is a deep neural network.
32. The method as in any of claims 30-31, wherein the neural network is a pre-trained neural network.

33. The method as in any of claims 30-32, wherein the amount of stochasticity is such that information content of the perturbed data retains essential pieces that enable the inference to be serviced accurately by the neural network.

34. The method as in any of claims 30-33, wherein the amount of stochasticity is determined by discovering the stochasticity via an offline gradient-based optimization that reformulates the neural network as an analytical function of the stochasticity.

35. The method as in any of claims 30-34, wherein weights of the neural network do not change their values.

36. A method of providing privacy for data, comprising:

- 10 determining, for a pre-trained deep neural network (DNN), an amount of stochastic perturbation,
 applying the amount of the stochastic perturbation to source data to obtain perturbed source data, and
 transmitting the perturbed source data to the DNN or to a partition of the DNN.

15 37. The method as in claim 36, wherein the amount of stochastic perturbation is such that information content of the perturbed source data retains essential pieces that enable an inference request to be serviced accurately by the DNN.

38. The method as in any of claims 36-37, wherein the amount of stochastic perturbation is determined by discovering the stochasticity via an offline gradient-based optimization problem that reformulates the pre-trained DNN as an analytical function of the stochastic perturbation.

39. The method as in any of claims 36-38, wherein weights of the pre-trained DNN do not change their values.

40. A method of providing privacy for data, comprising:

- 25 determining an amount of stochastic perturbation in a source data without accessing sensitive information or labels associated with the source data, and

transmitting, to a neural network or to a partition of the neural network, a perturbed data obtained by perturbing the source data using the amount of stochastic perturbation.

41. The method as in claim 40, wherein the neural network is a deep neural network.

5 42. The method as in any of claims 40-41, wherein the neural network is a pre-trained neural network.

43. The method as in any of claims 40-42, wherein the amount of stochastic perturbation is such that information content of the perturbed data retains essential pieces that enable an inference request to be serviced accurately by the neural network.

10 44. The method as in any of claims 40-43, wherein the amount of stochastic perturbation is determined by discovering the stochasticity via an offline gradient-based optimization problem that reformulates the neural network as an analytical function of the stochastic perturbation.

45. The method as in any of claims 40-44, wherein weights of the neural network do not change their values.

15 46. The method as in any of claims 30-45, wherein the neural network resides on a device which performs the method.

47. The method as in any of claims 30-45, wherein the neural network resides on a device other than the device which performs the method.

20 48. The method as in any of claims 30-45, wherein the neural network is a distributed neural network.

49. The method as in claim 48, wherein a part of the neural network resides on a device which performs the method and another part of the neural network resides on a device other than the device which performs the method.

50. A communication apparatus, comprising a memory and a processor, wherein the processor is configured to read code from the memory and implement a method as in any one of claims 1-49.

51. A non-transitory computer-readable medium storing instructions that, when executed by
5 a computer, cause the computer to perform a method as in any one of claims 1-49.

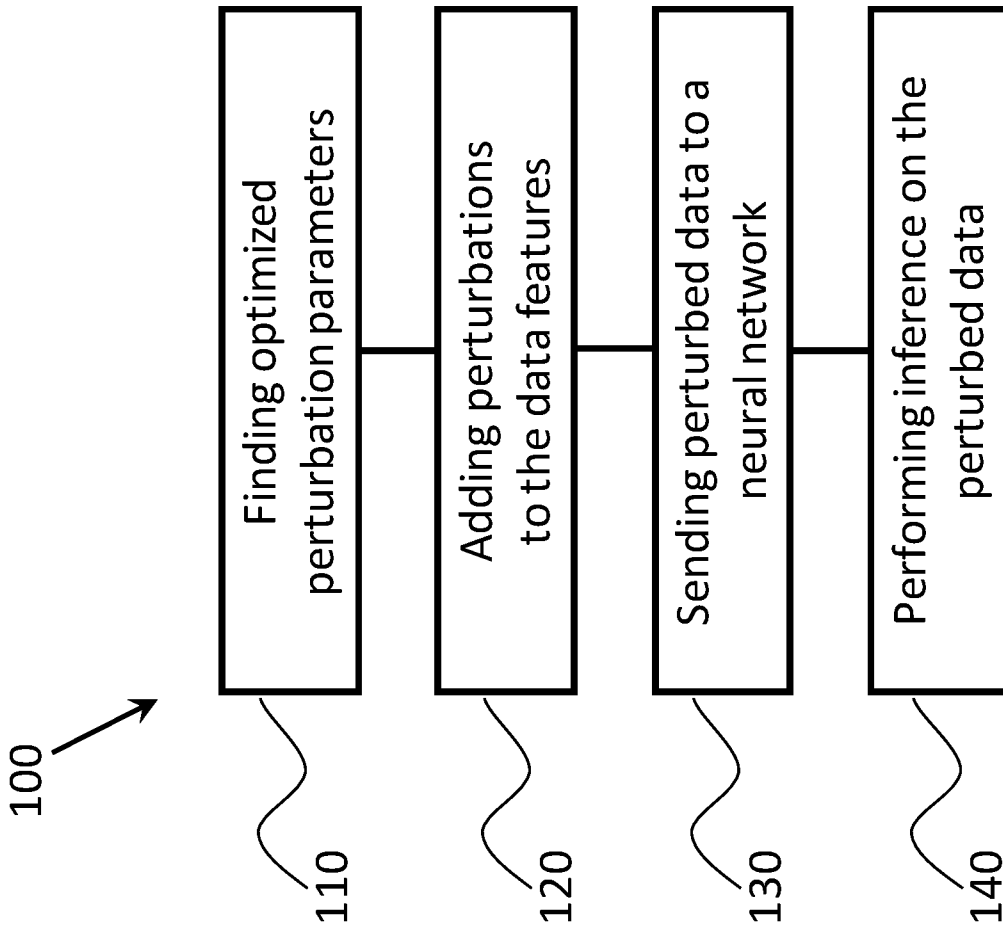


FIG. 1

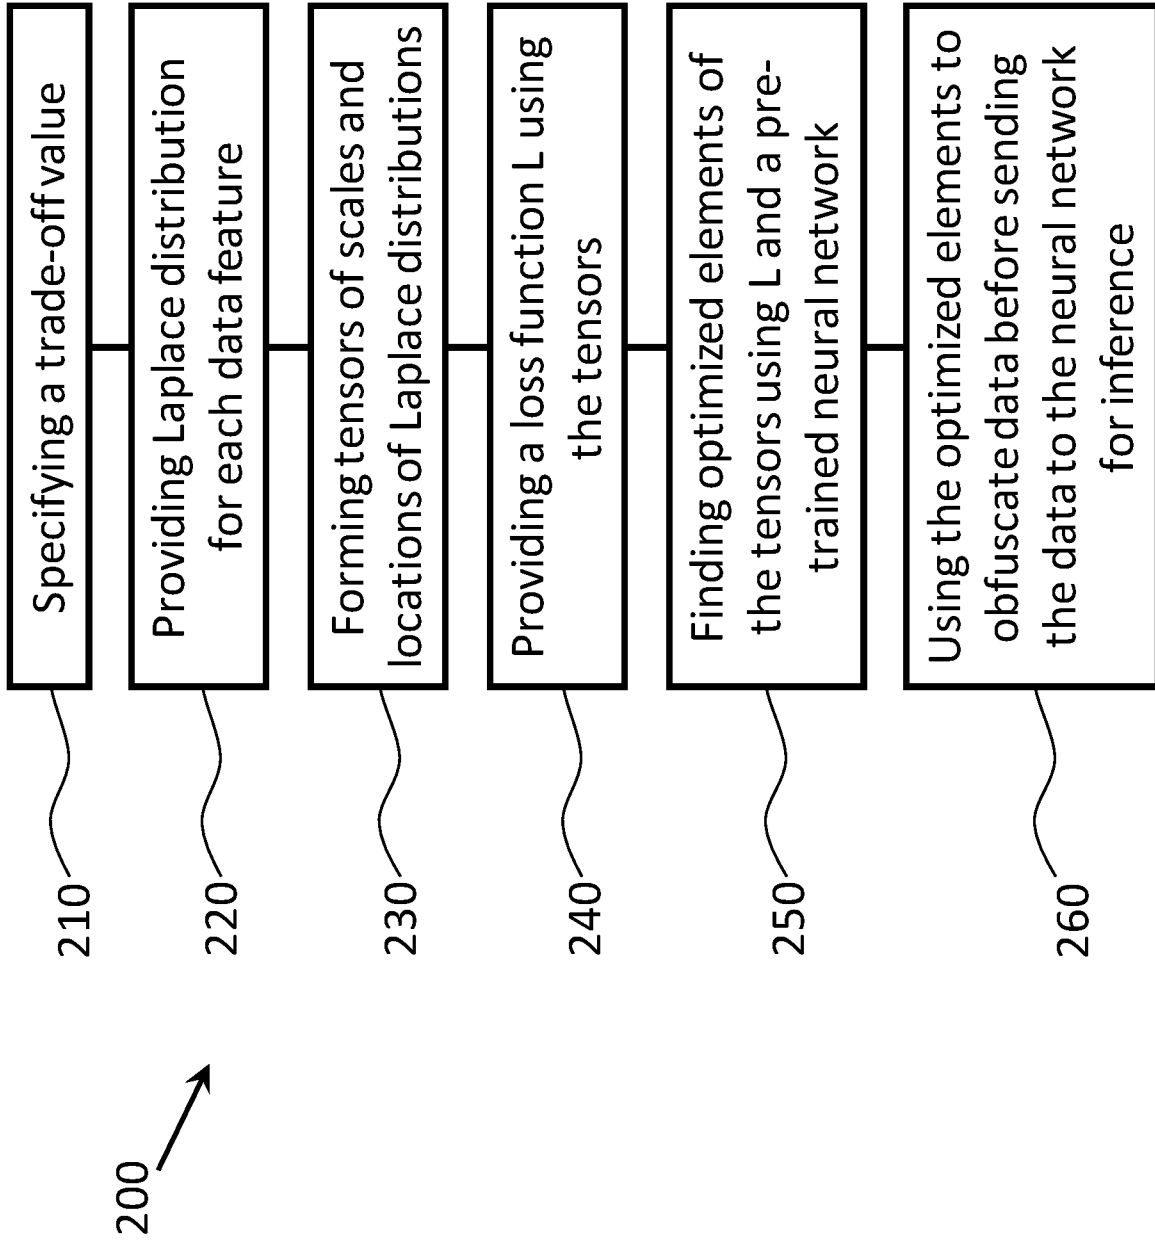


FIG. 2

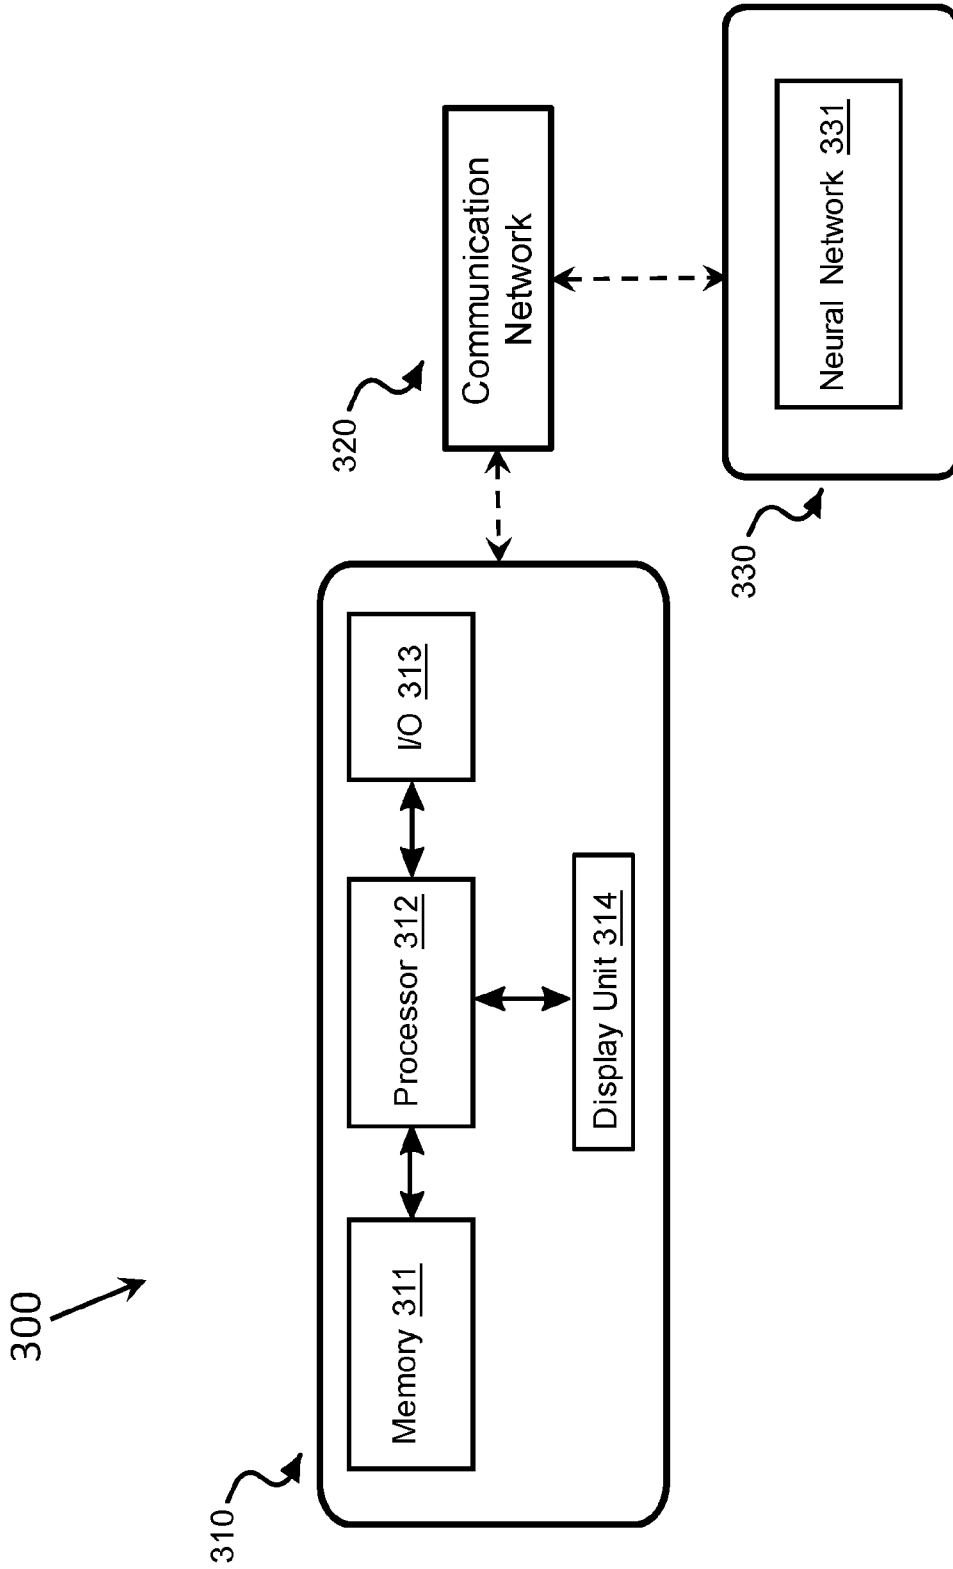


FIG. 3

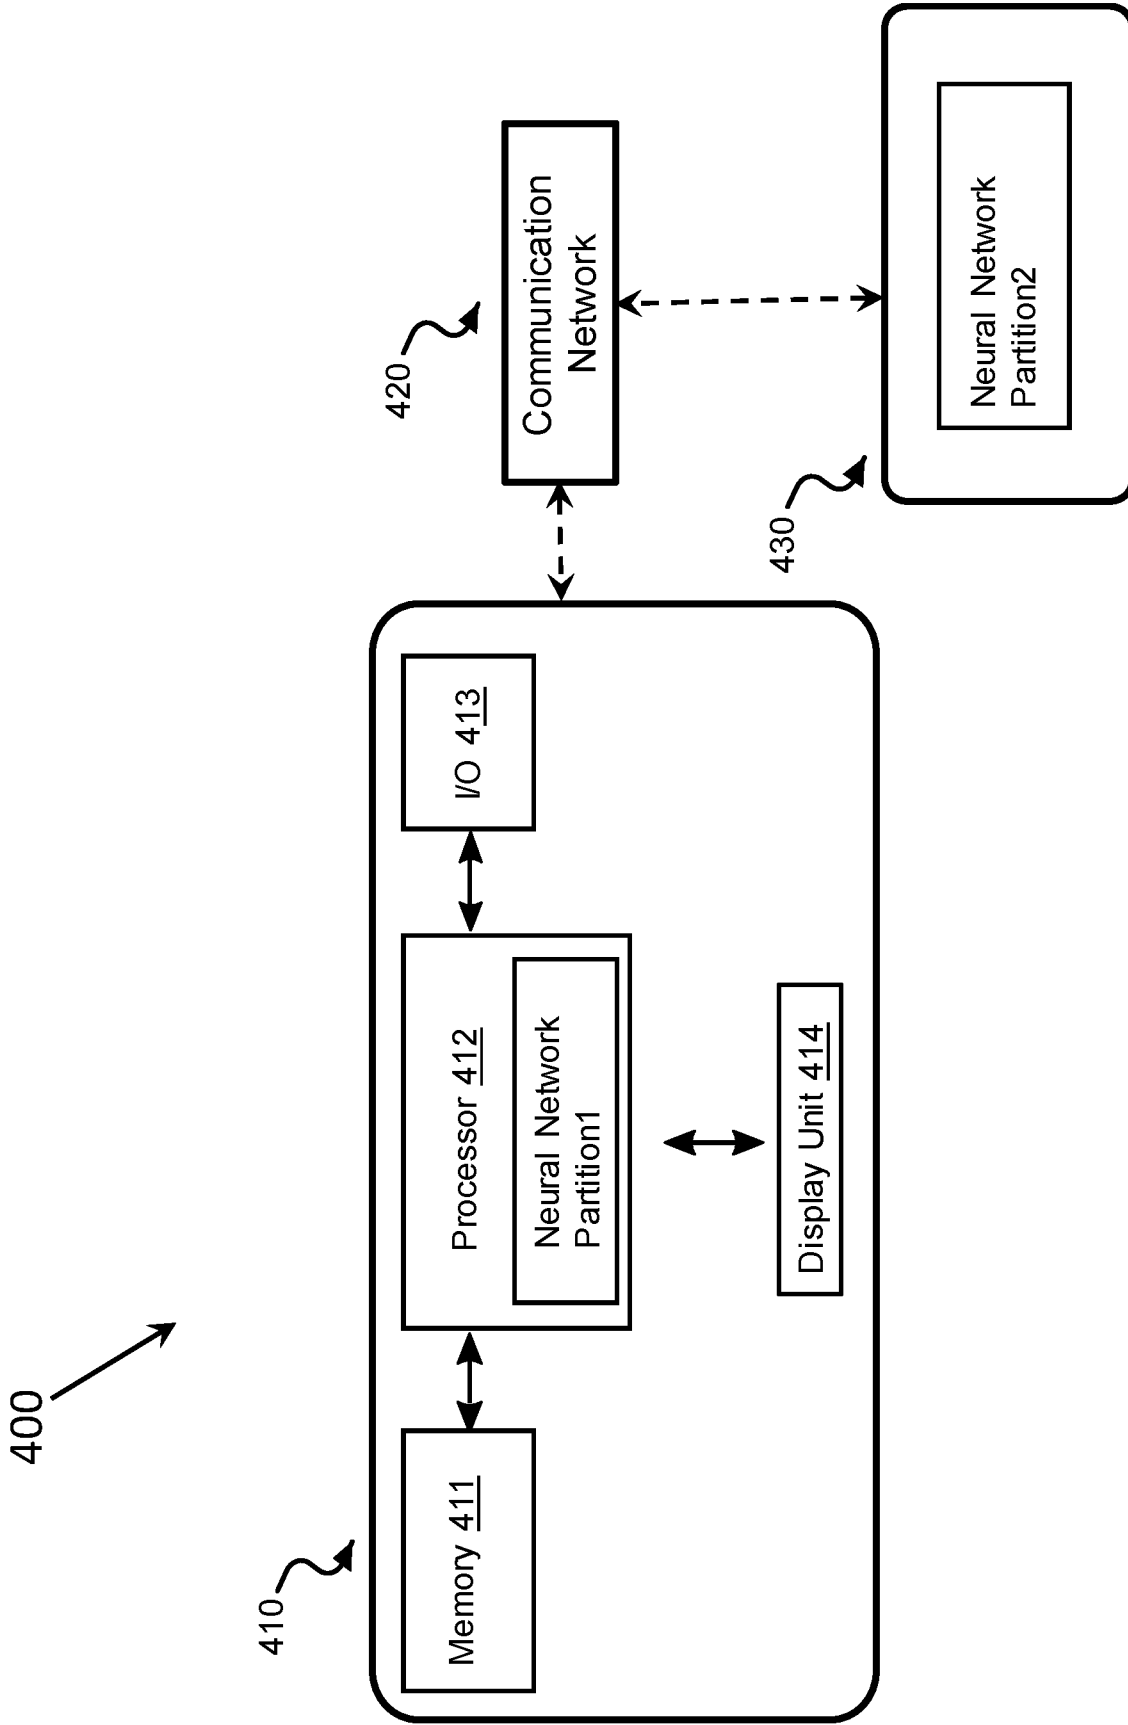


FIG. 4

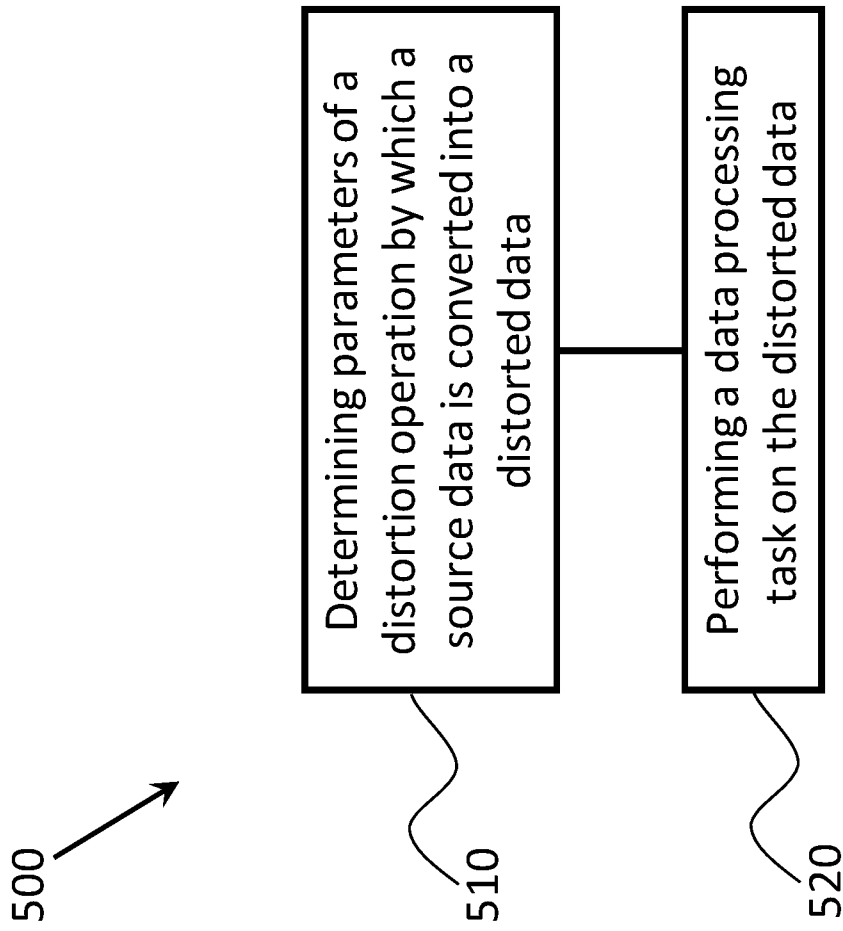


FIG. 5

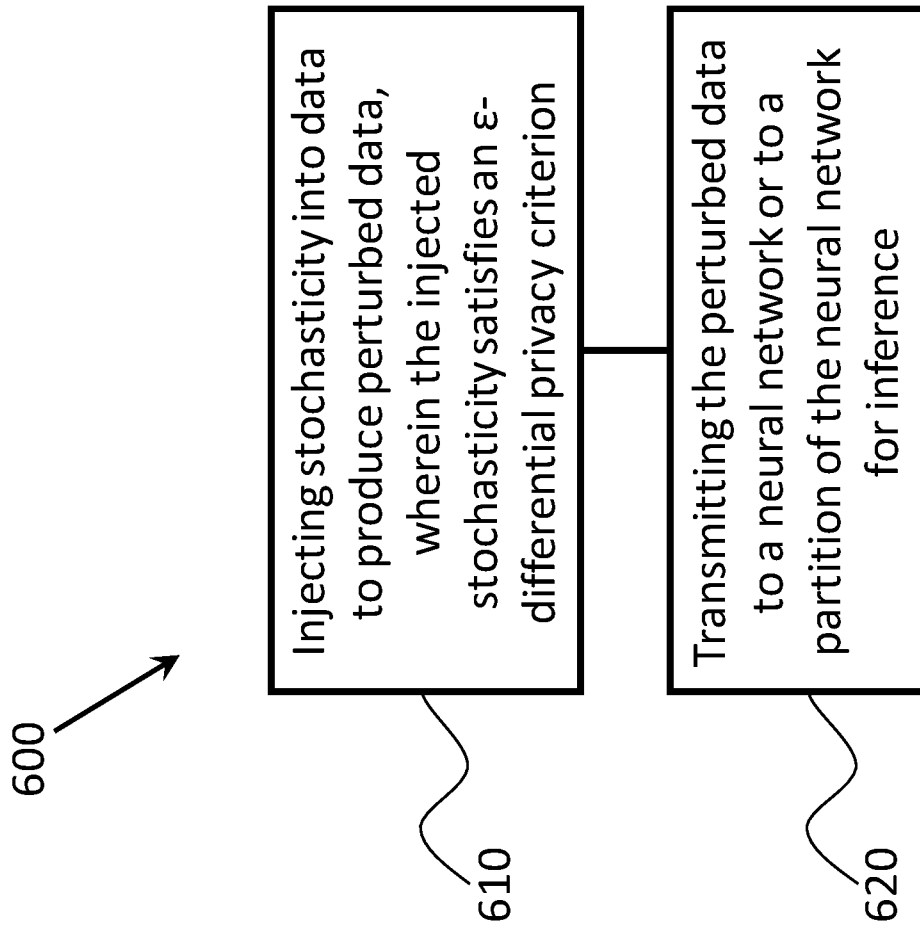


FIG. 6

700 ↗

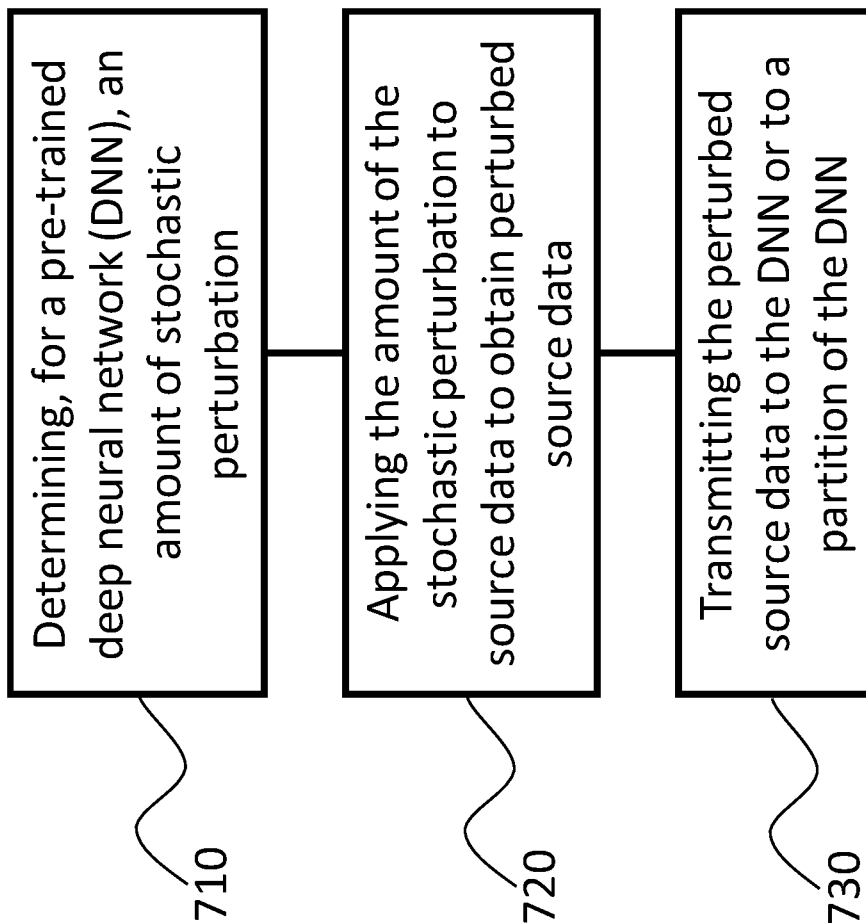


FIG. 7

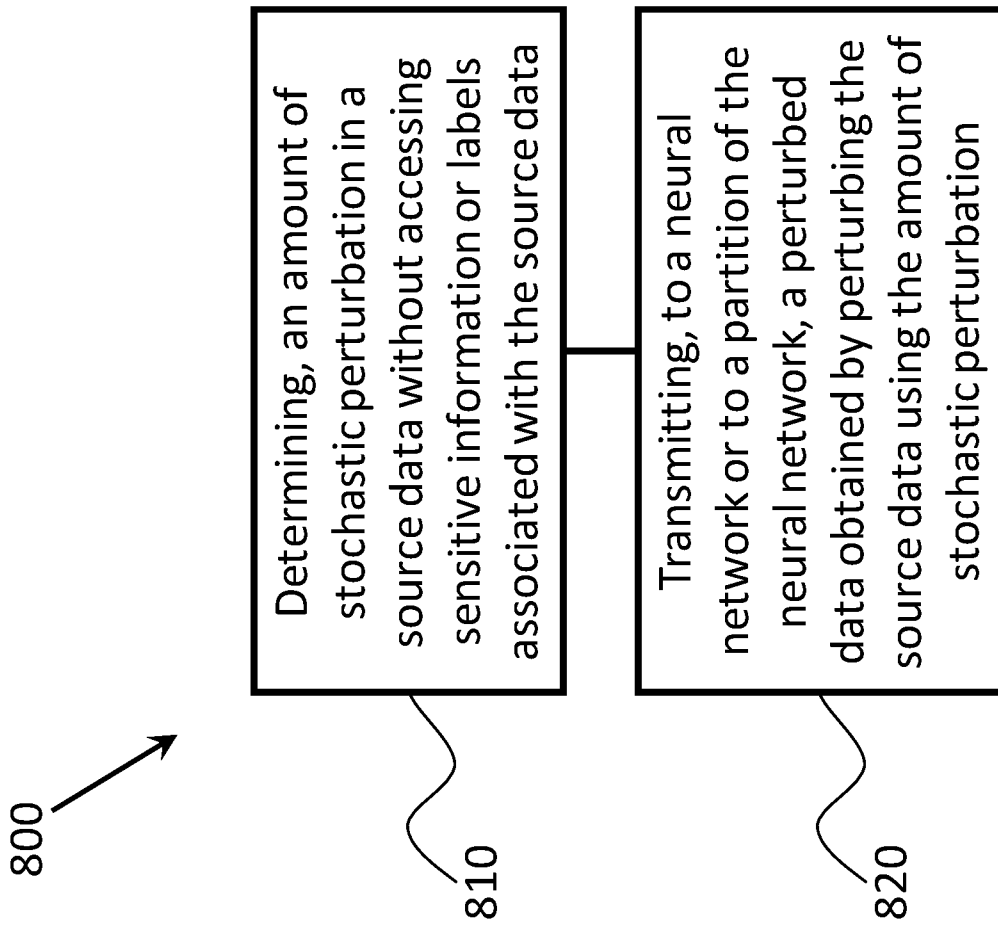


FIG. 8

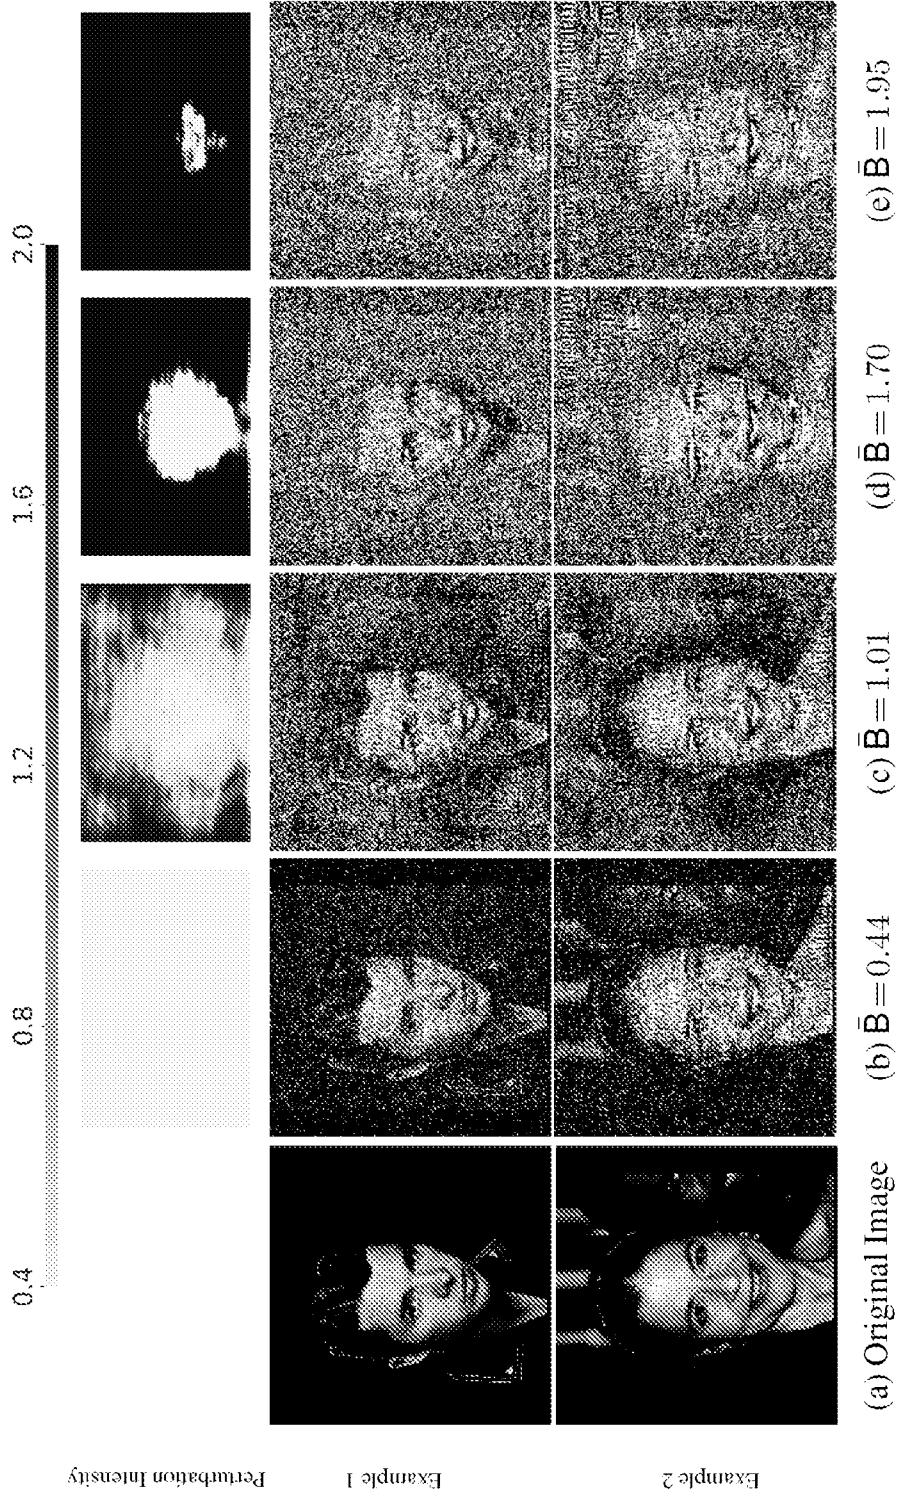


FIG. 9

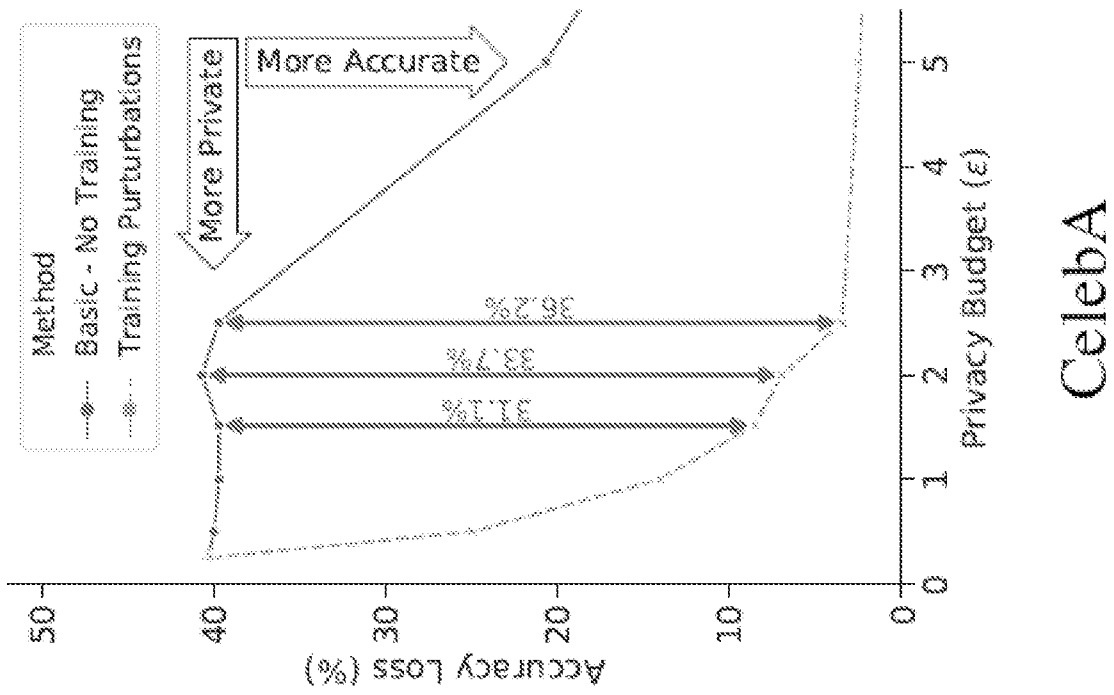


FIG. 10A

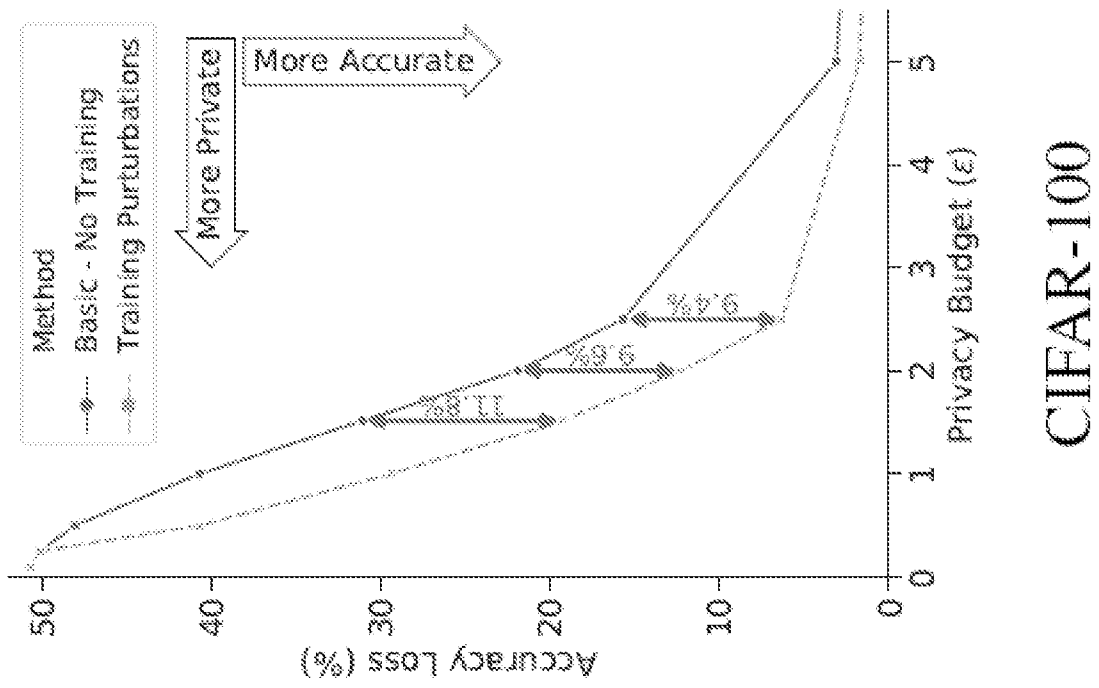


FIG. 10B

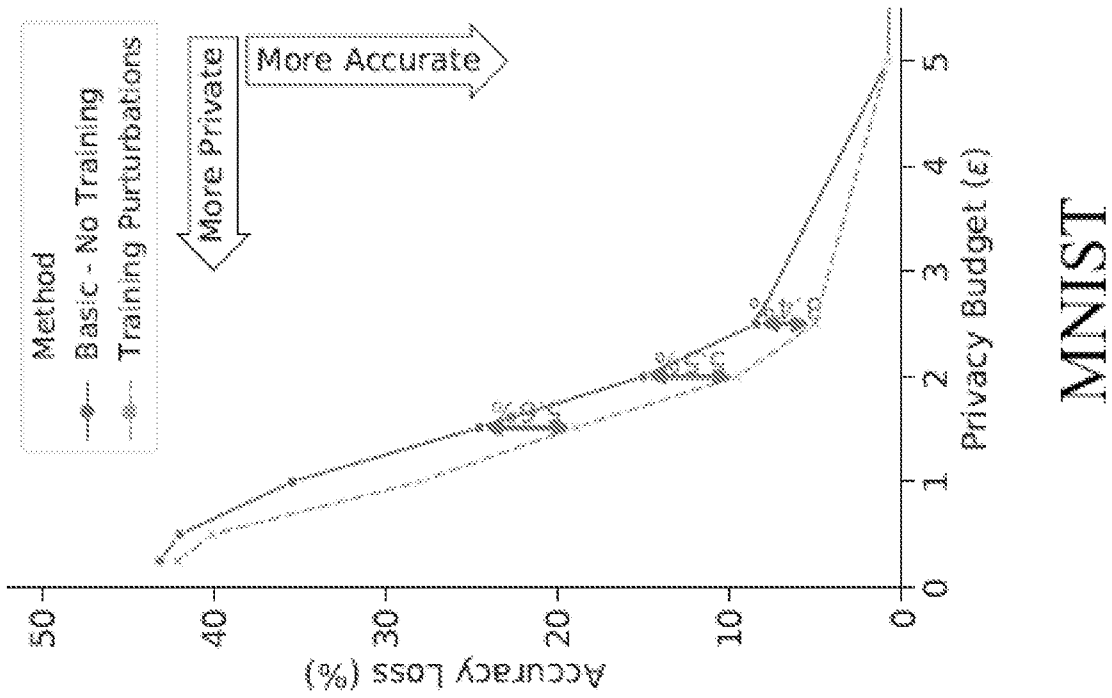


FIG. 10C

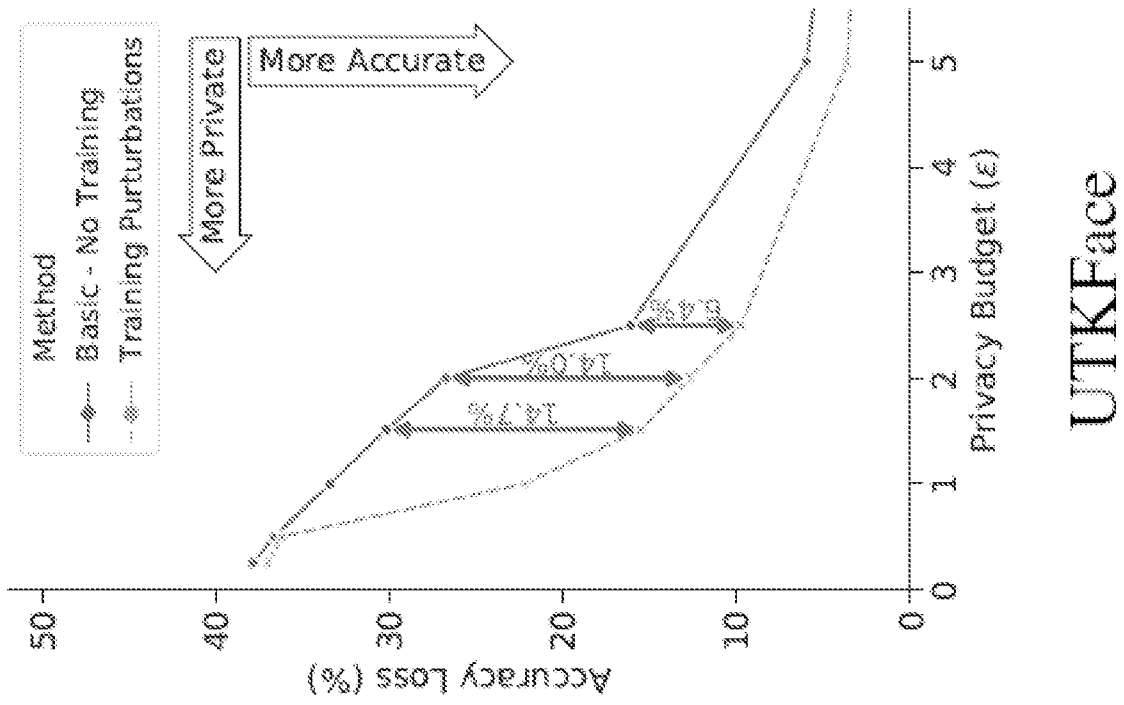


FIG. 10D

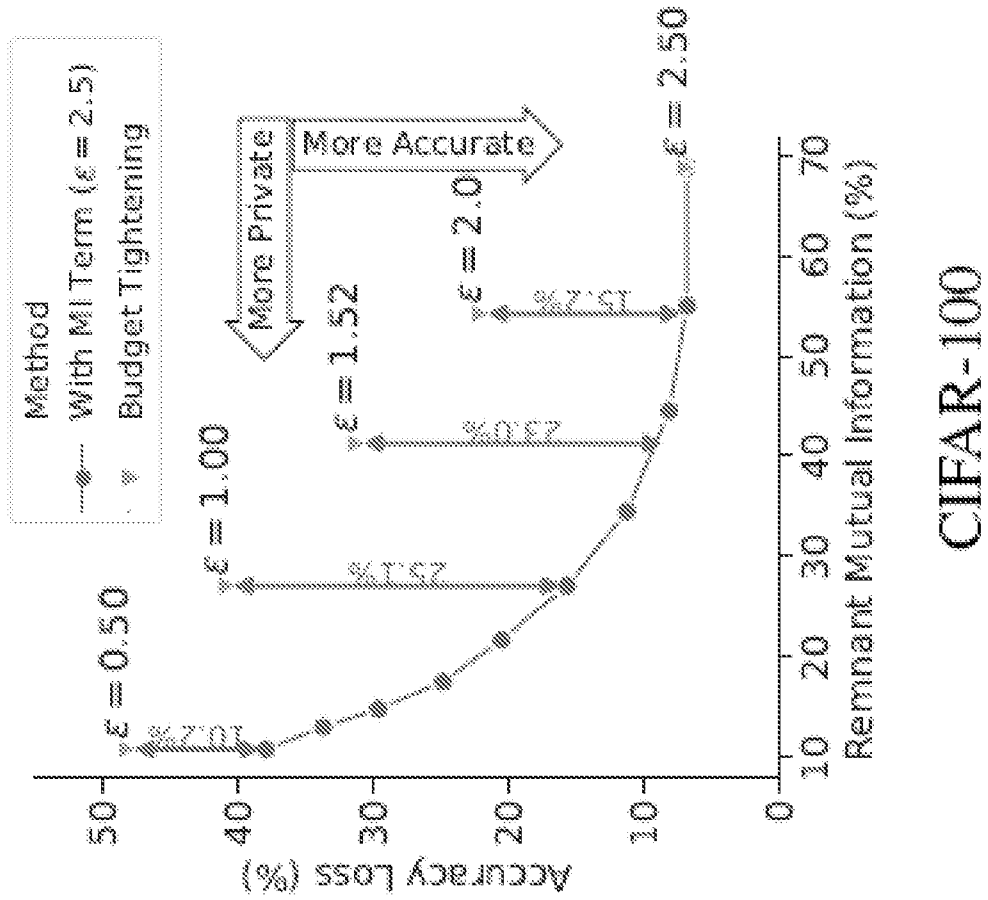
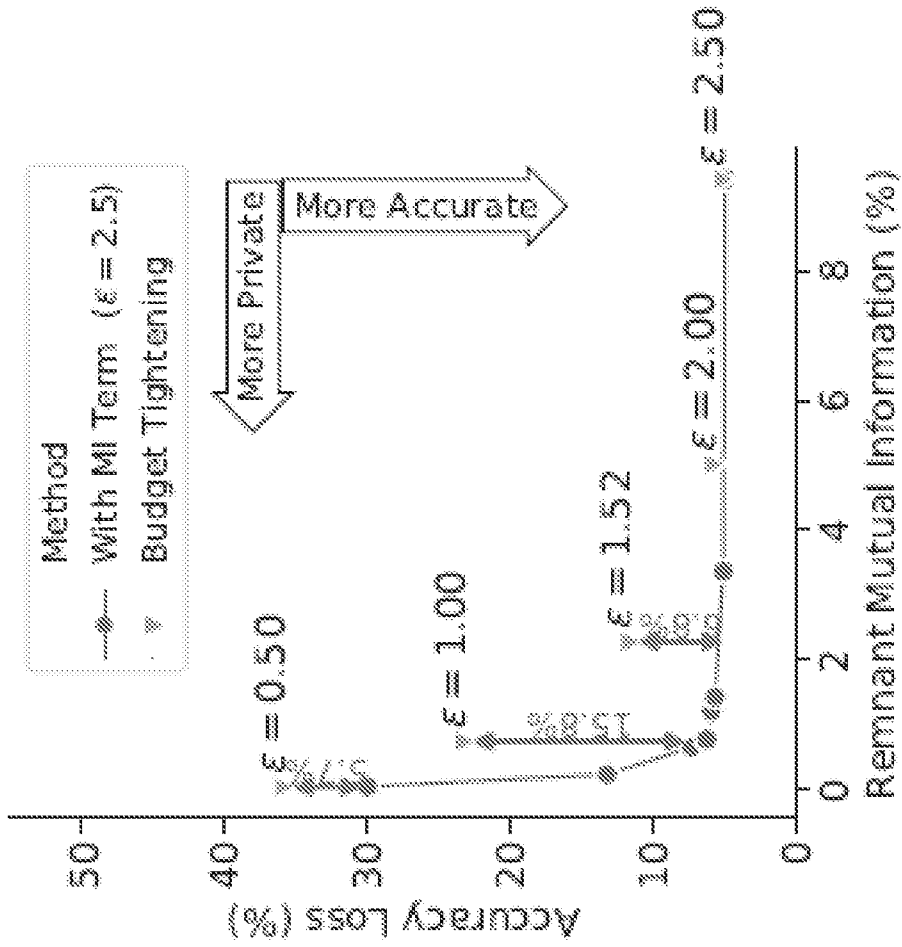
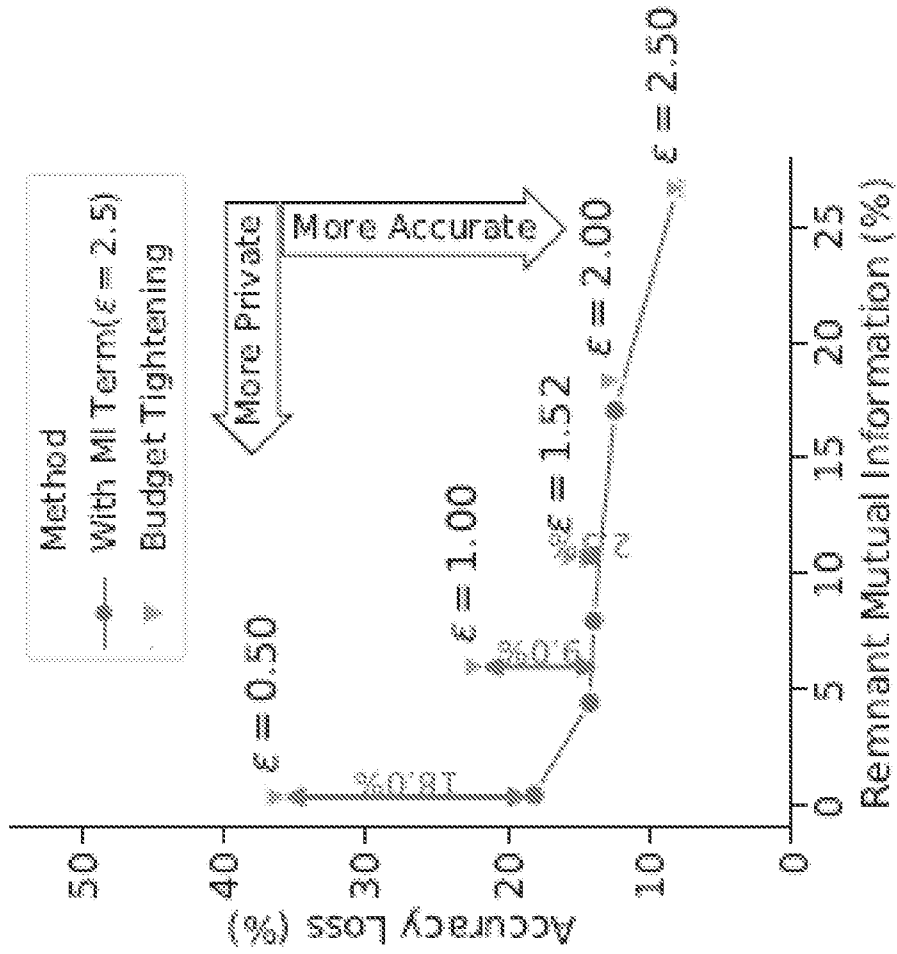


FIG. 11A



MNIST

FIG. 11B



UTKFace

FIG. 11C

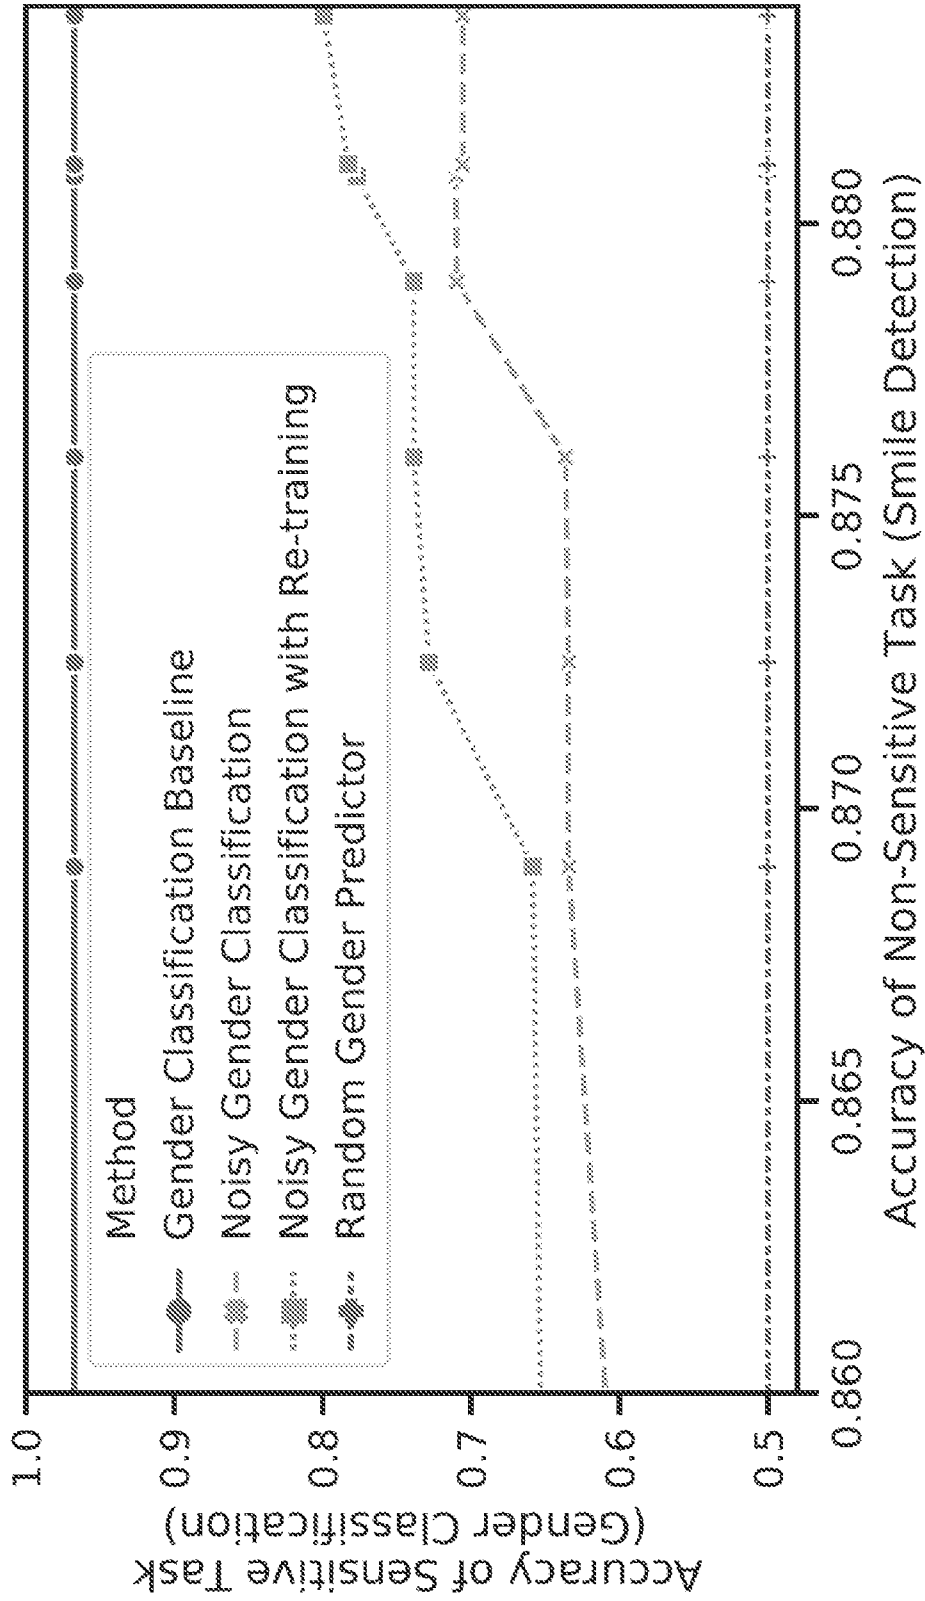


FIG. 12

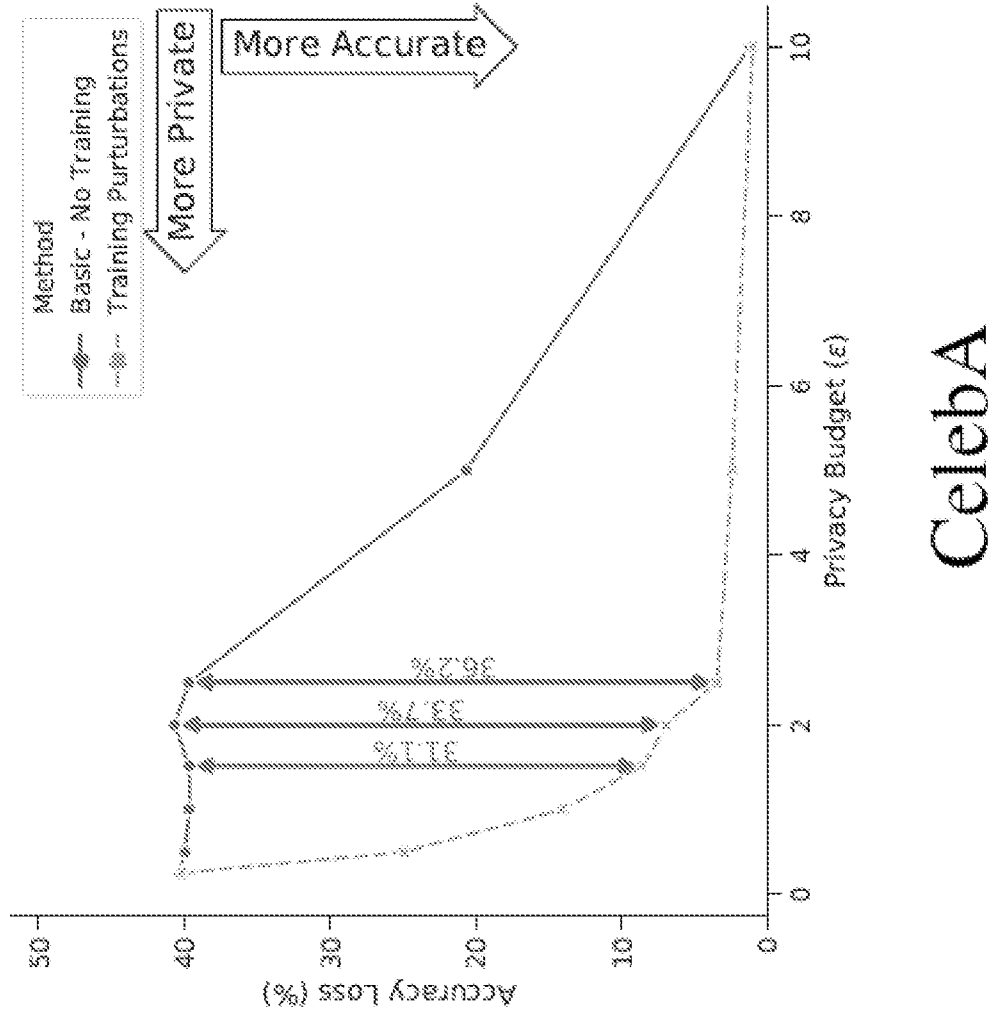
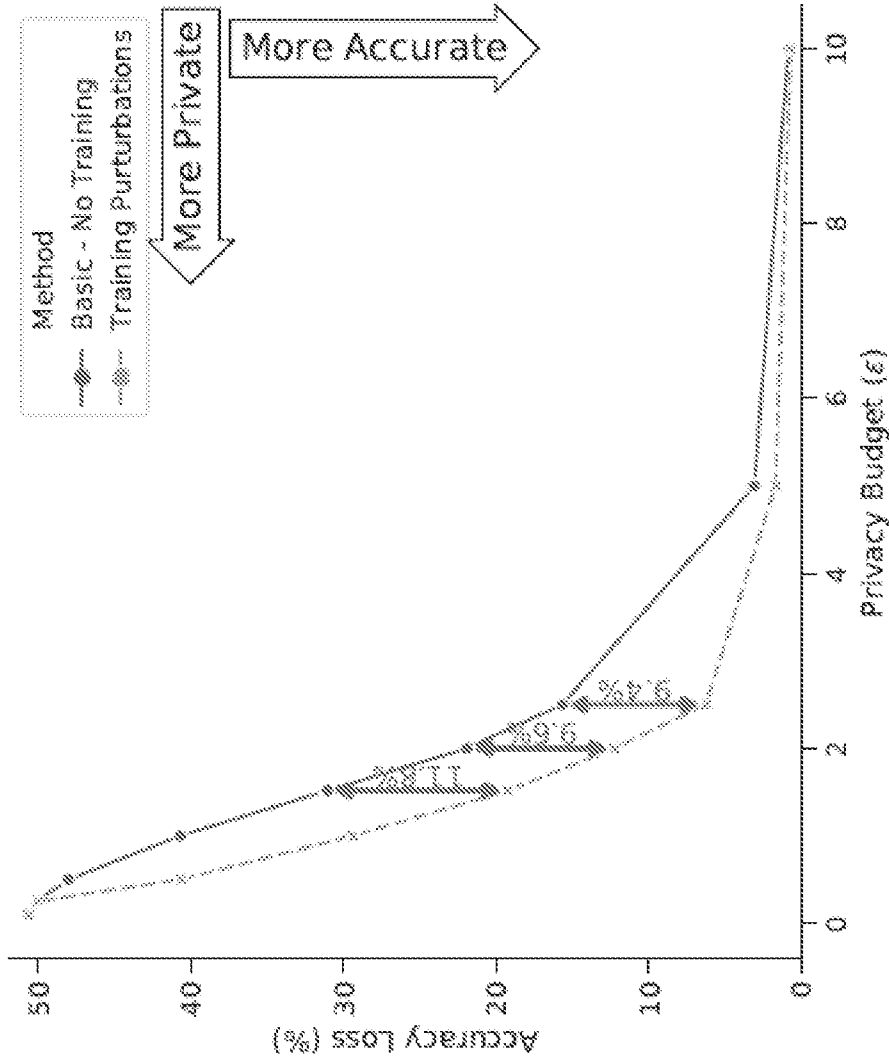
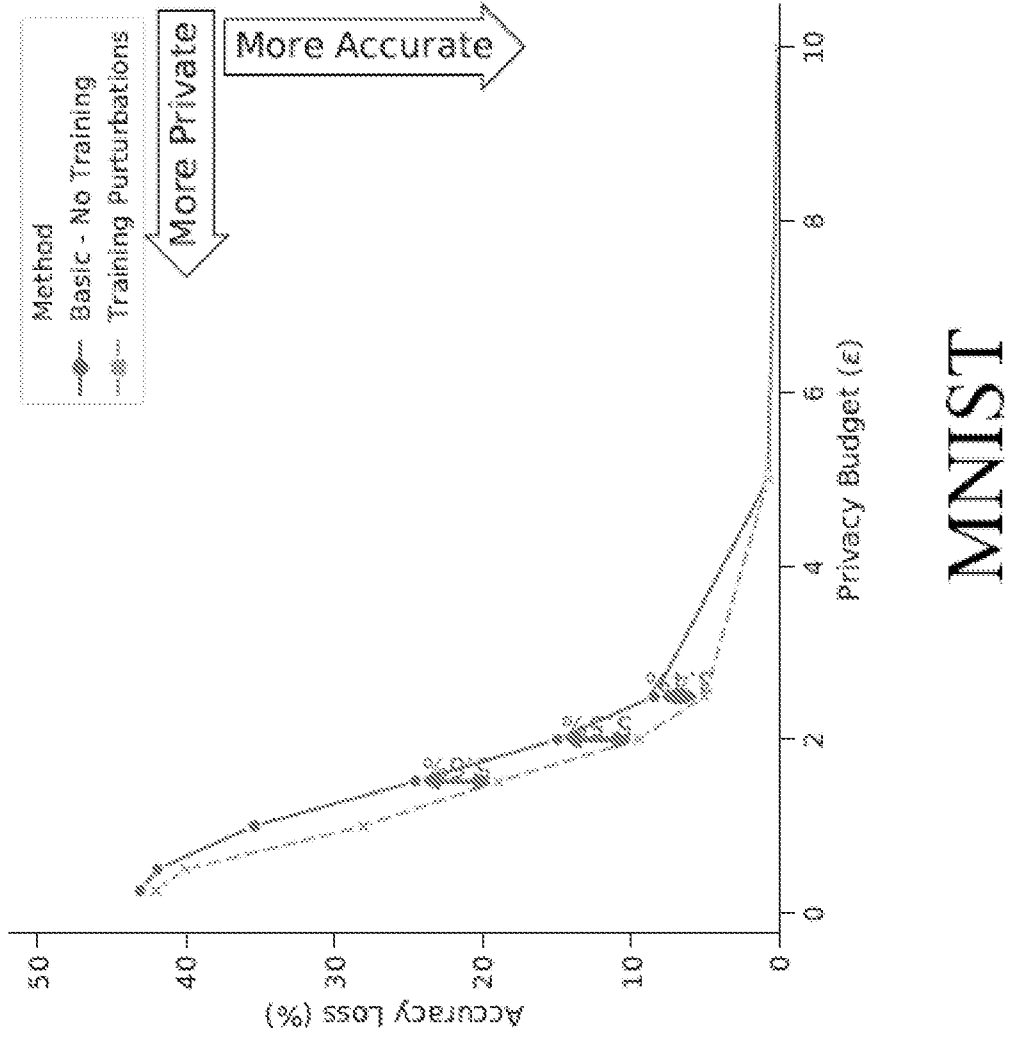


FIG. 13A



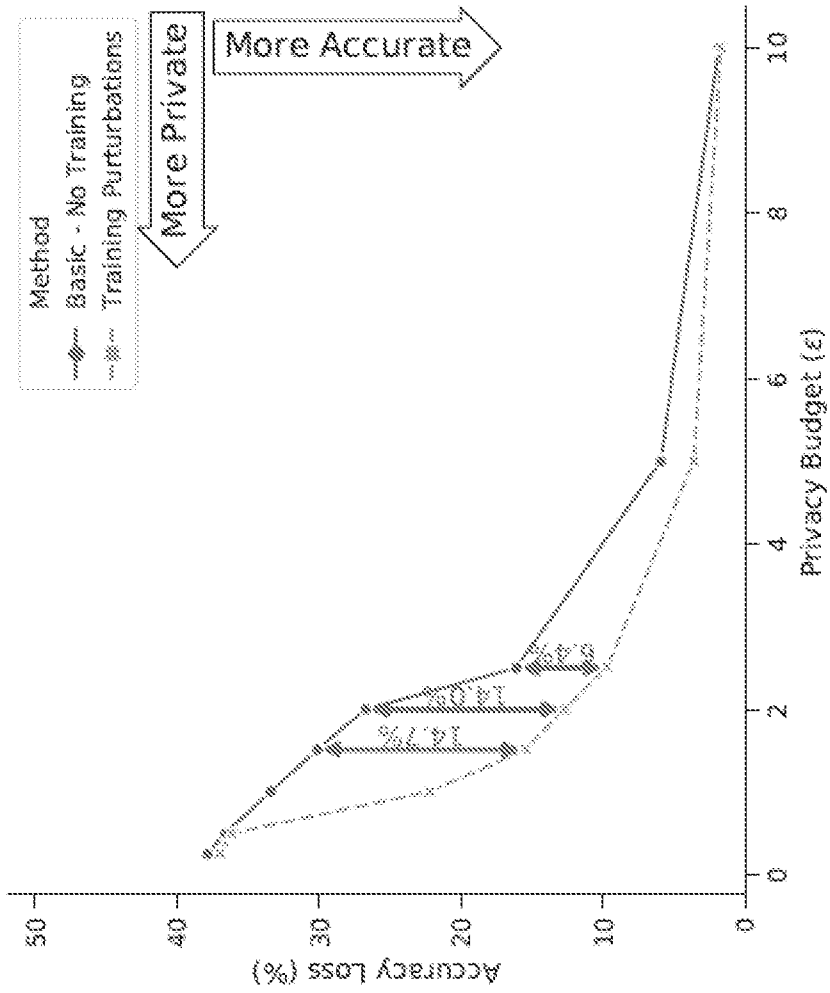
CIFAR-100

FIG. 13B



MNIST

FIG. 13C



UTKFace

FIG. 13D

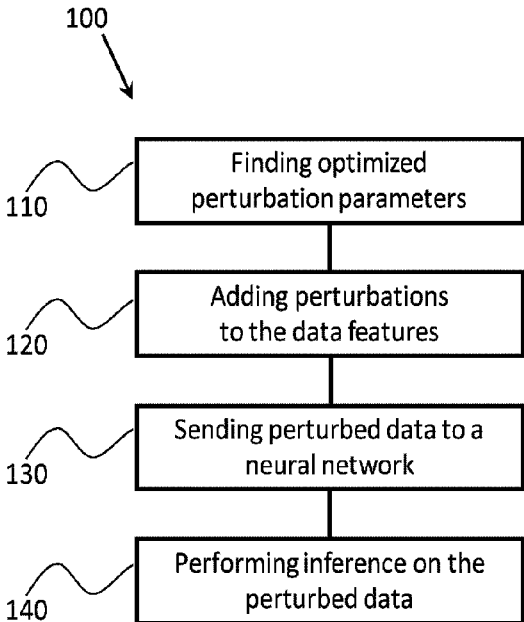


FIG. 1