



US 20180089587A1

(19) **United States**

(12) **Patent Application Publication**
Suresh et al.

(10) **Pub. No.: US 2018/0089587 A1**

(43) **Pub. Date: Mar. 29, 2018**

(54) **SYSTEMS AND METHODS FOR
COMMUNICATION EFFICIENT
DISTRIBUTED MEAN ESTIMATION**

Publication Classification

(51) **Int. Cl.**
G06N 99/00 (2006.01)
G06N 7/00 (2006.01)
(52) **U.S. Cl.**
CPC **G06N 99/005** (2013.01); **G06N 7/005**
(2013.01)

(71) Applicant: **Google Inc.**, Mountain View, CA (US)

(72) Inventors: **Ananda Theertha Suresh**, New York, NY (US); **Sanjiv Kumar**, Jericho, NY (US); **Hugh Brendan McMahan**, Seattle, WA (US); **Xinnan Yu**, Forest Hills, NY (US)

(21) Appl. No.: **15/676,076**


(22) Filed: **Aug. 14, 2017**

Related U.S. Application Data

(60) Provisional application No. 62/400,019, filed on Sep. 26, 2016.

(57) **ABSTRACT**

The present disclosure provides systems and methods for communication efficient distributed mean estimation. In particular, aspects of the present disclosure can be implemented by a system in which a number of vectors reside on a number of different clients, and a centralized server device seeks to estimate the mean of such vectors. According to one aspect of the present disclosure, a client computing device can rotate a vector by a random rotation matrix and then subsequently perform probabilistic quantization on the rotated vector. According to another aspect of the present disclosure, subsequent to quantization but prior to transmission, the client computing can encode the quantized vector according to a variable length coding scheme (e.g., by computing variable length codes).








FIG. 1




MNIST $k = 16$

FIG. 2A




MNIST $k = 32$

FIG. 2B




CIFAR $k = 16$

FIG. 2C




CIFAR $k = 32$

FIG. 2D




MNIST $k = 16$

FIG. 3A




MNIST $k = 32$

FIG. 3B



CIFAR $k = 16$

FIG. 3C



CIFAR $k = 32$

FIG. 3D




FIG. 4




FIG. 5




FIG. 6

SYSTEMS AND METHODS FOR COMMUNICATION EFFICIENT DISTRIBUTED MEAN ESTIMATION

FIELD

[0001] The present disclosure relates generally to distributed computing. More particularly, the present disclosure relates to systems and methods for communication efficient distributed mean estimation.

BACKGROUND

[0002] Given a number of vectors that reside on a number of different clients, the goal of distributed mean estimation is to estimate the mean of such vectors. This basic estimation problem is used as a subroutine in several learning and optimization tasks where data is distributed across several clients. For example, in Lloyd's algorithm for k-means clustering, if data is distributed across several clients, the server needs to compute the means of all clusters in each update step. Similarly, for principal components analysis (PCA), if data samples are distributed across several clients, then for the power-iteration method, the server needs to average the output of all clients in each step.

[0003] Recently, algorithms involving distributed mean estimation have been used extensively in training large-scale neural networks and other statistical models. In an example scenario of synchronized distributed learning, each client obtains a copy of a global model. The clients then update the model independently based on their local data. The updates (usually in the form of gradients) are then sent to a server, where they are averaged and used to update the global model. A critical step in all of the above algorithms is to estimate the mean of a set of vectors.

[0004] However, one of the main bottlenecks in distributed algorithms is the communication cost, which can be prohibitive for modern applications. For example, communication cost can be significant in example distributed computing systems where each client can be a low-power and/or low-bandwidth device such as, for example, a mobile phone, an embedded computing device, or other connected smart devices such as intelligent speakers, cameras, home appliances, vehicle computing systems, etc.

SUMMARY

[0005] Aspects and advantages of embodiments of the present disclosure will be set forth in part in the following description, or can be learned from the description, or can be learned through practice of the embodiments.

[0006] One aspect of the present disclosure is directed to a computing system to facilitate transmission of machine-learned model updates from client devices to a centralized server computing device. The computing system includes one or more client computing devices. Each client computing device includes one or more processors and one or more non-transitory computer-readable media that store instructions that, when executed by the one or more processors cause the client computing device to perform operations. The operations include determining an update to a machine-learned model based at least in part on a local dataset stored at the client computing device. The operations include rotating the update by a random rotation matrix to obtain a rotated update. The operations include performing probabilistic quantization of the rotated update to obtain a quantized

rotated update. The operations include transmitting the quantized rotated update to the centralized server computing device.

[0007] Another aspect of the present disclosure is directed to a computing system. The computing system includes one or more client computing devices. Each client computing device includes one or more processors and one or more non-transitory computer-readable media that store instructions that, when executed by the one or more processors cause the client computing device to perform operations. The operations include obtaining a vector. The operations include rotating the vector by a random rotation matrix to obtain a rotated vector. The operations include performing probabilistic quantization of the rotated vector to obtain a quantized rotated vector. The operations include transmitting the quantized rotated vector.

[0008] Another aspect of the present disclosure is directed to a computing system. The computing system includes one or more client computing devices. Each client computing device includes one or more processors and one or more non-transitory computer-readable media that store instructions that, when executed by the one or more processors cause the client computing device to perform operations. The operations include obtaining a vector. The operations include performing probabilistic quantization of the vector to obtain a quantized vector. Performing probabilistic quantization of the vector includes determining a value for each of a number of quantization levels based at least in part on a magnitude of the vector and a minimum coordinate value included in the vector. Performing probabilistic quantization of the vector includes quantizing each coordinate of the vector into one of the number of quantization levels.

[0009] Other aspects of the present disclosure are directed to various systems, apparatuses, non-transitory computer-readable media, user interfaces, and electronic devices.

[0010] These and other features, aspects, and advantages of various embodiments of the present disclosure will become better understood with reference to the following description and appended claims. The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate example embodiments of the present disclosure and, together with the description, serve to explain the related principles.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] Detailed discussion of embodiments directed to one of ordinary skill in the art is set forth in the specification, which makes reference to the appended figures, in which:

[0012] FIG. 1 depicts a graphical diagram of example results of distributed mean estimation on data generated from a Gaussian distribution according to example embodiments of the present disclosure.

[0013] FIGS. 2A-D depict graphical diagrams of example results for performance of Lloyd's algorithm with different types of quantizations according to example embodiments of the present disclosure.

[0014] FIGS. 3A-D depict graphical diagrams of example results for performance of power iteration with different types of quantizations according to example embodiments of the present disclosure.

[0015] FIG. 4 depicts a block diagram of an example computing system according to example embodiments of the present disclosure.

[0016] FIG. 5 depicts a swim lane flow diagram of an example method to perform a stochastic rotated quantization technique according to example embodiments of the present disclosure.

[0017] FIG. 6 depicts a swim lane flow diagram of an example method to perform a variable length coding technique according to example embodiments of the present disclosure.

[0018] Generally, the present disclosure provides systems and methods for communication efficient distributed mean estimation. In particular, the present disclosure addresses the need for distributed learning and optimization algorithms with low communication cost. However, unlike previous works, the systems and methods of the present disclosure make no probabilistic assumptions on the data.

[0019] Aspects of the present disclosure can be implemented by a system in which a number of vectors reside on a number of different client computing devices, and each client device seeks to transmit its respective vector to a centralized server computing device to enable the server device to estimate the mean of the vectors. As one example application, the techniques of the present disclosure can enable communication efficient uploads of local machine-learned model gradients from the client devices to the server device, where the server device aggregates the received gradients to update a global machine-learned model.

[0020] One aspect of the present disclosure is directed to a random rotation technique to improve communication efficiency. In particular, a client computing device can rotate a vector by a random rotation matrix and then subsequently perform probabilistic quantization on the rotated vector. For example, the present disclosure provides a stochastic k-level quantization technique. As demonstrated by the present disclosure, performing the random rotation prior to quantization can significantly reduce the quantization error, thereby leading to improved communication efficiency.

[0021] After quantization, the client device can transmit the quantized rotated vector to a centralized server computing device. The server computing device can receive multiple vectors from multiple client computing devices. The server computing device can determine a mean of the multiple received vectors and can de-rotate the mean vector using an inverse random rotation matrix.

[0022] In such fashion, the vectors can be communicated from the clients to the server in a highly efficient manner. In particular, in some implementations, this random rotation technique can reduce the mean square error of the mean estimation significantly by a factor of

$$\mathcal{O}\left(\frac{d}{\log d}\right).$$

[0023] Another aspect of the present disclosure is directed to a variable length coding technique to improve communication efficiency. In particular, subsequent to quantization but prior to transmission of the vector, the client computing device can encode the quantized vector according to a variable length coding scheme (e.g., by computing variable length codes). The server computing device can then decode the received vectors using the variable length coding scheme. This variable length coding technique can reduce the mean squared error by $\mathcal{O}(d)$ as compared to a naive approach that includes neither random rotation nor variable

length coding. In fact, the present disclosure mathematically demonstrates that this variable length coding approach is communication optimal.

[0024] The systems and methods of the present disclosure provide a number of technical effects and benefits. As a first example technical effect and benefit, the present disclosure enables transmission of information (e.g., machine-learned model updates) in a more efficient and lower bandwidth manner. For example, the random rotation technique and the variable length coding technique can reduce the number of bits required to represent the information, thereby enabling faster and lower bandwidth transmission of the information. As another example technical effect and benefit, by improving the ability to transmit vectors for aggregation at a central location, the present disclosure enables distributed training techniques in which machine-learned models are trained locally on-device and then a global model is generated or updated based on a mean of the updates resulting from the local training. Thus, the present disclosure enables and enhances such distributed scheme. The distributed training scheme results in better (e.g., more accurate) global models which have improved performance due to their exposure to additional training data. Furthermore, the distributed training scheme that is enhanced by the present disclosure improves user privacy as training data can be retained on the device and is not required to be sent to a central location. As yet another example technical effect and benefit, the present disclosure also provides enhanced security or encryption of information that is transmitted to a central location. For example, through the use of private rotation matrices the transmitted information can be rotated and only entities with the inverse of the private matrix can de-rotate to extract the information.

[0025] The present disclosure is structured as follows: first, the present disclosure demonstrates that for d dimensional data with n clients, a naive stochastic rounding approach yields a mean squared error (MSE) of $\Theta(d/n)$ and uses a constant number of bits per dimension per client. Next, this naive algorithm is extended in two ways: by demonstrating that application of a structured random rotation before quantization reduces the error to $\mathcal{O}((\log d)/n)$ and application of a better coding strategy further reduces the error to $\mathcal{O}(1/n)$. The present disclosure also demonstrates that the latter coding strategy is optimal up to a constant in the minimax sense. That is, it achieves the best MSE for a given communication cost. In addition, the present disclosure demonstrates the practicality of the algorithms described herein by applying them to distributed Lloyd's algorithm for k-means and power iteration for principal component analysis (PCA).

1.1 Example Discussion of Distributed Mean Aggregation

[0026] Given n vectors $X^n \stackrel{\text{def}}{=} X_1, X_2, \dots, X_n \in \mathbb{R}^d$ that reside on n clients, the goal of distributed mean estimation is to estimate the mean of the vectors:

$$\bar{X} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n X_i. \quad (1)$$

[0027] As described above, this basic estimation problem is used in several learning and optimization tasks where data

is distributed across several clients, including, for example, Lloyd's algorithm for k-means clustering (see, e.g., Lloyd, Stuart. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129-137, 1982); principal components analysis (PCA); and the training of large-scale neural networks and other statistical models. For example, in an example scenario of synchronized distributed learning, each client obtains a copy of a global model. The clients then update the model independently based on their local data. The updates (usually in the form of gradients) are then sent to a server, where they are averaged and used to update the global model. A critical step in all of the above algorithms is to estimate the mean of a set of vectors as in Eq. Error! Reference source not found.

[0028] More generally, however, the communication efficient systems, methods, and techniques of the present disclosure can involve or be performed in any configuration or scenario where data (e.g., in the form of vectors) resides on or is produced by one or more different client devices and is communicated to another device such as, for example, a central server device.

[0029] In particular, one of the main bottlenecks in distributed algorithms or other scenarios where data is communicated from a client device to a server device is the communication cost, which can be prohibitive for modern applications. For example, communication cost can be significant in example distributed computing systems where each client can be a low-power and/or low-bandwidth device such as, for example, a mobile phone or other mobile computing device, an embedded computing device, a worker device in a datacenter, or other connected smart devices such as intelligent speakers, cameras, home appliances, vehicle computing systems, etc.

[0030] Given such a wide set of applications, the systems and methods of the present disclosure address the basic problem of achieving the optimal minimax rate in distributed mean estimation with limited communication or other communication constraints (e.g., communication costs).

[0031] The systems, methods, models, and results of the present disclosure differ from previous works on mean estimation in a number of ways, including the following examples. As a first example, certain previous works assume that the data is generated independent and identically distributed (i.i.d.) according to some distribution. In contrast, the present disclosure does not make any distribution assumptions on data. As a second example difference, the objective in certain prior works is to estimate the mean of the underlying statistical model, while the goal in at least some implementations of the present disclosure is to estimate the empirical mean of the data.

1.2 Example Model

[0032] In at least some implementations, the communication algorithms provided herein are simultaneous and independent. That is, the clients independently send data to the server and they can transmit at the same time. In any independent communication protocol, each client transmits a function of X_i (say $f(X_i)$), and a central server estimates the mean by some function of $f(X_1), f(X_2), \dots, f(X_n)$. Let π be any such protocol and let $C_i(\pi, X_i)$ be the expected number of transmitted bits by the i -th client during protocol π , where throughout the present disclosure, expectation is over the randomness in protocol π .

[0033] The total number of bits transmitted by all clients with the protocol π is

$$C(\pi, X^n) \stackrel{\text{def}}{=} \sum_{i=1}^n C_i(\pi, X_i).$$

[0034] Let the estimated mean be $\hat{X}$. For a protocol π , the MSE of the estimate is

$$\mathcal{E}(\pi, X^n) = \mathbb{E}[\|\hat{X} - \bar{X}\|_2^2].$$

[0035] The systems and methods of the present disclosure can be implemented with the use of private and/or public randomness. Private randomness refers to random values that are generated by each machine separately, and public randomness refers to a sequence of random values that are shared among all parties.

[0036] As one example, in the absence of public randomness, the server can communicate a random seed that can be used by clients to emulate public randomness. More particularly, in some implementations, both the client and the server can have knowledge of a synchronized random seed. The client and the server can each use the random seed (e.g., along with a pseudo-random number generator) to separately produce random values (e.g., a random rotation matrix as will be discussed further below). In one example, the seed can be 32 bits long.

[0037] The algorithms provided by the present disclosure work for any X^n . However, to measure the minimax performance, without loss of generality, the following is provided with reference to the scenario where each $X_i \in S^d$, the ball of radius 1 in $\mathbb{R}^d$ i.e., $X \in S_d$ iff

$$\|X\|_2 \leq 1,$$

where $\|X\|_2$ denotes the l_2 norm of the vector X . For a protocol π , the worst case error for all $X^n \in S^d$ is

$$\mathcal{E}(\pi, S^d) \stackrel{\text{def}}{=} \max_{X^n: X_i \in S^d \forall i} \mathcal{E}(\pi, X^n).$$

[0038] Let $\Pi(c)$ denote the set of all protocols with communication cost at most c . The minimax MSE is

$$\mathcal{E}(\Pi(c), S^d) \stackrel{\text{def}}{=} \min_{\pi \in \Pi(c)} \mathcal{E}(\pi, S^d).$$

1.3 Example Results and Discussion

1.3.1 Example Algorithms

[0039] The MSE $\mathcal{E}(\pi, X^n)$ is first analyzed for three algorithms, when $C(\pi, X^n) = \Theta(nd)$, i.e., each client sends a constant number of bits per dimension.

[0040] Stochastic Uniform Quantization.

[0041] In Section 2.1, it is first shown that a naive stochastic binary quantization algorithm (denoted by π_{sb}) achieves an MSE of

$$\mathcal{E}(\pi_{sb}, X^n) = \mathcal{O}\left(\frac{d}{n} \cdot \frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2\right),$$

and $\mathcal{C}(\pi_{sb}, X^n) = n \cdot (d + \tilde{\mathcal{O}}(1))$, where $\tilde{\mathcal{O}}(1)$ is used to denote $\mathcal{O}(\log(dn))$. That is, each client sends one bit per dimension. It is further shown that this bound is tight. In many practical scenarios, d is much larger than n and the above error is prohibitive.

[0042] A natural way to decrease the error is to increase the number of levels of quantization. If k levels of quantization are used, in Theorem 2, the error decreases as

$$\mathcal{E}(\pi_{sk}, X^n) = \mathcal{O}\left(\frac{d}{n(k-1)^2} \cdot \frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2\right). \quad (2)$$

[0043] However, the communication cost would increase to $\mathcal{C}(\pi_{sk}, X^n) = n \cdot (d \lceil \log_2 k \rceil + \tilde{\mathcal{O}}(1))$ bits, which can be expensive, if the MSE is desired to be $\mathcal{O}(d/n)$.

[0044] In order to reduce the communication cost, the present disclosure proposes two approaches.

[0045] Stochastic Rotated Quantization:

[0046] The present disclosure demonstrates that preprocessing the data by a random rotation reduces the mean squared error. Specifically, in Theorem 3, it is shown that this new scheme (denoted by π_{srk}) achieves an MSE of

$$\mathcal{E}(\pi_{srk}, X^n) = \mathcal{O}\left(\frac{\log d}{n(k-1)^2} \cdot \frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2\right),$$

and has a communication cost of $\mathcal{C}(\pi_{srk}, X^n) = n \cdot (d \lceil \log_2 k \rceil + \tilde{\mathcal{O}}(1))$. Note that all logarithms are to base e , unless stated otherwise. Also note that the new scheme achieves much smaller MSE than naive stochastic quantization for the same communication cost.

[0047] Variable Length Coding:

[0048] A second approach provided by the present disclosure uses the same quantization as π_{sk} but encodes levels via variable length coding. Instead of using $\lceil \log_2 k \rceil$ bits per dimension, it is shown that using variable length encoding such as arithmetic coding to compress the data reduces the communication cost significantly. In particular, in Theorem 4 it is shown that there is a scheme (denoted by π_{svk}) such that

$$\mathcal{C}(\pi_{svk}, X^n) = \mathcal{O}\left(nd \left(1 + \log\left(\frac{k^2}{d} + 1\right)\right) + \tilde{\mathcal{O}}(n)\right), \quad (3)$$

and

$$\mathcal{E}(\pi_{svk}, X^n) = \mathcal{E}(\pi_{sk}, X^n).$$

[0049] Hence, setting $k = \sqrt{d}$ in Eqs. Error! Reference source not found. and Error! Reference source not found. yields

$$\mathcal{E}(\pi_{svk}, X^n) = \mathcal{O}\left(\frac{1}{n} \cdot \frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2\right),$$

and with $\Theta(nd)$ bits of communication i.e., constant number of bits per dimension per client. Of the three protocols, π_{svk} has the best MSE for a given communication cost. Note that π_{svk} uses k quantization levels but still uses $\tilde{\mathcal{O}}(1)$ bits per dimension per client for all $k \leq \sqrt{d}$.

[0050] Theoretically, while variable length coding has better guarantees, stochastic rotated quantization has several practical advantages: it uses fixed length coding and hence can be combined with encryption schemes for privacy preserving secure aggregation. It can also provide lower quantization error in some scenarios due to better constants (see, e.g., Section 7 for details).

1.3.2 Example Minimax MSE

[0051] In the above protocols, all of the clients transmit the data. According to an aspect of the present disclosure, these protocols can be augmented with a sampling procedure, where only a random fraction of clients transmit data. The present disclosure demonstrates that a combination of k -level quantization, variable length coding, and sampling can be used to achieve information theoretically optimal MSE for a given communication cost. In particular, combining Corollary 1 and Theorem 5 yields the following minimax result:

[0052] Theorem 1 There exists a universal constant $t < 1$ such that for communication cost $c \leq ndt$ and $n \geq 1/t$,

$$\mathcal{E}(\Pi(c), S^d) = \Theta\left(\min\left(1, \frac{d}{c}\right)\right).$$

[0053] This result shows that the product of communication cost and MSE scales linearly in the number of dimensions.

[0054] The remainder of the present disclosure is organized as follows. First, the stochastic uniform quantization technique is analyzed in Section 2. In Section 3, a novel stochastic rotated quantization technique is proposed, and in Section 4 arithmetic coding is analyzed. In Section 5, the above algorithms are combined with a sampling technique and the upper bound on the minimax risk is stated. In Section 6, the matching minimax lower bounds are stated. In Section 7 some practical considerations are discussed and the algorithms are applied on distributed power iteration and Lloyd's algorithm. Section 8 provides a proof of Lemma 7. Section 9 describes example computing systems that can implement the techniques described herein. Section 10 describes example methods to implement the techniques described herein and Section 11 provides some additional disclosure.

2. Example Stochastic Uniform Quantization

2.1 Example Stochastic Binary Quantization

[0055] For a vector X_i , let $X_i^{\max} = \max_{1 \leq j \leq d} X_i(j)$ and similarly let $X_i^{\min} = \min_{1 \leq j \leq d} X_i(j)$. In the stochastic binary quan-

tization protocol π_{sb} , for each client i , the quantized value for each coordinate j is generated independently with private randomness as

$$Y_i(j) = \begin{cases} X_i^{max} & \text{with probability } \frac{X_i(j) - X_i^{min}}{X_i^{max} - X_i^{min}}, \\ X_i^{min} & \text{otherwise.} \end{cases}$$

[0056] Observe $\mathbb{E} Y_i(j) = X_i(j)$. The server estimates $\bar{X}$ by

$$\hat{X}_{\pi_{sb}} = \frac{1}{n} \sum_{i=1}^n Y_i.$$

[0057] First, the communication cost of this protocol will be bounded as described below.

[0058] Lemma 1 There exists an implementation of stochastic binary quantization that uses $d + \tilde{O}(1)$ bits per client and hence $\mathcal{C}(\pi_{sb}, X^n) \leq n \cdot (d + \tilde{O}(1))$.

[0059] Proof. Instead of sending vectors Y_i , clients transmit two real values X_i^{max} and X_i^{min} (to a desired error) and a bit vector Y'_i such that $Y'_i(j) = 1$ if $Y_i(j) = X_i^{max}$ and 0 otherwise. Hence each client transmits $d + 2r$ bits, where r is the number of bits to transmit the real value to a desired error.

[0060] Let B be the maximum norm of the underlying vectors. To bound r , observe that using r bits, one can represent a number between $-B$ and B to an error of $B/2^{r-1}$. Thus using $3 \log_2(dn) + 1$ bits one can represent the minimum and maximum to an additive error of $B/(nd)^3$. This error in transmitting minimum and maximum of the vector does not affect the calculations and it is ignored for simplicity. Note that in practice, each dimension of X_i is often stored as a 32 bit or 64 bit float, and r should be set as either 32 or 64. In this case, using an even larger r does not further reduce the error.

[0061] End proof.

[0062] Next, the estimation error of this protocol is computed as described below.

[0063] Lemma 2 For any set of vectors X^n ,

$$\mathcal{E}(\pi_{sb}, X^n) = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^d (X_i^{max} - X_i(j))(X_i(j) - X_i^{min}).$$

Proof.

$$\begin{aligned} \mathcal{E}(\pi_{sb}, X^n) &= \mathbb{E} \|\hat{X} - \bar{X}\|_2^2 \\ &= \frac{1}{n^2} \mathbb{E} \left\| \sum_{i=1}^n (Y_i - X_i) \right\|_2^2 \\ &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \|Y_i - X_i\|_2^2, \end{aligned}$$

where the last equality follows by observing that $Y_i - X_i$, $\forall i$, are independent zero mean random variables. The proof follows by observing that for every i ,

$$\begin{aligned} \mathbb{E} \|Y_i - X_i\|_2^2 &= \sum_{j=1}^d \mathbb{E} [(Y_i(j) - X_i(j))^2] \\ &= \sum_{j=1}^d (X_i^{max} - X_i(j))(X_i(j) - X_i^{min}). \end{aligned}$$

[0064] End proof.

[0065] Lemma 2 implies the following upper bound.

[0066] Lemma 3 For any set of vectors X^n ,

$$\mathcal{E}(\pi_{sb}, X^n) \leq \frac{d-2}{2n} \cdot \frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2.$$

[0067] Proof. The proof follows by Lemma 2 observing that $\forall j$

$$(X_i^{max} - X_i(j))(X_i(j) - X_i^{min}) \leq \frac{(X_i^{max} - X_i^{min})^2}{4}, \quad (4)$$

and

$$(X_i^{max} - X_i^{min})^2 \leq 2 \|X_i\|_2^2.$$

[0068] End proof.

[0069] It can also be shown that the above bound is tight:

[0070] Lemma 4 There exists a set of vectors X^n such that

$$\mathcal{E}(\pi_{sb}, X^n) \geq \frac{d-2}{2n} \cdot \frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2.$$

[0071] Proof. For every i , let X_i be defined as follows. $X_i(1) = 1/\sqrt{2}$, $X_i(2) = -1/\sqrt{2}$, and for all $j > 2$, $X_i(j) = 0$. For every i ,

$$X_i^{max} = \frac{1}{\sqrt{2}} \text{ and } X_i^{min} = -\frac{1}{\sqrt{2}}.$$

Substituting these bounds in the conclusion of Lemma 2 (which is an equality) yields the theorem.

[0072] End proof.

[0073] Therefore, the algorithm proposed in this subsection gives MSE $\Theta(d/n)$. Such an error is too large for real-world use. For example, in the application of neural networks, d can be on the order of millions, yet n can be much smaller than that. In such cases, the MSE is even larger than the norm of the vector.

2.2 Example Stochastic k-Level Quantization

[0074] A natural generalization of binary quantization is k-level quantization. Let k be a positive integer larger than 2. The present disclosure proposes a k-level stochastic quantization scheme π_{sk} to quantize each coordinate. Recall that for a vector X_i , $X_i^{max} = \max_{1 \leq j \leq d} X_i(j)$ and $X_i^{min} = \min_{1 \leq j \leq d} X_i(j)$. For every integer r in the range [0, k), let

$$B_i(r) \stackrel{\text{def}}{=} X_i^{min} + \frac{rs_i}{k-1},$$

where s_i satisfies $X_i^{min} + s_i \geq X_i^{max}$. A natural choice for s_i would be $X_i^{max} - X_i^{min}$. However, as will be demonstrated in Section 4, a higher value of s_i and variable length coding has better guarantees. The algorithm quantizes each coordinate into one of $B_i(r)$ s stochastically. In π_{sk} , for the i-th client and j-th coordinate, if $X_i(j) \in [B_i(r), B_i(r+1))$,

$$Y_i(j) = \begin{cases} B_i(r+1) & \text{with probability } \frac{X_i(j) - B_i(r)}{B_i(r+1) - B_i(r)} \\ B_i(r) & \text{otherwise} \end{cases}$$

[0075] The server estimates $\bar{X}$ by

$$\hat{X}_{\pi_{sk}} = \frac{1}{n} \sum_{i=1}^n Y_i.$$

[0076] As before, the communication complexity of this protocol is bounded. The proof is similar to that of Lemma 1 and hence omitted.

[0077] Lemma 5 There exists an implementation of stochastic k-level quantization that uses $d \lceil \log(k) \rceil + \tilde{O}(1)$ bits per client and hence $\mathcal{C}(\pi_{sk}, X^n) \leq n \cdot (d \lceil \log_2 k \rceil + \tilde{O}(1))$.

[0078] The mean squared loss can be bounded as follows.

[0079] Theorem 2 If $X_i^{max} - X_i^{min} \leq s_i \leq \sqrt{2} \|X_i\|_2 \forall i$, then for any X^n , the π_{sk} protocol satisfies,

$$\mathcal{E}(\pi_{sk}, X^n) \leq \frac{d}{2n(k-1)^2} \cdot \frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2. \quad (5)$$

Proof.

$$\begin{aligned} \mathcal{E}(\pi_{sk}, X^n) &= \mathbb{E} \|\hat{X} - \bar{X}\|_2^2 = \frac{1}{n^2} \mathbb{E} \left\| \sum_{i=1}^n (Y_i - X_i) \right\|_2^2 \\ &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \|Y_i - X_i\|_2^2 \leq \frac{1}{n^2} \sum_{i=1}^n d \frac{s_i^2}{4(k-1)^2}, \end{aligned}$$

where the last equality follows by observing $Y_i(j) - X_i(j)$ is an independent zero mean random variable with

$$\mathbb{E}(Y_i(j) - X_i(j))^2 \leq \frac{s_i^2}{4(k-1)^2} \cdot s_i \leq \sqrt{2} \|X_i\|_2$$

completes the proof.

[0080] End proof.

[0081] This section is concluded by noting that $s_i = X_i^{max} - X_i^{min}$ satisfies the conditions for the above theorem by Eq. Error! Reference source not found.

3. Example Stochastic Rotated Quantization

[0082] Next, it is shown that the algorithm of the previous section can be significantly improved by a new protocol. The motivation comes from the fact that the MSE of stochastic binary quantization and stochastic k-level quantization is

$$\mathcal{O}\left(\frac{d}{n} (X_i^{max} - X_i^{min})^2\right)$$

(the proof of Lemma 3 and Theorem 2 with $s_i = X_i^{max} - X_i^{min}$). Therefore the MSE is smaller when X_i^{max} and X_i^{min} are close. For example, when X_i is generated uniformly on the unit sphere, with high probability,

$$X_i^{max} - X_i^{min} \text{ is } \mathcal{O}\left(\sqrt{\frac{\log d}{d}}\right).$$

See, e.g., Dasgupta, Sanjoy and Gupta, Anupam. An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Structures & Algorithms*, 22(1):60-65, 2003. In such case, $\mathcal{E}(\pi_{sk}, X^n)$ is

$$\mathcal{O}\left(\frac{\log d}{n}\right)$$

instead of

$$\mathcal{O}\left(\frac{d}{n}\right).$$

[0083] In this section, it is shown that even without any assumptions on the distribution of the data, $X_i^{max} - X_i^{min}$ can be “reduced” with a structured random rotation, yielding an

$$\mathcal{O}\left(\frac{\log d}{n}\right)$$

error. This method can be denominated as stochastic rotated quantization and is denoted herein by π_{srk} .

[0084] In some implementations of the stochastic rotated quantization technique, using public randomness, all clients and the central server generate a random rotation matrix (random orthogonal matrix) $R \in \mathbb{R}^{d \times d}$ according to some

known distribution. Let $Z_i = RX_i$ and $Z = R\bar{X}$. In the stochastic rotated quantization protocol $\pi_{srk}(R)$, clients quantize the vectors Z_i instead of X_i and transmit them similar to π_{srk} . The server estimates $\bar{X}$ by

$$\hat{\bar{X}}_{\pi_{srk}} = R^{-1} \hat{Z}, \hat{Z} = \frac{1}{n} \sum_{i=1}^n Y_i.$$

[0085] The communication cost is same as π_{sk} and is given by Lemma 5. Next the MSE is bounded.

[0086] Lemma 6 For any X^n , $\epsilon(\pi_{srk}(R), X^n)$ is at most

$$\frac{d}{2n^2(k-1)^2} \sum_{i=1}^n \mathbb{E}_R[(Z_i^{max})^2 + (Z_i^{min})^2],$$

where $Z_i = RX_i$ and for every i , let $s_i = Z_i^{max} - Z_i^{min}$.

[0087] Proof.

$$\begin{aligned} \mathcal{E}(\pi_{srk}, X^n) &= \mathbb{E}_\pi \|\hat{\bar{X}} - \bar{X}\|^2 = \mathbb{E}_\pi \|R^{-1} \hat{Z} - R^{-1} \bar{Z}\|^2 \stackrel{(a)}{=} \mathbb{E}_\pi \|\hat{Z} - \bar{Z}\|^2 \stackrel{(b)}{=} \\ &\mathbb{E}_R \mathbb{E}_\pi [\|\hat{Z} - \bar{Z}\|^2 | Z_i^n] \leq \frac{d}{4n^2(k-1)^2} \sum_{i=1}^n \mathbb{E}_R [(Z_i^{max} - Z_i^{min})^2], \end{aligned}$$

where the last inequality follows Eq. Error! Reference source not found. and the value of s_i . (a) follows from the fact that rotation does not change the norm of the vector, and (b) follows from the tower law of expectation. The lemma follows from observing that

$$(Z_i^{max} - Z_i^{min})^2 \leq 2(Z_i^{max})^2 + 2(Z_i^{min})^2.$$

[0088] End proof.

[0089] To obtain strong bounds, an orthogonal matrix R that achieves low $(Z_i^{max})^2$ and $(Z_i^{min})^2$ is beneficial. In addition, due to the fact that d can be huge in practice, a type of orthogonal matrix that permits fast matrix-vector products is also beneficial. Naive orthogonal matrices that support fast multiplication such as block-diagonal matrices often result in high values of $(Z_i^{max})^2$ and $(Z_i^{min})^2$. As such, the present disclosure provides a special type of orthogonal matrix $R = HD$, where D is a random diagonal matrix with i.i.d. Rademacher entries (± 1 with probability 0.5). H is a Walsh-Hadamard matrix. See Horadam, Kathy J. *Hadamard matrices and their applications*. Princeton university press, 2012. The Walsh-Hadamard matrix of dimension 2^m for $m \in \mathcal{N}$ is given by the recursive formula,

$$H(2^1) = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, H(2^m) = \begin{bmatrix} H(2^{m-1}) & H(2^{m-1}) \\ H(2^{m-1}) & -H(2^{m-1}) \end{bmatrix}.$$

[0090] Both applying the rotation and inverse rotation take $\mathcal{O}(d \log d)$ time and $\mathcal{O}(1)$ additional space (with an in-place algorithm). The next lemma bounds $\mathbb{E}(Z_i^{max})^2$ and $\mathbb{E}(Z_i^{min})^2$ for this choice of R . The lemma is similar to that of Ailon, Nir and Chazelle, Bernard. Approximate nearest

neighbors and the fast Johnson-Lindenstrauss transform. *STOC*, 2006. The proof is provided in Section 8 for completeness.

[0091] Lemma 7 Let $R = HD$, where D is a diagonal matrix with independent Radamacher random variables. For every i and every sequence X^n ,

$$\mathbb{E}[(Z_i^{min})^2] = \mathbb{E}[(Z_i^{max})^2] \leq \frac{\|X_i\|_2^2 (2 \log d + 2)}{d}.$$

[0092] Combining the above two lemmas yields a significant result.

[0093] Theorem 3 For any X^n , $\pi_{srk}(HD)$ protocol satisfies,

$$\epsilon(\pi_{srk}(HD), X^n) \leq \frac{2 \log d + 2}{n(k-1)^2} \cdot \frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2.$$

4. Example Variable Length Coding

[0094] Instead of preprocessing the data via a rotation matrix as in π_{srk} , in this section the present disclosure provides use of a variable length coding strategy to minimize the number of bits.

[0095] Consider the stochastic k-level quantization technique. A natural way of transmitting Y_i is sending the bin number for each coordinate, thus the total number of bits the algorithm sends per transmitted coordinate would be $d \lceil \log_2 k \rceil$. This naive implementation is sub-optimal. Instead, The present disclosure proposes to, in some implementations, further encode the transmitted values using universal compression schemes. See, e.g., Krichevsky, R and Trofimov, V. The performance of universal encoding. *IEEE Transactions on Information Theory*, 27(2):199-207, 1981; and Falahatgar, Moein and Jafarpour, Ashkan and Orlitsky, Alon and Pichapati, Venkatadheeraj and Suresh, Ananda Theertha. Universal compression of power-law distributions. *ISIT*, 2015.

[0096] As one example technique that can be performed: encode h_r , the number of times each quantized value r has appeared, and then use a variable length coding scheme (e.g., arithmetic or Huffman coding) corresponding to the distribution

$$p_r = \frac{h_r}{d}.$$

This technique is denoted herein by π_{svk} . Since the vectors are quantized the same way in π_{sk} and π_{svk} , the MSE of π_{svk} is also given by Theorem 2. The communication cost is now bounded.

[0097] Theorem 4 Let $s_i = \sqrt{2} \|X_i\|$. There exists an implementation of π_{svk} such that $\mathcal{C}(\pi_{svk}, X^n)$ is at most

$$n \left(d \left(2 + \log_2 \left(\frac{(k-1)^2}{2d} + \frac{5}{4} \right) \right) + k \log_2 \frac{(d+k)e}{k} + \tilde{\mathcal{O}}(1) \right).$$

[0098] Proof. As in Lemma 1, $\bar{\theta}(1)$ bits are used to transmit the s_i 's and X_i^{min} . Recall that h_r is the number of coordinates that are quantized into bin r , and r takes k possible values. Furthermore, $\sum_r h_r = d$. Thus the number of bits necessary to represent the h_r 's is

$$\left\lceil \log_2 \binom{d+k-1}{k-1} \right\rceil \leq k \log_2 \frac{(d+k)e}{k}.$$

[0099] Once the h_r 's have been compressed, variable length coding, such as, for example, arithmetic coding, corresponding to the distribution $p_r = h_r/d$ can be used to compress and transmit bin values for each coordinate. The total number of bits arithmetic coding uses is

$$d \sum_{r=0}^{k-1} \frac{h_r}{d} \log_2 \frac{d}{h_r} + 2.$$

See, MacKay, David J C. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.

[0100] Let $p_r = h_r/d$, $a = (k-1)X_i^{min}$, $b = s_i$, and $\beta = \sum_{r=0}^{k-1} 1$ $((a+br)^2 + \delta)$. Note that

$$\begin{aligned} \sum_r p_r \log_2 \frac{1}{p_r} &= \\ \sum_r p_r \log_2 \frac{1/((a+br)^2 + \delta)\beta}{p_r} + \sum_r p_r \log_2 (((a+br)^2 + \delta)\beta) &\leq \\ \sum_r p_r \log_2 (((a+br)^2 + \delta)\beta) &\leq \log_2 \left(\sum_r p_r (a+br)^2 + \delta \right) + \log_2 \beta, \end{aligned}$$

where the first inequality follows from the positivity of KL-divergence. Choosing $\delta = s_i^2$, yields $\beta \leq 4/s_i^2$ and hence,

$$\sum_r p_r \log_2 \frac{1}{p_r} \leq \log_2 \left(\sum_r p_r (a+br)^2 + s_i^2 \right) + \log_2 (4/s_i^2).$$

[0101] Note that if $Y_i(j)$ belongs to bin r , $(a+br)^2 = (k-1)^2 Y_i^2(j)$. Recall that h_r is the number of coordinates quantized into bin r . Hence $\sum_r h_r (a+br)^2$ is the scaled norm-square of Y_i , i.e.,

$$\sum_r h_r (a+br)^2 = (k-1)^2 \sum_{j=1}^d Y_i^2(j) = \sum_{j=1}^d ((X_i(j) + \alpha(j))(k-1))^2,$$

where the $\alpha(j) = Y_i(j) - X_i(j)$. Taking expectations on both sides and using the fact that the $\alpha(j)$ are independent zero mean random variables over a range of $s_i/(k-1)$, provides

$$\mathbb{E} \sum_r h_r (a+br)^2 = \sum_{j=1}^d \mathbb{E} (X_i(j)^2 + \alpha(j)^2)(k-1)^2 \leq \|X_i\|_2^2 \left((k-1)^2 + \frac{d}{2} \right).$$

[0102] Using Jensen's inequality yields the result.

[0103] End proof.

[0104] Thus if $k = \sqrt{d} + 1$, the communication complexity is $\mathcal{O}(nd)$ and the MSE is $\mathcal{O}(1/n)$.

5. Example Communication MSE Trade-off

[0105] In the above protocols, all the clients transmit and hence the communication cost scales linearly with n . However, the present disclosure demonstrates that any of the above protocols can be combined by client sampling to obtain trade-offs between the MSE and the communication cost. Note that similar analysis also holds for sampling the coordinates.

[0106] Let π be a protocol where the mean estimate is of the form:

$$\hat{X} = R^{-1} \frac{1}{n} \sum_{i=1}^n Y_i. \quad (6)$$

[0107] All three protocols that have been discussed are of this form. Let π_p be the protocol where each client participates independently with probability p . The server estimates $\hat{X}$ by

$$\hat{X}_{\pi_p} = R^{-1} \cdot \frac{1}{np} \sum_{i \in S} Y_i,$$

where Y_i s are defined in the previous section and S is the set of clients that transmitted.

[0108] Lemma 8 For any set of vectors X^n and protocol π of the form Equation Error! Reference source not found. its sampled version π_p satisfies

$$\varepsilon(\pi_p, X^n) = \frac{1}{p} \cdot \varepsilon(\pi, X^n) + \frac{1-p}{np} \sum_{i=1}^n \|X_i\|_2^2.$$

$$\mathcal{C}(\pi_p, X^n) = p \cdot \mathcal{C}(\pi, X^n).$$

[0109] Proof. The proof of communication cost follows from Lemma 5 and the fact that in expectation, np clients transmit. The MSE is now bounded. Let S be the set of clients that transmit. The error $\varepsilon(\pi_p, X^n)$ is

$$\begin{aligned} \mathbb{E} \left[\left\| \hat{X} - \bar{X} \right\|_2^2 \right] &= \mathbb{E} \left[\left\| \frac{1}{np} \sum_{i \in S} R^{-1} Y_i - \bar{X} \right\|_2^2 \right] = \\ &= \mathbb{E} \left[\left\| \frac{1}{np} \sum_{i \in S} X_i - \bar{X} \right\|_2^2 + \frac{1}{n^2 p^2} \left\| \sum_{i \in S} (R^{-1} Y_i - X_i) \right\|_2^2 \right], \end{aligned}$$

where the last equality follows by observing that $R^{-1}Y_i - X_i$ are independent zero mean random variables and hence for any i , $\mathbb{E}[(R^{-1}Y_i - X_i)^T (\sum_{i \in S} X_i - \bar{X})] = 0$. The first term can be bounded as

$$\mathbb{E} \left\| \frac{1}{np} \sum_{i \in S} X_i - \bar{X} \right\|_2^2 = \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left\| \frac{1}{p} X_i 1_{i \in S} - X_i \right\|_2^2 =$$

$$\frac{1}{n^2} \sum_{i=1}^n \left(p \frac{(1-p)^2}{p^2} \|X_i\|_2^2 + (1-p) \|X_i\|_2^2 \right) = \frac{1-p}{np} \cdot \frac{1}{n} \sum_{i=1}^n \|X_i\|_2^2.$$

[0110] Furthermore, the second term can be bounded as

$$\mathbb{E} \left[\frac{1}{n^2 p^2} \left\| \sum_{i \in S} (R^{-1} Y_i - X_i) \right\|_2^2 \right] \stackrel{(a)}{=} \frac{1}{n^2 p^2} \sum_{i \in S} \mathbb{E} [\| (R^{-1} Y_i - X_i) \|_2^2] =$$

$$\frac{1}{n^2 p^2} \sum_{i=1}^n \mathbb{E} [\| (R^{-1} Y_i - X_i) \|_2^2 1_{i \in S}] = \frac{1}{n^2 p} \sum_{i=1}^n \mathbb{E} [\| (R^{-1} Y_i - X_i) \|_2^2] =$$

$$\frac{1}{n^2 p} \mathbb{E} \left[\left\| \sum_{i=1}^n (R^{-1} Y_i - X_i) \right\|_2^2 \right] = \frac{1}{p} \mathcal{E}(\pi, X^n)$$

where the last equality follows from the assumption that π 's mean estimate is of the form Error! Reference source not found. (a) follows from the fact that $R^{-1}Y_i - X_i$ are independent zero mean random variables.

[0111] End proof.

[0112] Combining the above lemma with Theorem 4, and choosing $k = \sqrt{d} + 1$ results in the following.

[0113] Corollary 1 For every $c \leq nd(2 + \log_2(7/4))$, there exists a protocol π such that $\mathcal{C}(\pi, S^d) \leq c$ and

$$\mathcal{E}(\pi, S^d) = \mathcal{O} \left(\min \left(1, \frac{d}{c} \right) \right).$$

6. Example Lower Bounds

[0114] The lower bound relies on the lower bounds on distributed statistical estimation due to Zhang, Yuchen, Duchi, John, Jordan, Michael I, and Wainwright, Martin J. Information-theoretic lower bounds for distributed statistical estimation with communication constraints. *NIPS*, 2013.

[0115] Lemma 9 (Zhang et al., 2013 Proposition 2) There exists a set of distributions $\mathcal{F}_d$ supported on

$$\left[-\frac{1}{\sqrt{d}}, \frac{1}{\sqrt{d}} \right]^d$$

such that if any centralized server wishes to estimate the mean of the underlying unknown distribution, then for any independent protocol π

$$\max_{p_d \in \mathcal{P}_d} \mathbb{E} [\| \theta(p_d) - \hat{\theta}_\pi \|_2^2] \geq t \min \left(1, \frac{d}{\mathcal{C}(\pi)} \right),$$

where $\mathcal{C}(\pi)$ is the communication cost of the protocol, $\theta(p_d)$ is the mean of p_d and t is a positive constant.

[0116] Theorem 5 Let t be the constant in Lemma 9. For every $c \leq ndt/4$ and $n \geq 4/t$,

$$\mathcal{E}(\Pi(c), S^d) \geq \frac{t}{4} \min \left(1, \frac{d}{c} \right).$$

[0117] Proof. Given n samples from the underlying distribution where each sample belongs to S^d , it is easy to see that

$$\mathbb{E} [\| \theta(p_d) - \hat{\theta}_\pi \|_2^2] \leq \frac{1}{n},$$

where $\hat{\theta}(p_d)$ is the empirical mean of the observed samples. Let $\mathcal{F}_d$ be the set of distributions in Lemma 9. Hence for any protocol π there exists a distribution p_d such that

$$\mathbb{E} [\| \hat{\theta}(p_d) - \hat{\theta}_\pi \|_2^2] \stackrel{(a)}{\geq} \frac{1}{2} \mathbb{E} [\| \theta(p_d) - \hat{\theta}_\pi \|_2^2 - \mathbb{E} [\| \theta(p_d) - \hat{\theta}_\pi \|_2^2]] \stackrel{(b)}{\geq}$$

$$\frac{t}{2} \min \left(1, \frac{d}{\mathcal{C}(\pi)} \right) - \frac{1}{n} \stackrel{(c)}{\geq} \frac{t}{4} \min \left(1, \frac{d}{\mathcal{C}(\pi)} \right),$$

[0118] Inequality (a) follows from the fact that $2(a-b)^2 + 2(b-c)^2 \geq (a-c)^2$. Inequality (b) follows from Lemma 9 and (c) follows from the fact that $\mathcal{C}(\pi, S^d) \leq ndt/4$ and $n \geq 4/t$.

[0119] End proof.

[0120] Corollary 1 and Theorem 5 yield Theorem 1. Note that the above lower bound holds only for communication cost $c < \mathcal{O}(nd)$.

7. Example Practical Considerations and Applications

[0121] Based on the theoretical analysis, the variable-length coding method provides the lowest quantization error asymptotically when using a constant number of bits. However in practice, stochastic rotated quantization may be preferred due to (hidden) constant factors and the fact that it uses a fixed amount of bits per dimension. For example, considering quantizing the vector $[-1, 1, 0, 0]$, stochastic rotated quantization can use 1 bit per dimension and gives zero error, whereas the other two protocols do not. To see this, observe that the naive quantization will quantize 0 to either 1 or -1 and variable length coding cannot achieve 0 error with 1 bit per dimension due to its constant factors.

[0122] In addition, note that the rotated quantization is preferred when applied on “unbalanced” data, due to the fact that the rotation can correct the unbalancedness. This is demonstrated by generating a dataset where the value of the last feature dimension entry is much larger than others. As an example of such dataset, 1000 datapoints were generated each with 256 dimensions. The first 255 dimensions are generated i.i.d. from $N(0, 1)$, and the last dimension is generated from $N(100, 1)$. As shown in FIG. 1, the rotated

stochastic quantization has the best performance on this example dataset. The improvement is especially significant for low bit rate cases. In particular, FIG. 1 illustrates distributed mean estimation on data generated from a Gaussian distribution.

[0123] Two example applications are demonstrated in the rest of this section. The experiments are performed on the MNIST (d=2304) and CIFAR (d=512) datasets.

[0124] Distributed Lloyd's Algorithm.

[0125] In the distributed Lloyd's (k-means) algorithm, each client has access to a subset of data points. In each iteration, the server broadcasts the cluster centers to all the clients. Each client updates the centers based on its local data, and sends the centers back to the server. The server then updates the centers by computing the weighted average of the centers sent from all clients. In the quantized setting proposed by the present disclosure, the client compresses the new centers before sending to the server. This saves the uplink communication cost, which is often the bottleneck of distributed learning. Furthermore, in this setting, the downlink is a broadcast, and therefore its cost can be reduced by a factor of $O(n/\log n)$ without quantization, where n is the number of clients.

[0126] FIGS. 2A-D show example results in which both the number of centers and the number of clients are set to 10. In particular, FIGS. 2A-D illustrate Lloyd's algorithm with different types of quantizations. Two settings are illustrated as tested: 16 quantization levels and 32 quantizations levels. The x-axis is the average number of bits sent for each data dimensions, and the y-axis is the global objective of Lloyd's algorithm.

[0127] Distributed Power Iteration.

[0128] Power iteration is a widely used method to compute the top eigenvector of a matrix. In the distributed setting, each client has access to a subset of data. In each iteration, the server broadcasts the current estimate of the eigenvector to all clients. Each client then updates the eigenvector based on one power iteration on its local data, and sends the updated eigenvector back to the server. The server updates the eigenvector by computing the weighted average of the eigenvectors sent by all clients.

[0129] As one example of power iteration, given a set of vectors $A_1, A_2, \dots, A_m$, at each round t , compute:

$$v_{t+1} = \frac{1}{m} \sum_{i=1}^m A_i A_i^T v_t$$

[0130] Data points $A_1, A_2, \dots, A_m$ are distributed across n clients. At each round of power iteration t , a client sends:

$$\sum_{j \in \text{Client}_i} A_j A_j^T v_t$$

[0131] The server can average the received information to compute v_{t+1} .

[0132] Similar to the above distributed Lloyd's algorithm, in the quantized setting proposed by the present disclosure, the client compresses the estimated eigenvector before sending to the server. FIGS. 3A-D show the results. The dataset is distributed over 100 clients.

[0133] In particular, FIGS. 3A-D illustrate power iteration with different types of quantizations. Two settings were tested: 16 quantization levels and 32 quantization levels. The x-axis is the average number of bits sent for each data dimension and they y-axis is the l_2 distance between the computed eigenvector and the ground-truth vector.

[0134] For both distributed Lloyd's and distributed power iteration applications, variable-length coding achieves the lowest quantization error in most of the settings. Furthermore, for low-bit rate, stochastic rotated quantization is competitive with variable-length coding.

8. Proof of Lemma 7

[0135] The equality follows from the symmetry in HD. To prove the upper bound, observe that:

$$\mathbb{E}[(Z_i^{\max})^2] = \text{Var}(Z_i^{\max}) + (\mathbb{E}[Z_i^{\max}])^2.$$

[0136] Let $D(j)$ be the j^{th} diagonal entry of D . To bound the first term observe that $Z_i^{\max}$ is a function of d independent random variables $D(1), D(2), \dots, D(d)$. Changing $D(j)$ changes the $Z_i^{\max}$ by at most

$$\frac{2X_i(j)}{\sqrt{d}}.$$

Hence, applying Efron-Stein variance bound yields

$$-0.1 \text{ cmVar}(Z_i^{\max}) \leq \sum_{j=1}^d \frac{4X_i^2(j)}{2d} = \frac{2\|X_i\|_2^2}{d}.$$

[0137] To bound the second term, observe that for every $\beta > 0$.

$$-0.1 \text{ cm}\beta Z_i^{\max} = \log \exp(\beta Z_i^{\max}) \leq \log \left(\sum_{j=1}^d e^{\beta Z_i(j)} \right).$$

[0138] Note that

$$Z_i(k) = \frac{1}{\sqrt{d}} \sum_{j=1}^d D(j) H(k, j) X_i(j).$$

Since the $D(j)$'s are Radamacher random variables and $|H(k,j)|=1$ for all k, j , the distributions of $Z_i(k)$ is same for all k . Hence by Jensen's inequality,

$$\mathbb{E}[Z_i^{\max}] \leq \frac{1}{\beta} \mathbb{E} \left[\log \left(\sum_{j=1}^d e^{\beta Z_i(j)} \right) \right] \leq \frac{1}{\beta} \log \left(\sum_{j=1}^d \mathbb{E}[e^{\beta Z_i(j)}] \right) = \frac{1}{\beta} \log(d \mathbb{E}[e^{\beta Z_i(1)}]).$$

$$\text{Since } Z_i(1) = \frac{1}{\sqrt{d}} \sum_{j=1}^d D(j) X_i(j),$$

-continued

$$\mathbb{E}[e^{\beta Z_i(1)}] = \mathbb{E}\left[e^{\frac{\beta \sum_{j=1}^d D(j) X_i(j)}{\sqrt{d}}}\right] \stackrel{(a)}{=} \prod_{j=1}^d \mathbb{E}\left[e^{\frac{\beta D(j) X_i(j)}{\sqrt{d}}}\right] =$$

$$\prod_{j=1}^d \frac{e^{-\beta X_i(j)/\sqrt{d}} + e^{\beta X_i(j)/\sqrt{d}}}{2} \stackrel{(b)}{\leq} \prod_{j=1}^d e^{\beta^2 X_i^2(j)/2d} = e^{\beta^2 \|X_i\|_2^2/2d},$$

where (a) follows from the fact that the $D(i)$'s are independent and (b) follows from the fact that $e^a + e^{-a} \leq 2e^{a^2/2}$ for any a . Hence,

$$\mathbb{E}[Z_i^{\max}] \leq \min_{\beta \geq 0} \frac{\log d}{\beta} + \frac{\beta \|X_i\|_2^2}{2d} \leq \frac{2\|X_i\|_2 \sqrt{\log d}}{\sqrt{2d}}.$$

9. Example Computing Systems

[0139] FIG. 4 depicts an example system 200 for distributed computing. System 200 can include a server computing device 210 and a plurality of client computing devices 230.

[0140] In some implementations, and as one example application, the server computing device 210 can be configured to access a global machine-learned model and to provide the global model to the plurality of client devices 230. The model can be, for instance, a linear regression model, logistic regression model, a support vector machine model, a neural network (e.g. convolutional neural network, recurrent neural network, etc.), or other machine-learned model. In some implementations, the sever 210 can be configured to communicate with client devices 230 over one or more networks 280.

[0141] Client devices 230 can each be configured to determine one or more local updates associated with the model based at least in part on locally stored data 236. The locally stored data 236 can include audio files, image files, video files, log entries, and/or various other suitable data. In some implementations, the data 236 can be any data derived through a user interaction with a client device 230. Thus, the data 236 across the plurality of devices 230 includes data that is respectively stored at each device 230. Thus, in some implementations, the collective data 236 across all devices 230 is highly unbalanced and not independent and identically distributed.

[0142] Client devices 230 can be configured to provide the local updates to the server 210. The data 236 may be privacy sensitive. In this manner, the local updates can be performed and provided to server 210 without compromising the privacy of data 236. For instance, in such implementations, data 236 is not itself provided to server 210 since the local update does not include the actual data 236. In some implementations, one or more of encryption techniques, random noise techniques, and/or other security techniques can be added to the training process to obscure any inferable information from the local updates.

[0143] As indicated above, server 210 can receive each local update from client device 230, and can aggregate the local updates to determine a global update to the machine-learned model. In some implementations, server 210 can

determine a weighted average or other mean of the local updates and determine the global update based at least in part on the average.

[0144] In some implementations, scaling or other techniques can be applied to the local updates to determine the global update. For instance, a local step size can be applied for each client device 230, the aggregation can be performed proportionally to various data partition sizes of client devices 230, and/or one or more scaling factors can be applied to the local and/or aggregated updates. It will be appreciated that various other techniques can be applied without deviating from the scope of the present disclosure.

[0145] More generally, FIG. 4 depicts an example computing system 200 that can be used to implement the methods and systems of the present disclosure. The system 200 can be implemented using a client-server architecture that includes a server 210 that communicates with one or more client devices 230 over a network 280.

[0146] The system 200 includes a server 210, such as a web server. The server 210 can be implemented using any suitable computing device(s). The server 210 can have one or more processors 212 and one or more memory devices 214. The server 210 can be implemented using one server device or a plurality of server devices. In implementations in which a plurality of devices are used, such plurality of devices can operate according to a parallel computing architecture, a sequential computing architecture, or a combination thereof.

[0147] The server 210 can also include a network interface 236 used to communicate with the client devices 230 over the network 280. The network interface can include any suitable components for interfacing with one more networks, including for example, transmitters, receivers, ports, controllers, antennas, or other suitable components.

[0148] The one or more processors 212 can include any suitable processing device, such as a microprocessor, micro-controller, integrated circuit, logic device, or other suitable processing device. The one or more memory devices 214 can include one or more computer-readable media, including, but not limited to, non-transitory computer-readable media, RAM, ROM, hard drives, flash drives, or other memory devices. The one or more memory devices 214 can store information accessible by the one or more processors 212, including computer-readable instructions 218 that can be executed by the one or more processors 212.

[0149] The instructions 218 can be any set of instructions that when executed by the one or more processors 212, cause the one or more processors 212 to perform operations.

[0150] The server 210 can further include a de-rotator 220. The de-rotator 220 can de-rotate a vector that has been rotated by a client device 230. For example, the de-rotator can use an inverse random rotation matrix that is an inverse of a random rotation matrix used by the client device 230 to rotate the vector. Alternatively, a mean of several rotated vectors can be determined prior to de-rotation and the de-rotator can de-rotate the determined mean. In some implementations, the de-rotator 220 can determine one or more inverse random rotation matrices based at least in part on one or more seeds that are shared (e.g., respectively shared or universally shared) with the client devices 230.

[0151] The server 210 can further include a decoder 222. The decoder 222 can decode a vector that has been encoded by a client device 230 (e.g., according to one of the encoding techniques discussed above). For example, the decoder can

decode a vector that has been encoded according to variable length coding techniques such as, for example, Huffman coding or arithmetic coding.

[0152] The server 210 can further include a mean calculator 224. The mean calculator 224 can be configured to receive a plurality of vectors (e.g., decoded vectors, rotated vectors, and/or de-rotated vectors) and to determine a mean of the vectors (e.g., a mean vector). In some implementations in which the mean is determined prior to de-rotation, the mean vector can subsequently be de-rotated by the de-rotator 220.

[0153] As shown in FIG. 4, the one or more memory devices 214 can also store data 216 that can be retrieved, manipulated, created, or stored by the one or more processors 212. The data 216 can include, for instance, local updates, global parameters, and other data. The data 216 can be stored in one or more databases. The one or more databases can be connected to the server 210 by a high bandwidth LAN or WAN, or can also be connected to server 210 through network 280. The one or more databases can be split up so that they are located in multiple locales.

[0154] The server 210 can exchange data with client devices 230 over the network 280. Any number of client devices 230 can be connected to the server 210 over the network 280. Each of the client devices 230 can be any suitable type of computing device, such as a general purpose computer, special purpose computer, laptop, desktop, mobile device, navigation system, intelligent speaker or home assistant, home appliance, smartphone, tablet, computing device that is able to be worn, gaming console, worker device in a datacenter, a display with one or more processors, or other suitable computing device.

[0155] Similar to the server 210, a client device 230 can include one or more processor(s) 232 and a memory 234. The one or more processor(s) 232 can include, for example, one or more central processing units (CPUs), graphics processing units (GPUs) dedicated to efficiently rendering images or performing other specialized calculations, Tensor processing units (TPUs), and/or other processing devices. The memory 234 can include one or more computer-readable media and can store information accessible by the one or more processors 232, including instructions 238 that can be executed by the one or more processors 232 and data 236.

[0156] The client computing device 230 can include a vector calculator 240 that is implementable to determine one or more vectors (e.g., local updates) according to example aspects of the present disclosure. For example, in some example applications, the vector calculator 240 can perform one or more training techniques such as, for example, backwards propagation of errors to re-train or otherwise update a machine-learned model based on the locally stored data 236, thereby generating an update vector (e.g., a gradient). The vector calculator 240 can be included in an application or can be included in the operating system of the device 230. In other implementations, the vector calculator 240 can be any component or system that determines a vector to be transmitted to the server computing device 210.

[0157] The client computing device 230 can further include a rotator 242. The rotator 242 can rotate a vector by a random rotation matrix (e.g., by multiplying the vector by the matrix). In some implementations, the rotator 242 can determine the random rotation matrix based on a seed.

[0158] The client computing device 230 can further include a quantizer 246. The quantizer 246 can quantize a

vector. For example, the quantizer 246 can perform stochastic binary quantization or stochastic k-level quantization as described above to quantize the vector.

[0159] The client computing device 230 can further include an encoder 248. For example, the encoder 248 can perform one or more of the encoding techniques described above (e.g., variable length coding such as, for example, Huffman coding or arithmetic coding).

[0160] The data 236 can include data examples to be used in solving one or more optimization problems. In some applications, the data examples of each client device 230 can be distributed unevenly among the client devices, such that no client device 230 includes a representative sample of the overall distribution of the training data examples. The data 236 can further include a vector to be communicated to the server 210.

[0161] The client device 230 of FIG. 4 can include various input/output devices for providing and receiving information from a user, such as a touch screen, touch pad, data entry keys, speakers, and/or a microphone suitable for voice recognition.

[0162] The client device 230 can also include a network interface 250 used to communicate with one or more remote computing devices (e.g. server 210) over the network 280. The network interface 250 can include any suitable components for interfacing with one more networks, including for example, transmitters, receivers, ports, controllers, antennas, or other suitable components.

[0163] The network 280 can be any type of communications network, such as a local area network (e.g. intranet), wide area network (e.g. Internet), cellular network, or some combination thereof. The network 280 can also include a direct connection between a client device 230 and the server 210. In general, communication between the server 210 and a client device 230 can be carried via network interface using any type of wired and/or wireless connection, using a variety of communication protocols (e.g. TCP/IP, HTTP, SMTP, FTP), encodings or formats (e.g. HTML, XML), and/or protection schemes (e.g. VPN, secure HTTP, SSL).

[0164] The vector calculator 240, the rotator 242, the quantizer 246, the encoder 248, the de-rotator 220, the decoder 222, and the mean calculator 224 can include computer logic utilized to provide desired functionality. Thus, each of vector calculator 240, the rotator 242, the quantizer 246, the encoder 248, the de-rotator 220, the decoder 222, and the mean calculator 224 can be implemented in hardware, firmware and/or software controlling a general purpose processor. In some implementations, each of vector calculator 240, the rotator 242, the quantizer 246, the encoder 248, the de-rotator 220, the decoder 222, and the mean calculator 224 includes program code files stored on the storage device, loaded into memory and executed by a processor or can be provided from computer program products, for example, computer executable instructions that are stored in a tangible computer-readable storage medium such as, for example, a RAM disk or card or other computer-readable optical or magnetic media.

[0165] Furthermore, while the techniques of the present disclosure are discussed primarily with reference to vectors to be transmitted or uploaded, the techniques described herein can also be applied to other data structures as well. As an example, for any other type of data to be transmitted by the client computing device 230 (e.g., a matrix (2D), a tensor (3D and above), or other data types or structures), the client

computing device **230** (e.g., the vector calculator **240**) can first flatten the data structure to form a vector. The techniques described herein (e.g., random rotation, probabilistic quantization, and/or variable length coding) can then be applied to the vector. After the entire process (e.g., after de-rotation, decoding, and/or mean estimation), the server computing device **210** can re-shape the resulting vector (e.g., a mean vector) back to the original dimension(s).

[0166] Thus, the present disclosure can be generalized to other data structures through conversion (e.g., flattening) of the other data structure to vector format prior to rotation, quantization, encoding, etc. After de-rotation, decoding, etc., the de-rotated or decoded vector can be reshaped back into the original dimension(s).

10. Example Methods

[0167] FIG. 5 depicts a swim lane flow diagram of an example method **500** to perform a stochastic rotated quantization technique according to example embodiments of the present disclosure. In particular, the left-hand side of FIG. 5 illustrates operations performed by each of one or more client computing devices while the right-hand side of FIG. 5 illustrates operations performed by a server computing device.

[0168] At **502**, a client computing device obtains a vector. The vector can be any vector that is to be transmitted to the server computing device. Example vectors include a machine-learned model update vector that describes one or more parameters of a machine-learned model or one or more updates to the one or more parameters of the machine-learned model; a cluster vector that describes a plurality of cluster centers or a plurality of updates to the plurality of cluster centers; and a power iteration vector that describes an eigenvector. In some implementations, obtaining the vector can include computing the vector based on a local dataset that is stored locally at the client computing device.

[0169] At **504**, the client computing device rotates the vector by a random rotation matrix to obtain a rotated vector. As an example, in some implementations, the random rotation matrix can be a product of a Walsh-Hadamard matrix with a diagonal matrix. In some implementations, the diagonal matrix can include independent and identically distributed Rademacher entries.

[0170] In some implementations, the method **500** can further include obtaining a seed that is shared with the server computing device. The method **500** can include generating the random rotation matrix based at least in part on the seed.

[0171] At **506**, the client computing device performs probabilistic quantization of the rotated vector to obtain a quantized rotated vector. For example, in some implementations, performing probabilistic quantization at **506** can include performing stochastic binary quantization or performing stochastic k-level quantization.

[0172] In some implementations, performing probabilistic quantization at **506** can include determining a value for each of a number of quantization levels based at least in part on a magnitude of the rotated vector and a minimum coordinate value included in the rotated vector; and quantizing each coordinate of the rotated vector into one of the number of quantization levels.

[0173] At **508**, the client computing device transmits the quantized rotated vector to the server computing device. In some implementations, various additional encodings, cryp-

tography, and/or privacy preserving manipulations can be performed prior to transmission.

[0174] At **510**, the server computing device receives the quantized rotated vectors from the client computing devices. In some implementations, at **510**, the server computing device can decode or unquantize each quantized rotated vector to the extent possible. For example, this may include transforming references or codings into particular data entries to which the references refer.

[0175] At **512**, the server computing device determines a mean of all the quantized rotated vectors received from all of the client computing devices to obtain a mean rotated vector.

[0176] At **514**, the server computing device de-rotates the mean rotated vector by an inverse random rotation matrix to obtain a mean de-rotated vector. For example, the inverse random rotation matrix can be the inverse of the random rotation matrix used at **504**.

[0177] In some implementations, method **500** can include obtaining, by the server computing device, a seed, where the seed is shared with at least one of the client computing devices from which a quantized rotated vector is received. For example, each client computing device can have a different seed or seeds can be used by multiple client computing devices. The method **500** can include generating, by the server computing device, the inverse random rotation matrix based at least in part on the seed.

[0178] In some implementations, the operations illustrated at **512** and **514** can be performed in reverse order, such that the vectors are individually de-rotated prior to taking the mean.

[0179] In some implementations, the method **500** can further include performing, by the server computing device, a global update based on the mean de-rotated vector.

[0180] FIG. 6 depicts a swim lane flow diagram of an example method **600** to perform a variable length coding technique according to example embodiments of the present disclosure. In particular, the left-hand side of FIG. 6 illustrates operations performed by each of one or more client computing devices while the right-hand side of FIG. 6 illustrates operations performed by a server computing device.

[0181] At **602**, a client computing device obtains a vector. The vector can be any vector that is to be transmitted to the server computing device. Example vectors include a machine-learned model update vector that describes one or more parameters of a machine-learned model or one or more updates to the one or more parameters of the machine-learned model; a cluster vector that describes a plurality of cluster centers or a plurality of updates to the plurality of cluster centers; and a power iteration vector that describes an eigenvector. In some implementations, obtaining the vector can include computing the vector based on a local dataset that is stored locally at the client computing device.

[0182] At **604**, the client computing device performs probabilistic quantization of the vector to obtain a quantized vector. For example, in some implementations, performing probabilistic quantization at **604** can include performing stochastic binary quantization or performing stochastic k-level quantization.

[0183] In some implementations, performing probabilistic quantization at **604** can include determining a value for each of a number of quantization levels based at least in part on a magnitude of the vector to be quantized and a minimum

coordinate value included in the vector; and quantizing each coordinate of the vector into one of the number of quantization levels.

[0184] At 606, the client computing device encodes the quantized vector according to a variable length coding scheme to obtain an encoded quantized vector. For example, the variable length coding scheme can include Huffman coding or arithmetic coding.

[0185] At 608, the client computing device transmits the encoded quantized vector to the server computing device.

[0186] At 610, the server computing device receives the encoded quantized vectors from the client computing devices.

[0187] At 612, the server computing device decodes each encoded quantized vector according to the variable length coding scheme to obtain the quantized vector.

[0188] In addition, in some implementations, at 612 and after decoding with the variable length coding scheme, the server computing device can decode or unquantize each quantized rotated vector to the extent possible. For example, this may include transforming references or codings into particular data entries to which the references refer.

[0189] At 614, the server computing device determines a mean of all the quantized vectors received from all of the client computing devices to obtain a mean vector.

[0190] In some implementations, the method 600 can further include performing, by the server computing device, a global update based on the mean de-rotated vector.

11. Additional Disclosure

[0191] The technology discussed herein makes reference to servers, databases, software applications, and other computer-based systems, as well as actions taken and information sent to and from such systems. The inherent flexibility of computer-based systems allows for a great variety of possible configurations, combinations, and divisions of tasks and functionality between and among components. For instance, processes discussed herein can be implemented using a single device or component or multiple devices or components working in combination. Databases and applications can be implemented on a single system or distributed across multiple systems. Distributed components can operate sequentially or in parallel.

[0192] While the present subject matter has been described in detail with respect to various specific example embodiments thereof, each example is provided by way of explanation, not limitation of the disclosure. Those skilled in the art, upon attaining an understanding of the foregoing, can readily produce alterations to, variations of, and equivalents to such embodiments. Accordingly, the subject disclosure does not preclude inclusion of such modifications, variations and/or additions to the present subject matter as would be readily apparent to one of ordinary skill in the art. For instance, features illustrated or described as part of one embodiment can be used with another embodiment to yield a still further embodiment. Thus, it is intended that the present disclosure cover such alterations, variations, and equivalents.

What is claimed is:

1. A computing system to facilitate transmission of machine-learned model updates from client devices to a centralized server computing device, the computing system comprising:

one or more client computing devices, wherein each client computing device comprises one or more processors and one or more non-transitory computer-readable media that store instructions that, when executed by the

one or more processors cause the client computing device to perform operations, the operations comprising:

determining an update to a machine-learned model based at least in part on a local dataset stored at the client computing device;

rotating the update by a random rotation matrix to obtain a rotated update;

performing probabilistic quantization of the rotated update to obtain a quantized rotated update; and

transmitting the quantized rotated update to the centralized server computing device.

2. The computing system of claim 1, further comprising: the centralized server computing device, the centralized server computing device comprising one or more processors and one or more non-transitory computer-readable media that store instructions that, when executed by the one or more processors cause the centralized server computing device to perform operations, the operations comprising:

receiving the one or more quantized rotated updates respectively transmitted by the one or more client computing devices; and

determining a mean rotated update of all quantized rotated updates obtained for all of the one or more client computing devices;

de-rotating the mean rotated update by an inverse random rotation matrix to obtain a mean de-rotated update;

updating a global machine-learned model based at least in part on the mean de-rotated update.

3. The computing system of claim 1, wherein the update to the machine-learned model comprises a gradient associated with training of the machine-learned model.

4. The computing system of claim 1, wherein the random rotation matrix comprises a product of a Walsh-Hadamard matrix with a diagonal matrix.

5. The computing system of claim 4, wherein the diagonal matrix comprises independent and identically distributed Rademacher entries.

6. A computing system, comprising:

one or more client computing devices, wherein each client computing device comprises one or more processors and one or more non-transitory computer-readable media that store instructions that, when executed by the one or more processors cause the client computing device to perform operations, the operations comprising:

obtaining a vector;

rotating the vector by a random rotation matrix to obtain a rotated vector;

performing probabilistic quantization of the rotated vector to obtain a quantized rotated vector; and

transmitting the quantized rotated vector.

7. The computing system of claim 6, further comprising: a server computing device that comprises one or more processors and one or more non-transitory computer-readable media that store instructions that, when executed by the one or more processors cause the server computing device to perform operations, the operations comprising:

receiving the quantized rotated vectors transmitted by the client computing devices; and

- determining a mean of all quantized rotated vectors obtained for all of the one or more client computing devices; and
de-rotating the mean by an inverse random rotation matrix.
8. The computing system of claim 6, wherein the random rotation matrix comprises a product of a Walsh-Hadamard matrix with a diagonal matrix.
9. The computing system of claim 8, wherein the diagonal matrix comprises independent and identically distributed Rademacher entries.
10. The computing system of claim 6, wherein the operations further comprise:
obtaining a seed, wherein the seed is shared with a server computing device to which the quantized rotated vector is transmitted; and
generating the random rotation matrix based at least in part on the seed.
11. The computing system of claim 6, wherein performing probabilistic quantization of the rotated vector to obtain a quantized rotated vector comprises:
performing stochastic binary quantization of the rotated vector to obtain the quantized rotated vector; or
performing stochastic k-level quantization of the rotated vector to obtain the quantized rotated vector.
12. The computing system of claim 6, wherein performing probabilistic quantization of the rotated vector comprises:
determining a value for each of a number of quantization levels based at least in part on a magnitude of the rotated vector and a minimum coordinate value included in the rotated vector; and
quantizing each coordinate of the rotated vector into one of the number of quantization levels.
13. The computing system of claim 6, wherein the vector comprises one or more of:
a machine-learned model update vector that describes one or more parameters of a machine-learned model or one or more updates to the one or more parameters of the machine-learned model;
a cluster vector that describes a plurality of cluster centers or a plurality of updates to the plurality of cluster centers; and
a power iteration vector that describes an eigenvector.
14. The computing system of claim 6, wherein obtaining the vector comprises computing the vector based on a local dataset that is stored locally at the client computing device.
15. The computing system of claim 1, wherein each of the client computing devices comprises:
a mobile computing device;
a worker device in a datacenter;
an embedded computing device; or
a connected smart device.

16. A computing system, comprising:
one or more client computing devices, wherein each client computing device comprises one or more processors and one or more non-transitory computer-readable media that store instructions that, when executed by the one or more processors cause the client computing device to perform operations, the operations comprising:
obtaining a vector;
performing probabilistic quantization of the vector to obtain a quantized vector, wherein performing probabilistic quantization of the vector comprises:
determining a value for each of a number of quantization levels based at least in part on a magnitude of the vector and a minimum coordinate value included in the vector; and
quantizing each coordinate of the vector into one of the number of quantization levels.
17. The computing system of claim 16, wherein the operations further comprise:
encoding the quantized vector according to a variable length coding scheme; and
transmitting the encoded quantized vector.
18. The computing system of claim 17, further comprising:
a server computing device that comprises one or more processors and one or more non-transitory computer-readable media that store instructions that, when executed by the one or more processors cause the server computing device computing to perform operations, the operations comprising:
for each of the one or more client computing devices:
receiving the encoded quantized vector transmitted by the client computing device; and
decoding the encoded quantized vector according to the variable length coding scheme to obtain the quantized vector for such client computing device; and
determining a mean of all quantized vectors obtained for all of the one or more client computing devices.
19. The computing system of claim 17, wherein the variable length coding scheme comprises Huffman coding or arithmetic coding.
20. The computing system of claim 17, wherein the vector comprises one or more of:
a machine-learned model update vector that describes one or more parameters of a machine-learned model or one or more updates to the one or more parameters of the machine-learned model;
a cluster vector that describes a plurality of cluster centers or a plurality of updates to the plurality of cluster centers; and
a power iteration vector that describes an eigenvector.

* * * * *