CLAIMS

1. A method of privacy-preserving logistic regression training based on homomorphically encrypted ciphertexts, the method comprising:

obtaining a first packed ciphertext comprising at least a portion of a first training data sample packed into a first vector of slots thereof for training a privacy-preserving logistic regression model;

obtaining a second packed ciphertext comprising a plurality of weights of the privacy-preserving logistic regression model packed into a first vector of slots thereof;

determining at least a first output probability of the privacy-preserving logistic regression model based on the first packed ciphertext and the second packed ciphertext; and

updating the plurality of weights based on the first output probability, wherein

said determining at least the first output probability comprises performing a dot product between the first packed ciphertext and the second packed ciphertext, and

said performing the dot product comprises:

multiplying the first vector of slots of the first packed ciphertext and the first vector of slots of the second packed ciphertext to obtain a plurality of first multiplication results in a first vector of slots of a first resultant ciphertext; and

performing a first inter-slot summation of the first vector of slots of the first resultant ciphertext to obtain a first dot product result of the dot product between the first packed ciphertext and the second packed ciphertext, wherein

said performing the first inter-slot summation comprises:

obtaining a first binary representation corresponding to the number of slots in the first vector of slots of the first resultant ciphertext;

for each digit in the first binary representation having a first binary value from the digit being the second most significant digit to the digit being the least significant digit amongst all digits in the first binary representation having the first binary value,

generating a first new ciphertext based on rotating an immediately preceding ciphertext thereof by a number of slots based on a digit position of the digit in the first binary representation, wherein for a first rotation step, the immediately preceding ciphertext comprises the first resultant ciphertext, and for each subsequent rotation step, the immediately preceding ciphertext comprises the ciphertext generated in the immediately preceding rotation step, and

adding the first new ciphertext with a ciphertext having an index corresponding to said digit position to produce a second new ciphertext, wherein each of the ciphertext having an index that corresponds to a digit position were generated by rotations and additions of the first resultant

ciphertext based on increasing power-of-two steps from a 1^{st} step to a 2^{k-1} step where k is defined as $\log_2(n)$ and n is defined as a number of slots in the first resultant ciphertext; and

determining a summation of the first vector of slots of the first resultant ciphertext based on the second new ciphertext produced for said digit in the first binary representation having the first binary value and being the least significant digit, to produce the first dot product result.

2. The method according to claim 1, wherein the first binary value is binary value of 1.

3. The method according to claim 1, wherein the first new ciphertext is generated based on rotating the immediately preceding ciphertext thereof by the number of slots determined based on a power of two with an exponent thereof having a value corresponding to the digit position of the digit in the first binary representation.

4. The method according to claim 3, wherein said performing the first inter-slot summation further comprises: for each digit in the first binary representation from the digit being the least significant digit to the digit being the second most significant digit, generating a third new ciphertext based on rotating an immediately preceding ciphertext thereof by a number of slots based on a digit position of the digit in the first binary representation.

5. The method according to claim 4, wherein the third new ciphertext is generated based on rotating the immediately preceding ciphertext thereof by the number of slots determined based on a power of two with an exponent thereof having a value corresponding to the digit position of the digit in the first binary representation.

6. The method according to claim 5, wherein a plurality of new ciphertexts are generated, comprising the third new ciphertext for said each digit in the first binary representation from the digit being the least significant digit to the digit being the second most significant digit and the first and second new ciphertexts for said each digit in the first binary representation from the digit being the second most significant digit to the digital being the least significant digit, each of the plurality of new ciphertexts having an index corresponding to an order of the new ciphertext in the plurality of new ciphertexts.

7. The method according to claim 6, wherein

said generating the first new ciphertext for the digit in the first binary representation being the second most significant digit is based on rotating the third new ciphertext generated for the digit in the first binary representation being the second most significant digit as the immediately preceding ciphertext thereof.

8. The method according to claim 1, wherein

said determining the summation of the first vector of slots of the first resultant ciphertext comprises obtaining a summation result of the first vector of slots of the first resultant ciphertext from a predetermined slot of the second new ciphertext produced for said digit in the first binary representation having the first binary value and being the least significant digit.

9. The method according to claim 1, wherein

the first packed ciphertext comprises a plurality of training data samples packed into a plurality of vectors of slots, respectively, thereof, for training the privacy-preserving logistic regression model, the plurality of training data samples comprising the first training data sample and one or more second training data samples, and the plurality of vectors of slots of the first packed ciphertext comprising the first vector of slots and one or more second vectors of slots of the first packed ciphertext,

the second packed ciphertext comprises a plurality of vectors of slots, each vector of slots having packed therein the plurality of weights, the plurality of vectors of slots of the second packed ciphertext comprising the first vector of slots and one or more second vectors of slots of the second packed ciphertext,

said determining at least the first output probability comprises determining a plurality of output probabilities of the privacy-preserving logistic regression model based on the first packed ciphertext and the second packed ciphertext, the plurality of output probabilities comprising the first output probability and one or more second output probabilities of the privacy-preserving logistic regression model,

said updating the plurality of weights comprises updating the plurality of weights based on the plurality of output probabilities,

said performing the dot product further comprises, for each second vector of slots of said one or more second vectors of slots of the first packed ciphertext:

multiplying the second vector of slots of the first packed ciphertext and the corresponding second vector of slots of the second packed ciphertext to obtain a plurality of second multiplication results in a second vector of slots of the first resultant ciphertext; and

performing a second inter-slot summation of the second vector of slots of the first resultant ciphertext to obtain a second dot product result of the dot product between the first packed ciphertext and the second packed ciphertext, wherein

said performing the second inter-slot summation comprises:

obtaining a second binary representation corresponding to the number of slots in the second vector of slots of the first resultant ciphertext;

for each digit in the second binary representation having the first binary value from the digit being the second most significant digit to the digital being the least significant digit amongst all digits in the second binary representation having the first binary value, generating a fourth new ciphertext based on rotating an immediately preceding ciphertext thereof by a number of slots based on a digit position of the digit in the second binary representation, and adding the fourth new ciphertext with a ciphertext having an index corresponding to said digit position to produce a fifth new ciphertext; and

determining a summation of the second vector of slots of the first resultant ciphertext based on the fifth new ciphertext produced for said digit in the second binary representation having the first binary value and being the least significant digit, to produce the second dot product result.

10. The method according to claim 9, wherein said performing the first inter-slot summation of the first resultant ciphertext and said performing the second inter-slot summation of the first resultant ciphertext are performed in parallel.

11. The method according to claim 1, wherein

said determining at least the first output probability further comprises applying an activation function on the first dot product result to produce the first output probability,

the method further comprises:

determining a first loss gradient relating to the first training data sample based on the first output probability; and

determining, for each of the plurality of weights, a second loss gradient relating to the weight based on the first loss gradient and a slot of the first vector of slots of the first packed ciphertext corresponding to the weight to obtain a plurality of second loss gradients relating to the plurality of weights, and

said updating the plurality of weights comprises updating the plurality of weights based on the plurality of second loss gradients relating to the plurality of weights.

12. A system for privacy-preserving logistic regression training based on homomorphically encrypted ciphertexts, the system comprising:

a memory; and

at least one processor communicatively coupled to the memory and configured to:

obtain a first packed ciphertext comprising at least a portion of a first training data sample packed into a first vector of slots thereof for training a privacy-preserving logistic regression model;

obtain a second packed ciphertext comprising a plurality of weights of the privacy-preserving logistic regression model packed into a first vector of slots thereof;

determine at least a first output probability of the privacy-preserving logistic regression model based on the first packed ciphertext and the second packed ciphertext; and

update the plurality of weights based on the first output probability, wherein

said determine at least the first output probability comprises performing a dot product between the first packed ciphertext and the second packed ciphertext, and

said perform the dot product comprises:

multiplying the first vector of slots of the first packed ciphertext and the first vector of slots of slots of the second packed ciphertext to obtain a plurality of first multiplication results in a first vector of slots of a first resultant ciphertext; and

performing a first inter-slot summation of the first vector of slots of the first resultant ciphertext to obtain a first dot product result of the dot product between the first packed ciphertext and the second packed ciphertext, wherein

said performing the first inter-slot summation comprises:

obtaining a first binary representation corresponding to the number of slots in the first vector of slots of the first resultant ciphertext;

for each digit in the first binary representation having a first binary value from the digit being the second most significant digit to the digit being the least significant digit amongst all digits in the first binary representation having the first binary value, generating a first new ciphertext based on rotating an immediately preceding ciphertext thereof by a number of slots based on a digit position of the digit in the first binary representation, wherein for a first rotation step, the immediately preceding ciphertext comprises the first resultant ciphertext, and for each subsequent rotation step, the immediately preceding ciphertext comprises the ciphertext generated in the immediately preceding rotation step, and adding the first new ciphertext with a ciphertext having an index corresponding to said digit position to produce a second new ciphertext, wherein each of the ciphertext having an index that corresponds to a digit position were generated by rotations and additions of the first resultant ciphertext based on increasing power-of-two steps from a 1^{st} step to a 2^{k-1} step where k is defined as $\log_2(n)$ and n is defined as a number of slots in the first resultant ciphertext; and

determining a summation of the first vector of slots of the first resultant ciphertext based on the second new ciphertext produced for said digit in the first binary representation having the first binary value and being the least significant digit, to produce the first dot product result.

13. The system according to claim 12, wherein the first new ciphertext is generated based on rotating the immediately preceding ciphertext thereof by the number of slots determined based on a power of two with an exponent thereof having a value corresponding to the digit position of the digit in the first binary representation.

14. The system according to claim 13, wherein said performing the first inter-slot summation further comprises: for each digit in the first binary representation from the digit being the least significant digit to the digit being the second most significant digit, generating a third new ciphertext based on rotating an immediately preceding ciphertext thereof by a number of slots based on a digit position of the digit in the first binary representation.

15. The system according to claim 14, wherein the third new ciphertext is generated based on rotating the immediately preceding ciphertext by the number of slots determined based on a power of two with an exponent thereof having a value corresponding to the digit position of the digit in the first binary representation.

16. The system according to claim 15, wherein a plurality of new ciphertexts are generated, comprising the third new ciphertext for said each digit in the first binary representation from the digit being the least significant digit to the digit being the second most significant digit and the first and second new ciphertexts for said each digit in the first binary representation from the digit being the second most significant digit to the digital being the least significant digit, each of the plurality of new ciphertexts having an index corresponding to an order of the new ciphertext in the plurality of new ciphertexts.

17. The system according to claim 16, wherein

said generating the first new ciphertext for the digit in the first binary representation being the second most significant digit is based on rotating the third new ciphertext generated for the digit in the first binary representation being the second most significant digit as the immediately preceding ciphertext thereof.

18. The system according to claim 12, wherein

the first packed ciphertext comprises a plurality of training data samples packed into a plurality of vectors of slots, respectively, thereof, for training the privacy-preserving logistic regression model, the plurality of training data samples comprising the first training data sample and one or more second training data samples, and the plurality of vectors of slots of the first packed ciphertext comprising the first vector of slots and one or more second vectors of slots of the first packed ciphertext,

the second packed ciphertext comprises a plurality of vectors of slots, each vector of slots having packed therein the plurality of weights, the plurality of vectors of slots of the second packed ciphertext comprising the first vector of slots and one or more second vectors of slots of the second packed ciphertext,

said determine at least the first output probability comprises determining a plurality of output probabilities of the privacy-preserving logistic regression model based on the first packed ciphertext and the second packed ciphertext, the plurality of output probabilities comprising the first output probability and one or more second output probabilities of the privacy-preserving logistic regression model,

said update the plurality of weights comprises updating the plurality of weights based on the plurality of output probabilities,

said performing the dot product further comprises, for each second vector of slots of said one or more second vectors of slots of the plurality of slots of the first packed ciphertext:

multiplying the second vector of slots of the first packed ciphertext and the corresponding second vector of slots of the second packed ciphertext to obtain a plurality of second multiplication results in a second vector of slots of the first resultant ciphertext; and

performing a second inter-slot summation of the second vector of slots of the first resultant ciphertext to obtain a second dot product result of the dot product between the first packed ciphertext and the second packed ciphertext, wherein

said performing the second inter-slot summation comprises:

obtaining a second binary representation corresponding to the number of slots in the second vector of slots of the first resultant ciphertext;

for each digit in the second binary representation having the first binary value from the digit being the second most significant digit to the digital being the least significant digit amongst all digits in the second binary representation having the first binary value, generating a fourth new ciphertext based on rotating an immediately preceding ciphertext thereof by a number of slots based on a digit position of the digit in the second binary representation, and adding the fourth new ciphertext with a ciphertext having an index corresponding to said digit position to produce a fifth new ciphertext; and

determining a summation of the second vector of slots of the first resultant ciphertext based on the fifth new ciphertext produced for said digit in the second binary

representation having the first binary value and being the least significant digit, to produce the second dot product result.

19. The system according to claim 12, wherein

said determine at least the first output probability further comprises applying an activation function on the first dot product result to produce the first output probability,

the method further comprises:

determining a first loss gradient relating to the first training data sample based on the first output probability; and

determining, for each of the plurality of weights, a second loss gradient relating to the weight based on the first loss gradient and a slot of the first vector of slots of the first packed ciphertext corresponding to the weight to obtain a plurality of second loss gradients relating to the plurality of weights, and

said update the plurality of weights comprises updating the plurality of weights based on the plurality of second loss gradients relating to the plurality of weights.

20. A computer program product, embodied in one or more non-transitory computer readable storage mediums, comprising instructions executable by at least one processor to perform a method of privacy-preserving logistic regression training based on homomorphically encrypted ciphertexts, the method comprising:

obtaining a first packed ciphertext comprising at least a portion of a first training data sample packed into a first vector of slots thereof for training a privacy-preserving logistic regression model;

obtaining a second packed ciphertext comprising a plurality of weights of the privacy-preserving logistic regression model packed into a first vector of slots thereof;

determining at least a first output probability of the privacy-preserving logistic regression model based on the first packed ciphertext and the second packed ciphertext; and

updating the plurality of weights based on the first output probability, wherein

said determining at least the first output probability comprises performing a dot product between the first packed ciphertext and the second packed ciphertext, and

said performing the dot product comprises:

multiplying the first vector of slots of the first packed ciphertext and the first vector of slots of the second packed ciphertext to obtain a plurality of first multiplication results in a first vector of slots of a first resultant ciphertext; and

performing a first inter-slot summation of the first vector of slots of the first resultant ciphertext to obtain a first dot product result of the dot product between the first packed ciphertext and the second packed ciphertext, wherein

said performing the first inter-slot summation comprises:

obtaining a first binary representation corresponding to the number of slots in the first vector of slots of the first resultant ciphertext;

for each digit in the first binary representation having a first binary value from the digit being the second most significant digit to the digit being the least significant digit amongst all digits in the first binary representation having the first binary value, generating a first new ciphertext based on rotating an immediately preceding ciphertext thereof by a number of slots based on a digit position of the digit in the first binary representation, wherein for a first rotation step, the immediately preceding ciphertext comprises the first resultant ciphertext, and for each subsequent rotation step, the immediately preceding ciphertext comprises the ciphertext generated in the immediately preceding rotation step, and adding the first new ciphertext with a ciphertext having an index corresponding to said digit position to produce a second new ciphertext, wherein each of the ciphertext having an index that corresponds to a digit position were generated by rotations and additions of the first resultant ciphertext based on increasing power-of-two steps from a 1^{st} step to a 2^{k-1} step where k is defined as $\log_2(n)$ and n is defined as a number of slots in the first resultant ciphertext; and

determining a summation of the first vector of slots of the first resultant ciphertext based on the second new ciphertext produced for said digit in the first binary representation having the first binary value and being the least significant digit, to produce the first dot product result.