

Non-interactive privacy-preserving naïve Bayes classifier using homomorphic encryption

Jingwei Chen



CIGIT

Joint work with Y. Feng, Y. Liu, W. Wu & G. Yang

December, 2021 @ SPNCE 2021

Machine Learning as a Service (MLaaS)

- Learn a model from massive amounts of data
 - Data owner \implies model provider
- Infer predicted results for client's sample data
 - Client can easily obtain the predicted results.

Machine Learning as a Service (MLaaS)

- Learn a model from massive amounts of data
 - Data owner \implies model provider
- Infer predicted results for client's sample data
 - Client can easily obtain the predicted results.

Privacy concerns

- The model may be sensitive:
 - financial model, disease diagnosis, ...
- Sample data may be sensitive:
 - credit history, medical records, ...




Figure: Framework of privacy-preserving classifiers

- The model is already known \implies no learning
- Model provider is also service provider \implies 2PC
- Light-blue boxes are encrypted \implies secure

Threat model

Adversaries are passive (honest-but-curious).

Here are several privacy-preserving naïve Bayes classifiers based on HE:

- [Bost *et al.* '15]: 2PC, Quadratic Residuosity + Paillier
- [Li, *et al.* '16]: 4PC, Paillier
- [Kim, *et al.* '18]: 4PC, BGV
- [Yasumura, *et al.* '19]: MPC, BGV
- [Sun, *et al.* '20]: 2PC, BGV
-


Here are several privacy-preserving naïve Bayes classifiers based on HE:

- [Bost *et al.* '15]: 2PC, Quadratic Residuosity + Paillier
- [Li, *et al.* '16]: 4PC, Paillier
- [Kim, *et al.* '18]: 4PC, BGV
- [Yasumura, *et al.* '19]: MPC, BGV
- [Sun, *et al.* '20]: 2PC, BGV
-


Common feature

During the classification phase, interactions are needed:

- communication burden
- (potential) information leakage




We propose a privacy-preserving naïve Bayes classifier:



We propose a privacy-preserving naïve Bayes classifier:


- **non-interactive**

- Client sends an encryption of his sample data.
- Server evaluates the **model** and sends the encrypted **results**.
- Client decrypts the results to recovery the classification.



We propose a privacy-preserving naïve Bayes classifier:

- **non-interactive**
 - Client sends an encryption of his sample data.
 - Server evaluates the **model** and sends the encrypted **results**.
 - Client decrypts the results to recovery the classification.
- **post-quantum safe**
 - Based on BGV



We propose a privacy-preserving naïve Bayes classifier:

- **non-interactive**
 - Client sends an encryption of his sample data.
 - Server evaluates the **model** and sends the encrypted **results**.
 - Client decrypts the results to recovery the classification.
- **post-quantum safe**
 - Based on BGV

Main technique

An algorithm to compute **argmax** of an encrypted array.

1 Background

2 Building blocks

3 Proposed protocol

4 Experimental results

1 Background

2 Building blocks

3 Proposed protocol

4 Experimental results

Consider a data set with

- s categories $1, 2, \dots, s$
- n features X_1, \dots, X_n
- each feature X_k has at most t different values $1, 2, \dots, t$.

Then the classification of a sample $\mathbf{x} = (x_1, \dots, x_n)$ is

$$s^* = \underset{i=1, \dots, s}{\operatorname{argmax}} \Pr[Y = i] \prod_{k=1}^n \Pr[X_k = x_k | Y = i],$$

- prior probability: $\Pr[Y = i]$
- likelihood: $\Pr[X_k = x_k | Y = i]$

Consider a data set with

- s categories $1, 2, \dots, s$
- n features X_1, \dots, X_n
- each feature X_k has at most t different values $1, 2, \dots, t$.

Then the classification of a sample $\mathbf{x} = (x_1, \dots, x_n)$ is

$$s^* = \underset{i=1, \dots, s}{\operatorname{argmax}} \left\{ \log \Pr[Y = i] + \sum_{k=1}^n \log \Pr[X_k = x_k | Y = i] \right\}$$

- prior probability: $\Pr[Y = i]$
- likelihood: $\Pr[X_k = x_k | Y = i]$

Homomorphic encryption

A public key encryption scheme consists of

- **KeyGen**: $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$,
- **Enc**: $c \leftarrow \text{Enc}(pk, x)$ for $x \in \mathcal{P}$,
- **Dec**: $x \leftarrow \text{Dec}(sk, c)$ for $c \in \mathcal{C}$.

LHE/FHE

Let \mathcal{F} be a set of function in $\{0, 1\}^* \rightarrow \{0, 1\}$. A public key scheme is **\mathcal{F} -homomorphic** if there exists an evaluation algorithm **Eval** s.t.

$$\forall f \in \mathcal{F}, \forall x \in \mathcal{P}, \text{Dec}(sk, \text{Eval}(f, \text{Enc}(pk, x))) = f(x).$$

- **Leveled HE (LHE)**: $f \in \mathcal{F}$ has an *a priori* bound on the depth of its circuit.
 - **Fully HE (FHE)**: $f \in \mathcal{F}$ can be an arbitrary function.
- LHE + bootstrapping \implies FHE

BGV scheme [Brakerski-Gentry-Vaikuntanathan '12]

- Efficient LHE, supporting FHE, RLWE-based
- $\mathcal{P} = R_p := \mathbb{Z}[X]/(\Phi_m(X), p) \cong \mathbb{F}_p^\ell$
 - $d = \text{ord}(p)$ in \mathbb{Z}_m^* and $\ell = \phi(m)/d$
 - Support SIMD operations
- $\mathcal{C} = R_q := \mathbb{Z}[X]/(\Phi_m(X), q)$
- It evaluates L -level circuits with $O(\lambda \cdot L^3)$ per-gate computation.

HElib: an implementation of BGV¹

- Based on NTL², thread safe
- Batching based on SIMD operations
- “Assembly language” for HE:
 - Add, Mul, Rotate, Shift, TotalSum, AddConst, MulConst, ...

¹<https://github.com/homenc/HElib>

²<https://github.com/libntnl/ntl>

1 Background

2 Building blocks

3 Proposed protocol

4 Experimental results

Plain matrix-encrypted vector multiplication

Input: $\mathbf{c}' = (c'_i)_{i \leq t} \in R_q^t$ (c'_i encrypts the i th entry of $\mathbf{z} = (z_i)_{i \leq t}$), and pk ; $\mathbf{A} = (a_{i,j}) \in \mathbb{Z}^{s \times t}$.

Output: $\mathbf{c} = (c_i)_{i \leq s}$ with $c_i = \text{Enc}_{\text{pk}}\left(\sum_{j=1}^t a_{i,j} z_j\right)$.

- 1: For $i = 1, \dots, s$ do the following:
- 2: $c_i \leftarrow \text{Enc}_{\text{pk}}(0)$;
- 3: For $j = 1, \dots, t$, update $c_i := \text{Add}_{\text{pk}}(c_i, \text{MulConst}_{\text{pk}}(a_{i,j}, c'_j))$.
- 4: **return** $(c_i)_{i \leq s}$.

Algorithm 1: Naïve plaintext matrix-encrypted vector multiplication

Plain matrix-encrypted vector multiplication

Input: $\mathbf{c}' = (c'_i)_{i \leq t} \in R_q^t$ (c'_i encrypts the i th entry of $\mathbf{z} = (z_i)_{i \leq t}$), and pk ; $\mathbf{A} = (a_{i,j}) \in \mathbb{Z}^{s \times t}$.

Output: $\mathbf{c} = (c_i)_{i \leq s}$ with $c_i = \text{Enc}_{\text{pk}}\left(\sum_{j=1}^t a_{i,j} z_j\right)$.

- 1: For $i = 1, \dots, s$ do the following:
- 2: $c_i \leftarrow \text{Enc}_{\text{pk}}(0)$;
- 3: For $j = 1, \dots, t$, update $c_i := \text{Add}_{\text{pk}}(c_i, \text{MulConst}_{\text{pk}}(a_{i,j}, c'_j))$.
- 4: **return** $(c_i)_{i \leq s}$.

Algorithm 1: Naïve plaintext matrix-encrypted vector multiplication

- Only pt-ct multiplications and ct-ct additions are involved.
- It costs no ct-ct multiplication depth.

The less-than function over $S = [0, (p - 1)/2]$ is defined as

$$\text{LT}_S(x, y) = \begin{cases} 1, & \text{if } 0 \leq x < y \leq (p - 1)/2, \\ 0, & \text{if } 0 \leq y \leq x \leq (p - 1)/2. \end{cases}$$

The less-than function over $S = [0, (p-1)/2]$ is defined as

$$\text{LT}_S(x, y) = \begin{cases} 1, & \text{if } 0 \leq x < y \leq (p-1)/2, \\ 0, & \text{if } 0 \leq y \leq x \leq (p-1)/2. \end{cases}$$

It can be interpolated by the following polynomial over \mathbb{F}_p of degree $p-1$:

$$\frac{p+1}{2}(x-y)^{p-1} + \sum_{i=1, \text{odd}}^{p-2} \left(\sum_{a=1}^{\frac{p-1}{2}} a^{p-1-i} \right) \cdot (x-y)^i.$$

The less-than function over $S = [0, (p-1)/2]$ is defined as

$$\text{LT}_S(x, y) = \begin{cases} 1, & \text{if } 0 \leq x < y \leq (p-1)/2, \\ 0, & \text{if } 0 \leq y \leq x \leq (p-1)/2. \end{cases}$$

It can be interpolated by the following polynomial over \mathbb{F}_p of degree $p-1$:

$$\frac{p+1}{2}(x-y)^{p-1} + \sum_{i=1, \text{odd}}^{p-2} \left(\sum_{a=1}^{\frac{p-1}{2}} a^{p-1-i} \right) \cdot (x-y)^i.$$

- Evaluating the polynomial homomorphically costs at most

$$\sqrt{p-3} + \frac{3}{2} \log_2(p-3) + O(1)$$

ct-ct multiplication depth.

Argmax of an encrypted array

Given an array $\mathbf{z} = (z_0, \dots, z_{t-1})$,

$$\arg \max_i(\mathbf{z}) = \sum_{j=0}^{t-1} j \cdot \prod_{k=0, k \neq j}^{t-1} (1 - \text{LT}(z_j, z_k)).$$

Argmax of an encrypted array

Given an array $\mathbf{z} = (z_0, \dots, z_{t-1})$,

$$\mathop{\text{arg max}}_i(\mathbf{z}) = \sum_{j=0}^{t-1} j \cdot \prod_{k=0, k \neq j}^{t-1} (1 - \text{LT}(z_j, z_k)).$$

Evaluating $\mathop{\text{arg max}}$ costs:

- $t(t-1)/2$ comparisons and
- at most

$$\log_2 t + \sqrt{p-3} + \frac{3}{2} \log_2(p-3) + O(1)$$

ct-ct multiplication depth.

1 Background

2 Building blocks

3 Proposed protocol

4 Experimental results

Privacy-preserving naïve Bayes classifier

Input of client: A sample $x = (i_1, \dots, i_n)$, sk and pk .

Input of server: Likelihood and prior information: $(A_k)_{k \leq n}$, $(b_i)_{i \leq s}$, and pk .

- 1: The client encode x to a matrix $(e_{i_1}, \dots, e_{i_n}) \in \{0, 1\}^{t \times n}$.
- 2: The client encrypts e_{i_k} and sends the ciphertexts to the server.
- 3: The server does the following:
 - 4: For $i = 1, \dots, s$, set $c_i \leftarrow \text{Enc}_{pk}(0)$ and $c_i := \text{AddConst}_{pk}(c_i, b_i)$.
 - 5: For $k = 1, \dots, n$, calling Algorithm 1 with input as the ciphertexts of e_{i_k} , A_k and pk outputs $(c'_i)_{i \leq s}$. Update $c_i := \text{Add}_{pk}(c_i, c'_i)$ for $i = 1, \dots, s$.
- 6: Calling **arg max** with input as $c = (c_i)_{i \leq s}$ and pk returns c .
- 7: The server sends c to the client.
- 8: The client decrypts c to $y = \text{Dec}_{sk}(c)$.

Protocol 1: Privacy-preserving naïve Bayes classifier

Privacy-preserving naïve Bayes classifier

Input of client: A sample $x = (i_1, \dots, i_n)$, sk and pk .

Input of server: Likelihood and prior information: $(\mathbf{A}_k)_{k \leq n}$, $(b_i)_{i \leq s}$, and pk .

- 1: The client encode x to a matrix $(\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_n}) \in \{0, 1\}^{t \times n}$.
- 2: The client encrypts \mathbf{e}_{i_k} and sends the ciphertexts to the server.
- 3: The server does the following:
 - 4: For $i = 1, \dots, s$, set $c_i \leftarrow \text{Enc}_{\text{pk}}(0)$ and $c_i := \text{AddConst}_{\text{pk}}(c_i, b_i)$.
 - 5: For $k = 1, \dots, n$, calling Algorithm 1 with input as the ciphertexts of \mathbf{e}_{i_k} , \mathbf{A}_k and pk outputs $(c'_i)_{i \leq s}$. Update $c_i := \text{Add}_{\text{pk}}(c_i, c'_i)$ for $i = 1, \dots, s$.
- 6: Calling **arg max** with input as $\mathbf{c} = (c_i)_{i \leq s}$ and pk returns c .
- 7: The server sends c to the client.
- 8: The client decrypts c to $y = \text{Dec}_{\text{sk}}(c)$.

Protocol 1: Privacy-preserving naïve Bayes classifier

- No interaction during classification phase (Step 3–6)
- Secure classification without sacrificing privacy, assuming secure HE

1 Background

2 Building blocks

3 Proposed protocol

4 Experimental results

- Data set:
 - $s = 3, n = 4, t = 5$
 - 150 samples = 120 for training (80%) + 30 for testing
- HElib parameters:
 - $p = 37, m = 14539$ ($\implies \ell = 1980$ slots), $\log q \approx 387$
 - The security parameter $\lambda \approx 100$
- Performance (Ubuntu 20.04/Intel i7-10750H/16GB):
 - 5.42s for 30 testing samples
 - ≈ 0.18 s for each
 - Supporting at most 1980 testing samples
 - Amortized costs ≈ 2.7 ms per sample
- Accuracy: $\approx 97\%$

³<http://archive.ics.uci.edu/ml>

- Data set:
 - $s = 2, n = 9, t = 10$
 - 683 samples = 478 for training (70%) + 205 for testing
- HElib parameters:
 - $p = 113, m = 12883$ ($\implies \ell = 3960$ slots), $\log q \approx 382$
 - The security parameter $\lambda \approx 100$
- Performance (Ubuntu 20.04/Intel i7-10750H/16GB):
 - 4.75s for 205 testing samples
 - ≈ 23 ms for each
 - Supporting at most 3960 testing samples
 - Amortized costs ≈ 1.2 ms per sample
- Accuracy: $\approx 97\%$

³<http://archive.ics.uci.edu/ml>

Conclusion

- We proposed a privacy-preserving naïve Bayes classifier based on BGV.
- R-LWE based LHE \implies secure against quantum attackers
- Argmax of encrypted arrays \implies non-interactive
- Batching \implies efficient

Conclusion

- We proposed a privacy-preserving naïve Bayes classifier based on BGV.
- R-LWE based LHE \implies secure against quantum attackers
- Argmax of encrypted arrays \implies non-interactive
- Batching \implies efficient

Future work

- More efficient
- More classifiers

Conclusion

- We proposed a privacy-preserving naïve Bayes classifier based on BGV.
- R-LWE based LHE \implies secure against quantum attackers
- Argmax of encrypted arrays \implies non-interactive
- Batching \implies efficient

Future work

- More efficient
- More classifiers

The full version is available at

https://www.arcnl.org/jchen/download/ppnb_full.pdf

THANKS