

Privacy Preserving Machine Learning



Angel Navia Vázquez
Machine Learning for Data Science (ML4DS) Group

MLG, may 2018

Telling stories: Alice and Bob



Alice



Bob

- *Alice*: Take this suitcase and please count the money
- *Bob*: Sure! But it is locked. Please, give me key.
- *Alice*: I am sorry. I can't trust you. You have to count it without opening it.
- *Bob*: Then, it can't be done!

What is Privacy Preserving Machine Learning (PPML)?

- Some users have some data
- A computation has to be carried out without revealing the data



- *Alice:* Take these encrypted files. Please sum all the numbers in them.
- *Bob:* Sure! Please, give me decryption key.
- *Alice:* I am sorry. I can't trust you. You have to sum them without decrypting the files.
- *Bob:* Hummmmm... let me think...

3

How to learn while protecting privacy?

- PPML approaches:
 - **Noise perturbation techniques:** an additive or multiplicative noise is included in the data aiming at hiding the original values. **WE'RE NOT INTERESTED**
 - **Differential privacy:** probabilistic approach that aims at maximizing the information provided by queries from a database while minimizing the chance of identifying individual records. Relies on sanitization techniques:
 - **Input/output perturbation:** add noise to inputs or computation results. **WE'RE NOT INTERESTED.**
 - **k-anonymization:** verify that at least k elements provide the same attributes. **WE'RE NOT INTERESTED.**
 - **Federated Machine Learning:** updating a master model with incremental updates from every client owning a subset of the training set. **Warning! The model is shared...**
 - **Distributed Machine Learning:** the exchanged information is not raw data and does not reveal any information concerning individual users. **Warning! The model can be shared...**
 - **Partially Homomorphic Cryptography:** Homomorphic encryption schemes guarantee some equivalences between operations performed on the clear and encrypted domain. **Warning! some operations may not be supported! Trusted Third Party may be necessary.**
 - **Fully Homomorphic Cryptography:** schemes that guarantee the equivalence between any operation performed on the clear and encrypted domain. **Complexity? Feasibility?**
 - **Secure Multiparty Computation:** every user only gets as information the output of a function on its own data. **Warning! Very restrictive, some operations can be very complex / time consuming.**

4

Differential privacy

- Privacy definition: With or without including me in the database, my privacy risk should not change much
- Privacy is lost if:
 - An individual record is leaked
 - If a cross-analysis with public data allows to identify individuals
 - A published statistics S allows to better estimate a private value, e.g.: *Alice's height is 2 cm above average height in the population.*

Hospital Patient Data

DOB	Sex	Zipcode	Disease
1/21/76	Male	53715	Heart Disease
4/13/86	Female	53715	Hepatitis
2/28/76	Male	53703	Brochitis
1/21/76	Male	53703	Broken Arm
4/13/86	Female	53706	Flu
2/28/76	Female	53706	Hang Nail

Vote Registration Data


Name	DOB	Sex	Zipcode
Andre	1/21/76	Male	53715
Beth	1/10/81	Female	55410
Carol	10/1/44	Female	90210
Dan	2/21/84	Male	02174
Ellen	4/19/72	Female	02237

Information leak!!!: Andre has heart disease!

5

Federated Machine Learning

- [Google's Federated Averaging algorithm](#) (MacMahan et al.):
 1. Construct some master model with weights \mathbf{w} .
 2. Download your master model to each of your clients.
 3. For each client, compute an updated, personalized set of weights \mathbf{w}_k using local training data — this is done in parallel and offline. **Data does not leave the client.**
 4. Now, update \mathbf{w} with the weighted average of all \mathbf{w}_k .
 5. (Return to step #2)



6

Federated Machine Learning (II)


- Useful when:
 - Users have Privacy Protected Data
 - Data cannot leave the users' facilities
 - Computation is decentralized
- The algorithms must be robust against:
 - Non IID data
 - Asynchronous updates
 - Unbalanced data
 - Massive distribution
 - Dynamic datasets
 - Other restrictions: *in mobiles, only operate when plugged in, connected to wifi, etc.*
- We must minimize:
 - The amount of information to be transmitted
 - The communication cost
 - The local computation (low cost processors)

7

Federated Machine Learning (III)

- Complex models better preserve the privacy:
 - *E.g.: in linear language model on sparse data, the nonzero gradients reveal the typed words... the less explicable, the better...*
- Deep Learning experiments [Shokri'15]
 - SSGD Selective Stochastic Gradient Descent: helps to better generalize
 - Even sharing 1% of the parameters improves with respect to standalone learning

	SGD	0.1	0.01	0.001	Standalone
MNIST, CNN	0.9917	0.9914	0.9871	0.9645	0.9316
SVHN, CNN	0.9299	0.9312	0.8986	0.7481	0.8182
	SGD	0.1	0.01	0.001	Standalone
MNIST, MLP	0.9810	0.98	0.9707	0.9171	0.8832
SVHN, MLP	0.8476	0.8394	0.7833	0.6542	0.5136




8


Distributed Machine Learning

- Data is located in different nodes
- Partial computations are locally done and shared ($K^T K$, $K^T y$)
- The final model is adjusted. e.g: Distributed SVM training [Navia'06]

Local training




Distributed training



9

Alice and Bob are back!


- Symmetric key cryptography
 - A single (secret) key is used to encrypt and decrypt
 - The only method known until 1976
 - Key management is a problem:
 - How to securely send the key to Alice?
 - An increasing number of keys is needed (square with the number of network members)
- Asymmetric or public key cryptographic system is proposed in 1976: W. Diffie and M. Hellman. New Directions in Cryptography. IEEE Transactions on Information Theory. IT-22 (6): 644–654. 1976.



10

Public Key cryptography

- The encryption key is public
- The decryption key is private
- Both keys are related, but it is unfeasible to obtain the private key from the public key
- The public key can be freely distributed
- Diffie and Hellman did not find the complete system, only the protocol for key sharing
- 1978: Rivest, Shamir and Adleman defined a working system, currently known as the RSA algorithm.
- Used in most popular digital signature systems: SSL/TLS, VPNs, etc.



← → ↻ Es seguro | https://

11

Homomorphisms

- Algebra definition:
 - structure-preserving map between two algebraic structures of the same type (two groups, two rings, etc.)
- Example:
 - G_1 : Real numbers are a group for addition
 - G_2 : Positive real numbers are a group for multiplication
 - The exponential function $f(x) = e^x$ is a homomorphism between them:
 - $f(x+y) = e^{x+y} = e^x e^y$ (addition in G_1 is equivalent to multiplication in G_2)
- Ring homomorphism: map between rings that preserves the addition, multiplication, and the multiplicative identity
 - $f(a + b) = f(a) + f(b)$ for all a and b in R
 - $f(ab) = f(a) f(b)$ for all a and b in R
 - $f(1_R) = 1_S$

12

Homomorphic encryption

- **Partially homomorphic encryption PHE** (either sum or product are supported)
 - Unpadded RSA (multiplicative)
 - ElGamal (multiplicative)
 - Goldwasser–Micali (additive)
 - Benaloh (additive)
 - Paillier (additive)
- **Somewhat homomorphic encryption SHE:**
 - Also known as (XOR, AND)-homomorphic encryption
 - (XOR, AND) defines a Turing-complete system, we can compute any function on encrypted bits
 - Allows many additions and **some** multiplications
- **Fully homomorphic encryption FHE**
 - Supports both sum and product
 - It preserves ring structure $(A, +, \bullet)$
 - Examples:
 - Gentry's cryptosystem [Gentry09]
 - Cryptosystem over the integers
 - 2nd generation

13

Paillier cryptosystem

- Paillier cryptosystem exploits the fact that certain discrete logarithms can be computed easily
- Powers of a multiplicative subgroup $G = \{..., b^{-3}, b^{-2}, b^{-1}, 1, b^1, b^2, b^3, ...\}$ allow an easy computation of $\log_b a$ (e.g. $\log_b b^3 = 3$)
- *Binomial Theorem:*

$$(1 + n)^x = \sum_{k=0}^x \binom{x}{k} n^k = 1 + nx + \binom{x}{2} n^2 + \text{higher powers of } n$$

- *Then:*

$$(1 + n)^x \equiv 1 + nx \pmod{n^2}$$

$$y = (1 + n)^x \pmod{n^2}$$

$$(y = 1 + nx \pmod{n^2})$$

$$x \equiv \frac{y - 1}{n} \pmod{n}.$$

$$L((1 + n)^x \pmod{n^2}) \equiv x \pmod{n}, \quad L(u) = \frac{u - 1}{n}$$

14

Paillier cryptosystem (II)

- **Key generation:**
 - Choose two large prime numbers p and q randomly and independently of each other such that $\gcd(pq, (p-1)(q-1))=1$
 - Compute $n=pq$ and $\lambda = \text{lcm}(p-1, q-1)$
 - Select random integer g in Z^*n^2
 - check the existence of the following modular multiplicative inverse: $\mu = (L(g^\lambda \bmod n^2))^{-1} \bmod n$
 - The public (encryption) key is (n, g)
 - The private (decryption) key is (λ, μ)
- **Encryption:**
 - Let m be a message to be encrypted where $0 \leq m < n$
 - Select random r where $0 < r < n$ (nondeterministic encryption)
 - Compute ciphertext as: $c=(g^m \times r^n) \bmod n^2$
- **Decryption:**
 - Let c be the ciphertext to decrypt, c in Z^*n^2
 - Compute the plaintext message as: $m=L(c^\lambda \bmod n^2) \times \mu \bmod n$

15

Paillier homomorphic properties

- The product of two ciphertexts will decrypt to the sum of their corresponding plaintexts

$$D(E(m_1, r_1)) \times E(m_2, r_2) \bmod n^2 = m_1 + m_2 \bmod n$$
- The product of a ciphertext with a plaintext raising g will decrypt to the sum of the corresponding plaintexts

$$D(E(m_1, r_1)) \times g^{m_2} \bmod n^2 = m_1 + m_2 \bmod n$$
- An encrypted plaintext raised to the power of another plaintext will decrypt to the product of the two plaintexts,

$$D(E(m_1, r_1)^{m_2} \bmod n^2) = m_1 m_2 \bmod n$$

$$D(E(m_2, r_2)^{m_1} \bmod n^2) = m_1 m_2 \bmod n$$
- An encrypted plaintext raised to a constant k will decrypt to the product of the plaintext and the constant,

$$D(E(m_1, r_1)^k \bmod n^2) = k m_1 \bmod n$$
- Given the Paillier encryptions of two messages there is no known way to compute an encryption of the product of these messages without knowing the private key.

16

Python example

```

from paillier.paillier import *
priv, pub = generate_keypair(512)

<PrivateKey:
798118547958500245259981069765619625628486136783124026485990453392220872813728033186211995518324090070051355953664166
615424331878953264099307183676192082857769585333462739006841720237665186879921732227444670044974240474743832343622275
3280574578702592416342330563340833663544646150158205288741878048629243651>

<PublicKey:
7981185479585002452599810697656196256284861367831240264859904533922208728137462594507789989881575029791719981405
347528323252888280515201802857785430853213>

x = 3
Y = 5
cx = encrypt(pub, x)
cy = encrypt(pub, y)

cx
2464122010767311540380010188701068415658146069097357176786693581125378203706108869610339087992140267694030516037
0503889413254858877138494539573422865180279650182282544366625778908110318857458598413139857309548184312343072739
32401660599914756279900682003314734411563525193213131616818273732455261263949141090L

cy
1639654100165493230200739114632710118806968788608610405613162502531479834237311119281938506869201185488535481898
2267759640795904619083164846017637756406419696406888252423648247526196464801232597847488454528793339210436792733
354386977678729706093166702429762725145482584069220181876314585322278921760606900715L

```

17

Python example (II)

```

cz = e_add(pub, cx, cy)
cz
30874372195658557163990323289550203652837363336018643416848853002649985289658553293037575252852126966976976519883033
282451433902046910165235673573984275386592437390893419636046407181933075983964631717130163163829787926871637375676873
20245157858852416749369364787331164090887233101422348241479577652127455245L

z = decrypt(priv, pub, cz)
z
8L

cq = e_add_const(pub, cz, 2)
cq
336866228544107405684372345216953267869146674574444052817365983881326621398350755863305216417867294096836266102120915
829686316609393300868837963363959995842092163275303619734513509576783126851594846626072762482030478546844578908922523
90852642739418417770605903466796699997142267348199023010959007761169456220L

q = decrypt(priv, pub, cq)
q
10L

cp = e_mul_const(pub, cq, 11)
cp
646813598018216077262570768607977562224086941654680088210333908266028369732940115977126870678406078633281335833371196213614
484790268596019355640438681709592894838979695433796669271542226265175393254432926105809266114318412216089596729249064309602
5880250340072226736382371845551931564628826536637320907225854L

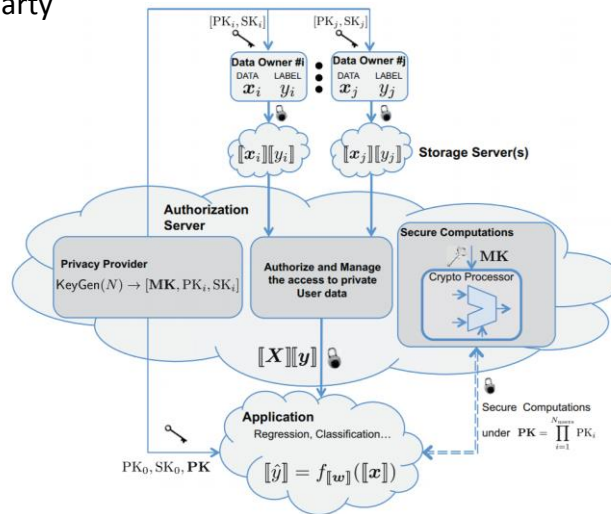
p = decrypt(priv, pub, cp)
p
110L

```

18

Proposed approach #1

- A trusted third party provides the unsupported operations (González17):



19

Towards FHE: (XOR, AND)-homomorphic encryption with private keys (DGHV'10)

- Choose large odd number p (n^2 bits) \rightarrow decryption key (fixed)
- To encrypt bit m :
 - Choose large q (n^5 bits) and small r (n bits), random numbers
 - $c = p \cdot q + 2 \cdot r + m$ ($2r$ is "noise"!). r is needed to avoid recovering the secret key using two encryptions of "0" as $\text{GCD}(q_1 p, q_2 p)$
 - $2 \cdot r + m$ much smaller than p , s.t. $2r + m \pmod{2} = m \pmod{2}$
 - ciphertext is close to a multiple of p , s.t. $c \pmod{p} = (2r + m) \pmod{p}$
- Decryption:
 - $m \equiv (c \pmod{p}) \pmod{2}$
- Adding (XOR) two encrypted bits ($2(r_1 + r_2)$ is "noise"!):

$$c_1 + c_2 = (q_1 + q_2) \cdot p + 2 \cdot (r_1 + r_2) + (m_1 + m_2)$$

\hookrightarrow if $2 \cdot (r_1 + r_2) + (m_1 + m_2)$ much smaller than p

$\hookrightarrow (c_1 + c_2 \pmod{p}) \pmod{2} \equiv m_1 + m_2 \pmod{2}$

20

(XOR, AND)-homomorphic encryption (DGHV'10)

- Multiply (AND) two encrypted bits ($2(r_1r_2)$ is "noise"!):

$$c_1 \cdot c_2 = \dots = (q_1q_2p + 2q_1r_2 + q_1m_2 + 2q_2r_1 + q_2m_1)p + 2(2r_1r_2 + r_1m_2 + m_1r_2) + m_1m_2$$

$$c_1q_2 = q_1q_2p + 2q_2r_1 + q_2m_1$$

$$c_2q_1 = q_1q_2p + 2q_1r_2 + q_1m_2$$

$$c_1 \cdot c_2 = (c_1q_2 + c_2q_1 - q_1q_2p)p + 2(2r_1r_2 + r_1m_2 + m_1r_2) + m_1m_2$$

↪ if $2(2r_1r_2 + \dots) + m_1m_2$ much smaller than p

$$\hookrightarrow (c_1 \cdot c_2 \bmod p) \bmod 2 \equiv m_1 \cdot m_2 \pmod{2}$$

- But...
 - Ciphertext grows with each operation: $|c_1 + c_2| = n + 1$ and $|c_1 \cdot c_2| = 2 \cdot n$
 - Noise doubles on addition and squares on multiplication... ciphertext may become unrecoverable after **some** operations.
 - This approach is **Some**what Homomorphic Encryption (SHE)

21

Fully Homomorphic Encryption

- Gentry's "bootstrapping" theorem:
 - "If an encryption scheme can evaluate its own decryption circuit, then it can evaluate everything"
 - Problem: operations increase noise
 - Decryption reduces noise, but the key is private. How to reduce noise without knowing the private key?
- Steps:
 - Construct a SHE scheme that supports evaluation of D -degree polynomials
 - Squash the decryption procedure and express it as a low degree polynomial ($D/2$)
 - Apply a bootstrapping transformation to obtain the FHE scheme

22

Bootstrapping


RSA&friends:

MANY mult

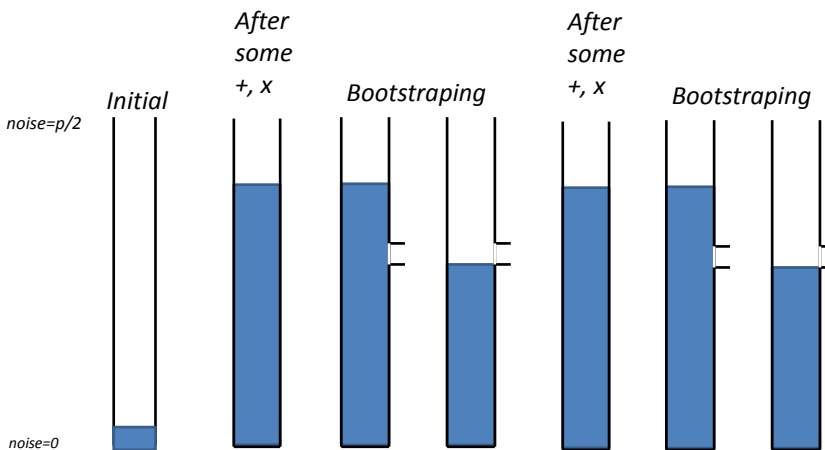
ZERO add



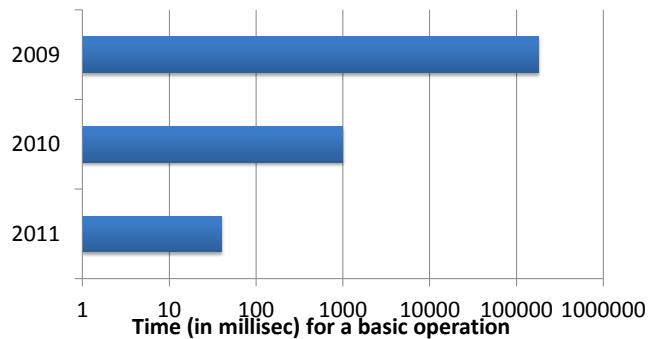
- How to reduce the noise?:
 - Decryption $D(\text{Enc}(m), p)$ reduces noise to zero (**but secret key is needed!!!**)
 - Use $\text{Enc}(p)$ (circular encryption) and homomorphically evaluate the decryption $\text{Enc}(m)^* = D(\text{Enc}(m), \text{Enc}(p))$
 - $\text{Enc}(m)^*$ and $\text{Enc}(m)$ have different noise levels
 - Regardless noise level in $\text{Enc}(m)$, the noise level in $\text{Enc}(m)^*$ is fixed!!!



Every Bootstrap stage reduces noise to a fixed level



Efficiency improvement over time



25

FHE practical applications

- **Advertising:**
 - a consumer wants product recommendations but without revealing its identity [Jeckmans]
 - A consumer receives encrypted recommendations from a recommender [Armknrecht]
 - A consumer wants to receive contextual ads, but without revealing its position [Naehrig2011]
- **Medical applications:**
 - The patient data is uploaded (encrypted) to a system that analyses it and tracks the users health [Naehrig2011]
- **Data Mining:**
 - Customer data analysis [Yang]
- **Financial privacy:**
 - Outsource computation on private data with private algorithms [Naehrig2011]
- **Image recognition:**
 - Detect illegal images on a hard drive based on a Police database [Bösch]
- **Computation delegation:**
 - Outsource computation to a third party while checking that it was done correctly [Chung]
- **Private internet search:**
 - Retrieve information without revealing the query
- **Spam filtering:**
 - Blacklisting encrypted emails
- **DNA analysis:**
 - Search for DNA markers without revealing DNA itself
- **Electronic voting:**
 - Compute the result of a voting without decrypting the ballots

26

References

- Frederik Armknecht and Thorsten Strufe. **An efficient distributed privacy preserving recommendation system**. In The 10th IFIP Annual Mediterranean Ad Hoc Networking Workshop, Med-Hoc-Net 2011, pages 65–70. IEEE, 2011.
- Christoph B^osch, Andreas Peter, Pieter H. Hartel, and Willem Jonker. **SOFIR: securely outsourced forensic image recognition**. In IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2014, pages 2694–2698. IEEE, 2014.
- Kai-Min Chung, Yael Tauman Kalai, and Salil P. Vadhan. **Improved delegation of computation using fully homomorphic encryption**. In Rabin [52], pages 483–501.
- W. Diffie and M. Hellman. **New Directions in Cryptography**. IEEE Transactions on Information Theory. IT-22 (6): 644–654. 1976.
- González-Serrano F.J, Navia-Vázquez A. and Amor-Martín A. **Training Support Vector Machines with privacy-protected data**. Pattern Recognition. Vol. 72, pp. 93-107, 2017
- Hayes, Brian. (2012). **Alice and Bob in Cipherspace**. American Scientist. 100. 362-367.
- Arjan Jeckmans, Andreas Peter, and Pieter H. Hartel. **Efficient privacy enhanced familiarity-based recommender system**. In Jason Crampton et al., editors, Computer Security – ESORICS 2013, volume 8134 of Lecture Notes in Computer Science, pages 400–417. Springer, 2013.
- McMahan et al. **Communication-Efficient Learning of Deep Networks from Decentralized Data**. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017, Fort Lauderdale, Florida, USA. JMLR: W&CP volume 54.
- A. Navia-Vázquez, D. Gutiérrez-González, E. Parrado-Hernández, J.J Navarro-Abellán, **Distributed Support Vector Machines**. *IEEE Tr. Neural Networks*, Vol. 17, No. 4, pp. 1091-1097, 2006.
- Michael Naehrig, Kristin E. Lauter, and Vinod Vaikuntanathan. **Can homomorphic encryption be practical?** In Christian Cachin and Thomas Ristenpart, editors, Proceedings of the 3rd ACM Cloud Computing Security Workshop, CCSW, pages 113–124. ACM, 2011.
- R. Shokri, V. Shmatikov. **Privacy-Preserving Deep Learning**. CCS'15 Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pp. 1310-1321
- C. Gentry. **A fully homomorphic encryption scheme**. PhD Dissertation. Stanford University, 2009
- Dijk, Gentry, Halevi, Vaikuntanathan. **Fully homomorphic encryption over the integers**. EUROCRYPT'10 Proceedings of the 29th Annual international conference on Theory and Applications of Cryptographic Techniques. pages 24-43, 2010.
- Zhiqiang Yang et al. **Privacy-preserving classification of customer data without loss of accuracy**. In Hillol Kargupta et al., editors, Proceedings of the 2005 SIAM International Conference on Data Mining, SDM 2005, pages 92–102. SIAM, 2005.

27



Bob



Alice

Thanks for your attention!

28