Private and Secure

A short overview of confidential anomaly detection

Mihail-Iulian Pleșa

orange™

Who am I?

- Passionate about research and dissemination
- Interests: applied cryptography, privacy preserving machine learning
- Security Researcher at Orange Services

What You Will Know

1. We can
2. How to design and implement privacy preserving anomaly detection

Contents

Anomaly detection in cybersecurity

User behavioral analytics	Login times, number of trials
Network traffic analysis	Spikes, connection to known malicious domains
Intrusion detection systems	Sudden access to critical server
Endpoint security	Sensible files modifications, registry changes

Anomaly Detection

K-Means

- Clustering algorithm
- Classical machine learning
- Requires the knowledge of the number of clusters in advance
- Compute the distance to the nearest cluster and check if it is below a threshold

Autoencoders

- Compress and reconstruct data
- Neural networks
- It does not require any input other than the data itself
- Compute the reconstruction error and check if it is below a threshold

Philosophies

Homomorphic encryption (Input Privacy)

- Protect the confidentiality of the input data
- Perform computations directly over encrypted data
- Suitable for a any type of computation

Differential Privacy (Output Privacy)

- Protect the privacy of the individual
- Suitable for machine learning tasks

Libraries

OpenFHE

Homomorphic encryption

Supports multiple encryption schemes e.g. CKKS, BGV, BFV, TFHE, etc.

TensorFlow

Differential privacy

Yes, TensorFlow does support training with differential privacy

K-Means

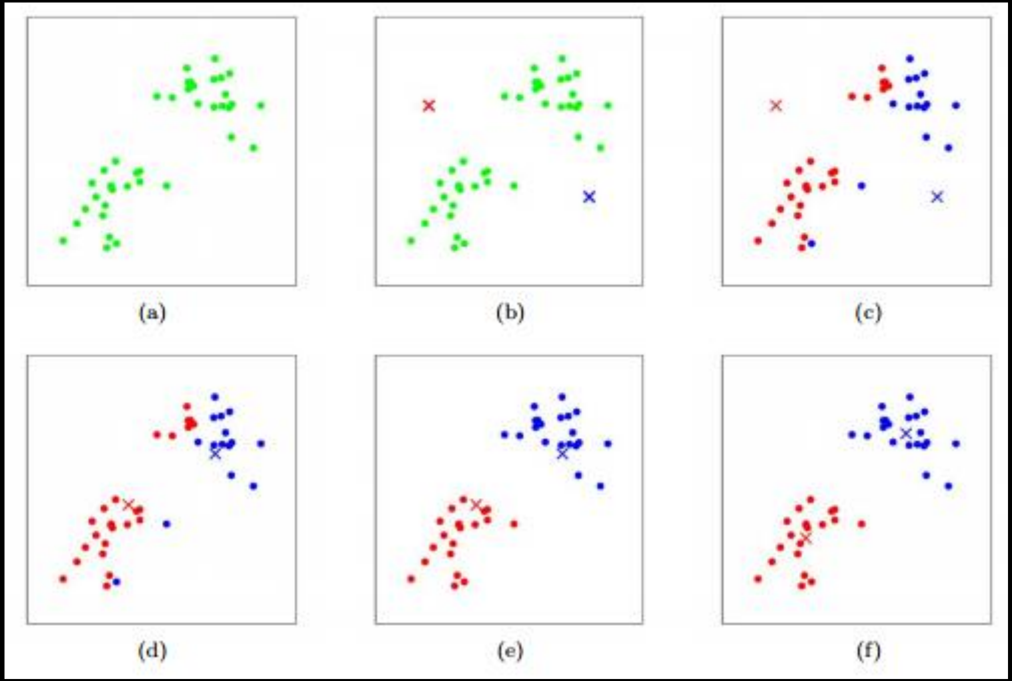

K-Means

1. Initialize **cluster centroids** $\mu_1, \mu_2, \ldots, \mu_k \in \mathbb{R}^n$ randomly.

2. Repeat until convergence: {

   For every i, set

   $$c^{(i)} := \arg\min_j ||x^{(i)} - \mu_j||^2.$$

   For each j, set

   $$\mu_j := \frac{\sum_{i=1}^m 1\{c^{(i)} = j\} x^{(i)}}{\sum_{i=1}^m 1\{c^{(i)} = j\}}.$$

   }

Fully Homomorphic Encryption

$$Enc_{PK}(m_1) \boxplus Enc_{PK}(m_2) = Enc_{PK}(m_1 + m_2)$$
$$Enc_{PK}(m_1) \odot Enc_{PK}(m_2) = Enc_{PK}(m_1 \cdot m_2)$$

Cheon-Kim-Kim-Song

C

K

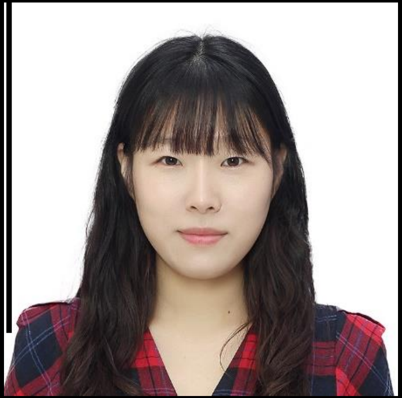

K

S

CKKS

- Leveled Homomorphic Encryption
- Approximate results
- Real numbers (suitable for ML tasks)

CKKS

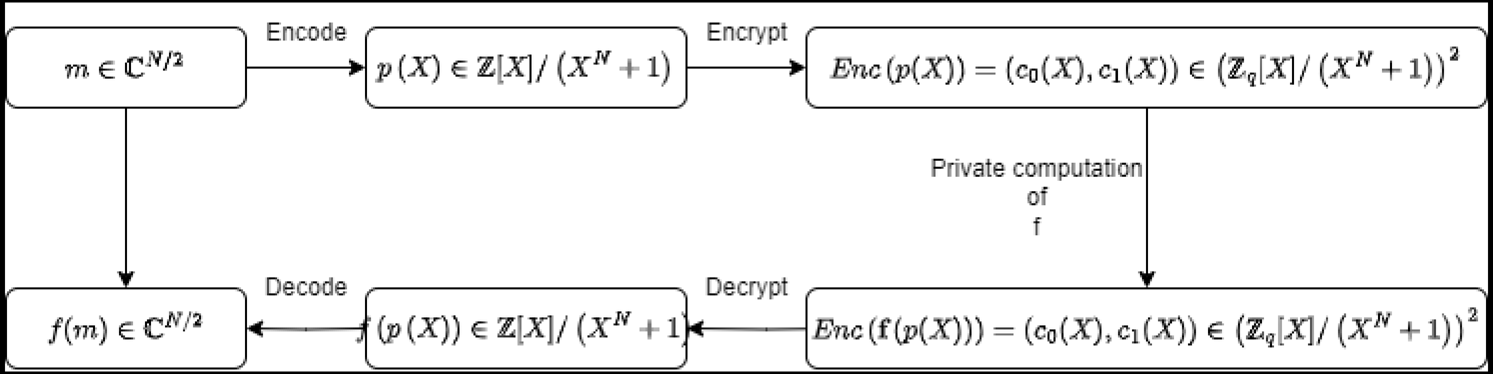
$m \in \mathbb{C}^{N/2}$
Encode
$p(X) \in \mathbb{Z}[X]/(X^N+1)$
Encrypt
$Enc(p(X)) = (c_0(X), c_1(X)) \in (\mathbb{Z}_q[X]/(X^N+1))^2$
Private computation
of
f
$Enc(f(p(X))) = (c_0(X), c_1(X)) \in (\mathbb{Z}_q[X]/(X^N+1))^2$
Decrypt
$f(p(X)) \in \mathbb{Z}[X]/(X^N+1)$
Decode
$f(m) \in \mathbb{C}^{N/2}$

Let’s practice

Autoencoders

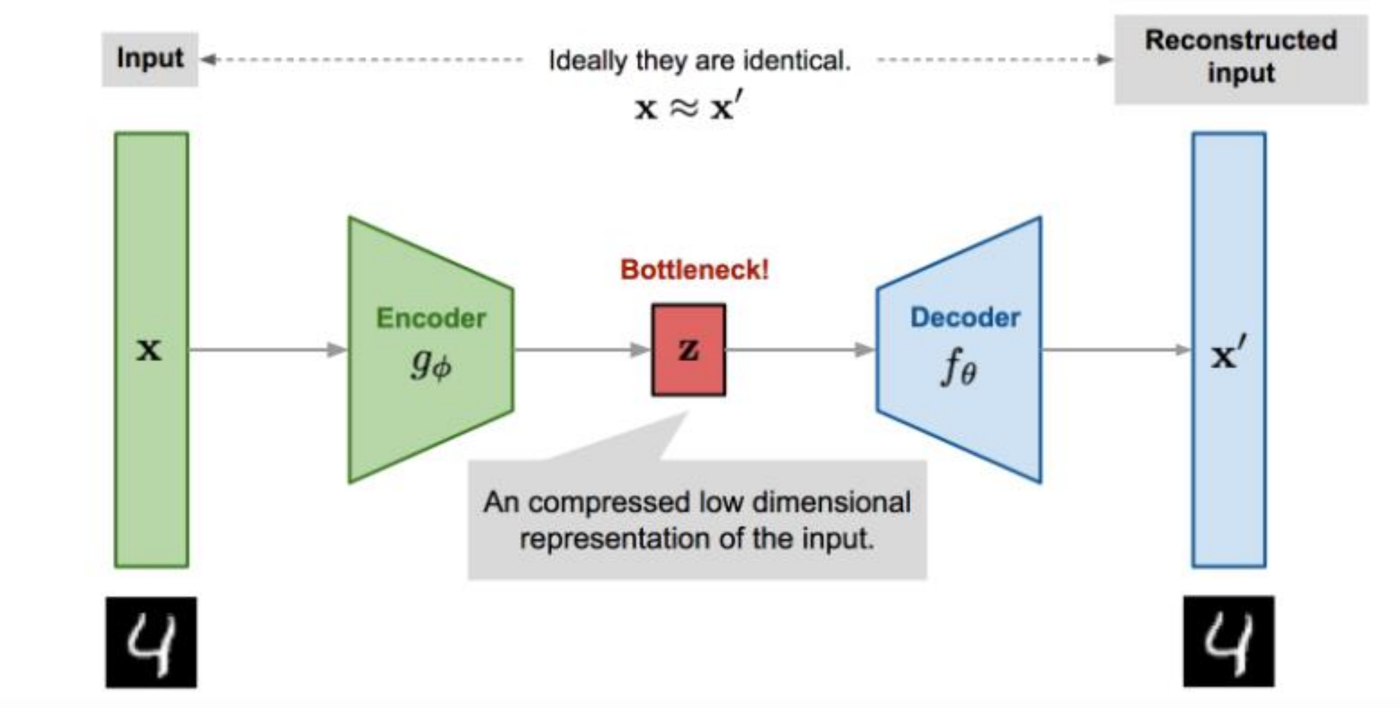

Autoencoders

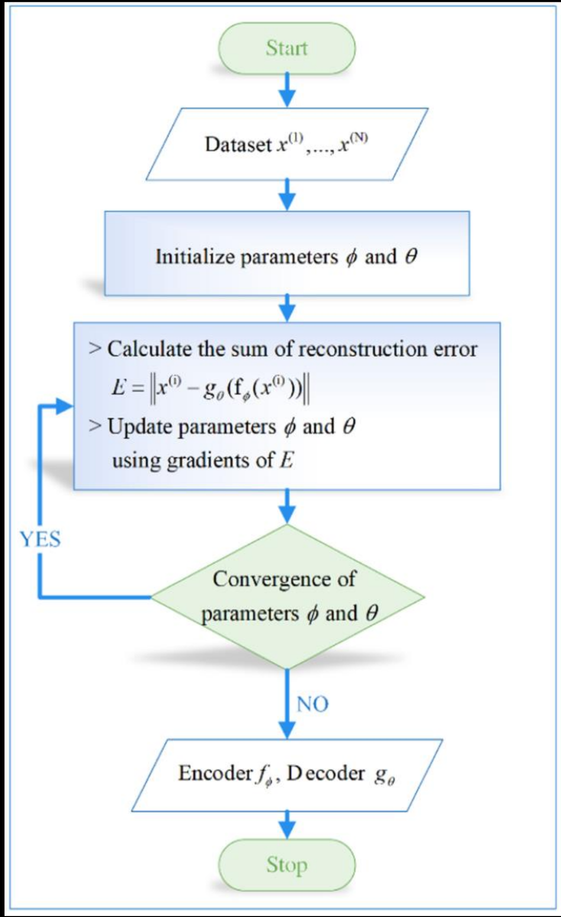

Autoencoders

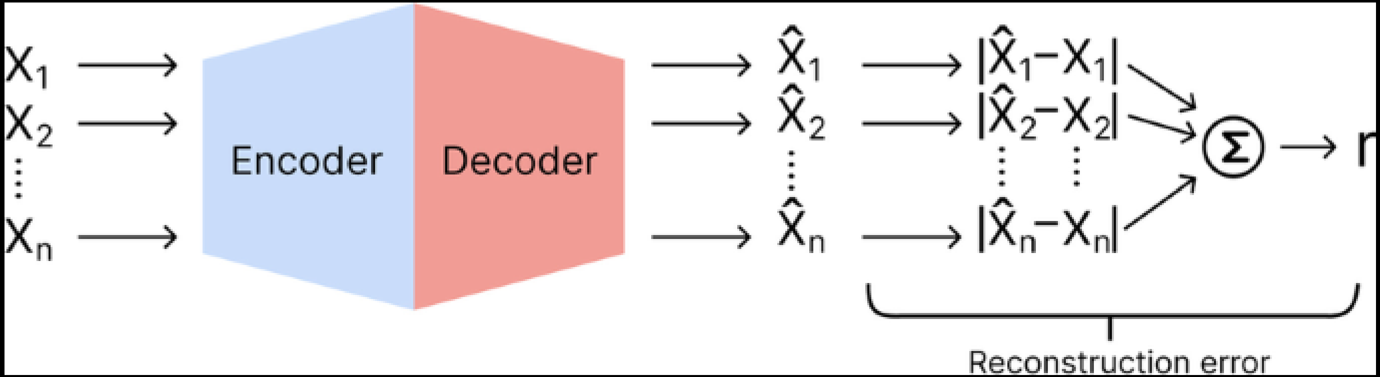

Differential Privacy

Differential Privacy

- Noise addition
- Privacy budget
- Trade-off between privacy and utility

Differential Privacy

A satisfies $\epsilon - DP$ if and only if

$$\Pr[A(D) \in T] \leq e^{\epsilon} \Pr[A(D') \in T] \,\forall\, T \subseteq range(A)$$

For any D and D' that differ on one element

NetFlix Cancels Recommendation Contest After Privacy Lawsuit

Netflix is canceling its second $1 million Netflix Prize to settle a legal challenge that it breached customer privacy as part of the first contest's race for a better movie-recommendation engine. Friday's announcement came five months after Netflix had announced a successor to its algorithm-improvement contest. The company at the time said it intended to [...]

Automatic Network Intrusion Detection

- Goal: detect buffer overflow attacks (KDD99 subset)
- Train with differential privacy an autoencoder on normal traffic
- Use the reconstruction error to detect anomalies

Let’s practice

Summary

- We can provide security and privacy
- Implementations with FHE are not simple translations (yet)
- Differential privacy is only a few hyperparameters away
- Privacy is a property of both input and output

It’s possible!

Security & Privacy

Thank you

orange™

References

- Cheon, J.H., Kim, A., Kim, M. and Song, Y., 2017. Homomorphic encryption for arithmetic of approximate numbers. In Advances in Cryptology–ASIACRYPT 2017: 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I 23 (pp. 409-437). Springer International Publishing.
- Li, N., Lyu, M., Su, D. and Yang, W., 2017. Differential privacy: From theory to practice. Morgan & Claypool.
- Torabi, H., Mirtaheri, S.L. and Greco, S., 2023. Practical autoencoder based anomaly detection by using vector reconstruction error. *Cybersecurity*, *6*(1), p.1.
- Piech, C. 2013. K Means. Stanford University. Available at: https://stanford.edu/~cpiech/cs221/handouts/kmeans.html .
- Li, N., Lyu, M., Su, D. and Yang, W., 2017. *Differential privacy: From theory to practice*. Morgan & Claypool.
- Al Badawi, A., Bates, J., Bergamaschi, F., Cousins, D.B., Erabelli, S., Genise, N., Halevi, S., Hunt, H., Kim, A., Lee, Y. and Liu, Z., 2022, November. Openfhe: Open-source fully homomorphic encryption library. In *Proceedings of the 10th Workshop on Encrypted Computing & Applied Homomorphic Cryptography* (pp. 53-63).
- Pang, B., Nijkamp, E. and Wu, Y.N., 2020. Deep learning with tensorflow: A review. *Journal of Educational and Behavioral Statistics*, *45*(2), pp.227-248.