Privacy by design

SHAPING ETHICAL AI USING FULLY HOMOMORPHIC ENCRYPTION

Oscar Licciardi

SFCON 2025

Table of Contents

Background

1. What is Machine Learning as a Service (MLaaS)?
2. An example of MLaaS: CHATGPT
3. Which is the cost?

What is MLaaS?

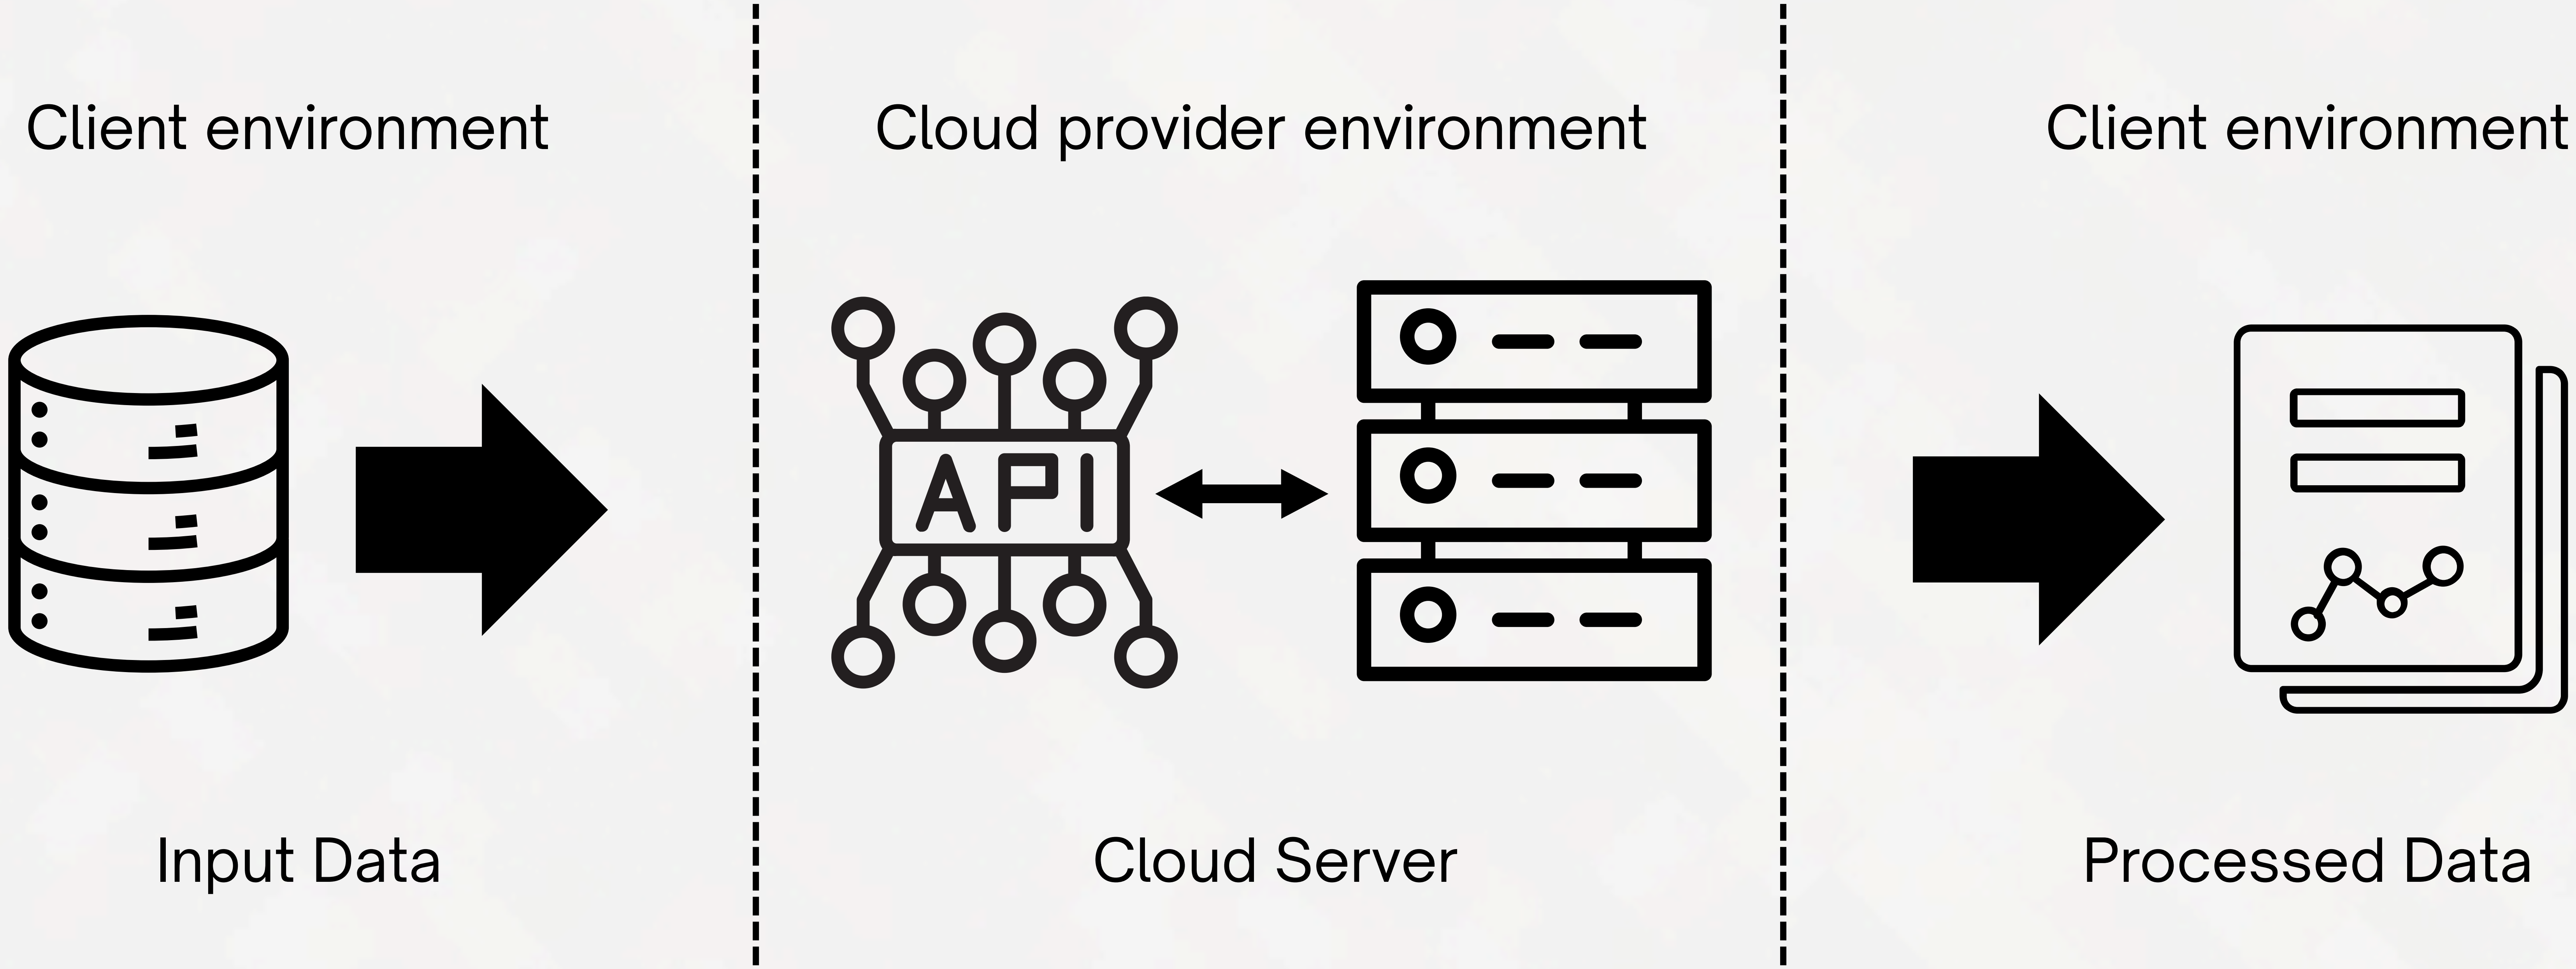

An example: CHATGPT

Common use cases

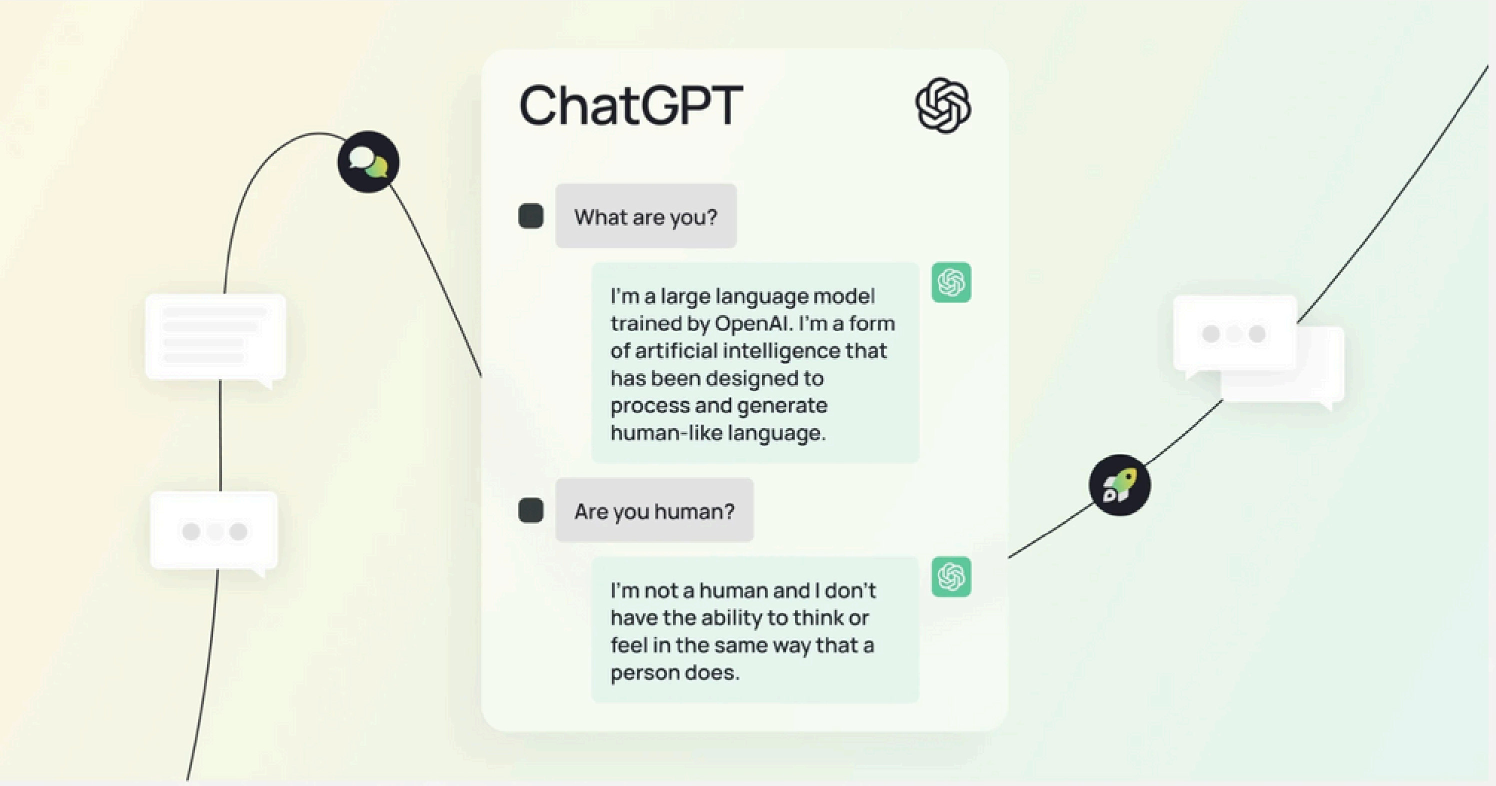

Writing

Content creation

Debugging code

Which is the cost?

1. Privacy Leakage due to **Public Data Exploitation**
2. Privacy Leakage due to **Personal Input Exploitation**
3. Privacy Leakage due to **Unauthorized Access**

Problems of MLaaS

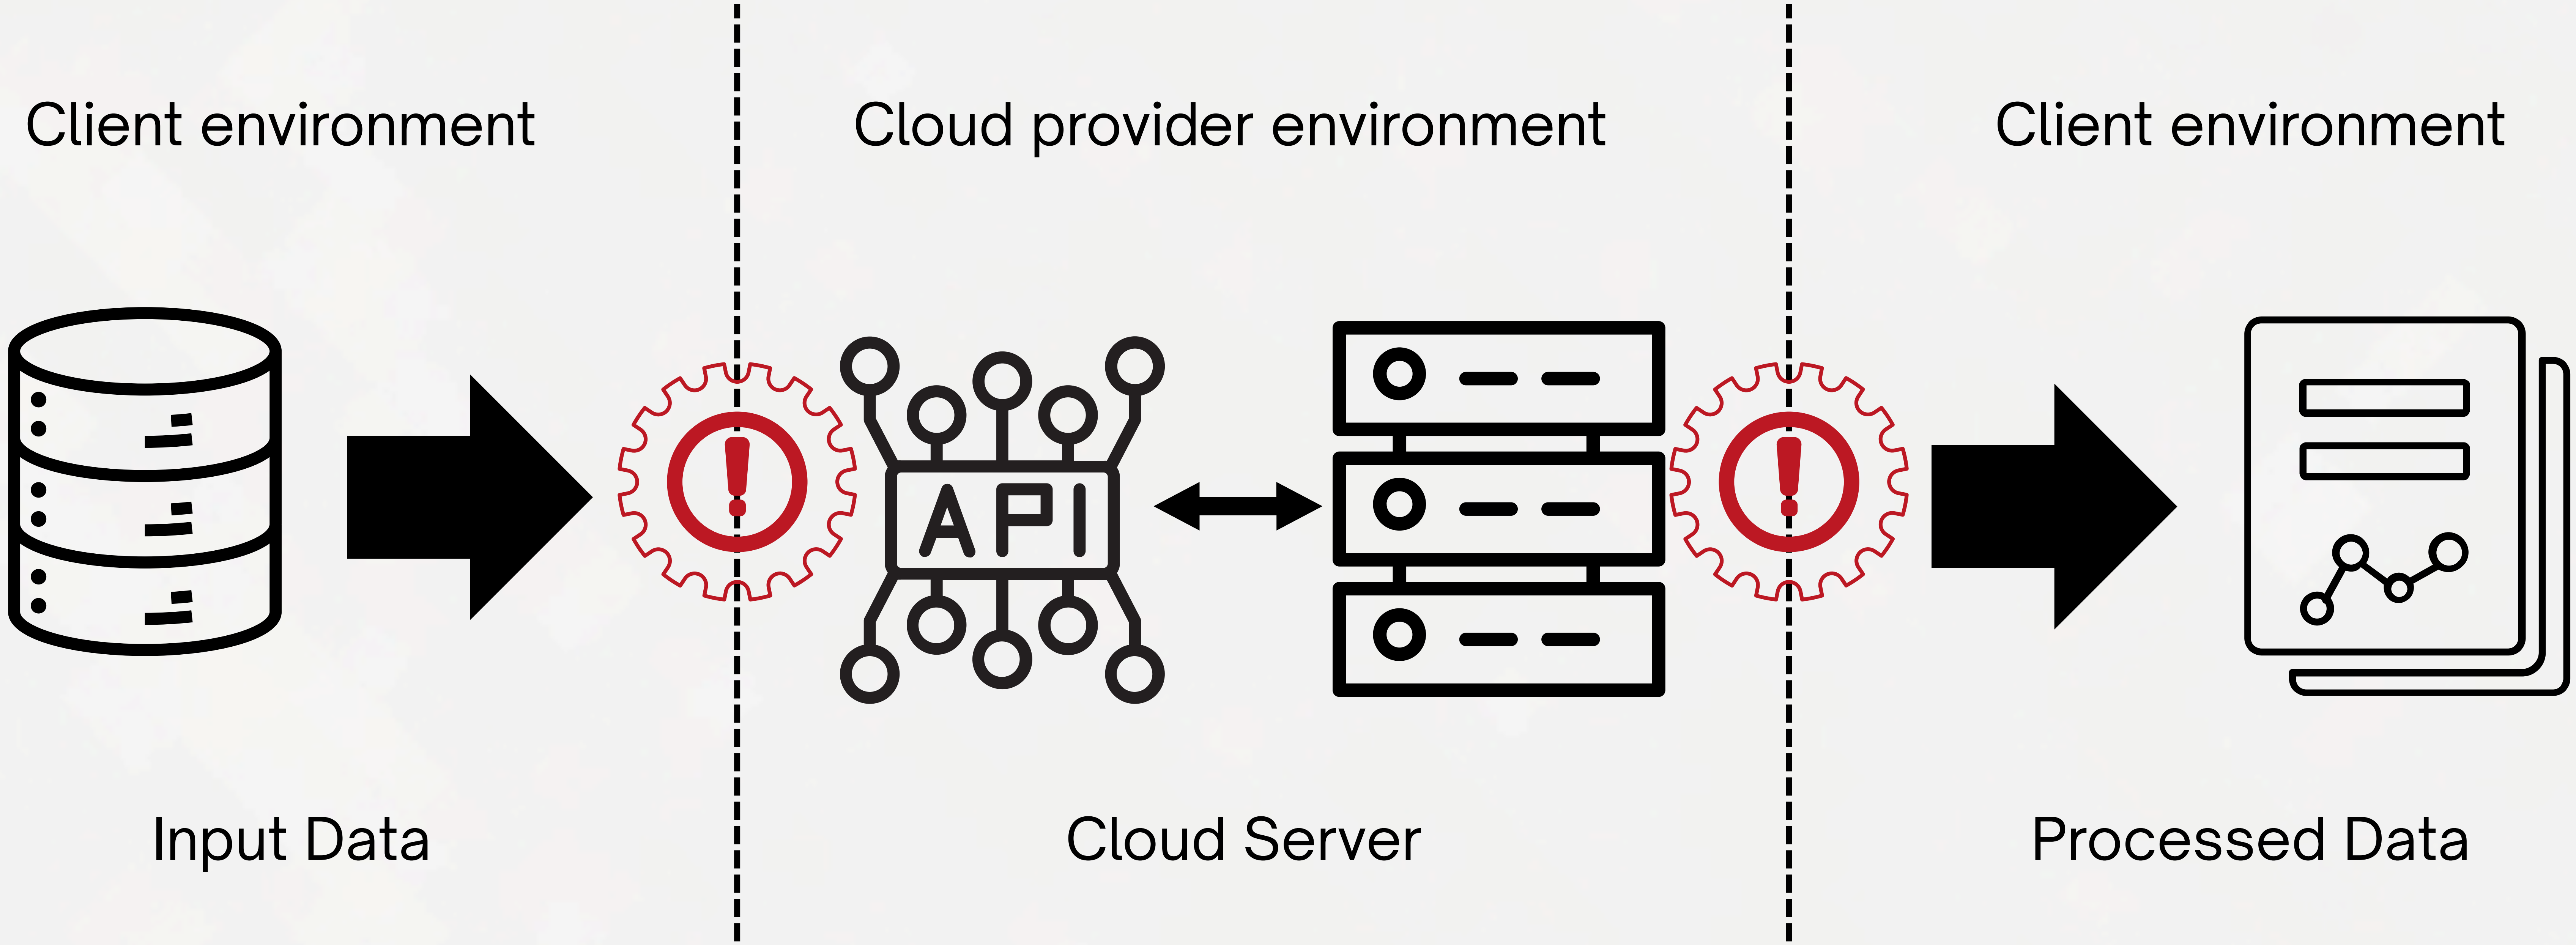

How can we do better?

Privacy and usability are inversely related as we strengthen one, we weaken the other.

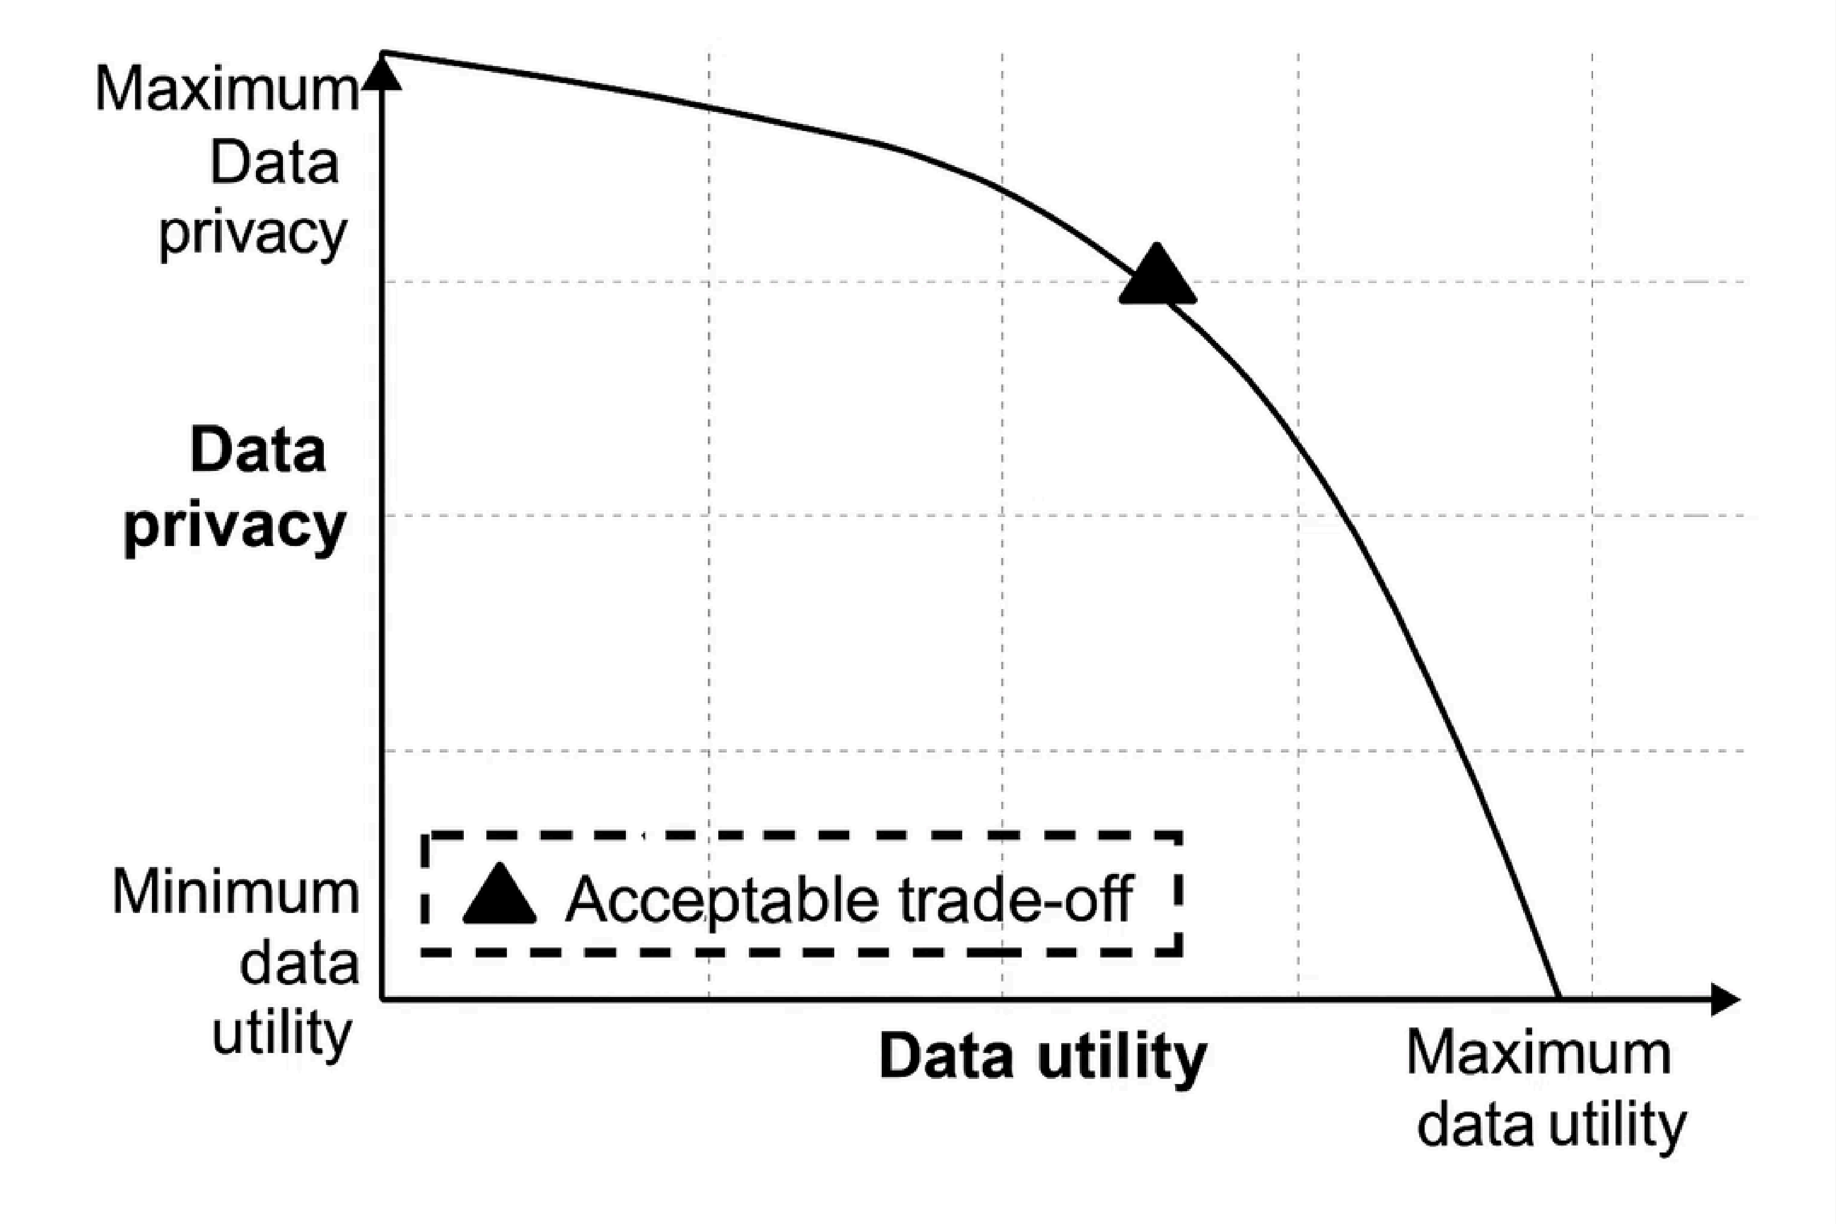

Goal: Increase privacy without crippling usability. Avoid access to the plaintext by using Homomorphic Encryption

How does FHE work?

- Inputs encrypted end-to-end MLaaS **never sees raw data**
- Inference done directly on ciphertexts
- Stored and logged data remain unreadable
- Balances high privacy with *growing efficiency*

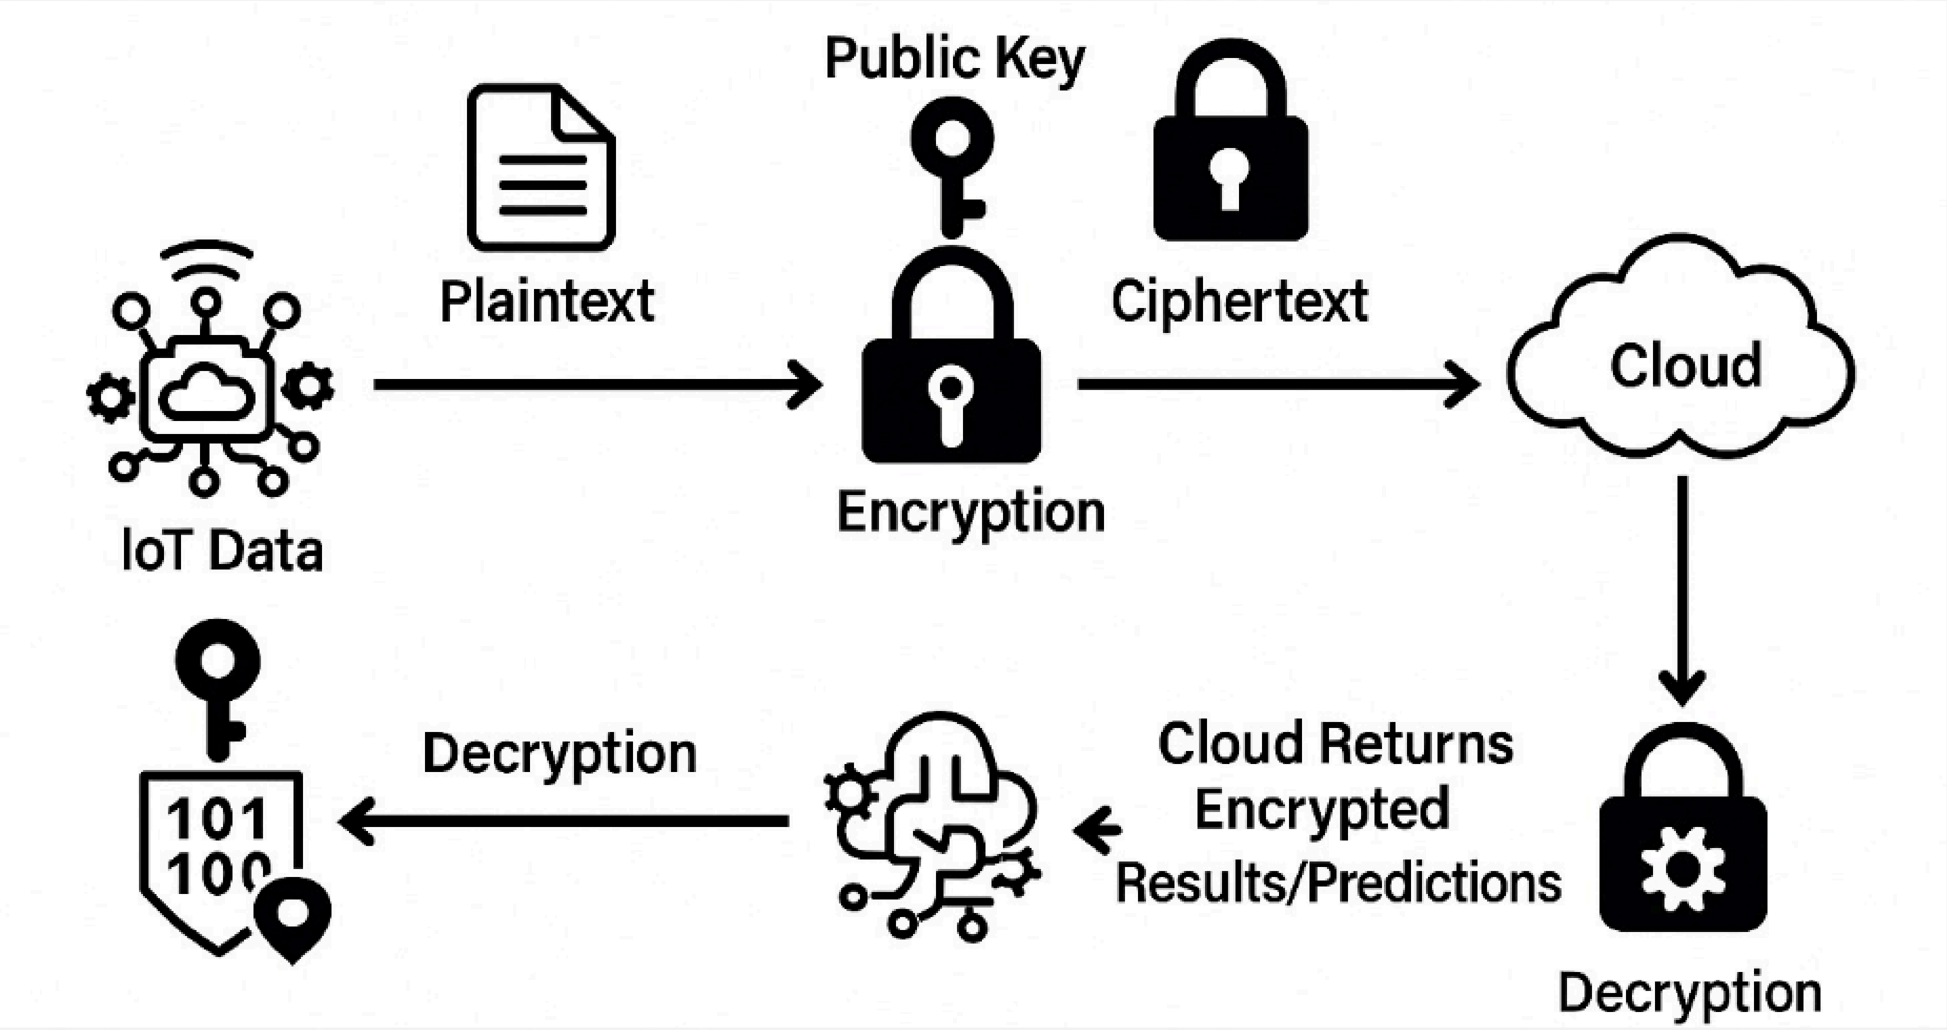

Analysis of applied FHE

PROS

- Complete protection against **Personal Input Exploitation**
- Complete protection against **Unauthorized access**

CONS

- Partial protection against **Public Data Exploitation**
- Introduction of prohibitive **overhead** in computation (x 10 - x 100) and in communication (x 10)

FHE-validator

- **Purpose of the framework:** Comparing standard scikit-learn implementations with Concrete-ML encrypted versions

- Measuring the computational cost and performance impact of using encryption on common ML models

README GPL-3.0 license

FHE Model Evaluator Library

A comprehensive Python library for evaluating and comparing Fully Homomorphic Encryption (FHE) models against traditional machine learning models. This library provides automated hyperparameter tuning, performance evaluation, and visualization capabilities for FHE implementations using concrete-ml.

Academic Context

This project is part of the **Undergraduate Research Opportunity Programme** at the **Polytechnic of Turin**, under the supervision of **Prof. Pelusi**. The research aims to contribute to the academic community by exploring the practical applications of Fully Homomorphic Encryption in machine learning and data analysis.

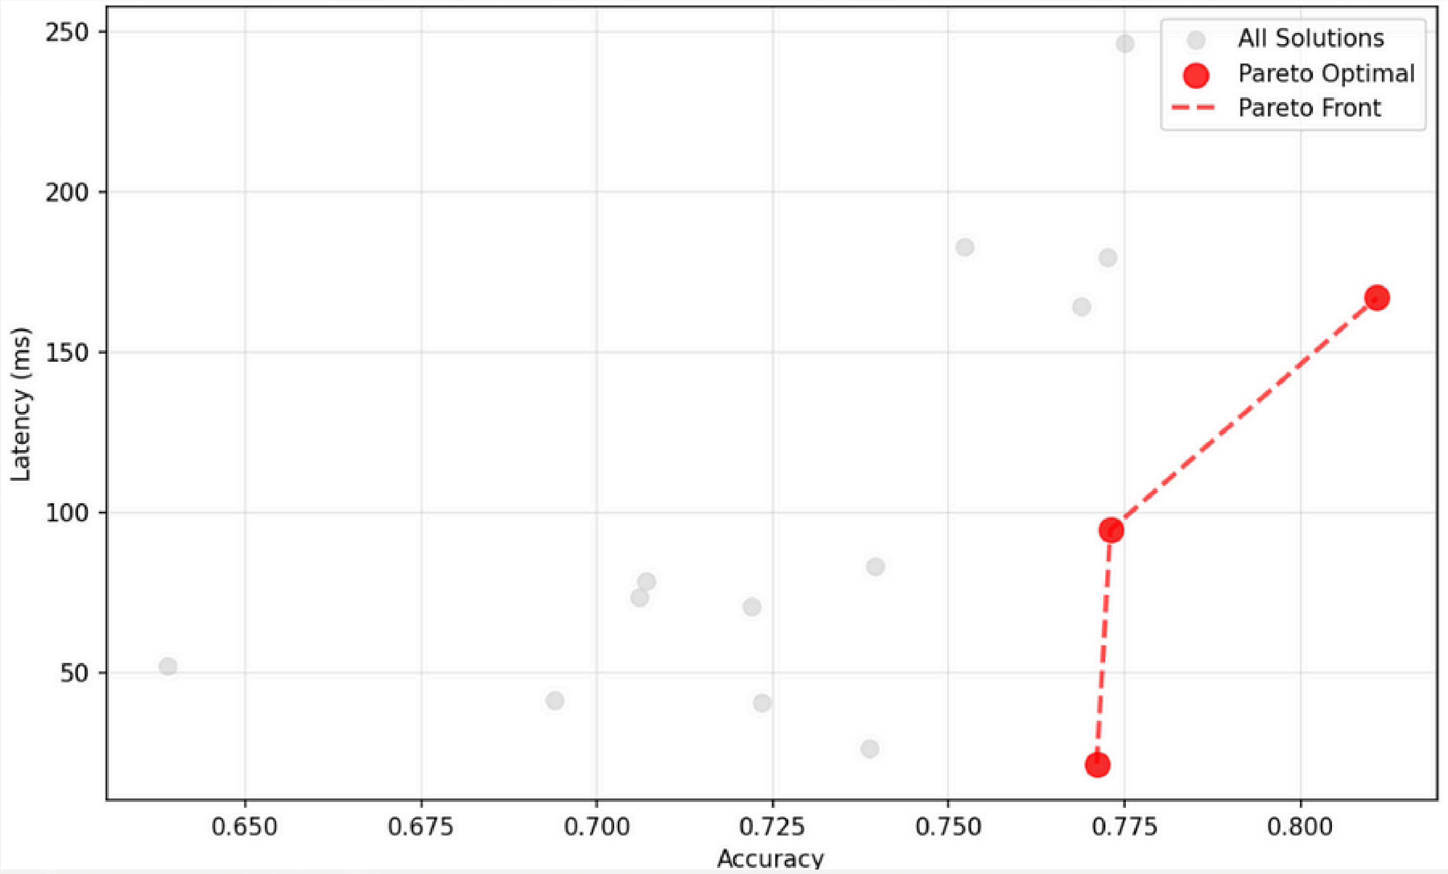

Conclusion

- **Takeaway:** Protecting privacy in ML is not only a technical challenge, it is a human responsibility

- *Collaboration* is key to a safer, more trustworthy internet

Thank you!

DO YOU HAVE ANY QUESTION?

Oscar Licciardi

SFCON 2025

References

[1] A. K. R. M. Al-Hussein, “Attack framework – data collection, processing, and sampling,” ResearchGate, 2023. [Online]. Available: https://www.researchgate.net/figure/Attack-framework-1-The-data-we-collected-will-be-processed-and-some-samples-will-be_fig4_367969786.
[2] J. Doe, “The trade-off between data privacy and data utility,” ResearchGate, 2022. [Online]. Available: https://www.researchgate.net/figure/The-trade-off-between-data-privacy-and-data-utility_fig1_363889143.
[3] Zama AI, “Concrete-ML repository,” GitHub, 2024. [Online]. Available: https://github.com/zama-ai/concrete-ml.
[4] M. Smith et al., “Costs of encrypted computation: Why fully homomorphic computations are slow,” ResearchGate, 2022. [Online]. Available: https://www.researchgate.net/publication/354842186_Costs_of_Encrypted_Computation_Why_Fully_homomorphic_computations_are_Slow.
[5] P. Kumar (Ed.), Privacy-Preserving Machine Learning, Springer, 2023, ISBN 978-981-13-6393-1. [Online]. Available: https://link.springer.com/book/10.1007/978-981-13-6393-1.