Privacy-Centric Digital Surveillance through Homomorphic Encryption and Deep Learning

Johannes Unruh, Dorian Przetakiewicz, Oscar H. Ramírez-Agudelo, Michael Karl

Deutsches Zentrum für Luft- und Raumfahrt e.V.

Introduction

The proliferation of digital surveillance technologies has intensified the need for robust data privacy safeguards—particularly when processing sensitive visual or biometric information. In response, we introduce **LYNX**[1], an open-source web platform designed to enable secure, privacy-preserving computation through **homomorphic encryption (HE)**. LYNX enables encrypted data to be processed directly—eliminating the need for decryption at any stage of computation. At the core of this work is our **Secure Inference Layer**, a modular component that supports the privacy-preserving deployment of deep learning models. This module supports **ONNX**-based neural networks and automatically translates their inference operations into HE-compatible form using the **TenSEAL** library. Our architecture facilitates modular and scalable encrypted inference pipelines, enabling practical privacy-centric applications in domains such as facial recognition, object detection, and behavioral analytics—all without exposing raw data. This poster presents the technical foundation and capabilities of the LYNX platform and demonstrates its potential to make secure machine learning both accessible and impactful.

Platform Architecture

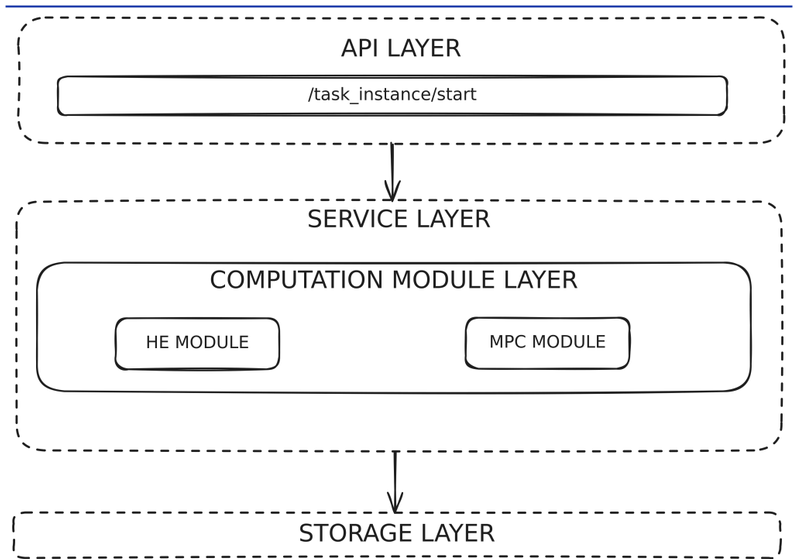

Architecture of LYNX.

Methods

We designed a modular web-based framework that enables privacy-preserving inference via homomorphic encryption:

- **Model Import:** ONNX models are loaded into the system.
- **HE Translation:** Inference layers are transformed into TenSEAL-compatible operations; extendable to high-resolution image inputs via ciphertext sharding.
- **Encrypted Execution:** Inference is performed directly on encrypted input.
- **Client-Side Interface:** A browser-based GUI facilitates user interaction.

All computations remain encrypted throughout; no server-side decryption is required.

HE Inference Pipeline

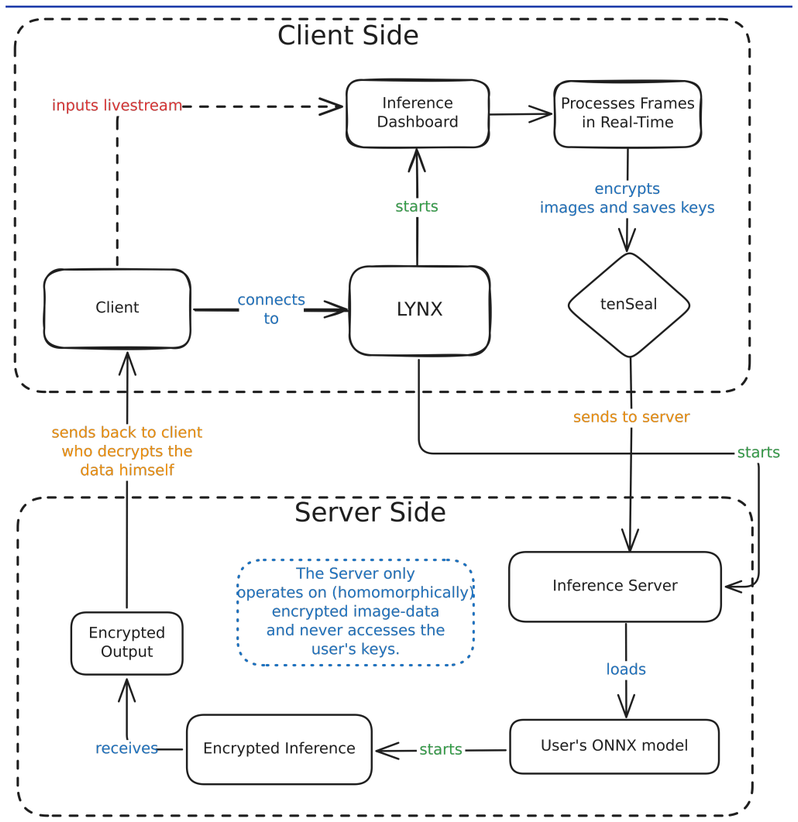

Pipeline for real-time homomorphically encrypted inference of frames taken from a livestream.

Results

We validated our HE inference module on benchmark image classification tasks using encrypted inputs:

- **Model:** Pre-trained on a human detection dataset[2].
- **Inference:** Inputs are fully encrypted on the client side; results are decrypted by the client after computation.
- **Accuracy:** Comparable to plaintext inference using the same model.
- **Latency:** Real-time capable for lightweight models.

These results demonstrate the feasibility of privacy-preserving deep learning in surveillance scenarios.

Conclusions

- We present LYNX, a web-based platform enabling homomorphic encrypted inference.
- The ONNX-to-TenSEAL pipeline makes deep learning models compatible with secure computation.
- Our results validate the practical feasibility of encrypted inference, with performance trade-offs acceptable for small to medium models.

References

- Benaissa et al., *TenSEAL: A Library for Encrypted Tensor Operations Using Homomorphic Encryption*
- ONNX Community, *Open Neural Network Exchange*
- Maloney, V., Obrecht, R. F., Saraph, V., Rama, P., Tallaksen, K. (2023). *High-Resolution Convolutional Neural Networks on Homomorphically Encrypted Data via Sharding Ciphertexts.* arXiv:2306.09189.

1 Layered privacY eNhancing eXchange

2 https://www.kaggle.com/datasets/constantinwerner/human-detection-dataset

This work is part of the LYNX project at DLR Institute for AI Safety and Security.