TENSORLOCK

PATENTED ENCRYPTION FOR MACHINE LEARNING

Train AI directly on encrypted data

Tensorlock allows neural networks to learn directly on encrypted datasets, with no decryption step.

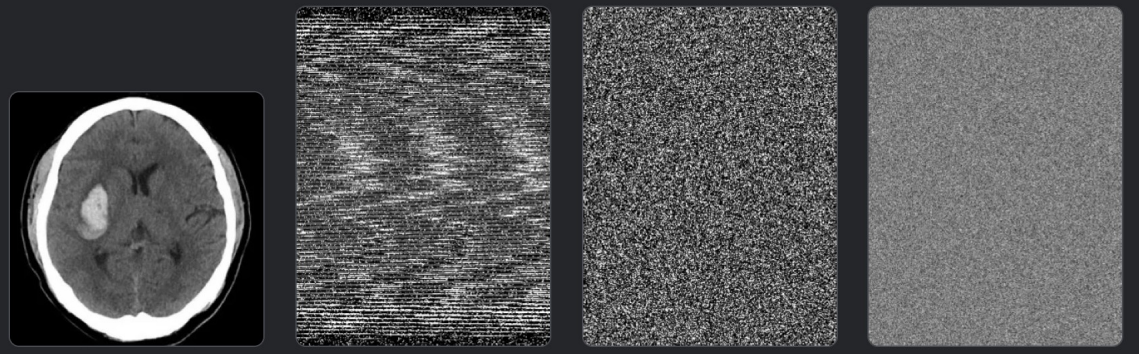

Original **Padded** **Shuffled** **Transformed**

Encryption of a CT image by interlaced padding, random index shuffle, and random orthogonal matrix transformation.

tensorlock.ai

info@tensorlock.ai

THE PROBLEM

Most high-value data never reaches model training

Data owners need to protect privacy, commercial value, and intellectual property.

97%

of the data hospitals produce goes unused

World Economic Forum, 2019

~70%

of health system leaders call privacy and security a major barrier to AI adoption

Sage Growth Partners, 2025

22%

of banking leaders name data privacy the main obstacle to generative AI progress

EY-Parthenon, 2025

62%

of government leaders say privacy and security concerns hold back digital tools

EY, 2025

THE SOLUTION

Neural networks train and run inferenece on encrypted data

No decryption step

No example is decrypted before any machine-learning operation.

Standard architectures

Networks operate on the encrypted vectors with little modification.

One key, end to end

The same transformation covers training, test, and inference data.

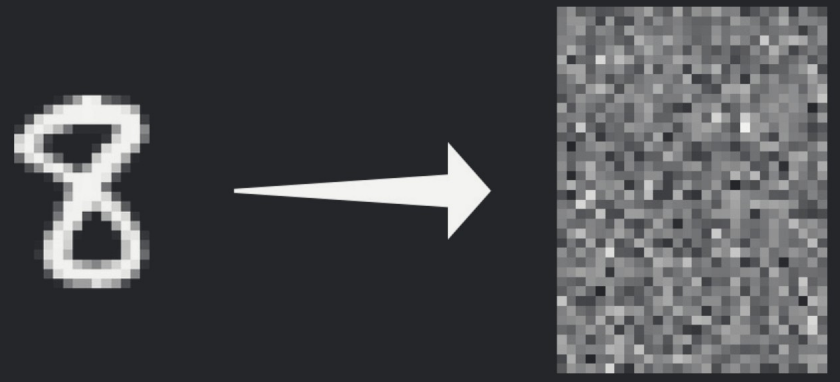

Plaintext

Encrypted, as trained on

HOW IT WORKS

Four steps encrypt each vector

01

Padding

Adds dimensions to each vector and conceals its original size.

02

Perturbation

Alters values within set bounds and keeps the data useful for learning.

03

Fixed index shuffle

Applies one secret permutation of coordinates to every example.

04

Orthogonal transformation

Rotates each vector into a secret basis with a random orthogonal matrix.

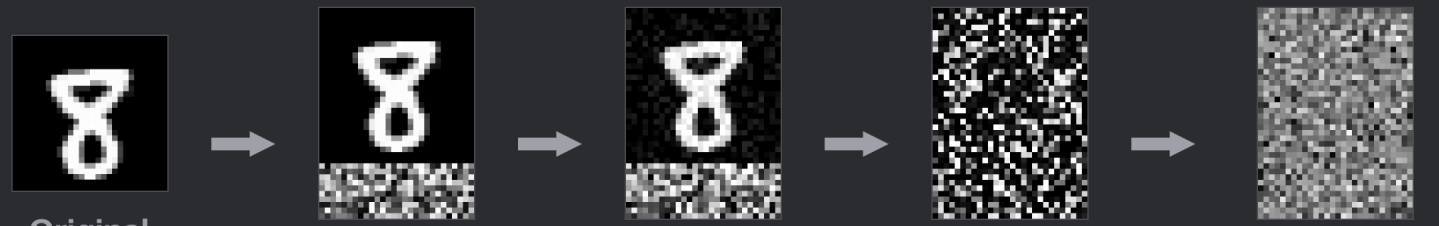

Distances and inner products survive. The original coordinates do not.

RESULTS

Encrypted training is very close to plaintext training

Dataset	Accuracy difference
MNIST	0.0012
Fashion-MNIST	0.0032
CIFAR-10	0.0227
EMNIST Digits	0.0002
EMNIST Letters	0.0028

Absolute difference in validation accuracy, encrypted vs plaintext.

+1.62%

training time with ciphertext compared to plaintext

Mixing hides plaintext coordinates

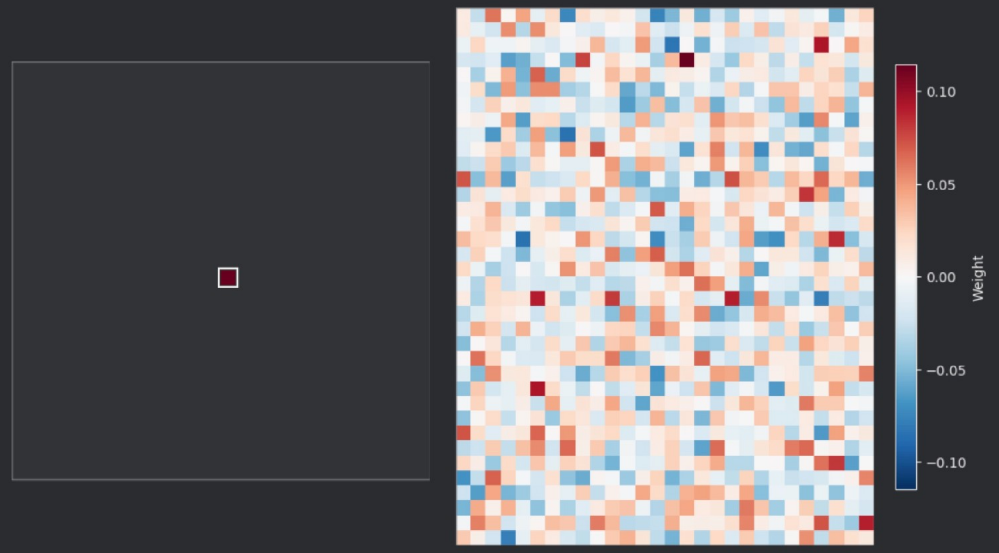

Each individual plaintext pixel [here, (14,14)] contributes a weight to the value in every ciphertext dimension.

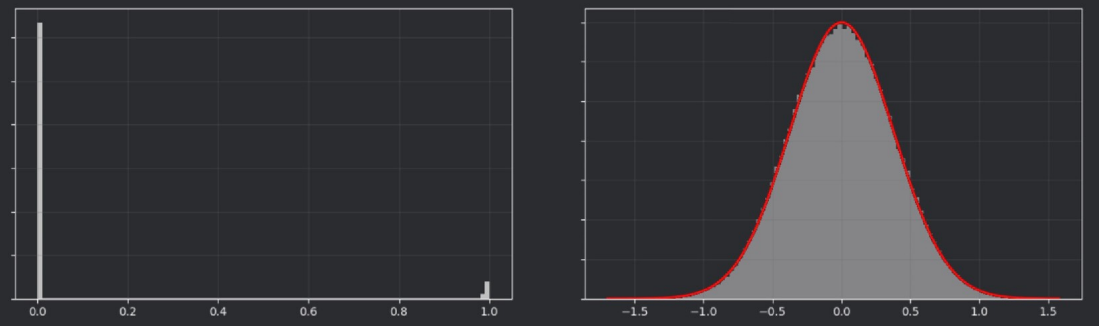

Plaintext pixels cluster at 0 and 1 for handwritten digits. Ciphertext values follow a Gaussian distribution.

The plaintext and the key stay with the data owner

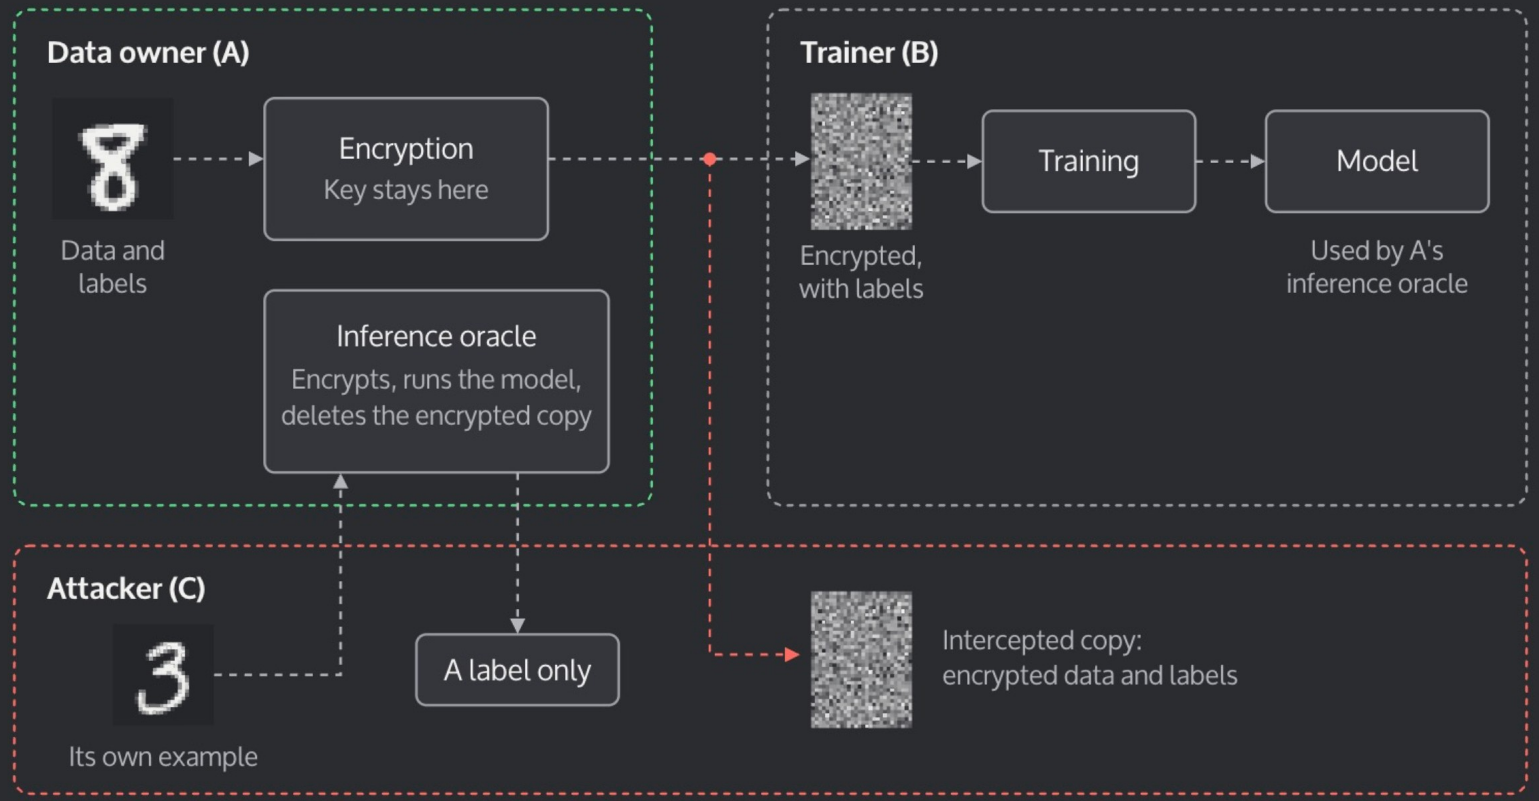

Training

The attacker obtains a copy of the encrypted data and labels.

Inference

The attacker can query the inference oracle and receive a label only.

Never shared:

The plaintext, the key, and the encrypted form of the attacker's own example.

SECURITY

Modeled attacks recover examples only with exact or near-exact reference data (outside the threat model)

Original

Attacker's reconstruction



0.241

attacker's mean absolute error with a near-exact reference (vs 0.173 for a blank guess)

$\approx 2^{36\text{ million}}$

number of orthogonal matrices at double precision that are consistent with the encrypted geometry

$\approx 10^{354}$

possible plaintext coordinate sets for one MNIST-sized example encrypted to 1,176 values

The data owner controls the data and the inference oracle

Standard networks

Train directly on encrypted data, with minimal changes to architecture.

Plaintext-level performance

Accuracy and training time stay close to those of plaintext training.

Configurable security

The owner sets the extent and type of padding and perturbation.

Owner-controlled inference

The trained network accepts only inputs encrypted with the owner's key.

Tensorlock differs in these respects from federated learning, homomorphic encryption, trusted execution environments, and differential privacy.

APPLICATIONS

Built for data too sensitive to share in plaintext

Healthcare

Medical images and patient records.

- Flag scans for review
- Predict outcomes and patient risk

Finance and business

Transactions, customer data, and business intelligence.

- Detect fraud
- Assess credit risk
- Forecast demand

Frontier models

Private datasets for training large models.

- License datasets without releasing plaintext
- Train on encrypted embeddings

Government

Personal records, imagery, and other sensitive data.

- Classify satellite and aerial imagery
- Work with contractors who see only ciphertext

Data types: images and video, tables and records, time series, and text embeddings.

One core method extends across network types

Network type	Summary
Fully connected	Manuscript results on five image datasets
Convolutional	Fixed convolution before encryption, or a key-holding service for trainable kernels
Language models	Classification of encrypted documents; training on encrypted token embeddings
Generative models	Adversarial and diffusion models trained on encrypted latent representations
Unified networks	One network for plaintext data and data encrypted under one or more keys

INTELLECTUAL PROPERTY

Granted US patents protect the method

I.P.	Scope	Status
US Patent 12,731,025 B2	Machine learning applications	Granted
US Patent 12,190,240 B2	Search and data storage applications	Granted
US Application 17/652,743	Additional machine learning applications	Pending
US Trademark 99931353	TENSORLOCK	Pending



info@tensorlock.ai

tensorlock.ai