RUB

RUHR-UNIVERSITÄT BOCHUM

Fast Homomorphic Evaluation of Deep Neural Networks

FHE–DiNN — Privacy-Preserving Image Classification in the Cloud

Chair for Cryptology and IT Security @ Ruhr-Uni Bochum, 7. January 2019

Matthias Minihold
ECRYPT-NET Early Stage Researcher

hgi
Horst Görtz Institut
für IT-Sicherheit

1 Quantum Computing Threatens IT Infrastructure

2 Privacy-Preserving Predictions in the Cloud
- Law Perspective
- Technical Perspective
- Machine Learning as a Service (MLaaS)
- Recapitulation: Homomorphisms and FHE
- Machine Learning & Neural Network Basics
- FHE-friendly Discretized Neural Networks (DiNNs)

3 Experiments: Digit Classification with FHE–DiNN
- MNIST Digit Recognition & Classification

Impact of Quantum Computing on IT Security—Overview

Goals of Cryptography within IT Security

- Confidentiality (A speaks in private with B)
- Authenticity (A knows it is B she speaks with)
- Integrity (A can verify that B signed data)
- Non-repudiation (A cannot undo signature on previously signed data)

Effects of Grover's and Shor's quantum algorithms in cryptanalysis

- Symmetric Ciphers (AES, ...): security level halved by Grover's algorithm;
  $\exists c \in \mathbb{R}\ \forall n \in \mathbb{N} : \mathcal{O}\left(c^n\right) \xrightarrow{\text{Grover}} \mathcal{O}\left(c^{\frac{n}{2}}\right) = \mathcal{O}\left(\sqrt{c^n}\right)$,
- Encryption (RSA, ECC) and signatures (RSA, (EC)DSA): broken by Shor's algorithm;
  $\exists c \in \mathbb{R}\ \forall n \in \mathbb{N} : \mathcal{O}\left(c^n\right) \xrightarrow{\text{Shor}} \mathcal{O}\left(n^c\right)$.

Implementation and integration issues lead to delayed migration to post-quantum crypto.

Computing on Encrypted Data Practice—Law Perspective

≈ 50 Years Data Protection Regulations: Timeline for the EU

1970 *Hessian Data Protection Regulation* privacy law (Hesse),
1986 Overhauled 2^{nd} version for public authorities (in Germany),
1995 Adapt & blue-print natural person's EU Data Protection Directive,
2016 Superseded by EU's General Data Protection Regulation (GDPR),
2018 GDPR is enforceable since May 2018 granting basic protection,
2020 GDPR is prominently covered in the media (known as DSGVO in Germany).

Any *'free'* Cloud-service means *user data* is the *product*.

Computing on Encrypted Data Theory—Theoretical Perspective

Let $n \in \mathbb{N}$ denote the security parameter. Typically $>$ 80 bit post-quantum security level.

(Public-Key) Encryption Scheme S

Given an encryption (resp. decryption) function $\mathsf{Enc}_{\mathsf{pk}} : \mathcal{M} \to \mathcal{C}$ (resp. $\mathsf{Dec}_{\mathsf{sk}} : \mathcal{C} \to \mathcal{M}$) with secret-key–public-key pair $(\mathsf{sk}, \mathsf{pk}) \xleftarrow{\$} \mathsf{Gen}(1^n)$; we call it private-key, if $\mathsf{sk} = \mathsf{pk}$, and require all algorithms to be efficiently computable ($\textsc{ppt}$).

For all plaintexts $m \in \mathcal{M}$, and all key-pairs $(\mathsf{sk}, \mathsf{pk}) \in \mathcal{K}$ we have

$\Pr[\mathsf{Dec}_{\mathsf{sk}}(\mathsf{Enc}_{\mathsf{pk}}(m)) = m] = 1 - \mathrm{negl}(n)$, holds with overwhelming probability ('w.o.p.').

Evaluating a Function f on Encrypted Data

Let $S = (\mathsf{Gen}(1^n), \mathsf{Enc}(.), \mathsf{Dec}(.))$ be a (public-key) encryption scheme:

$\mathsf{Eval}(f, \mathsf{Enc}_{\mathsf{pk}}(m)) = c \in \mathcal{C}$, such that w.o.p. $\mathsf{Dec}_{\mathsf{sk}}(c) = f(m)$ holds.

Machine Learning as a Service (MLaaS)

User submits $\mathsf{Enc}(x)$ and recovers $\mathsf{Enc}(\mathcal{M}(x))$; the encrypted prediction.

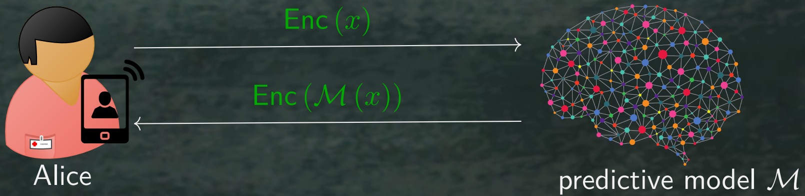

- ✓ Privacy — input & output data is encrypted (user has only key)
- ◆ Efficiency — is a central practical issue

Goal of PhD-Thesis: FHE–DiNN — fast homomorphic evaluation of neural networks ✓

Recapitulation: Homomorphisms and Fully Homomorphic Encryption (FHE)

Remarkably, FHE can evaluate any function f on encrypted inputs c.

FHE means "$\forall f:\ f \circ \mathsf{FHE.Enc}_{\mathsf{pk}} \mathrel{\widehat{=}} \mathsf{FHE.Enc}_{\mathsf{pk}} \circ f$"

Let $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Dec}, \mathsf{FHE.Eval})$ be an (IND-CPA–secure public-key) encryption scheme with compact ciphertexts $\mathcal{C}$.
If for any computable function $f \in \mathcal{F}$ and all plaintexts $m_1, m_2 \in \mathcal{M}$,

$$\begin{aligned}(f \circ \mathsf{FHE.Enc}_{\mathsf{pk}})(m_1, m_2) &= \overbrace{f([m_1]_{\mathsf{pk}}, [m_2]_{\mathsf{pk}})}^{f(c_1,c_2)=c} \overset{!}{=} \overbrace{[f(m_1, m_2)]_{\mathsf{pk}}}^{c' \in \mathcal{C}} \\ &= (\mathsf{FHE.Enc}_{\mathsf{pk}} \circ f)(m_1, m_2),\end{aligned}$$

holds with $f(m_1, m_2) = m_3 \in \mathcal{M} \subseteq \mathcal{C}$, then it is an FHE scheme.

Actually, w.o.p. $\mathsf{FHE.Dec}_{\mathsf{sk}}(c) = \mathsf{FHE.Dec}_{\mathsf{sk}}(c') \in \mathcal{M}$ must match!

FHE — 'The Holy Grail of Cryptography' [Mic10]

≈ 40 Years of FHE: Timeline

1978 Adleman, Dertouzos, and Rivest mention private homomorphisms

2009 Gentry's *theoretical breakthrough* construction: 1^{st} generation

2012 Brakerski, Gentry, and Vaikuntanathan (BGV)'s *simpler* 2^{nd} gen.

2013 Gentry, Sahai, and Waters (GSW)'s *efficient*: 3^{rd} generation

2016 Chillotti, Gama, Georgieva, and Izabachène (CGGI)'s *efficient implementation*: (TFHE)

2020 FHE schemes & applications' *practical breakthrough*?

Deep Feed-Forward Neural Network with $n_{\mathcal{I}} : n_1 : \ldots : n_d : n_{\mathcal{O}}$–topology

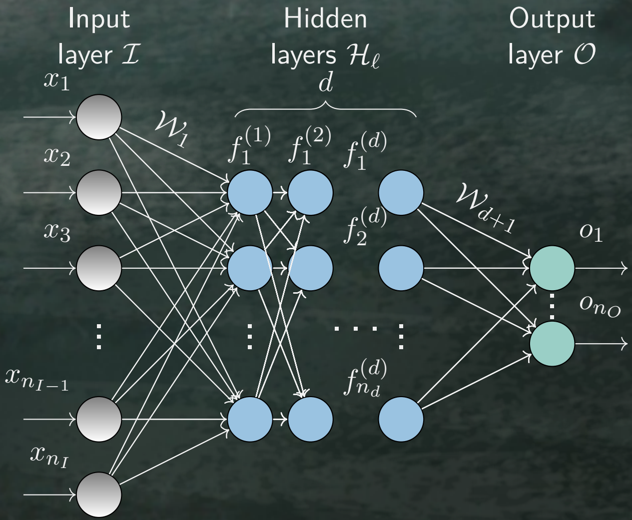

Close-up on Neuron

Computation for every neuron:

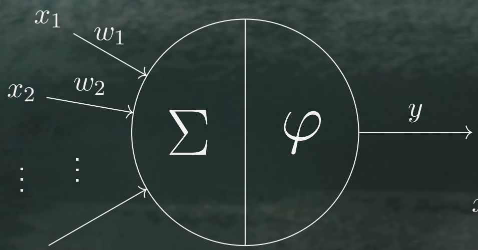

$x_i, w_i, y \in \mathbb{R}$

$$y = \varphi\left(\sum_i w_i\, x_i\right),$$

where φ is an *activation function*.

FHE-friendly Discretized Neural Networks

Goal: *FHE-friendly* model of neural network: $x_i, w_i, y \in \mathbb{Z}$.

Definition (DiNN)

A neural network whose layers have inputs in $\{-I, \ldots, I\} \subseteq \mathbb{Z}$, weights in $\{-W, \ldots, W\} \subseteq \mathbb{Z}$, for $I, W, O \in \mathbb{N}$, and each neuron's activation function maps the weighted sum to integer values in $\{-O, \ldots, O\} \subseteq \mathbb{Z}$.

1. Not restrictive as it seems as, e.g., binarized NNs perform well;
2. trade-off between size and performance;
3. conversion is straight-forward.

Main impediment: non-linear functions

Applying the non-linear activation function after linear layer.

Digit Recognition & Classification in the Cloud

We showcase a solution to the problem of *digit recognition*.

Digit Recognition & Classification in the Cloud

We showcase a solution to the problem of **blind** *digit recognition*.

Dataset: MNIST (60 000 images in training set + 10 000 in test set).

FHE–DiNN: Overview [BMMP18]

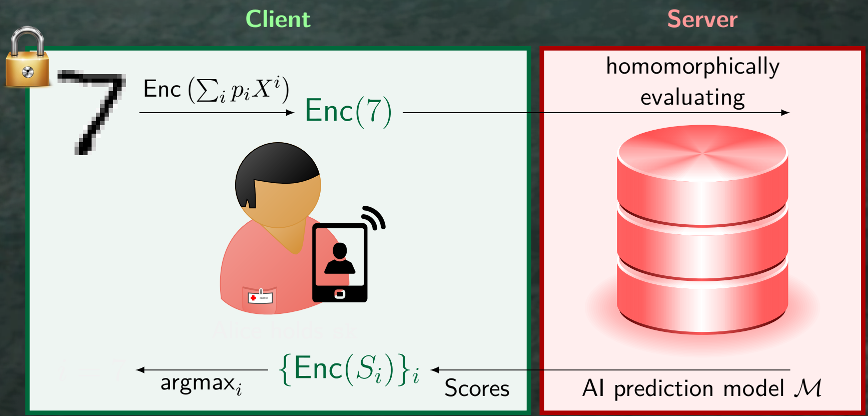

FHE–DiNN: Input Image and 784:100:10–Neural Network

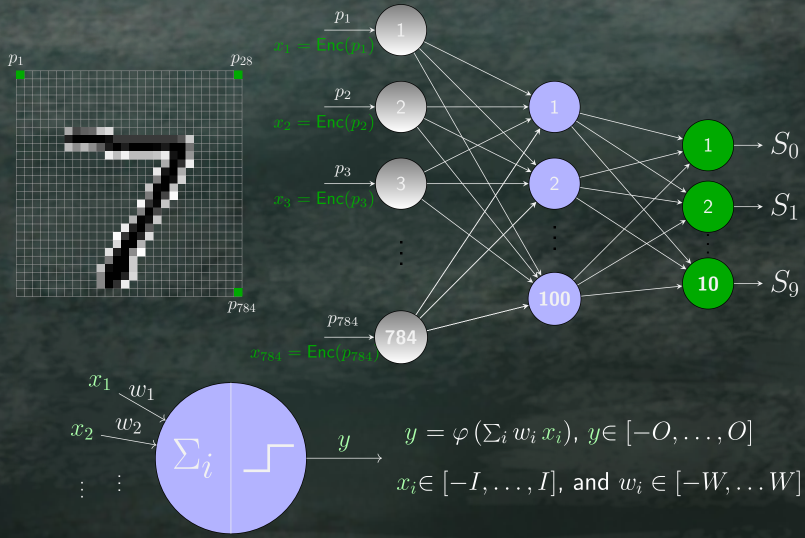

$y = \varphi\left(\Sigma_i\, w_i\, x_i\right)$, $y \in [-O, \ldots, O]$

$x_i \in [-I, \ldots, I]$, and $w_i \in [-W, \ldots W]$.

Hidden Neuron (zoomed)

FHE–DiNN: Algorithmic Overview [BMMP18]

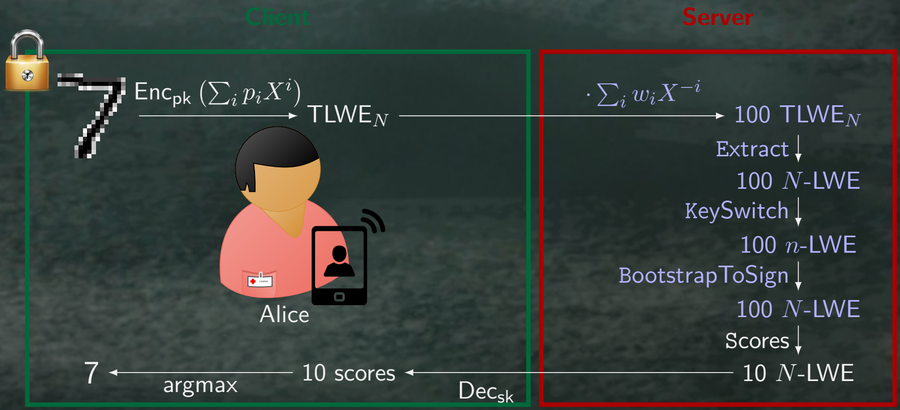

FHE–DiNN: Evaluation Formula of our 784:100:10–network

We assume a neural network trained on $D_{\text{train}} = \{(\mathbf{x}^{(i)}, L^{(i)})_i\}$.

$\mathcal{M}_{\text{FHE–DiNN}}$ models a weighted recomposition of a TLWE encryption $\mathbf{c}_0$;

$$\begin{cases} \mathbb{T}_N[X]^k \longrightarrow (\mathbb{T}_N[X]^k)^{10} \\ \mathbf{c}_0 \mapsto \vec{\mathbf{c}}_2 = \sum_{\ell_2=1}^{100} \left(\underbrace{\varphi_1 \left(\sum_{\ell_1=1}^{784} (\mathbf{c}_0)_{\ell_1} \cdot (\widehat{\mathbf{w}_{0\to 1}})_{\ell_1} \right)}_{\vec{\mathbf{c}}_1} \right)_{\ell_2} \cdot (\widehat{\mathbf{W}_{1\to 2}})_{\ell_2}. \end{cases}$$

The homomorphic evaluation yields 10 samples $\overrightarrow{\mathbf{c}_{\mathcal{O}}}$ as output, encrypting the perceptrons' predicted label likelihoods of an encrypted input digit $\mathbf{c}_{\mathcal{I}}$.

Label $L = \operatorname{argmax}_i \left(\mathrm{Dec}_{\mathsf{sk}}\left(\overrightarrow{\mathbf{c}_{\mathcal{O}}}\right)\right)_i$ is how the model sees the input's depicted digit: $L = \mathcal{M}_{\text{FHE–DiNN}}(\mathbf{c}_{\mathcal{I}})$, with $\mathrm{Dec}_{\mathsf{sk}}(\mathbf{c}_{\mathcal{I}}) \approx \mathbf{x}^{(\mathcal{I})} \in (D_{\text{train}})_{\mathbf{x}}$.

Main Result of the PhD-Thesis—Scalability

The analysis shows how to bootstrap the most expensive layer, then repeat for arbitrary many hidden neurons arranged in various layers.

FHE–DiNN Experiments: Practical Performance Neural Networks

Performance metrics on (clear) inputs x:

	Original NN	DiNN + hard_sigmoid	DiNN + sign
FHE–DiNN 30	94.76%	93.76% (-1 %)	93.55% (-1.21%)
FHE–DiNN 100	96.75%	96.62% (-0.13%)	96.43% (-0.32%)

Performance metrics on (encrypted) inputs $\mathsf{Enc}_{\mathsf{pk}}(x)$:

	Acc.	Disagreements	Total wrong BS	when dis.	Time
30	93.71%	273 (105–121)	3 383/300 000	196/273	0.515 s
100	96.26%	127 (61–44)	9 088/1 000 000	105/127	1.679 s
30 w	93.46%	270 (119–110)	2 912/300 000	164/270	0.491 s
100 w	96.35 %	150 (66–58)	7 452/1 000 000	99/150	1.640 s

window size w = 2

Performance Comparison with Microsoft Cryptonets [DGBL+16]

	Overall Network		per Image			
	$n_{\mathcal{H}}$	Accuracy	Eval [s]	$\lvert c\rvert$ [B]	Enc [s]	Dec [s]
Cryptonets	945	98.95 %	570	586 M	122	5
Cryptonets*	945	98.95 %	0.07	73.3 k	0.015	0.0006
FHE–DiNN30	30	93.71 %	0.49	$\approx$ 8.2 k	0.000168	0.0000106
FHE–DiNN100	100	96.35 %	1.64	$\approx$ 8.2 k	0.000168	0.0000106

Cryptonets* is amortized per image (accumulating 8192 inferences)

Experimental Results

Timing/Image on Intel Core i7-4720HQ CPU @ 2.60GHz: 1.64 [sec].

References

Horst Görtz Institut für IT-Sicherheit

RUB

Questions?

- Land Hessen Datenschutzgesetz. Gesetz- und Verordnungsblatt, 7. 10. 1970. http://starweb.hessen.de/cache/GVBL/1970/00041.pdf.
- Florian Bourse, Michele Minelli, Matthias Minihold, and Pascal Paillier. Fast homomorphic evaluation of deep discretized neural networks. Lecture Notes in Computer Science, 10993:483–512, 2018. DOI: 10.1007/978-3-319-96878-0_17.