EPIC: Efficient Private Image Classification*

Eleftheria Makri, Dragos Rotaru, Nigel Smart, Frederik Vercauteren

Dept. of Electrical Engineering, KU Leuven

Problem Statement

Design an *efficient* privacy-preserving system for image classification that does not leak any information about the images to be classified, nor about the classifier parameters, while being secure in the *active security* model.

Motivation

The images need to remain private, *and* the model parameters as well. For instance [1]:

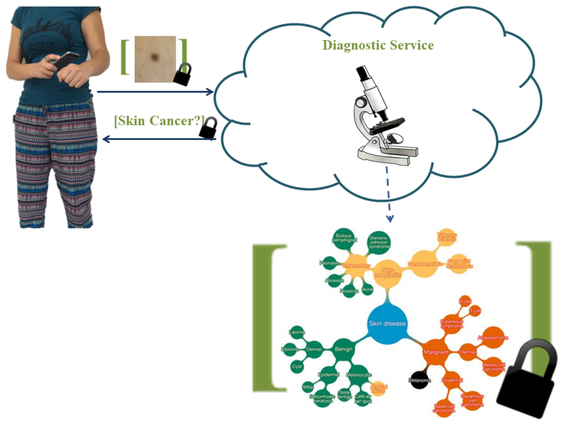

Contributions

- We enable full outsourcing of private image classification to a third independent party.
- Our solution does not leak any information about the private images, nor the classifier, while being the first to provide active security.
- We show how an effective, data-independent feature extraction technique can be deployed to alleviate the privacy-preserving computations.
- We demonstrate the practicality of our approach, both in terms of efficiency, and in terms of accuracy, by conducting experiments on realistic datasets.

*https://eprint.iacr.org/2017/1190

EPIC

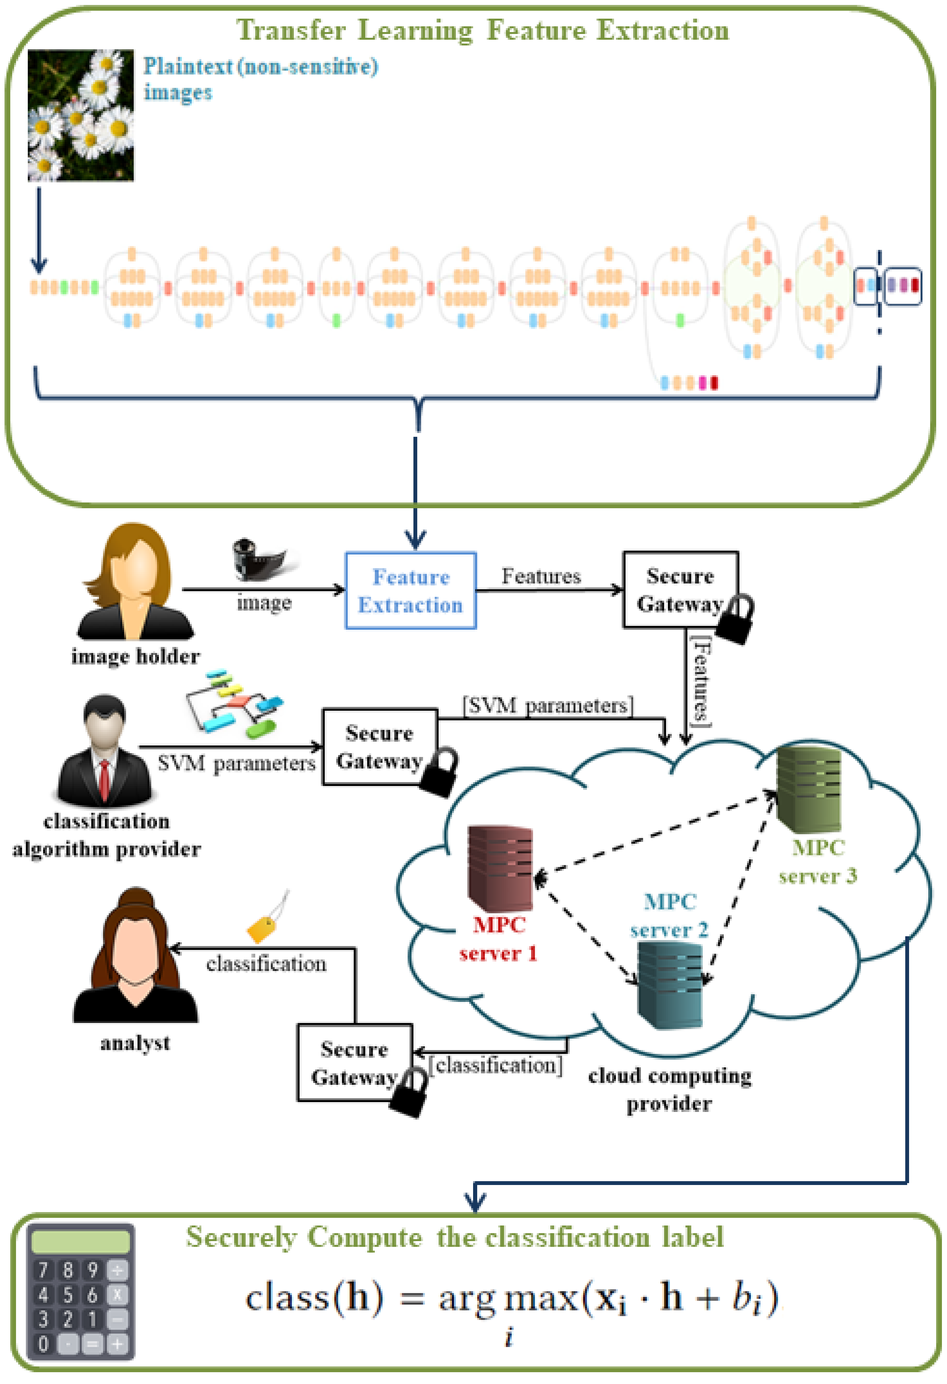

Efficiency Gains over the State-of-the-Art

EPIC vs. Gazelle [2] on CIFAR-10:

- 34 times faster runtime;
- 50 times improvement of communication cost;
- 7% higher classification accuracy.

EPIC vs. Gazelle [2] with the same accuracy:

- 700 times faster runtime;
- 500 times improvement of communication cost.

Results

Framework	Time (s)	Comm.(MB)	Accuracy (%)
MiniONN [3]	544	9272	81.61
Gazelle [2]	12.9	1236	81.61
EPIC	**0.37**	**24.33**	**88.8**

Table: 1 Gbps LAN timings for CIFAR-10 dataset.

Dataset	Time (s)	Comm.(MB)	Accuracy (%)
CIFAR-10	0.37	24.33	88.8
MIT-67	2.45	164.18	72.2
Caltech-101	3.74	250	91.4

Table: 1 Gbps LAN timings for EPIC (Linear SVM)

Dataset	Time (s)	Comm.(MB)	Accuracy (%)
CIFAR-10	0.037	2.5	81.74
MIT-67	0.261	17.4	64.4
Caltech-101	0.399	26.543	85.56

Table: 1 Gbps LAN timings for EPIC (RBF-SVM; 128 features)

References

[1] Andre Esteva, Brett Kuprel, Roberto A Novoa, Justin Ko, Susan M Swetter, Helen M Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115, 2017.

[2] Chiraag Juvekar, Vinod Vaikuntanathan, and Anantha Chandrakasan. GAZELLE: A low latency framework for secure neural network inference. In *27th USENIX Security Symposium (USENIX Security 18)*, Baltimore, MD, 2018. USENIX Association.

[3] Jian Liu, Mika Juuti, Yao Lu, and N Asokan. Oblivious Neural Network Predictions via MiniONN Transformations. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 619–631. ACM, 2017.