

Privacy-preserving Surveillance Methods using Homomorphic Encryption

William Bowditch¹, Will Abramson¹, William J Buchanan¹, Nikolaos Pitropakis¹ and Adam J Hall¹

¹*Blockpass ID Lab, Edinburgh Napier University, Edinburgh*

Keywords: Cryptography; SEAL; Machine Learning; Homomorphic Encryption; FV (Fan and Vercauteren)

Abstract: Data analysis and machine learning methods often involve the processing of cleartext data, and where this could breach the rights to privacy. Increasingly, we must use encryption to protect all states of the data: *in-transit*, *at-rest*, and *in-memory*. While tunnelling and symmetric key encryption are often used to protect data in-transit and at-rest, our major challenge is to protect data within memory, while still retaining its value. Homomorphic encryption, thus, could have a major role in protecting the rights to privacy, while providing ways to learn from captured data. Our work presents a novel use case and evaluation of the usage of homomorphic encryption and machine learning for privacy respecting state surveillance.

1 Introduction

User privacy has always been a second priority behind revenue for technology companies, and for valid reasons; the revenue of these tech giants relies on selling behavioural futures predicted from analyzing vast quantities of data (Zuboff, 2015). Companies like Facebook and Google digest raw data that need to be normalized, sorted, and trained on in order to produce effective machine learning models. While this data can be protected with SSL tunnels on transit, it must be decrypted and stored in plain-text on corporate servers in order to provide any value. A parallel exists in civilian privacy and national security; government agencies like the NSA rely on Internet surveillance programs that search plain-text data in order to detect threats of national security. On the surface, free internet services and effective terrorism countermeasures seem like a reasonable trade in exchange for one's personal data. However, as is often the case in information security, humans are the weakest link in the chain.

Unfettered access to personal data has facilitated cases where this access has been abused by both government agencies and industry (Selyukh, 2019)(Pymnts, 2019). With these examples the predicament is clear; we wish to provide data for these models such that they continue to subsidize free internet services and protect homeland security, but we do not trust the human users that inevitably gain access to this data. Fortunately, the cryptographic community has been working on a solution for forty years but it was not until recently that implementations became practical.




Figure 1: Homomorphic Encryption

Homomorphic encryption is a cryptographic scheme that allows one to perform arbitrary functions on a ciphertext without the need to decrypt the ciphertext in advance. Furthermore, the decrypted ciphertext is equivalent to the output of the same arbitrary functions performed on the plain-text. Figure 1 illustrates this relationship. For example, fully homomorphic encryption would allow users to send encrypted data to government agencies and technology companies such that models can train and act on this data without knowing or needing to store the plain-text itself. A cryptographic scheme is partially homomorphic (PHE) if it allows unlimited operations to be performed but with one particular function, while a scheme is somewhat homomorphic if it allows limited operations of any arbitrary function.

Up until now, homomorphic encryption research has been mainly focused on the development of new methods and performance analysis. In order for it to be adopted into the mainstream, we need strong use cases. This paper thus presents a novel use case of homomorphic encryption that utilises machine learning and applied into surveillance. It uses (scikit-learn) and Python implementations of Pailler and FV

schemes, in order to create a homomorphic machine learning classification technique that allows model owners to classify data without jeopardizing user privacy. The paper aims to provide a review of modern HE schemes for non-cryptography specialists, and gives simple examples of the usage of homomorphic encryption. Code samples are used in order to illustrate the application of the methods, and support readability.

While the state-of-the-art homomorphic methods proposed today are impractical for computationally complex tasks like machine learning without substantial delay (Laine et al., 2018), the schemes reviewed below are capable of handling machine learning evaluation. We construct a hypothetical scenario, solved with homomorphic encryption, such that a government agency wishes to use machine learning in order to identify pro-ISIS messages without (a) collecting the messages of citizens and (b) allowing users to reverse engineer the model. Our implementation differs from previous approaches (Trask, 2019) because it utilizes the machine learning library scikit-learn (Pedregosa et al., 2011), the Github repository python-paillier (python-paillier, 2012), and a Python port (pySEAL, 2017) of Microsoft SEAL 2.3 (Laine, 2017) in order to benchmark and evaluate the parameters of the cryptosystems proposed by Paillier (python-paillier, 2012) and Fan Vercauteren (Fan and Vercauteren, 2012).

2 Background and Related Work

The literature review will investigate fully homomorphic encryption schemes, starting with Craig Gentry’s 2009 seminal paper, which was the first paper to describe a credible fully homomorphic encryption scheme (Gentry et al., 2009), followed by Brakerski and Vaikuntanathan’s work (Brakerski and Vaikuntanathan, 2014) as well as Junfeng Fan and Frederik’s Vercauteren’s (VF) suggestion (Fan and Vercauteren, 2012), both approaches that build off of Gentry’s work and are implemented by Microsoft in the C++ library SEAL (Laine, 2017). Melchor (Melchor et al., 2018) evaluated HELib, SEAL and FV-NFLlib, and found that SEAL V2.3 performed best for multiplicative homomorphic encryption.

2.1 Gentry

Craig Gentry broke new ground in the field of homomorphic encryption with his seminar paper, “Fully Homomorphic Encryption Using Ideal Lattices” (Gentry et al., 2009). Gentry’s method re-

lies on a somewhat homomorphic lattice-based crypto scheme; the scheme is limited in the number of operations that can be performed on a ciphertext before *noise*, a by-product of the probabilistic nature of the scheme, grows so large such that the plain-text mapping is inaccurate. The monumental insight gained from Gentry’s work was the concept of bootstrapping, a technique that refreshes the noise of a ciphertext by decrypting the ciphertext with a new key without revealing the plain-text. While strictly following Gentry’s algorithm was unrealistic due to Big-O complexity, his method was the foundation for practical implementations such as HELib and SEAL, the latter of which is utilized in Section 3.

2.1.1 Lattice Based Cryptography

In linear algebra, a basis of a vector space is a set of n independent vectors such that any coordinate point on said space is a linear combination of these basis vectors. The lattice of a vector space is the set of basic linear combinations with integer coefficients; for example, all (x,y) points where $x,y \in \mathbb{Z}$ on a Euclidean vector space make up the lattice. Ideal lattices are, “lattices corresponding to ideals in rings of the form $\mathbb{Z}[x]/(f)$ for some irreducible polynomial of degree n ” (Wikipedia contributors, 2004). Ideal lattices are essential to the semantic security of Gentry’s FHE method due to the intractable nature of the closest vector problem - given a vector v outside of any lattice points, which lattice point is closest to v ? The closest vector problem forces one to perform lattice basis reduction in order to be solved, but at the cost of exponential time.

When the vector without error is known by a party, this closest vector problem allows this party to “hide” an encoded message m_1 with an error if the message space is $(\text{mod } p)$ for some integer p , the cipher space is $(\text{mod } q)$ for some integer $q \gg p$, and the error is divisible by p , allowing simple future removal of the error. Consequently, the error is calculated by randomly generating e from a uniform distribution and multiplying e by p , thus ensuring this divisibility and clean error removal. Furthermore, it is essential that the chosen p is much less than q since all operations in the scheme are performed $(\text{mod } q)$ (Raynal,).

Due to the algebraic properties of vector addition and multiplication, it is possible to calculate the sum and product of two cipher-texts with the respective sum and products of the error. When the error is removed after the operation $F(c_1, c_2)$, via decryption, the output is equivalent to $F(m_1, m_2)$. However, this growth in the error is why lattice-based cryptography is somewhat homomorphic; if the error grows too large then the closest lattice vector during decryption

is no longer accurate. For example, on a Euclidean space if the correct lattice vector is $(1, 1)$ but the error is $x = .3$ and $y = 0.6$, then the decryption will incorrectly decrypt to $(1, 2)$.

2.1.2 Bootstrapping

The solution Craig Gentry proposes to counter noise growth in lattice-based cryptography is a "bootstrapping" technique, thus transforming the scheme from partially homomorphic to full homomorphic cryptography. The technique is based on the intuitive notion that the only way to remove noise is to decrypt the cipher-text; therefore if the decryption operations are performed with the private key k_1 encrypted with a new private key k_2 as well as the ciphertext c_1 encrypted with the new key k_2 , then the error is reduced to the noise added by the homomorphic decryption operation. Due to the circuit complexity of the decryption operation, Gentry invents a squashing technique to the decryption function that in a sense provides a "hint" to the decryption process for the evaluator, but relies on the intractability of the subset sum problem; given a set of integers find a non empty subset with a sum of 0 (Gentry et al., 2009). The majority of homomorphic research in the past decade has been built upon Gentry's proposal, mainly focusing on (1) reducing the homomorphic operation cost of decryption and (2) reducing the resources necessary to encrypt an already encrypted cipher-text. The bootstrapping method, though, has been criticised for its requirement to refresh noisy ciphertexts. Ducas et al (Ducas and Micciancio, 2015) present a new method (FHEW) and which homomorphically computes simple bit operations. This reduces the overhead to around half a second and is based on the worst-case hardness of lattice problems.

2.2 Brakerski and Vaikuntanathan

In 2011, Zvika Brakerski and Vinod Vaikuntanathan published a fully homomorphic encryption scheme that improved on Gentry's scheme in both efficiency and simplicity (Brakerski and Vaikuntanathan, 2014). This paper introduces two key concepts: a re-linearization technique that removes the need for intractability assumptions regarding ideal lattices, and a dimension-modulus technique that removes the need for squashing and the above mentioned intractability assumption of the subset sum problem.

2.2.1 Re-Linearization

Braverski and Vaikuntanathan migrated the assumption from ideal lattice cryptography to general lattices

by utilizing learning with errors, which states that given a basis, a linear combination of the basis vector with small error, and a lattice point, finding the latter vector from the former is computationally difficult. Since ideal lattices are a relatively new field of study in the mathematics community as opposed to general lattices, Brakerski and Vaikuntanathan claim that the community has, "a much better understanding of the complexity of lattice problems (thanks to [LLL82, Ajt98, Mic00] and many others), compared to the corresponding problems on ideal lattices" (Brakerski and Vaikuntanathan, 2014), thus providing more confidence to this scheme's semantic security. By using learning with errors, a homomorphic multiplication optimization called re-linearization can be performed, which utilizes a new key to decrease the degree of cipher-text. This optimization prevents noise from growing exponentially during multiplication, replacing the growth factor with a constant dependent on the initial security parameter λ .

2.2.2 Dimension-Modulus Reduction

As mentioned above, Gentry's 2009 paper utilizes a "squashing" technique in order to ensure that homomorphic operations were possible on the decryption circuits by relying on the hardness of the subset sum problem. Braverski and Vaikuntanathan demonstrate that a learning with error homomorphic scheme with a dimension-modulus reduction only requires a relatively small decryption, thus making any squashing of the decryption circuit unnecessary. Dimension-modulus reduction is the process of converting a ciphertext with dimension n and cipher modulo q and mapping this ciphertext with new parameters of dimension k where $k \ll n$ and modulo $\log(p)$ where p is the plain-text modulus. The semantic security of this dimension-modulus reduction relies on the hardness of learning with errors for dimension k modulo $\log(p)$. Consequently, this reduction allows Braverski and Vaikuntanathan's scheme to be boot-strappable, thus fully homomorphic, without assumptions beyond the intractability of the learning with errors problem.

2.3 Fan Vercauteran Scheme

Jufeng Fan and Frederik Vercauteran's "Somewhat Practical Fully Homomorphic Encryption" directly builds off of Braverski's learning with errors homomorphic scheme by introducing a ring variant of the learning with error problem (Fan and Vercauteran, 2012). The 2012 paper optimizes Braverski's re-linearization with the aid of smaller re-linearization keys, as well as a modulus switching trick in order to simplify bootstrapping.

2.3.1 Ring Learning with Errors

In mathematics, a ring R is a set with two binary operations that allows generalization from normal arithmetic to other frames like polynomials and functions. Thus, a polynomial ring can be $R = \mathbb{Z}/f(x)$ where $f(x) \in \mathbb{Z}[x]$ is a monic irreducible polynomial of degree d . Fan and Vercauteren utilize polynomial rings in creating the hardness of their scheme: Definition 1 (Decision-RLWE):

For security parameter λ , let $f(x)$ be a cyclotomic polynomial $\omega_m(x)$ with $\deg(f) = \sigma(m)$ depending on λ and set $R = \mathbb{Z}[x]/(f(x))$. Let $q = q(\lambda) \geq 2$ be an integer. For a random element $s \in R_q$ and a distribution $\chi = \chi(\lambda)$ over R , denote with $A(a)$ the distribution obtained by choosing a uniformly random element $a \leftarrow \chi$ and outputting $(a, [a \cdot s + e]_q)$. The Decision-RLWE problem is to distinguish between the distribution $A_{s,\chi}^q$ and the uniform distribution $U(R_q^2)$. (Fan and Vercauteren, 2012)

3 Terrorist Surveillance Use Case

As mentioned the Introduction, internet technology companies and government agencies each have an imperative that require data analytics. However, the requirement for data analytics does not imply a need for data collection; these organizations do not relish data silo maintenance, data breach countermeasures, and the damage control against inevitable internal employee misuse. Homomorphic cryptography can diminish these vulnerabilities by separating the data evaluators from the data owners. For example, imagine a government agency who wishes to detect messages related to terrorist activity exchanged on a public network; we will refer to this agency as Big Brother. But unlike Orwell’s dystopic counterpart, our Big Brother has regulations in place that prevent the collection of plain-text messages. Our Big Brother requires probable cause before being granted a warrant for a citizen’s data, thus ensuring the honest citizen’s data privacy. Big Brother’s necessity for surveillance can be met by using a homomorphic scheme akin to a metal detector in an airport. Rather than forcing a body search of every passenger, airport security use metal detectors to single out the potentially dangerous passengers. Furthermore, since dangerous passengers cannot experiment with an airport metal detector from home, the ability to reverse engineer or trick the airport detector is severely limited.

In our implementation of homomorphic cryptography, our metal detector is a homomorphically en-

crypted logistic regression model trained to detect pro-ISIS tweets. In our hypothetical scenario, Big Brother trained his model using pro-ISIS messages collected from previous investigations and synthetic ISIS-related data; in reality, the pro-ISIS ($pro_isis = 1$) and ISIS-related data ($pro_isis = 0$) was collected from a Kaggle dataset (Tribe, 2016). After training, Big Brother encrypts the weights and y-intercept of his model using either the Paillier scheme or Fan Vercauteren scheme, and sends the encrypted model with any necessary evaluation parameters to the messaging devices of his citizens. A generalisation of this is defined in Figure 2.

When an arbitrary citizen, who we will refer to as Winston Smith, sends a message from his device to his friend Julia, the message is evaluated by the encrypted model resulting in an encrypted prediction that is sent to Big Brother. Winston can not discern the result nor purpose of the encrypted prediction, and it is not possible for Winston to disable this feature. While it is feasible that Big Brother supplemented the encrypted model with an encrypted identity function in order to obtain the plain-text message, an assumption of our use case is that Big Brother has no desire to collect user data due to aforementioned regulations and a general lack of trust in his own employees. Upon receiving the encrypted prediction, Big Brother decrypts the prediction m with his private key and inputs this value into the logistic function $p = \frac{1}{1+e^{-m}}$ and ignores any probability p below an arbitrary threshold γ . However, a probability greater than γ is equivalent to a *beep* in airport security and can be used as probable cause for a warrant in order to lawfully obtain Winston’s plain-text messages (as well as more ground truth to improve the accuracy of Big Brother’s model). Consequently, two major goals have been achieved: lawful citizens of the state are ensured data privacy, and the state can protect national security without collecting personal data nor revealing the source code of their surveillance to potential criminals.

4 Methods

4.1 Logistic Regression

Logistic regression is a statistical technique used to estimate the parameters of a binary model, binary in that there are two possible outputs for the dependent variable discerned from independent variables. The estimated parameters, or weights, along with the y-intercept and logistic function allow us to calculate the probability for either one of the dependent target




Figure 2: Homomorphic Surveillance

variables. A logistic model is trained using a maximum likelihood estimate and labeled data, with the goal being to find the model parameters that most minimize the training error. Logistic regression is a fundamental machine learning method, and as such is implemented in the open-source Python machine learning library scikit-learn (Pedregosa et al., 2011). Before submitting text data for training in natural language processing, one must decide how they wish to quantify the representation of text. For the sake of simplicity, we will use a bag-of-words approach that tokenizes each message into a vector where each index represents a distinct word from the "vocabulary" of the training set. Since there are 44,000 unique words in our 37,000 message dataset, the training matrix has 44,000 columns and 37,000 rows. We can reduce this sparse column vector to any size we choose for the sake of optimization at the cost of increasing hash collisions and accuracy; the number of features used in our implementation is benchmarked in Section 5.

4.2 Python Paillier

The Paillier implementation used in this paper is from a publicly available Github repository owned by N1 Analytics (python-paillier, 2012). The initialization of the scheme is trivial; the API key-pair generation function is given a requested key size in bits and returns a randomly generated public and private key pair. The precision argument is used during the encryption process for rounding floating point integers, as the Paillier scheme only works on integer values. When Big Brother uses his public key to encrypt the model parameters, an EncryptedNumber object is in-

stantiated containing the ciphertext along with Big Brother's public key; no further information needs to be provided to Winston in order to evaluate the ciphertext.

4.3 PySEAL

The PySEAL library is a Python port of the C++ library SEAL, developed by a homomorphic research team at Microsoft (pySEAL, 2017) (Laine, 2017). SEAL is an implementation of the Fan Vercauteren homomorphic scheme, including further optimizations introduced by the 2016 paper "A Full RNS Variant of FV like Somewhat Homomorphic Encryption Schemes" (Bajard et al., 2016). Due to the fully homomorphic nature of the Fan Vercauteren scheme, the initialization parameters are more complex relative to the Paillier scheme.

The polynomial modulus is parameter required to be a power-of-two cyclotomic polynomial of the form $1x^{\text{power-of-two}} + 1$. The size of the polynomial modulus is proportional to the security level of the homomorphic scheme and inversely proportional to the computation time because a larger polynomial modulus increases the size of cipher-texts.

The second parameter chosen is the coefficient modulus, which is directly proportional to the allowance of noise accumulated on a ciphertext before the plain-text message is unrecoverable. However, increasing the coefficient modulus will also decrease the security level of the scheme. The SEAL documentation advises the use of the helper function illustrated on lines 11 to 16, where $\text{bit_strength} = n$ denotes that it would take 2^n operations to break the cipher. Furthermore, the plain-text modulus is configured in or-

der to determine the size of the plain-text data, and is inversely proportional to the noise budget. The importance of these parameters on the remaining noise budget and computation time will be investigated in Section 5. Ultimately, the noise budget of an encrypted ciphertext can be estimated with:

$$\text{noise_budget} \approx \log_2\left(\frac{\text{coefficient_modulus}}{\text{plain_modulus}}\right) \quad (1)$$

The fractional encoder allows the encoding of a fixed-precision rational number into a plain-text polynomial. Given a base b , the fractional encoder maps a rational number to a polynomial where $x = b$, and will limit the number of fractional and integral coefficients by `fractional_coeff` and `fractional_base`, respectively. The fractional encoder will also move any fractional component x of the number to the degree of the polynomial modulus minus the degree of the fractional component multiplied by $-1; 0.75$ where $b = 2$ and the polynomial modulus degree is 12 would be encoded as $-1x^{11} - 1x^{10}$. To further illustrate this encoding method, a fraction encoded with polynomial degree $n = 2048$, base $b = 2$, `fractional_coeff` = 1, and `fractional_base` = 64 will be encoded as:

$$24.2351 = 1x^4 + 1x^3 - 1x^{-2047} \quad (2)$$

4.4 Big Brother and Winston Smith

4.4.1 Big Brother

Our Big Brother is initialized with a classifier object, a vectorizer object, a cryptographic scheme, a test set, and a training set. The classifier is the machine learning model used for surveillance, while the vectorizer is used to extract independent variables from raw text. The cryptographic scheme is a 'Homomorphic-Cryptography' object with the decrypt and encrypt implementations abstracted from Big Brother, allowing simple substitution of any homomorphic cryptographic method.

4.4.2 Winston Smith

Winston Smith is initialized with test data, a homomorphic evaluator, a vectorizer, and the encrypted model. The test data is used to simulate original messages written by Winston Smith. The evaluator contains any parameters necessary to evaluate his messages. The vectorizer is used to transform his raw message to a token vector; while this reveals the number of features used in the encrypted model, it does not reveal the importance or weight of any feature. The encrypted model is supplied by Big Brother, and

Table 1: PySEAL Benchmark

Parameter	Default Value
nfeatures	1000
polymodulus	1024
plainmodulus	16
integralcoeffs	32
fractionalcoeffs	64

is a tuple of the encrypted weights and encrypted y-intercept.

5 Evaluation

A truly practical homomorphic cryptographic scheme needs to be able to calculate arbitrary calculations without developers needing to fine tune their evaluation code or wait substantial time for calculations that are relatively fast on plain-text data. Fortunately, the computational requirements for a logistic regression prediction are small; multiplication against plain-text data is not only possible in the Fan Vercauteren and Paillier scheme, but also has no affect on the noise budget of cipher-texts in the former. A logistic regression prediction, where n represents the number of features used in the model, x_i represents the data token at index i , w_i represents the weight coefficient at index i , and y_o represents the intercept is calculated as such:

$$Y = w_0 \cdot x_0 + w_1 \cdot x_1 + \dots + w_{n-1} \cdot x_{n-1} + y_0 \quad (3)$$

This formula illustrates that addition will be the primary culprit for any noise accumulated during Winston's evaluation. Furthermore, the expansion and exponentiation of the plain-text model to the ciphertext model will be time intensive for both the Paillier and FV scheme.

5.1 PySEAL Benchmark

The PySEAL benchmark was performed by timing different components of our exchange while changing one exchanging parameter. The different components are: the scheme initialization, the model encryption, the homomorphic evaluation, and the prediction decryption. The benchmark was performed on the same message throughout the experiment, and using a 2.2GHz, 6core 8thgeneration Intel Core i7 processor with 32GB RAM. Unless otherwise stated, the default parameters for this experiment are defined in Table 1.

Table 2: Polynomial Modulus

Polynomial Modulus	Noise Budget
1024	10
2048	36
4096	90
8192	198

Table 3: Bit Strength

Bit Strength	Noise Budget
128	10
192	1

5.2 Polynomial Modulus

For the FV scheme of PySEAL, the importance of the polynomial modulus degree on the computation time and remaining noise budget of the final ciphertext $E(Y)$ was evaluated for 2^n where $10 \leq n \leq 14$. The polynomial modulus increases the size of the ciphertext dramatically; although not recorded precisely in this paper, the benchmark scripts memory usage grew linearly with the modulus degree such that a degree of 2^{10} required 400MB of RAM while a degree of 2^{14} required 20G of RAM. It can be seen from the below table how the degree of the polynomial modulus has a positive relationship with the available noise budget; this increased noise budget allows us to evaluate the ciphertext further without needing to re-linearize or otherwise refresh the noise, an expensive operation (Table 2).

Figure 3 demonstrates the cost that increasing the polynomial modulus and thus the ciphertext size has on the computation time.

5.3 Bit Strength

As mentioned in Section 3, the bit strength denotes the level of security for our scheme. While Figure 4 illustrates that while increasing the bit strength from 128 to 192 has little effect on the computation time for our model with 1000 features, the bit strength table demonstrates that it does have a negative effect on the remaining noise budget of the evaluated ciphertext (128-bit strength gives a noise budget of 10, and 192-bit strength gives a noise budget of 1).

(Table 3).

5.4 Plain Modulus

The effect of the plain modulus degree on the noise budget is evident by the negative relationship demonstrated in Table 4. Furthermore, Figure 5 illustrates




Figure 3: Polynomial Modulus




Figure 4: Bit Strength

that the plain modulus is uncorrelated with the time complexity of the FV cryptographic scheme.

5.5 Number of Features

The number of features used by the logistic regression model has a direct correlation with the accuracy of the model. Figure 6 shows that the number of features is also correlated with the model encryption and evaluation time because more weights need to be encrypted and evaluated (Table 5).

Due to the partially homomorphic nature of the Paillier cryptosystem, there is no noise budget to manage as such there does not exist a noticeable difference in time required to perform logistic regression evaluation. Figure 7 demonstrates the private key lengths lack of correlation with the encryption and evaluation time of Paillier cryptography. The benchmark of Paillier was performed using a count vectorizer with 44,000 features, as opposed to 1000 features for SEAL; this difference demonstrates the light-weight efficiency of the Paillier cryptosystem for this addition-intensive specific evaluation.

Table 4: Plain Modulus

Plain Modulus	Noise Budget	Plain Modulus	Noise Budget
64	8	128	7
256	6	512	5
1024	4	2048	3
4096	2	8192	1
16384	0		




Figure 5: Plain Modulus




Figure 6: Features




Figure 7: Features

Table 5: Number of Features

Features	Accuracy	Features	Accuracy
1000	0.9304	2000	0.9396
3000	0.9445	4000	0.9459
5000	0.9502	6000	0.9534
7000	0.9548	8000	0.9517
9000	0.9548	10000	0.9545

6 Conclusions

This paper has outlined and evaluated a method which preserves privacy within a surveillance infrastructure. We increasingly live in a world where the rights of citizens to privacy are core to fundamental rights. Homomorphic encryption can thus provide a foundation element in building surveillance systems that respect these rights.

REFERENCES

- Bajard, J.-C., Eynard, J., Hasan, M. A., and Zucca, V. (2016). A full rns variant of fv like somewhat homomorphic encryption schemes. In *International Conference on Selected Areas in Cryptography*, pages 423–442. Springer.
- Brakerski, Z. and Vaikuntanathan, V. (2014). Efficient fully homomorphic encryption from (standard) lwe. *SIAM Journal on Computing*, 43(2):831–871.
- Ducas, L. and Micciancio, D. (2015). Fhe: bootstrapping homomorphic encryption in less than a second. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 617–640. Springer.
- Fan, J. and Vercauteren, F. (2012). Somewhat practical fully homomorphic encryption. *IACR Cryptology ePrint Archive*, 2012:144.
- Gentry, C. et al. (2009). Fully homomorphic encryption using ideal lattices. In *Stoc*, volume 9, pages 169–178.
- Laine, K. (2017). Simple encrypted arithmetic library 2.3. 1. *Microsoft Re-*

search <https://www.microsoft.com/en-us/research/uploads/prod/2017/11/sealmanual-2-3-1.pdf>.

- Laine, K., Gilad-Bachrach, R., Chen, H., Lauter, K., Han, K., Huang, Z., and Jalali, A. (2018). Logistic regression over encrypted data from fully homomorphic encryption.
- Melchor, C. A., Kilijian, M.-O., Lefebvre, C., and Ricosset, T. (2018). A comparison of the homomorphic encryption libraries helib, seal and fv-nflib. In *International Conference on Security for Information Technology and Communications*, pages 425–442. Springer.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Pymnts (2018 (accessed April 20, 2019)). Lyft accused of giving access to rider data, including zuck’s phone number.
- pySEAL (2017). Pyseal. Lab41.
- python-paillier (2012). python-paillier. n1analytics.
- Raynal, F. A brief survey of fully homomorphic encryption, computing on encrypted data.
- Selyukh, A. (2013 (accessed April 20, 2019)). Nsa staff used spy tools on spouses, ex-lovers: Watchdog.
- Trask, A. (2017 (accessed April 20, 2019)). Safe crime detection.
- Tribe, F. (2016). How isis uses twitter. We scraped over 17,000 tweets from 100+ pro-ISIS fanboys from all over the world since the November 2015 Paris Attacks.
- Wikipedia contributors (2004). Ideal lattice cryptography — Wikipedia, the free encyclopedia.
- Zuboff, S. (2015). Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30(1):75–89.