Doing Statistics Blindfolded

Louis J. M. Aslett (aslett@stats.ox.ac.uk), Pedro M. Esperança, and Chris C. Holmes
Department of Statistics, University of Oxford

i-like.org.uk

UNIVERSITY OF OXFORD

1. Security in statistics

The extensive use of private and personally identifiable information in modern statistical applications, especially in biomedical applications, can present serious privacy concerns.

Indeed, industry is on the brink on embarking on biomedical applications on a scale never before witnessed via the impending wave of so-called 'wearable devices' such as smart watches, which can monitor vital health signs round the clock, perhaps fitting classification models to alert on different health conditions. Such constrained devices will almost certainly leverage cloud services, uploading reams of private health diagnostics to corporate servers.

Is there any hope of honouring people's desire for security while still performing statistical analyses?

2. Homomorphic encryption : the blindfold

Traditional encryption schemes (AES, SSL, . . .) secure data for archive or communication using a key k or keypair (k_p, k_s). Encrypt a *message*, $m \in M$, to a *ciphertext*, $c \in C$, with public key:

$$c \leftarrow \texttt{Enc}(k_p, m)$$

Decrypt with secret key:

$$m = \texttt{Dec}(k_s, c)$$

But if we want to compute, have to decrypt first because they are 'brittle':

$$\texttt{Dec}(k_s, f(c)) \neq f(m) \quad \forall f(\cdot) \neq \texttt{Id}(\cdot)$$

Homomorphic encryption

Rivest *et al.* (1978) hypothesised $\exists$ schemes allowing blindfolded computation. Not until Gentry (2009) was it shown to be possible for arbitrary numbers of additions & multiplications. Homomorphic if:

$$\texttt{Dec}(k_s, \texttt{Enc}(k_p, m_1) \diamond \texttt{Enc}(k_p, m_2)) = m_1 \circ m_2$$

for a set of operations $\circ \in \mathcal{F}_M$ acting in M that have corresponding operations $\diamond \in \mathcal{F}_C$ acting in C. "Fully homomorphic" $\implies \mathcal{F}_M = \{+, \times\}$.

Fully homomorphic exciting if $M = GF(2)$ because $+ \equiv \veebar$ and $\times \equiv \wedge$, so can reproduce arbitrary boolean logic (arbitrary computation).

Nirvana? Perhaps purgatory . . .

Flurry of excitement, followed by dose of reality.

- C usually complex (e.g. polynomial ring)
  - very slow computation
  - size of $c \gg$ size of m
- $\therefore$ $M = GF(2)$ impractical, but
  - $M = \mathbb{R}$ impossible
  - $M = \mathbb{Z}/n\mathbb{Z}$ for large n best
- . . . but if integers not boolean circuits we'd like
  - division
  - comparisons $(<, \leq, >, \geq, =)$

  which are not possible!

Quick reality check: can only evaluate polynomials of integers (in practise of limited degree).

The challenge: fit meaningful statistical models within these constraints.

3. Fan and Vercauteren (2012)

Notation

$\mathbb{Z}_q = \{n : n \in \mathbb{Z}, -q/2 < n \leq q/2\}$
$[a]_q \in \mathbb{Z}_q$ st $[a]_q = a \mod q$
$\mathbb{Z}[x], \mathbb{Z}_q[x]$ polynomials with coeff $\in \mathbb{Z}$ and $\in \mathbb{Z}_q$
$\Phi_n(x)$ n^{th} cyclotomic poly, $\Phi_{2^d}(x) = x^{2^{d-1}} + 1$
$R = \mathbb{Z}[x]/\Phi_{2^d}(x)$ and $R_q = \mathbb{Z}_q[x]/\Phi_{2^d}(x)$
$a(x) = \underline{a} \in R_q$ polynomial ring elements
$[\underline{a}]_q \implies$ centred reduction of coeff in $\mathbb{Z}_q$
$\cdot \sim \chi$ random poly with discrete Gaussian coeff
$\cdot \sim R_q$ random poly uniformly from R_q

Parameters

- d, degree of polynomial rings M and C;
- t, q, magnitude of coefficient sets of M, C;
- σ, magnitude of injected noise.

Encryption scheme

$M = R_t, C = R_q \times R_q$

Keys: $\underline{k}_s \sim R_2$ and $k_p := ([-(\underline{a} \cdot \underline{k}_s + \underline{e})]_q, \underline{a})$
where $\underline{a} \sim R_q$ and $\underline{e} \sim \chi$.

Encrypt: first map $m \in \mathbb{Z} \to \mathring{m}(x) \in R_t$

$$m = \sum_n a_n 2^n \to \mathring{m}(x) = \sum_{n=0}^{2^{d-1}-1} a_n x^n \in R_t$$
$$c := ([\underline{k}_{p1} \cdot \underline{u} + \underline{e}_1 + \Delta \cdot \underline{\mathring{m}}]_q, [\underline{k}_{p2} \cdot \underline{u} + \underline{e}_2]_q)$$

where $\underline{u}, \underline{e}_1, \underline{e}_2 \sim \chi$ and $\Delta = \left\lfloor \frac{q}{t} \right\rfloor$.

Add/mult: $c_1 + c_2 = ([\underline{c}_{11} + \underline{c}_{21}]_q, [\underline{c}_{12} + \underline{c}_{22}]_q)$

$$c_1 \times c_2 = \left(\left[\left\lfloor \frac{t(\underline{c}_{11} \cdot \underline{c}_{21})}{q} \right\rceil\right]_q, \left[\left\lfloor \frac{t(\underline{c}_{11} \cdot \underline{c}_{22} + \underline{c}_{12} \cdot \underline{c}_{21})}{q} \right\rceil\right]_q, \left[\left\lfloor \frac{t(\underline{c}_{12} \cdot \underline{c}_{22})}{q} \right\rceil\right]_q \right)$$

Decrypt c: $\underline{\mathring{m}} = \left[\left\lfloor \frac{t[\underline{c}_1 + \underline{c}_2 \cdot \underline{k}_s]_q}{q} \right\rceil\right]_t$

Then $m = \mathring{m}(2)$.

4. High performance R package

`HomomorphicEncryption` R package (Aslett, 2014) provides easy to use interface which hides all the complexity of homomorphic encryption.

Implementation is mostly high performance C++, with many operations setup to utilise multi-core parallelism without any end-user intervention.

Native support for vectors/matrices and all operators and common functions overloaded to run encrypted.

```
p <- parsHelp("FandV", lambda=80, L=8)
k <- keygen(p)
c <- enc(k$pk, matrix(1:9, nrow=3))
cres <- c[,1] %*% c
dec(k$sk, cres)
```

5. Completely Random Forests (CRF)

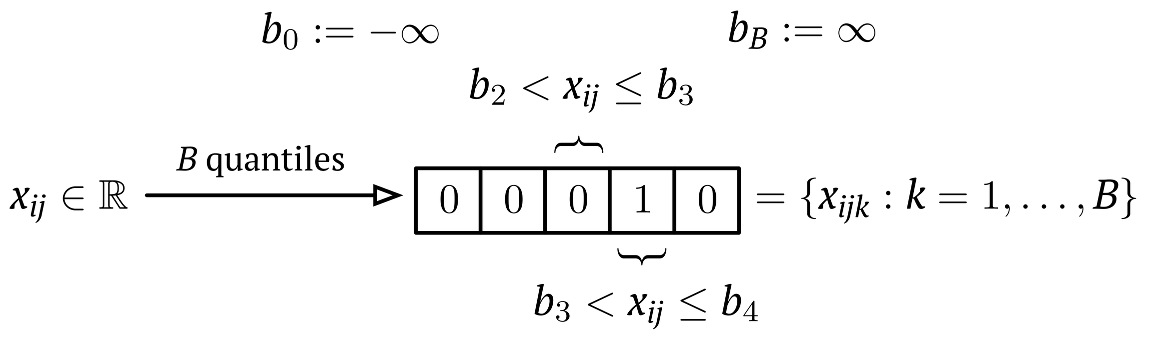

Then,

$$\mathbb{I}(x_{ij} \leq b_l) = \sum_{k=1}^{l} x_{ijk} \quad \text{and} \quad \mathbb{I}(x_{ij} > b_l) = \sum_{k=l+1}^{B} x_{ijk}$$

Also encode response category c, $y_i \to y_{ic} \in \{0, 1\}$ and build a decision tree selecting variable j and split point b_l *completely* at random to a fixed depth.

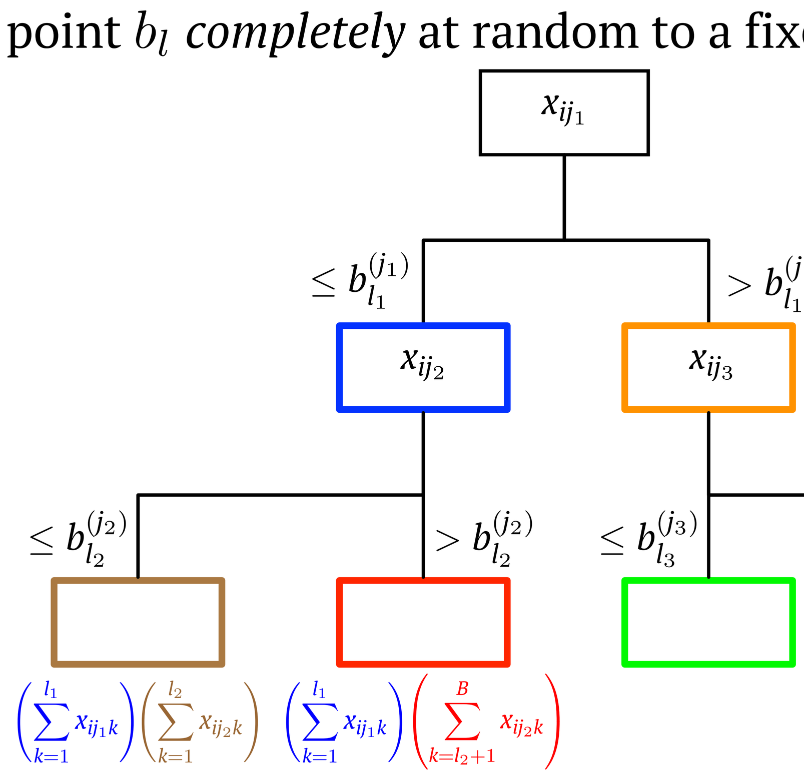

Count votes for class c in leaf, $\nu_c = \sum_i \nu_{ic}$

$$\sum_{i=1}^N \nu_{ic} = \sum_{i=1}^N y_{ic} \left(\sum_{k=1}^{l_1} x_{ij_1k} \right) \left(\sum_{k=1}^{l_2} x_{ij_2k} \right)$$

Stochastic fraction estimate

Relative class frequency should be $\frac{\nu_c}{\sum_c \nu_c}$.

But, $\in [0, 1]$ and involves $\div$. Target equivalently:

$$\nu_c \left\lfloor \frac{N}{\sum_c \nu_c} \right\rceil$$

By construction $\sum_c \nu_c \leq N$, so $0 \leq \frac{\sum_c \nu_c}{N} \leq 1$.

Simple fact: $X \sim \text{Geometric}(p) \implies \mathbb{E}[X] = p^{-1}$.

$\therefore$ unbiased approximation to fraction is a draw from Geometric distribution with $p = \frac{\sum_c \nu_c}{N}$. *Helping?!*

NB: $\{\sum_c \nu_{ic} : i = 1, \ldots, N\}$ is binary vector with exactly right proportion of 1's $\implies$ blind sample with replacement to get latent Bernoulli process underlying Geometric$\left(p = \frac{\sum_c \nu_c}{N}\right)$. *Still not helping?!*

Let $\xi_1, \ldots, \xi_M$ be resampled vector, M a power of 2. Then, for $l \in \{0, \ldots, \log_2(M) - 1\}$ set: $\xi_i = \xi_i \vee \xi_{i-2^l} = \xi_i + \xi_{i-2^l} - \xi_i \xi_{i-2^l} \quad \forall 2^l + 1 \leq i \leq M$

CPU hardware algorithm for renormalising mantissa of IEEE floating point number expressed with $+, \times$.

$$\implies \left\lfloor \frac{N}{\sum_c \nu_c} \right\rceil \approx M - \sum_{i=1}^{M} \xi_i + 1$$

6. Results

Other new crypto methods (see arXiv preprints)

- Semi-parametric naïve Bayes with linear logistic decision boundaries (SNB)
- One-step logistic regression (LR-onestep)

Tested on 20 different data sets from UCI repository.

Performance practicalities?

Full bcw_o data set, 100 tree SF with 3 levels deep fitted on Amazon EC2 cluster of 1152 CPU cores in 1 hour 36 minutes fully encrypted.

Total cost: less than US$ 24.

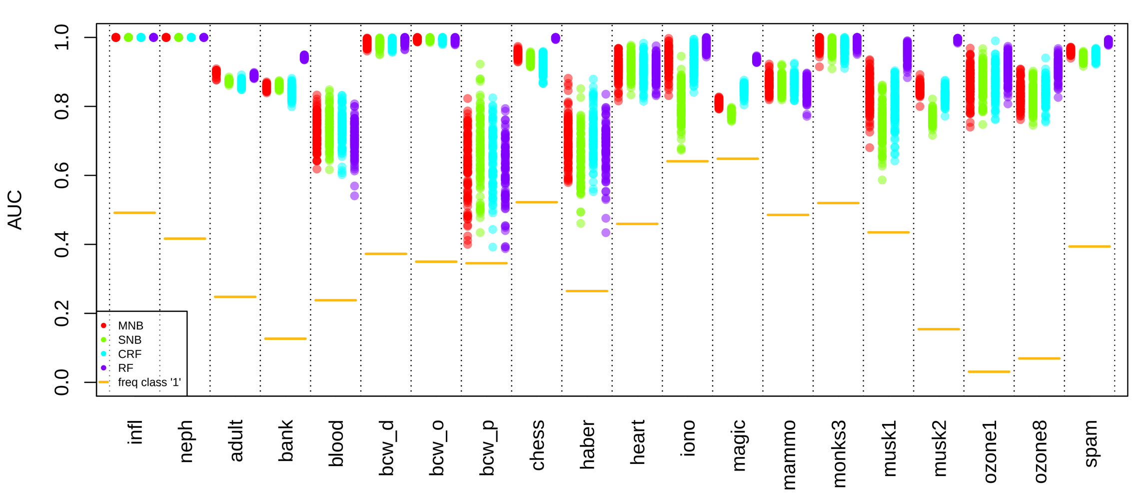

AUC for 100 randomisations of the train/test sets.

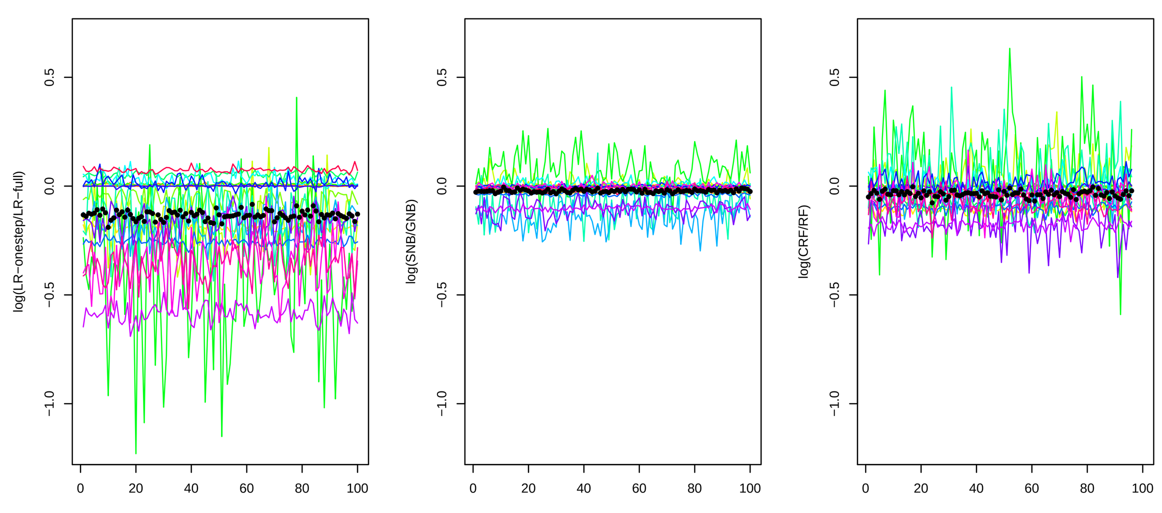

Ratio of encrypted method to traditional method AUC

References

Aslett, L. J. M. (2014), *HomomorphicEncryption: Fully Homomorphic Encryption*. R package version 0.2. www.louisaslett.com/HomomorphicEncryption/.
Aslett, L. J. M., Esperança, P. M. and Holmes, C. C. (2015*a*), `Encrypted statistical machine learning: new privacy preserving methods'. http://arxiv.org/abs/1508.06845arXiv:1508.06845 [stat.ML].
Aslett, L. J. M., Esperança, P. M. and Holmes, C. C. (2015*b*), A review of homomorphic encryption and software tools for encrypted statistical machine learning, Technical report, University of Oxford. http://arxiv.org/abs/1508.06574arXiv:1508.06574 [stat.ML].
Fan, J. and Vercauteren, F. (2012), `Somewhat practical fully homomorphic encryption', *IACR Cryptology ePrint Archive* .
Gentry, C. (2009), A fully homomorphic encryption scheme, PhD thesis, Stanford University.
Rivest, R. L., Adleman, L. and Dertouzos, M. L. (1978), `On data banks and privacy homomorphisms', *Foundations of Secure Computation* **4**(11), 169--180.

Funding

EPSRC programme grant EP/K014463/1
www.i-like.org.uk