Variational Bayes in Private Settings (VIPS)

James Foulds,[1*] Mijung Park,[2,3*] Kamalika Chaudhuri,[4] Max Welling[5]
UMBC,[1] MPI-IS,[2] U Tübingen,[3] UCSD,[4] UvA[5]
[*]Equal Contribution

Overview

- Motivation: data science applications in privacy-sensitive domains
  - Health informatics, MOOCs, social media data …
  - Many of these applications use Bayesian models
  - We need general privacy-preserving Bayesian inference algorithms!

- We propose a general privacy-preserving framework for variational Bayes: **Variational Bayes in Private Settings (VIPS)**
- To evaluate VIPS, we apply it to LDA topic models
  - In our full JAIR paper, we also study Bayesian logistic regression, sigmoid belief networks

Background

- ***Variational Bayes***
  - Optimization-based approach for approximate Bayesian inference
  - Make approximating distribution *Q* as similar as possible to target posterior distribution *P*
  - Equivalent to maximizing the *evidence lower bound (ELBO)*:

$$\log p(\mathcal{D}) = \log\Big(\int d\boldsymbol{l}\, d\boldsymbol{m}\; p(\boldsymbol{l},\boldsymbol{m},\mathcal{D})\Big) = \log\Big(\int d\boldsymbol{l}\, d\boldsymbol{m}\; p(\boldsymbol{l},\boldsymbol{m},\mathcal{D})\frac{q(\boldsymbol{l},\boldsymbol{m})}{q(\boldsymbol{l},\boldsymbol{m})}\Big)$$
$$= \log\Big(\mathbb{E}_q\Big[\frac{p(\boldsymbol{l},\boldsymbol{m},\mathcal{D})}{q(\boldsymbol{l},\boldsymbol{m})}\Big]\Big) \geq \mathbb{E}_q\Big[\log p(\boldsymbol{l},\boldsymbol{m},\mathcal{D}) - \log q(\boldsymbol{l},\boldsymbol{m})\Big] = \mathcal{L}(q)$$

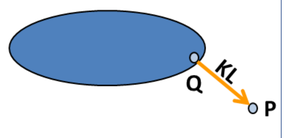

- ***Differential Privacy***
  - Gold-standard privacy definition for data-driven algorithms
  - Algorithm has similar behavior (outcome probabilities) if you change one data point

$\mathcal{M}(\mathcal{D})$ is said to be (ϵ,δ)-differentially private if

$$P(\mathcal{M}(\mathcal{D}) \in \mathcal{S}) \leq \exp(\epsilon) P(\mathcal{M}(\mathcal{D}') \in \mathcal{S}) + \delta$$

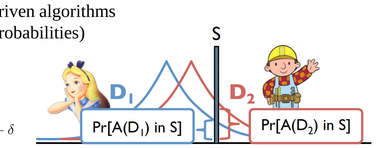

VIPS Technical Details

Mean Field VB

- Mean field assumes variational distribution factorizes: $q(\boldsymbol{l},\boldsymbol{m}) = q(\boldsymbol{l})q(\boldsymbol{m}) = q(\boldsymbol{m})\prod_{n=1}^{N} q(\boldsymbol{l}_n)$
- Mean field VB simplifies to two step VBEM procedure (E- and M-steps) for the conjugate-exponential (CE) family:

(1) The complete-data likelihood is in the exponential family:
$p(\mathcal{D}_n, \boldsymbol{l}_n|\boldsymbol{m}) = g(\boldsymbol{m}) f(\mathcal{D}_n, \boldsymbol{l}_n) \exp(\mathbf{n}(\boldsymbol{m})^\top \mathbf{s}(\mathcal{D}_n, \boldsymbol{l}_n))$,
(2) The prior on model parameters $\boldsymbol{m}$ is conjugate to the complete-data likelihood:
$p(\boldsymbol{m}|\tau,\boldsymbol{\nu}) = h(\tau,\boldsymbol{\nu}) g(\boldsymbol{m})^\tau \exp(\boldsymbol{\nu}^\top \mathbf{n}(\boldsymbol{m}))$,

VBEM

- E-step: Compute expected sufficient statistics
- M-step: Update expected natural parameters

- VIPS algorithm: we perturb the output of the E-step for stochastic VBEM

Algorithm 1 Private VIPS for CE family distributions

Require: Data $\mathcal{D}$. Define $\rho_t = (\tau_0 + t)^{-\kappa}$, noise variance σ^2, mini-batch size S, and maximum iterations J.
Ensure: Perturb expected natural parameters $\tilde{\mathbf{n}}$ and expected sufficient statistics $\tilde{\bar{\mathbf{s}}}$.
Compute the L2-sensitivity Δ of the expected sufficient statistics.
for $t = 1, \ldots, J$ **do**
Draw a minibatch of S datapoints, without replacement.
(1) E-step: Given the expected natural parameters $\bar{\mathbf{n}}$, compute $q(\boldsymbol{l}_n)$ for $n = 1, \ldots, S$. Perturb each coordinate of $\bar{\mathbf{s}} = \frac{1}{S}\sum_{n=1}^{S} \langle \mathbf{s}(\mathcal{D}_n, \boldsymbol{l}_n)\rangle_{q(\boldsymbol{l}_n)}$ by adding $\mathcal{N}(0, \sigma^2\Delta^2 I)$ noise, and output $\tilde{\bar{\mathbf{s}}}$.
(2) M-step: Given $\tilde{\bar{\mathbf{s}}}$, compute $q(\boldsymbol{m})$ by $\tilde{\boldsymbol{\nu}}^{(t)} = \boldsymbol{\nu} + N\tilde{\bar{\mathbf{s}}}$. Set $\tilde{\boldsymbol{\nu}}^{(t)} \leftarrow (1-\rho_t)\tilde{\boldsymbol{\nu}}^{(t-1)} + \rho_t \tilde{\boldsymbol{\nu}}^{(t)}$. Output the expected natural parameters $\tilde{\mathbf{n}} = \langle \mathbf{n}(\boldsymbol{m})\rangle_{q(\boldsymbol{m})}$.
end for
Compute the privacy loss $(\epsilon_{tot}, \delta_{tot})$ using the analytical moments account method (Wang *et al.*, 2019).

VIPS for LDA

Latent Dirichlet Allocation

Documents have **distributions over topics** $\theta^{(d)}$
Topics are **distributions over words** $\varphi^{(k)}$

Assumed generative process: (full model includes Dirichlet priors on θ, φ)
- For each **document** d
  - For each **word** $w_{d,n}$
    - Draw a **topic assignment** $z_{d,n} \sim \text{Discrete}(\theta^{(d)})$
    - Draw a **word** from the chosen topic $w_{d,n} \sim \text{Discrete}(\varphi^{(z}{}_{d,n)})$

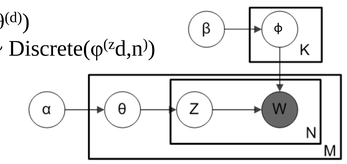

VIPS for LDA

- Preprocessing: resample all documents to have a fixed length, e.g., N = 500
- The expected sufficient statistics are the expected word-topic counts
- In each iteration, privatize by adding Gaussian noise to each word-topic count

$$\tilde{\bar{\mathbf{s}}}_k^v = \bar{\mathbf{s}}_k^v + Y_k^v, \text{ where } Y_k^v \sim \mathcal{N}(0, \sigma^2(\Delta\bar{\mathbf{s}})^2)$$

- The sensitivity is: $\Delta\bar{\mathbf{s}} \leq \frac{N}{S}$
- We use a clipping strategy to reduce the sensitivity, thereby reducing the amount of noise. Our approach leverages the fact that the norm for most documents is much smaller than the worst case $\bar{\mathbf{s}}^d := \frac{aN}{S}\frac{\bar{\mathbf{s}}^d}{|\bar{\mathbf{s}}^d|}$

a is a clipping hyperparameter. We use a = 0.1, for a ten-fold reduction in sensitivity, while clipping ¾ of the documents.

Key Ideas

- **Challenges**
  - Statistical efficiency when privatizing with latent variables
  - Iterative algorithms such as VB cumulatively increase privacy cost, hence increase noise
  - How to generalize beyond conjugate-exponential (CE) family models?

- **Our approach: VIPS**
  - *Perturb expected sufficient statistics*
    - Effective use of the privacy budget per iteration

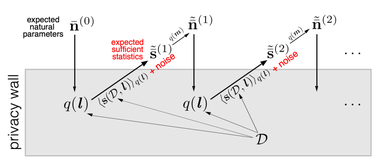

By perturbing expected sufficient statistics, we avoid having to perturb latent variables, variational parameters

  - *(Analytical) Moments Accountant (Abadi et al., 2016; Wang et al., 2019)*
    - Refined composition analysis, increase the privacy budget per iteration
    - Keeps track of a special quantity, the *log-moment function*, per iteration Easily composes across iterations. Use it to compute final privacy loss

  - *Privacy Amplification from Subsampling*
    - Stochastic VB scales to large datasets
    - Subsampling improves privacy guarantees

  - *Data Augmentation for non-CE family models*
    - Pólya-Gamma data augmentation brings a broad class of models into CE

Experimental Results

- VIPS with the moments accountant successfully learns a private topic model on Wikipedia, beats baselines
- For more details and experiments on Bayesian logistic regression, sigmoid belief networks, please see our full journal paper at JAIR:
  - *M. Park, J. R. Foulds, K. Chaudhuri, and M. Welling. Variational Bayes in Private Settings (VIPS). Journal of Artificial Intelligence Research (JAIR) 68:109-157, 2020.*

Non-private	Moments Accountant	Strong Composition	Moments Acc. (no clipping)
station	station	station	station
line	line	line	line
railway	railway	railway	french
opened	opened	opened	railway
services	services	services	opened
located	closed	stations	services
closed	code	closed	republic
owned	country	section	closed
stations	located	platform	stations
platform	stations	republic	country

Example topic from private LDA ($\epsilon = 2.38$)

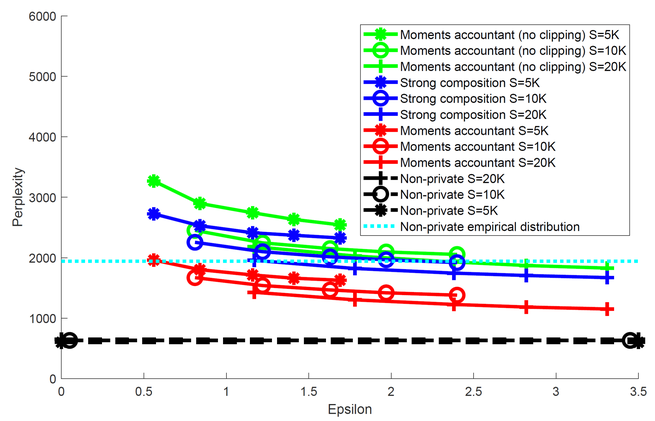

Email: jfoulds@umbc.edu, mijung.park@tuebingen.mpg.de, kamalika@cs.ucsd.edu, m.welling@uva.nl