Differentially Private and Efficient Sequential Learning Algorithms

Yu Wang, Zhenqi Huang, Sayan Mitra, Geir E. Dullerud

Introduction

A major concern in machine learning application is the privacy of sample data In this work, we adopt the concept of ε-differential privacy to study the privacy issue in a sequential empirical risk minimization setup where the number of samples needed is determined in the process of learning. Using exponential algorithm, we design a differentially private sequential learning algorithm.

Formulation

Advantages of sequential setup

- Fewer samples
- Less computation
- Probably Approximately Correct (PAC)

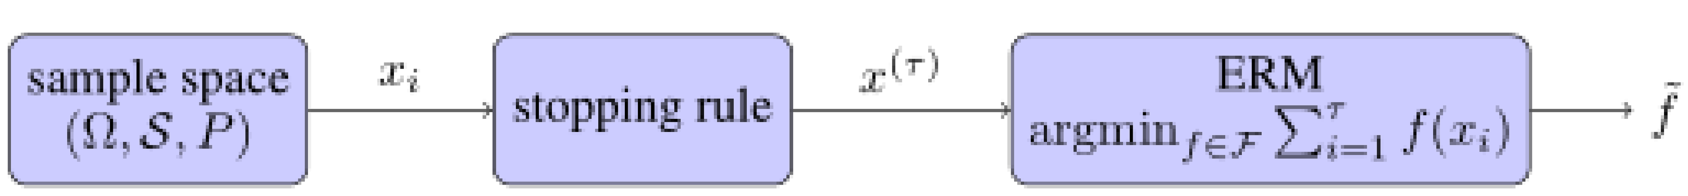

Learn from a class of binary functions F of finite VC-dimension using a sequence of samples $X = (x_1, x_2, \ldots)$

Ideal minimizer

$$f^* = \mathrm{argmin}_{f \in F} \int_{\Omega} f \mathrm{d}P$$

Empirical minimizer

$$f^{\mathrm{erm}} = \mathrm{argmin}_{f \in F} \sum_{i=1}^{\tau} f(x_i)$$

Efficiency

The stopping rule is (α, β)-*useful* if

$$\Pr[|P(f^*) - P(f^{\mathrm{erm}})| > \alpha] < \beta$$

It is (k_1, k_2, k_3)-*strongly efficient* if for any $(k_1 \alpha, \beta)$-useful stopping rule ν

$$\sup_P \Pr[\nu(k_2 \alpha, \beta) > \tau(\alpha, \beta)] < k_3 \beta$$

Differential Privacy

Adjacency: two sequences of samples differ in only one entry.

ε-differential privacy: for any stopping rule and adjacent samples X and X', $\Pr[(\tau_X, f_X^{\mathrm{erm}}) \in O] < e^{\varepsilon} \Pr[(\tau_{X'}, f_{X'}^{\mathrm{erm}}) \in O]$.

Algorithm

The algorithm is ε-differentially private and $(5 + \frac{3\varepsilon}{\alpha N}, 6 + \frac{3\varepsilon}{\alpha N}, 1)$-strongly efficient.

Algorithm 1 ε-differentially private sequential learning algorithm

Input $\alpha > 0, \beta \in (0,1), \varepsilon > 0, \tau = 1, r_{\mathcal{F}} = 0$ and $N(\alpha, \beta) = \left\lceil \frac{2}{\alpha^2} \ln \frac{2}{\beta(1 - e^{-\frac{\alpha^2}{2}})} \right\rceil$.
draw $\delta_\tau \sim \mathrm{Laplace}(1/\varepsilon)$
repeat
  draw $X_\tau \sim (\Omega, \mathcal{S}, P)$
  draw $\sigma_\tau \sim \mathrm{Bernoulli}(-1, 1)$
  $r_{\mathcal{F}} \leftarrow (\tau r_{\mathcal{F}} + X_\tau \sigma_\tau)/(\tau + 1)$
  $\tau \leftarrow \tau + 1$
until $\tau > N(\alpha, \beta)$ and $r_{\mathcal{F}} < \alpha + \delta/\tau$

$f^{\mathrm{erm}} = \mathrm{argmin}_{f \in \mathcal{F}} \sum_{i=1}^{\tau} f(X_i)$

Output $\mathrm{Exponential}(f^{\mathrm{erm}}, \varepsilon\tau)$

Conclusion

In this work, we designed a differentially private and strongly effective sequential learning algorithm, whose efficiency converges to non-differentially private case for large sample size.

Acknowledgement

This material is based upon work supported by the Maryland Procurement Office under Contract No. H98230-14-C-0141."