Preserving patient privacy in dynamic treatment regimes: Private outcome-weighted learning (PrOWL)

Dylan Spicker Erica E.M. Moodie Susan M. Shortreed

Department of Mathematics and Statistics
University of New Brunswick

Sunday December 17, 2023

Individualized Treatment Rules: Mathematical Statement

We want to estimate a decision function,

$$d\colon \mathcal{H} \longrightarrow \mathcal{A} = \{0, 1\},$$

where $H \in \mathcal{H}$ is the patient history and $A \in \mathcal{A}$ is the treatment decision.

We call this function a individualized treatment rule (ITR).

ITR Estimation as Classification

An ITR, d, has value

$$V(d) = E\{E[R|A = d(H)]\}.$$

Optimal ITRs maximize the value.

ITR Estimation as Classification

An ITR, d, has value

$$V(d) = E\{E[R|A = d(H)]\}\,.$$

Optimal ITRs maximize the value. Optimal ITRs minimize

$$E[R|A = 1] + E[R|A = -1] - V(d) = E\left[\frac{R}{P(A|H)} I(A \neq d(H))\right].$$

Outcome-Weighted Learning (OWL)

Outcome-Weighted Learning (OWL) estimates optimal ITRs by minimizing a regularized, empirical version of this error.

Support Vector Machines (SVM)

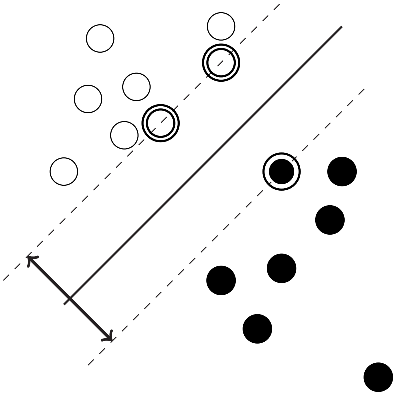

SVMs use hyperplanes to solve classification problems.

The resulting classifier exists as

$$f(H) = \sum_{i \in \mathcal{SV}} \alpha_i A_i K(H_i, H).$$

Privacy Implications of SVMs

Generally, the resulting decision function requires the **direct release** of the support vectors.

$$f(H) = \sum_{i \in \mathcal{SV}} \alpha_i A_i \exp\left(-\sigma^2 \|H_i - H\|\right)$$

Differential Privacy

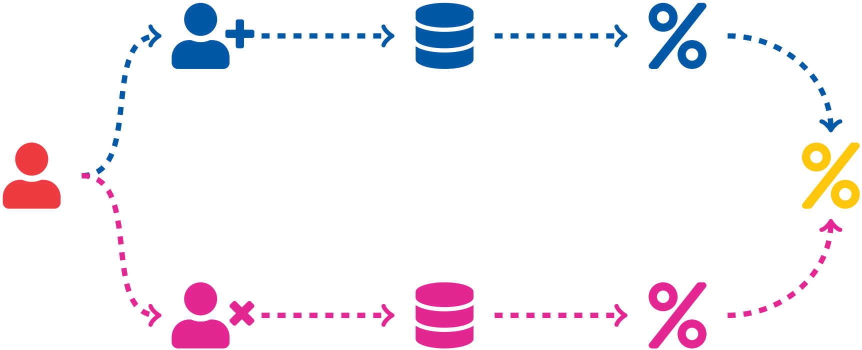

Differential Privacy

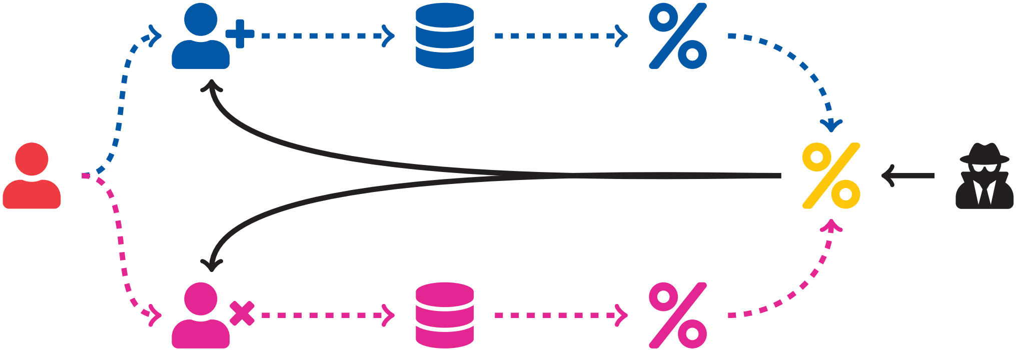

Differential Privacy: Mathematical Statement

We say that an estimator, $\mathcal{M}$, is ϵ-differentially private if for **all** neighbouring datasets, $\mathbb{X}$ and $\mathbb{X}^\dagger$, we have:

$$\frac{P(\mathcal{M}(\mathbb{X}) \in \mathcal{Y})}{P(\mathcal{M}(\mathbb{X}^\dagger) \in \mathcal{Y})} \leq e^{\epsilon}.$$

Private Outcome-Weighted Learning (PrOWL)

We propose a differentially private implementation of OWL, called PrOWL.

1. Approximate the kernel in finite dimensions.

Spicker, D., Moodie, E. E. M., Shortreed, S. M. Differentially private outcome-weighted learning for optimal dynamic treatment regime estimation. Stat. 2023; e641. https://doi.org/10.1002/sta4.641 *(Forthcoming)*.

Private Outcome-Weighted Learning (PrOWL)

We propose a differentially private implementation of OWL, called PrOWL.

1. Approximate the kernel in finite dimensions.
2. Compute the standard OWL estimator.

Spicker, D., Moodie, E. E. M., Shortreed, S. M. Differentially private outcome-weighted learning for optimal dynamic treatment regime estimation. Stat. 2023; e641. https://doi.org/10.1002/sta4.641 *(Forthcoming)*.

Private Outcome-Weighted Learning (PrOWL)

We propose a differentially private implementation of OWL, called PrOWL.

1. Approximate the kernel in finite dimensions.
2. Compute the standard OWL estimator.
3. Perturb the vector with Laplace distributed errors.

Spicker, D., Moodie, E. E. M., Shortreed, S. M. Differentially private outcome-weighted learning for optimal dynamic treatment regime estimation. Stat. 2023; e641. https://doi.org/10.1002/sta4.641 *(Forthcoming)*.

PrOWL Guarantees

- Quantifiable privacy-accuracy tradeoffs.
- Agreement on meaningful treatments w.h.p.
- Agreement on optimal value w.h.p.

Theorem (Noise Requirements for Privacy)

Suppose that we observe a dataset, $\mathbb{X}$, with n observations, such that $|\widetilde{Y}_i| \leq \xi$ (i.e., the modified rewards are bounded). Further, suppose that propensity scores are estimated with bounded sensitivity, $\|\pi(\mathbf{x}, \alpha) - \pi(\mathbf{x}, \alpha')\|_\infty \leq \zeta$, and are such that the estimated $P(A = 1|\mathbf{x}; \alpha) \in (c_L, c_H)$. Take ℓ to be an L-Lipschitz loss function, which is convex. Under regularity conditions, using kernel K, and loss ℓ, **the proposed private-WSVM run on $\mathbb{X}$ with kernel K is ϵ-DP, provided the noise parameter λ is such that**

$$\boldsymbol{\lambda \gtrsim \frac{C\xi\kappa\sqrt{F}}{\epsilon n}.}$$

Theorem (Clinically Meaningful Accuracy)

Suppose that we observe a dataset, $\mathbb{X}$, with n observations, and an F-dimensional feature mapping, $\varphi(\cdot)$. **Define an indifference parameter $\Delta > 0$, take $\beta \in (0.5, 1)$**, *and consider PrOWL with noise level $\lambda \lesssim \frac{-\Delta}{\log(2(1-\sqrt{\beta}))}$. Under certain regularity conditions,* **there is agreement between OWL and PrOWL with probability at least β for all individuals with a true effect size greater than Δ.**

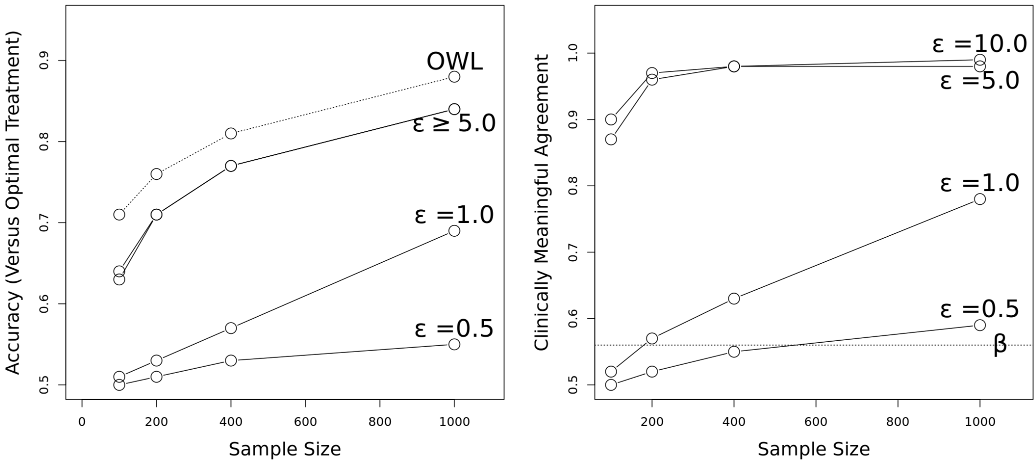

Accuracy (Versus Optimal Treatment)
0.5
0.6
0.7
0.8
0.9
0
200
400
600
800
1000
Sample Size
OWL
ε ≥ 5.0
ε =1.0
ε =0.5
Clinically Meaningful Agreement
0.5
0.6
0.7
0.8
0.9
1.0
200
400
600
800
1000
Sample Size
ε =10.0
ε =5.0
ε =1.0
ε =0.5
β

Summary

Privacy should be a major concern within precision medicine and beyond.

Differential privacy provides one framework for addressing these concerns, with promising results thus far.

Thank You!

www.dylanspicker.com | dylan.spicker@unb.ca