

Improving the detection of behavioral health conditions through positive and unlabeled learning: self-harm and opioid use disorder

Praveen Kumar, PhD¹; Jonathan K. Tsosie¹; Christophe G. Lambert, PhD¹

¹Division of Translational Informatics, Department of Internal Medicine, University of New Mexico Health Sciences Center, Albuquerque, New Mexico, USA



Abstract

Accurate detection and estimation of behavioral health conditions, such as self-harm and opioid use disorder (OUD), is crucial for identifying at-risk individuals, determining treatment needs, tracking prevention and intervention efforts, and finding treatment-naïve individuals for clinical trials. Despite the underdiagnosis and undercoding of these conditions in electronic health records (EHRs), our work aims to accurately estimate both the probability of a given patient having these conditions and the overall population prevalence.

We have developed a novel machine learning algorithm, “Positive Unlabeled Learning Selected Not At Random (PULSNAR)”, to estimate the prevalence of undiagnosed or unrecorded behavioral health conditions. Positive unlabeled learning differentiates between labeled positive instances and a mix of positive and negative instances (unlabeled). Our algorithm addresses the limitations of traditional methods, which do not accurately reflect the true prevalence of behavioral health conditions due to the fact that known, coded cases are not representative of undetected cases. Cases are generally not selected at random, for example, because more serious cases are more likely to generate a healthcare encounter.

In a study of 6,037,479 commercially insured patients with major mental illness (MMI) and 1,329,120 veterans, our PULSNAR algorithm estimates 3.97% visit-level self-harm among patients with MMI and 10.46% lifetime self-harm among Veterans, compared to the 0.453% and 1.85% coded in their EHR data, respectively. Chart review of 97 unlabeled individuals among the Veteran population confirmed that PULSNAR provides well-calibrated classification.

In a study of 1,000,000 patients with at least one opioid prescription fill, PULSNAR estimated 5.3% (53,144) of patients have OUD, compared to the 2.0% (20,079) that have a recorded diagnosis of OUD.

PULSNAR accurately estimates the prevalence of undiagnosed/unrecorded behavioral health conditions, including self-harm and OUD. This has the potential to inform public health, guide screening efforts, identify health disparities, and reduce the negative impacts of these conditions.

Acknowledgment: Research reported in this poster was supported by the National Institute of Mental Health of the National Institutes of Health under award numbers R01MH129764 and R56MH120826.

Background

Opioid use disorder (OUD) is a chronic behavioral health condition marked by prolonged opioid use that leads to significant distress or impairment of brain structure and function¹. The opioid crisis continues to be a significant public health problem worldwide². More than 16 million individuals globally, including over 2.1 million in the United States³, are impacted by opioid use disorders. In 2020, 91,799 drug overdose deaths occurred in the US, with opioids contributing to 74.8% of all those deaths^{4,5}.

The accurate detection and estimation of behavioral health conditions are essential for identifying at-risk individuals and determining treatment needs. However, underdiagnosis and undercoding of these conditions in EHRs are common^{6,7}. We developed novel Positive and Unlabeled (PU) learning algorithms, “Positive Unlabeled Learning Selected Completely At Random” (PULSCAR) and “Positive Unlabeled Learning Selected Not At Random” (PULSNAR).⁸ These algorithms aim to estimate the proportion of cases among unlabeled (undiagnosed and/or uncoded) samples. PULSNAR addresses the shortcomings of traditional PU learning methods, which rely on the assumption that *cases are selected completely at random* (SCAR). This assumption is often not valid in patient healthcare data, for instance because more severe cases are more likely to be diagnosed (among other sources of bias). Currently under review, and available in preprint form, is a full manuscript detailing our PULSNAR and PULSCAR methods and their performance on simulated and benchmark datasets.⁸ Figure 1 and 2 gives some intuition behind the approach, and in this poster we present unpublished work focusing on its application to imputing self-harm and opioid use disorder.

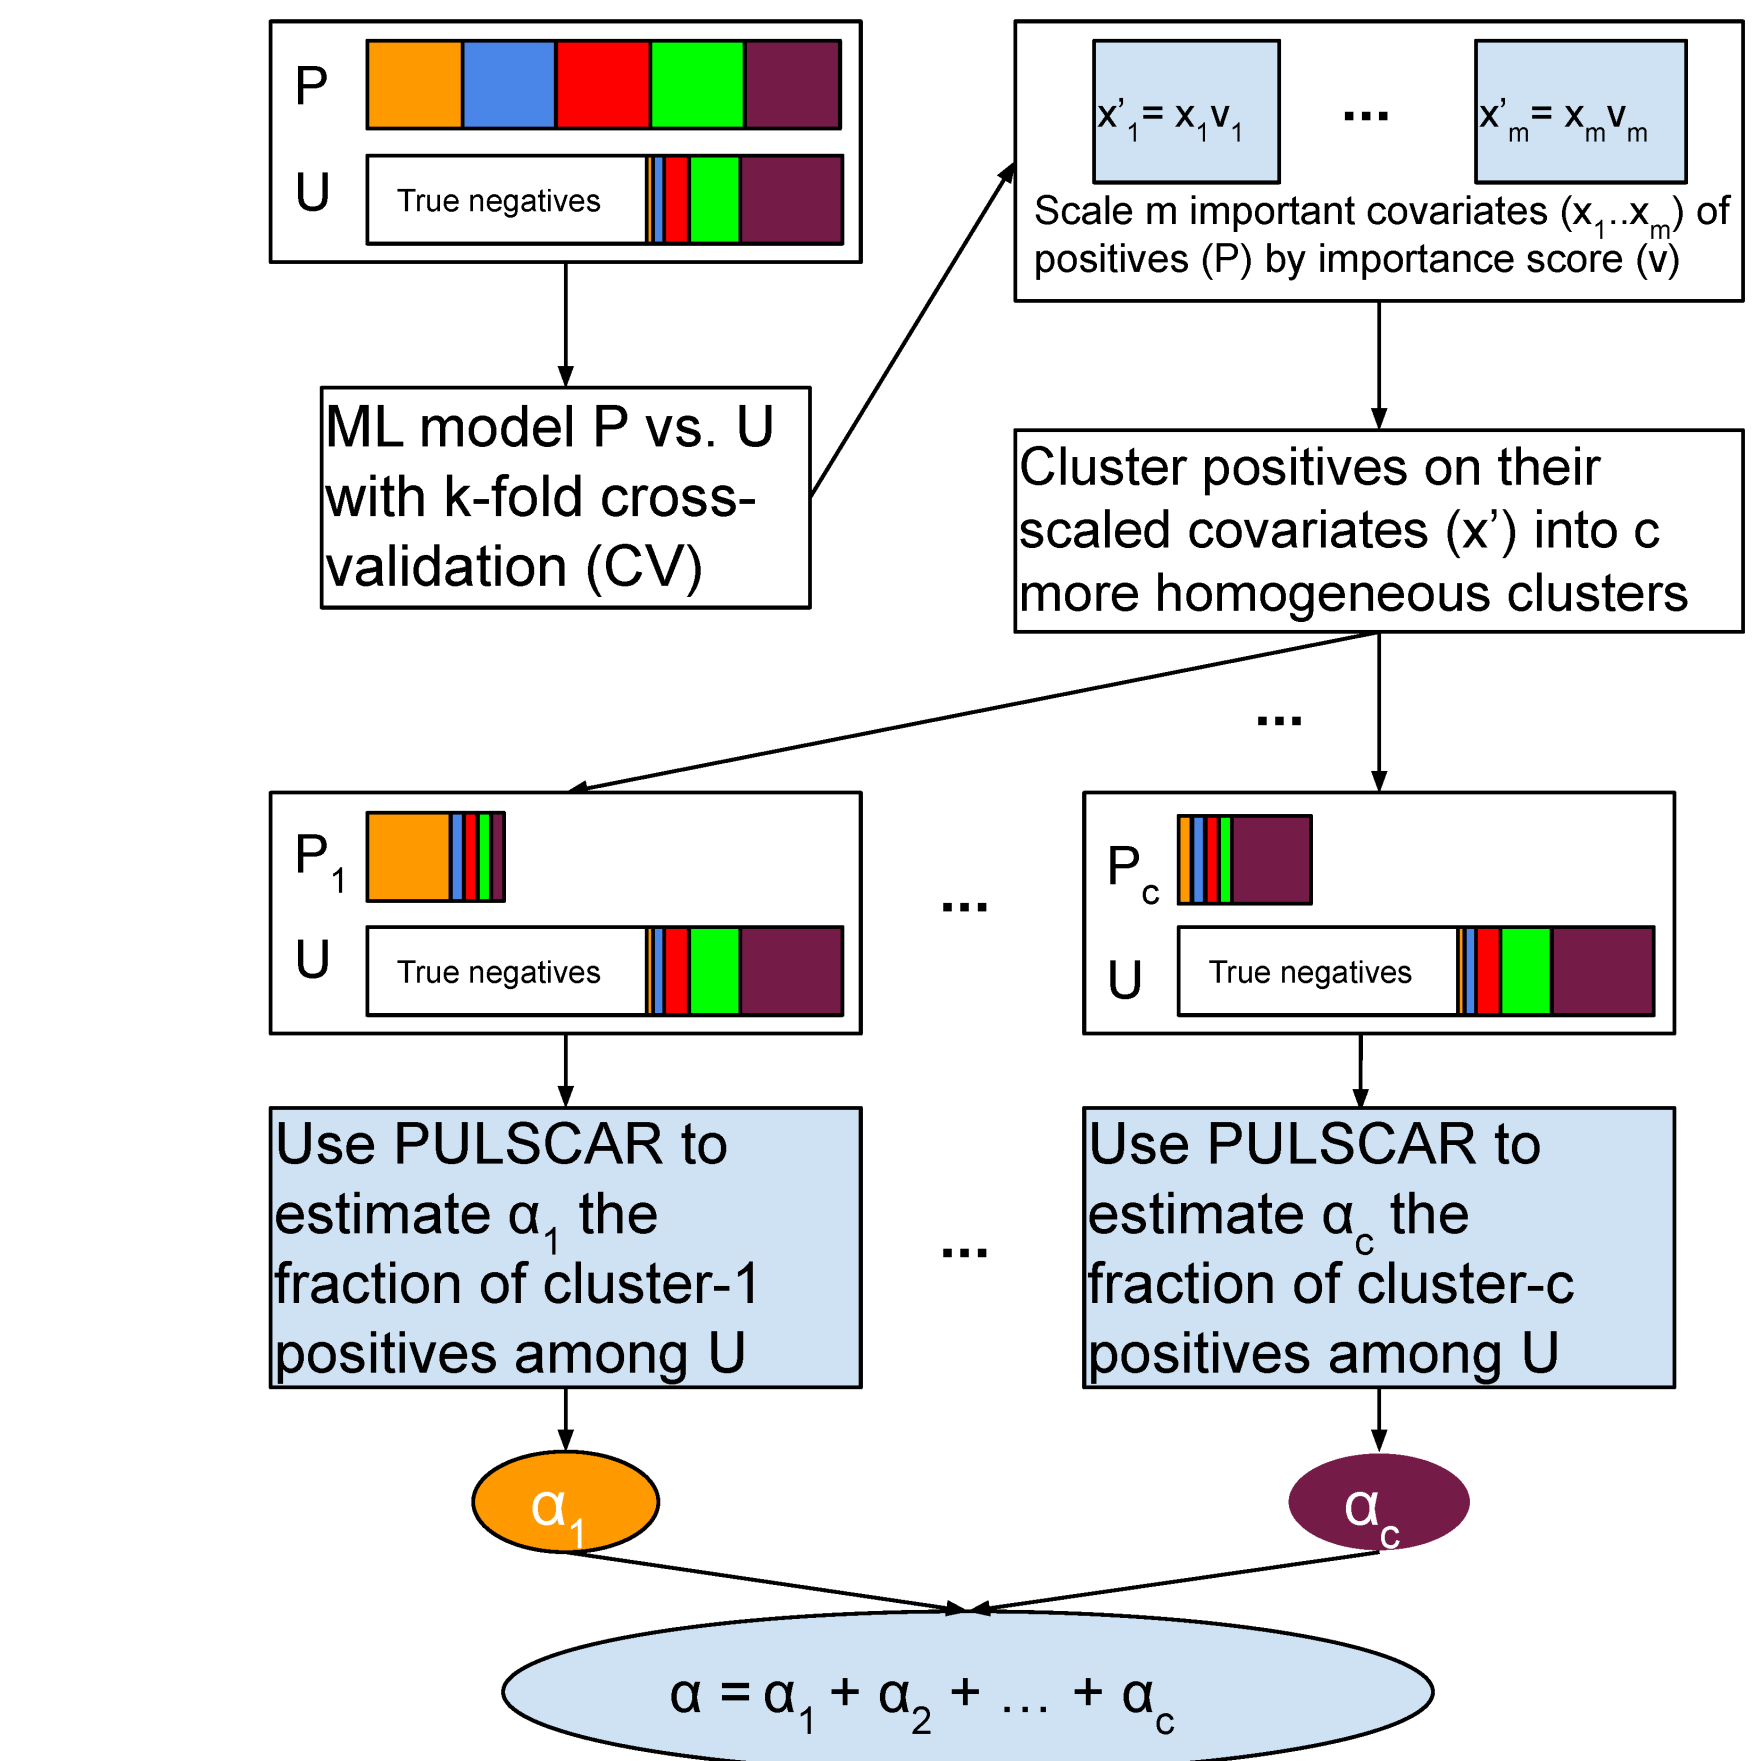


Figure 1. Schematic of PULSNAR algorithm. An ML model is trained and tested with 5-fold cross-validation (CV) on all positive and unlabeled examples. The important covariates that the model used are scaled by their importance value. Positives are divided into c clusters using the scaled important covariates. c ML models are trained and tested with 5-fold CV on the records from a cluster and all unlabeled records. We estimate the proportions ($\alpha_1, \dots, \alpha_c$) of each subtype of positives in the unlabeled samples using PULSCAR. The sum of those estimates gives the overall fraction of positive samples in the unlabeled set. P = positive examples, U= Unlabeled examples.

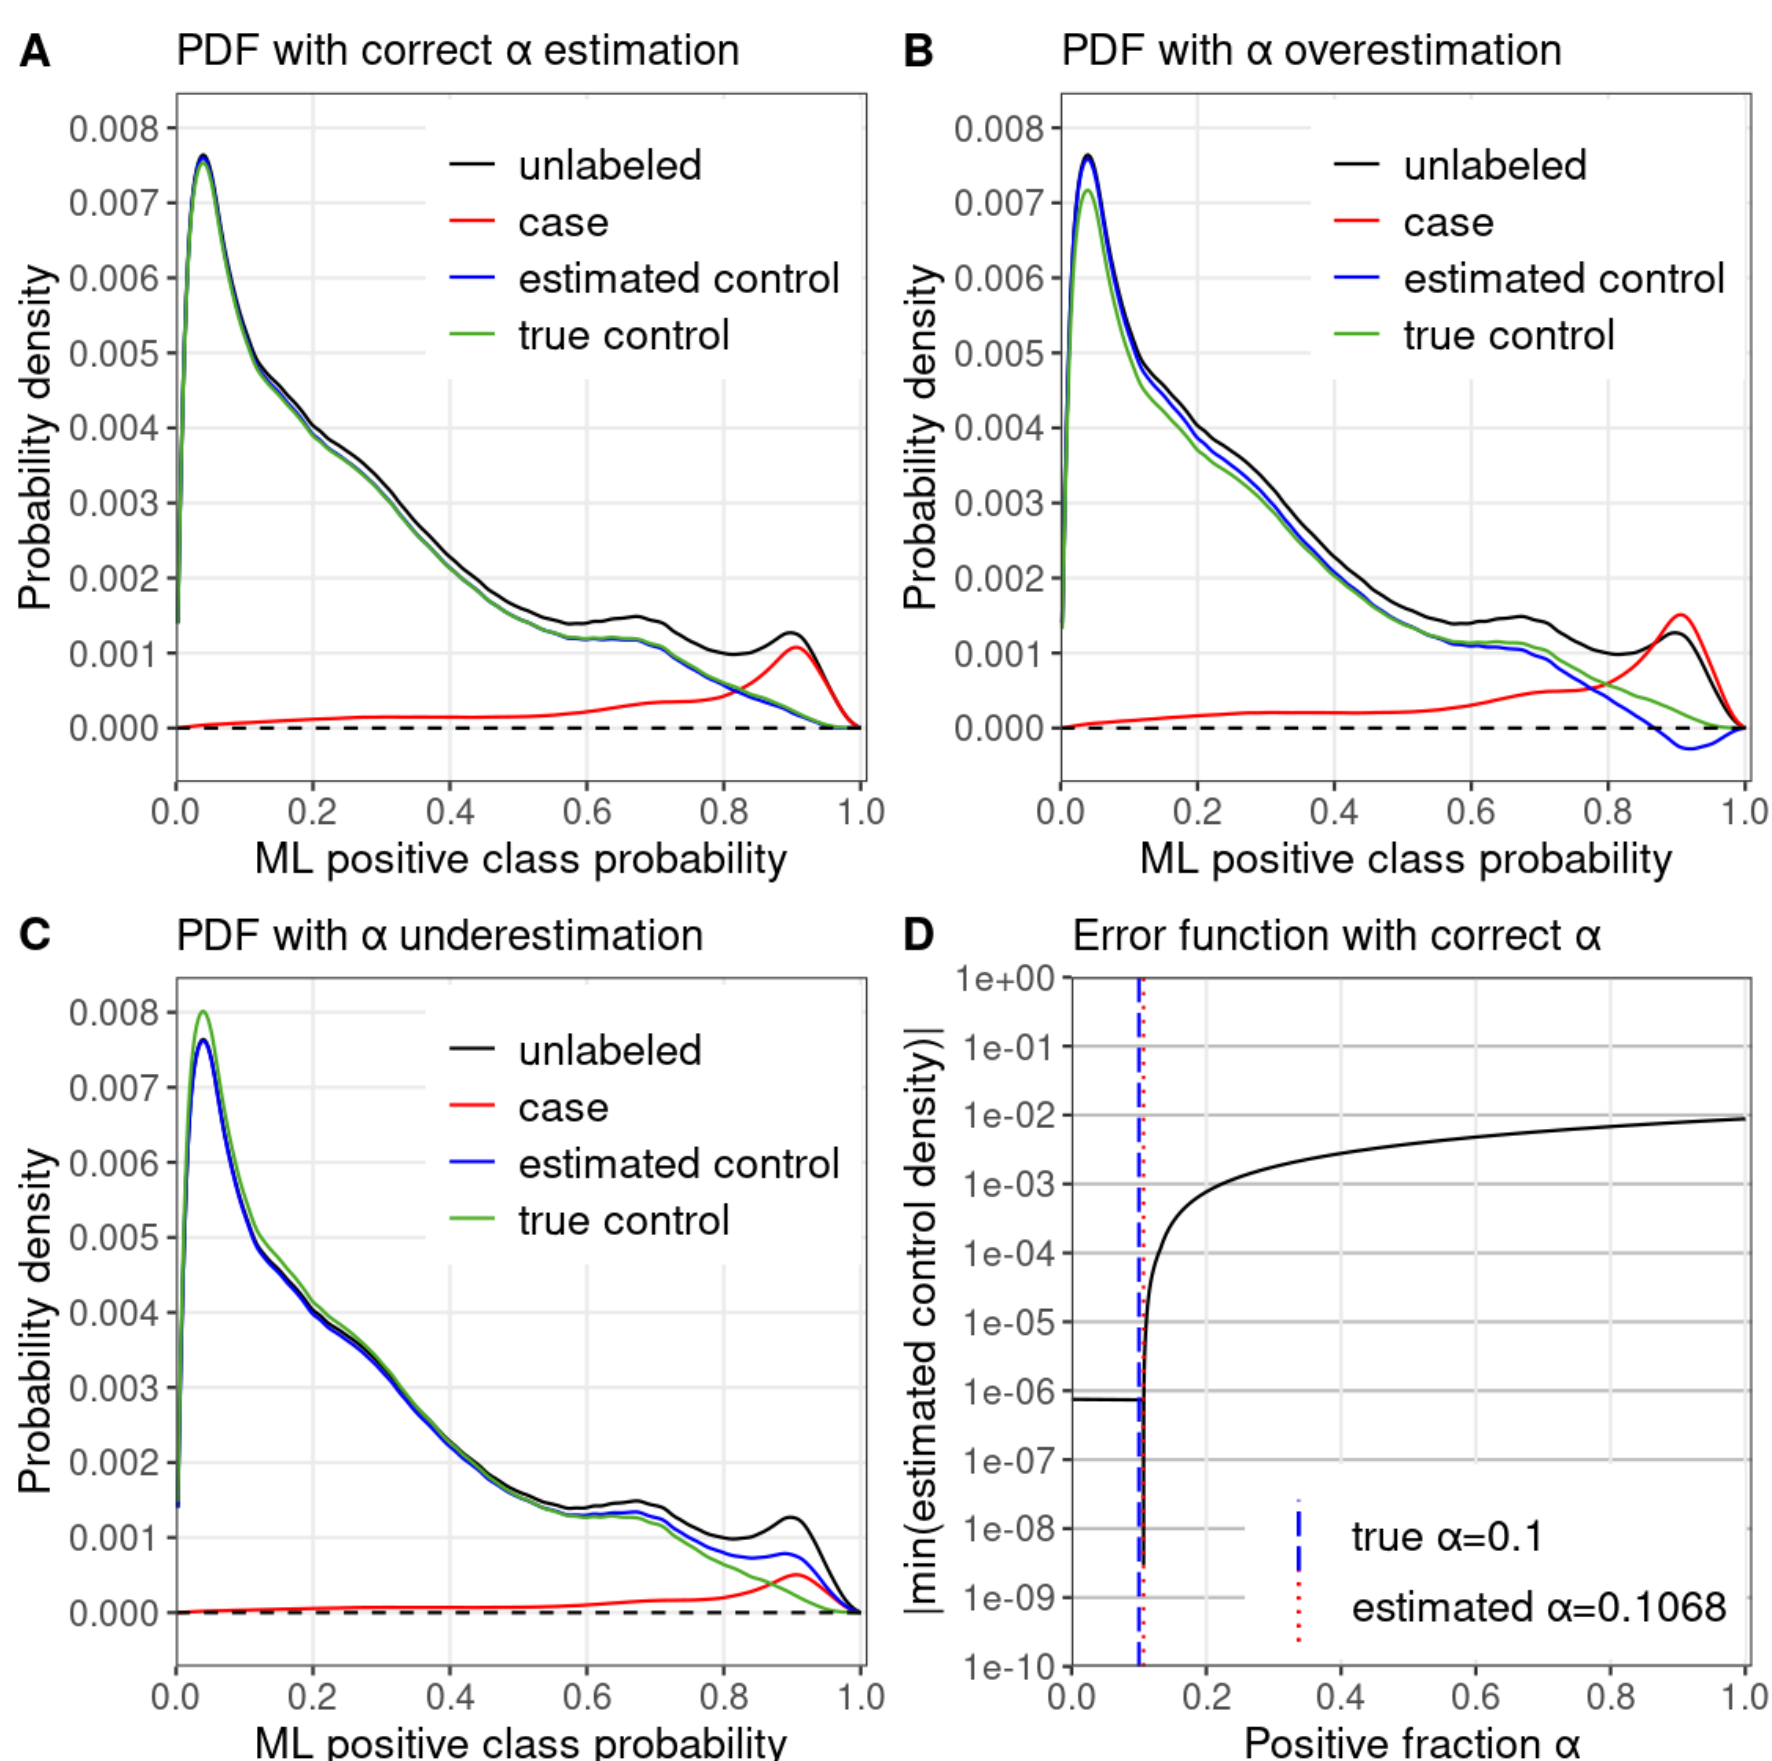


Figure 2: PULSNAR algorithm visual intuition. PULSCAR finds the smallest α such that $f_u(x) - \alpha f_p(x)$ is everywhere positive for x in $[0, \dots, 1]$. A) Kernel density estimates for simulated data with $\alpha = 10\%$ positives in the unlabeled set – estimated negative density (blue) nearly exactly equals the ground truth (green). B) Overweighting the positive density by $\alpha = 15\%$ results in the estimated negative density (blue), $f_u(x) - \alpha f_p(x)$ dropping below zero. C) Underweighting the positive density by $\alpha = 5\%$ results in the estimated negative density (blue) being higher than the ground truth (green). D) Error function with estimated $\alpha = 10.68\%$ selected where the change is largest - very close to ground truth $\alpha = 10\%$

Materials and Methods

Data sources: i) US Department of Veterans Affairs (VA) national EHR patient data for the self-harm analysis (2000-2020); and ii) IBM Health Analytics MarketScan commercial claims and encounters database for both self-harm and OUD analyses (2003-2019).

MarketScan Self-harm: We used emergency room/inpatient visit-level data of 6,037,479 unique patients with ≥ 2 diagnostic codes for major mental illness (MMI) during the observation period, resulting in 20,783,244 visits. See Table 2 for visit attributes.

VA Self-harm: We used person-level data instead of visit-level data. A total of 1,329,120 distinct patients were selected with their attributes (presence or absence of non-self-harm conditions, procedures, and observations) based on their interactions with the VA over their enrolled period.

Coded self-harm was defined by the presence of ICD-10-Clinical Modification (ICD-10-CM) and ICD-9-CM codes (e.g., E95{0-9}*, X7{1-9}*, etc.).

MarketScan Opioid use disorder: We selected a random sample of 1,000,000 individuals who had been exposed to at least one of 36 opioids during their coverage period (e.g., morphine, oxycodone, fentanyl, etc.). An OUD phenotype was defined by the presence of at least one of the ICD-10-CM codes F11*, T40.2*, T50.7* or the ICD-9-CM codes 304.0*, 304.7*, 305.5* over the covered period for each patient. Attributes were conditions and medications.

Analyses: Self-harm data were processed using both PULSCAR and PULSNAR, while only PULSNAR was used for the OUD data. To address class imbalance in both datasets, we created k balanced datasets by including all labeled positives and an equal number of unlabeled examples, where $k = \lfloor \text{unlabeled} / \text{positive} \rfloor$. The PU learning algorithms were assessed in terms of estimating the proportion of phenotypes (self-harm and OUD) in the uncoded records. We also conducted a chart review of 97 unlabeled individuals within the veteran population, selecting one from each probability bin in the range $[0, 0.01, \dots, 1]$. For OUD we assessed both the coded and estimated OUD fraction per state, as well as the fraction coded out of the estimated.

Results

- In a study of 1,000,000 commercially insured patients with at least one opioid prescription fill, PULSNAR estimated 5.3% (53,144) of patients have OUD, compared to the 2.0% (20,079) that have a recorded diagnosis of OUD.
- 37.5% of females with OUD had it diagnosed vs. 43.0% of males.
- Less than 20% of self-harm is recorded in EHR data

Concept name	Domain id
naloxone	Drug
Chronic pain	Condition
Chronic pain syndrome	Condition
buprenorphine	Drug
Drug-related disorder	Condition
Mental disorder	Condition
Drug withdrawal	Condition
Backache	Condition
Disorder of back	Condition
Low back pain	Condition
Mood disorder	Condition
Psychoactive substance-induced organic mental disorder	Condition
Substance abuse	Condition
Drug abuse	Condition
Hypnotic or anxiolytic dependence	Condition

Table 1: Top 15 covariates used by the OUD positive unlabeled learning model.

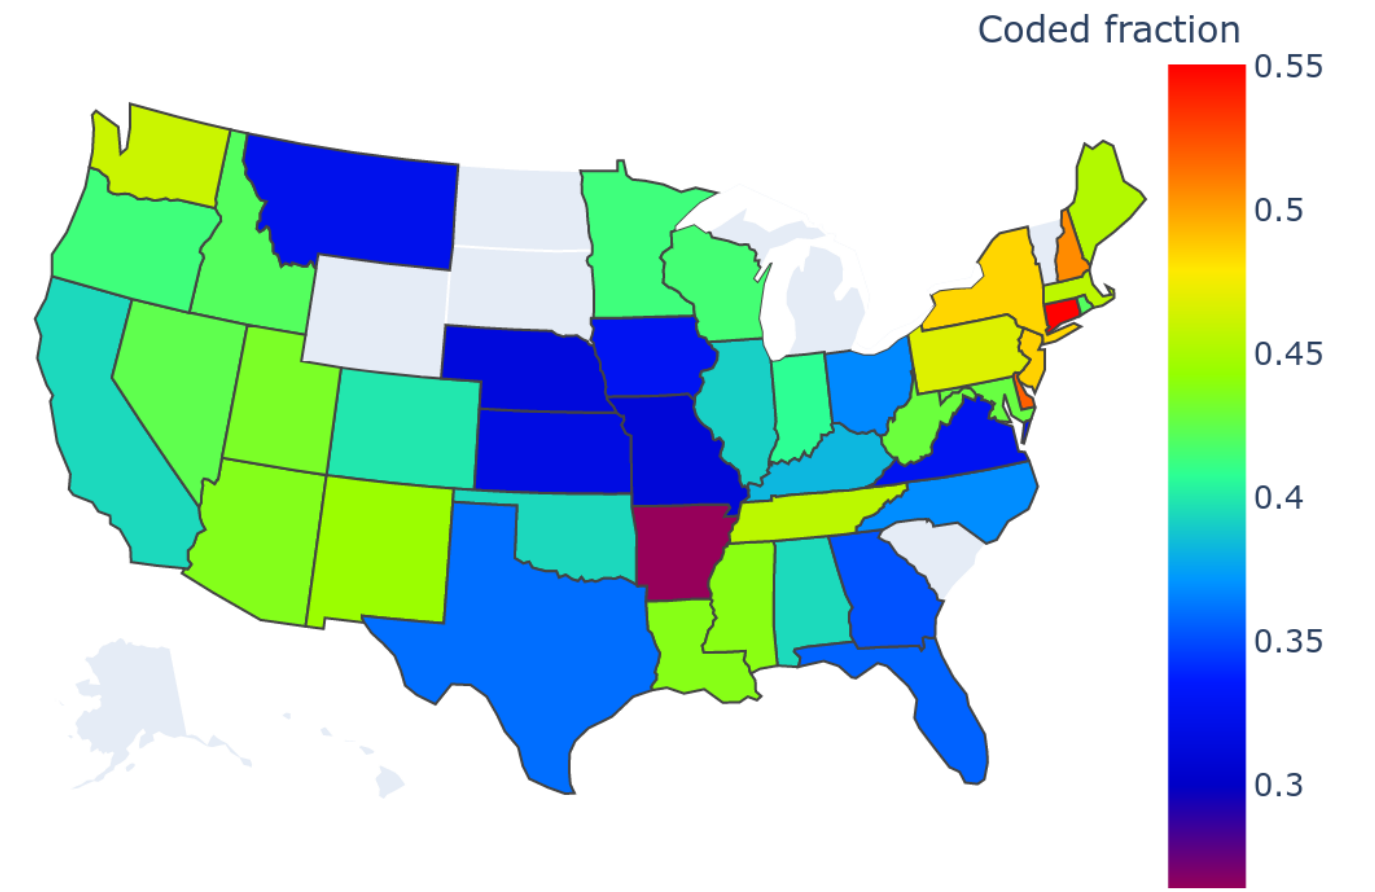


Figure 3: Fraction of coded opioid use disorder by state. Due to MarketScan license restrictions, data for South Carolina were excluded from the figure. Also, data for states PR, HI, VT, ND, DC, AK, WY, and SD were not included due to smaller sample size. Coded fraction=coded/(coded+imputed). State-level diagnosis of OUD ranges from 26.4-55.0%.

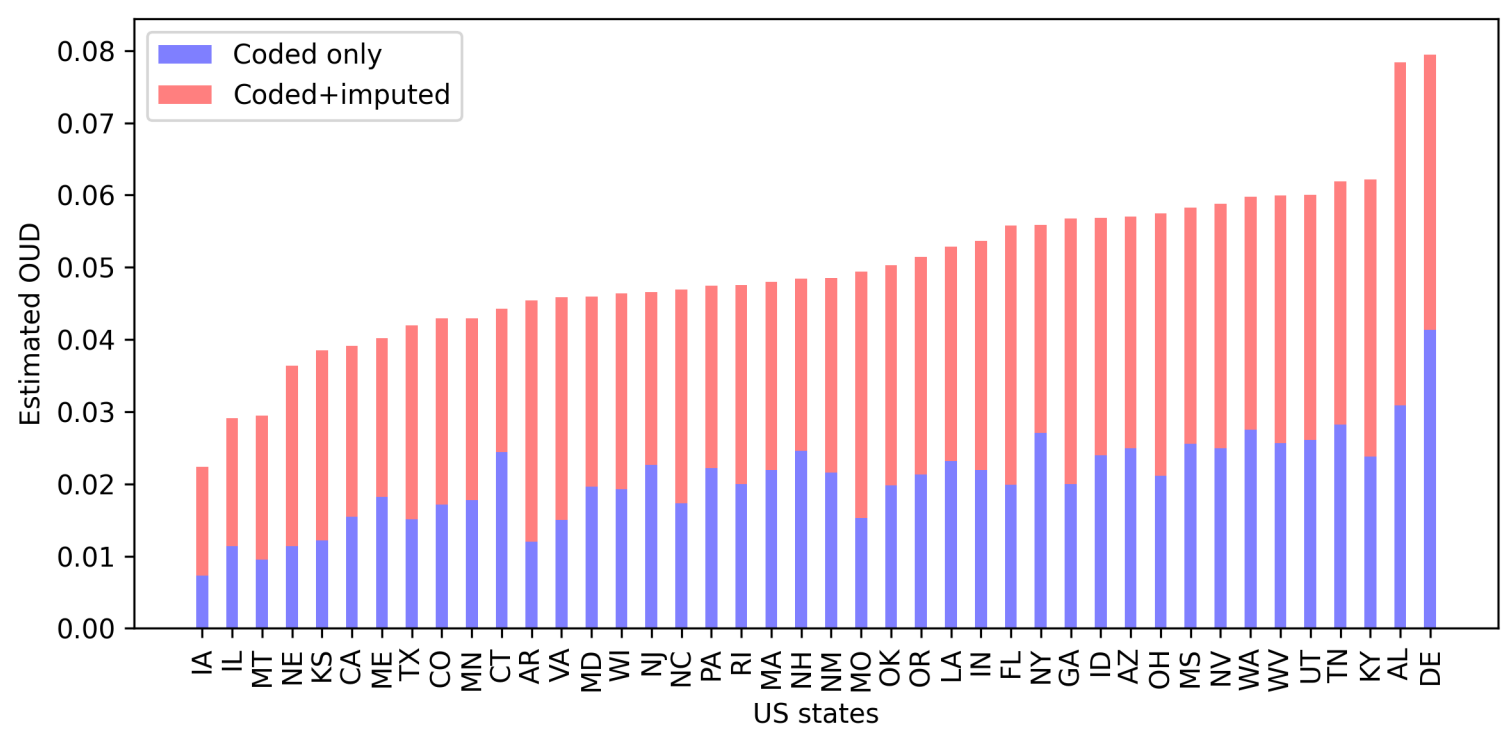


Figure 4: Estimated OUD among opioid users (ever). Coded plus imputed OUD fraction among those who have at least one opioid prescription fill ranges from 2.2%-7.9% across US states. Some states were excluded as described in Figure 3.

	MarketScan (visit-level self-harm)	VA (lifetime self-harm)
# of records	20,783,244	1,329,120
Coded self-harm	94,475 (0.454%)	24,625 (1.85%)
# of covariates	189,294	159,049
Covariate types	7 (procedure, condition, drug, device, observation, measurement, and ancestors)	3 (Condition, procedure, observation)
# of important covariates	2164	1302
α using PULSCAR	0.72%	1.65%
α using PULSNAR	3.53%	8.77%
coded+imputed self-harm using PULSCAR	1.17%	3.47%
coded+imputed self-harm using PULSNAR	3.97% (15 clusters)	10.46% (14 clusters)

Table 2: Self-harm imputation using PULSCAR and PULSNAR on MarketScan & VA data.

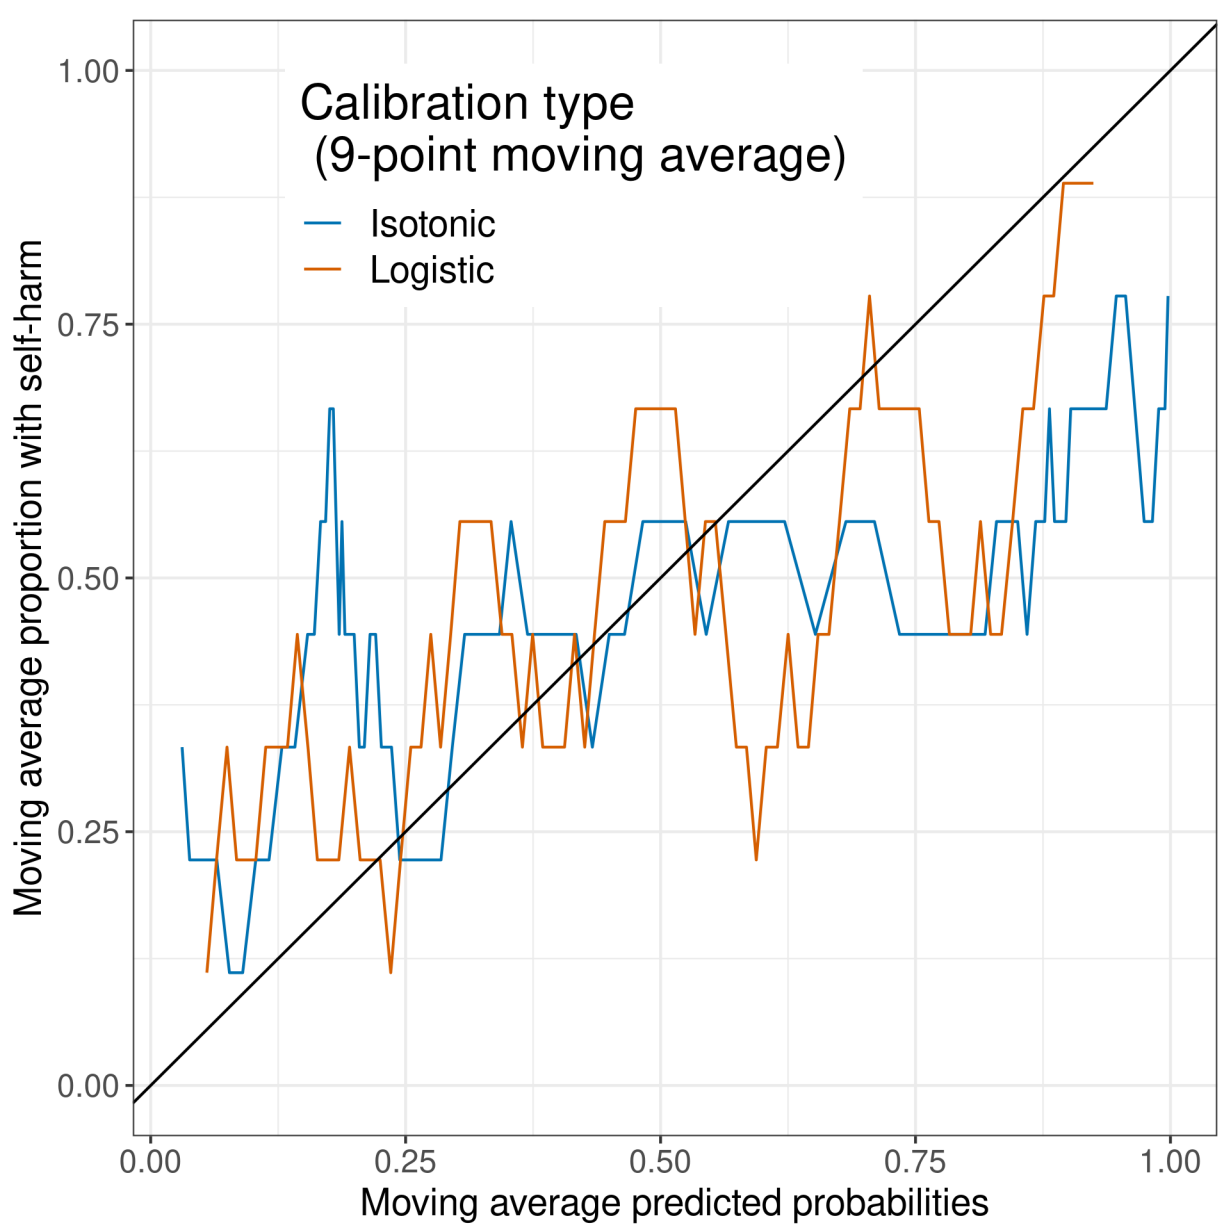


Figure 5: VA self-harm: calibration curve using 9-point moving average. The curves are based on the calibrated probabilities of 97 unlabeled individuals (no coded history of self-harm) and their true labels identified using chart review of patient notes. Out of the 97, our model projected that 48 would have uncoded self-harm, and 46 individuals were confirmed to have self-harm. *blue*: probabilities calibrated using isotonic regression, *red*: probabilities calibrated using sigmoid regression.

Conclusions

Accurately estimating the prevalence of undiagnosed/unreported behavioral health conditions, such as self-harm and OUD, can have significant implications for public health, screening efforts, identifying health disparities, and mitigating the negative impacts of these conditions. The PULSNAR algorithm has demonstrated its value in estimating the proportion of underreported phenotypes. Notably, OUD is more likely missed in females than males. It is also sobering that out of 1M randomly selected individuals across the US with some exposure to opioids, 2% have a coded OUD diagnosis, and an estimated 3.3% have unrecognized OUD, for a total of 1 in 19 people exposed to opioids. Exploration of why different states are better at detecting and diagnosing OUD (Figure 3) could have important public health implications. Detecting the 80+% of Veterans with uncoded self-harm could guide proactive intervention efforts.

References

- Dydyk AM, Jain NK, Gupta M. Opioid use disorder. InStatPearls [Internet] 2022 Jun 21. StatPearls Publishing.
- Leung, K., Xu, E., Rosic, T., Worster, A., Thabane, L. and Samaan, Z., 2021. Sensitivity and specificity of self-reported psychiatric diagnoses amongst patients treated for opioid use disorder. BMC psychiatry, 21, pp.1-8.
- National Academies of Sciences, Engineering, and Medicine, 2019. Medications for opioid use disorder save lives. National Academies Press.
- Death Rate Maps & Graphs <https://www.cdc.gov/drugoverdose/deaths/index.html>
- Hedegaard, H., Miniño, A.M., Spencer, M.R. and Warner, M., 2021. Drug overdose deaths in the United States, 1999–2020.
- Haight, S.C., Ko, J.Y., Tong, V.T., Bohm, M.K. and Callaghan, W.M., 2018. Opioid use disorder documented at delivery hospitalization—United States, 1999–2014. Morbidity and Mortality Weekly Report, 67(31), p.845.
- Kumar, P., Nestsiarovich, A., Nelson, S.J., Kerner, B., Perkins, D.J. and Lambert, C.G., 2020. Imputation and characterization of uncoded self-harm in major mental illness using machine learning. Journal of the American Medical Informatics Association, 27(1), pp.136-146.
- Kumar P and Lambert CG (2023) PULSNAR – Positive unlabeled learning selected not at random: class proportion estimation when the SCAR assumption does not hold. arXiv [cs.LG]. Available at: <http://arxiv.org/abs/2303.08269>.