

Pharmacokinetics-Informed Neural Network for Predicting Opioid Administration Moments with Wearable Sensors

Bhanu Teja Gullapalli¹, Stephanie Carreiro², Brittany P Chapman², Eric L Garland³, Tauhidur Rahman¹

¹University of California San Diego

²UMass Chan Medical School

³University of Utah

bgullapalli@ucsd.edu, stephanie.carreiro@umassmemorial.org, brittany.chapman@umassmed.edu, eric.garland@socwk.utah.edu, trahman@ucsd.edu

Abstract

Long-term and high-dose prescription opioid use places individuals at risk for opioid misuse, opioid use disorder (OUD), and overdose. Existing methods for monitoring opioid use and detecting misuse rely on self-reports, which are prone to reporting bias, and toxicology testing, which may be infeasible in outpatient settings. Although wearable technologies for monitoring day-to-day health metrics have gained significant traction in recent years due to their ease of use, flexibility, and advancements in sensor technology, their application within the opioid use space remains underexplored. In the current work, we demonstrate that oral opioid administrations can be detected using physiological signals collected from a wrist sensor. More importantly, we show that models informed by opioid pharmacokinetics increase reliability in predicting the timing of opioid administrations. Forty-two individuals who were prescribed opioids as a part of their medical treatment in-hospital and after discharge were enrolled. Participants wore a wrist sensor throughout the study, while opioid administrations were tracked using electronic medical records and self-reports. We collected 1,983 hours of sensor data containing 187 opioid administrations from the inpatient setting and 927 hours of sensor data containing 40 opioid administrations from the outpatient setting. We demonstrate that a self-supervised pre-trained model, capable of learning the canonical time series of plasma concentration of the drug derived from opioid pharmacokinetics, can reliably detect opioid administration in both settings. Our work suggests the potential of pharmacokinetic-informed, data-driven models to objectively detect opioid use in daily life.

Introduction

Opioid use and overdoses continue to fuel the opioid epidemic in the United States. In 2019, the National Survey on Drug Use and Health (NSDUH) showed that approximately 10.1 million United States residents aged 12 years or older misused opioids, with the majority (9.3 million) citing prescription opioid misuse. This number has further increased during the COVID-19 pandemic ((Khatri and Perrone 2020)), highlighting the need for immediate action. While opioids can provide temporary pain relief, frequent and repetitive opioid use can be considered an early indicator of opioid misuse and developing opioid use disorder.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Current methods for passively monitoring opioid use including subjective self-reports (based on the Diagnostic and Statistical Manual (DSM-5) criteria (Hasin et al. 2012)) and toxicology testing are susceptible to recall bias, intentional concealment, and short detection periods. Diagnostic methods that offer objective and real-time measures of opioid use are urgently needed. Such an automatic and passive opioid use tracking system can also help to detect treatment adherence and opioid relapse. Different digital health technologies including wearables and smartphone applications have been developed for OUD monitoring and treatment (Senyurek et al. 2020; Natarajan et al. 2013). In a recent work (Gullapalli et al. 2021), physiological signals continuously and passively collected from a wrist sensor was used to build a purely data-driven model with supervised learning to detect intravenous (IV) opioid administrations in a clinical setting. While the work showed the promise in using wearable devices for opioid use detection, the reliability of these approaches for oral opioid administrations in inpatient or outpatient settings remains unknown.

Moreover, models built using a purely data-driven approach, without learning any underlying pharmacological mechanisms of opioids, limit the generalizability of the model to different types of opioids. One candidate group of properties intrinsic to opioids that may be very useful to inform models are the drugs' pharmacokinetic (PK) parameters. Pharmacokinetics is the branch of pharmacology concerned with how the body handles a drug once it is administered (Gibaldi, Perrier et al. 1982). This includes the time it takes for the drug to be absorbed, metabolized, distributed, and eliminated. The plasma concentration of a drug defined as the amount of drug present in each volume of a blood plasma, changes over time in a predictable fashion.

Plasma drug concentration as a function of time for many drugs (including the opioids considered in our study) is described by a single-compartment model (Equation 1). To the best of our knowledge, models combining Opioid pharmacokinetic knowledge with wearable sensor-derived physiologic data to detect opioid use have not been proposed previously. Our current work aims to leverage pharmacokinetics to improve our opioid detection model. Specifically, we aim to develop a pharmacokinetics-informed machine learning framework that jointly learns the pharmacokinetics of opi-

oids while estimating opioid administration moments using multimodal physiological signals collected from a wearable wrist sensor in inpatient and outpatient settings.

Building data-driven models in healthcare is often challenging due to the scarcity of data and the costs associated with annotating it. These challenges become even more pronounced in substance use research, where privacy concerns and social stigma can lead people to conceal or withhold information about their substance use or provide inaccurate self-reports of their usage patterns. In our current work, we at first take advantage of the vast unlabeled wearable data and a simple self-supervised learning task to teach the model how to detect sudden changes in the multimodal wearable signals. We then teach the model opioid pharmacokinetics with the help of a canonical single-compartment model (of opioids) while the model is also tasked to detect opioid use and predict the precise timing of use with the input multimodal sensor data. In this paper, we will demonstrate that the proposed Pharmacokinetics-informed model training approach outperforms a purely data-driven supervised learning approach and can generate much improved predictions of opioid use moments.

Related Work

Sensor-based applications to opioid use have focused on the detection of physiology surrounding opioid use, including therapeutic effects, overdose, and withdrawal. A recent study demonstrated it is possible to predict stress and drug craving ninety minutes in the future with passively collected GPS data (Epstein et al. 2020). Previous work has demonstrated the ability to detect the use of prescription opioids administered by mouth and intravenously in an inpatient setting using physiological signals from a wrist sensor (Mahmud et al. 2018; Gullapalli et al. 2021), and has explored the impact of individual-level features on model performance (Chapman et al. 2022). Other investigators have leveraged similar non-invasive sensors from wristwatches and mobile devices to detect the most dangerous complication of opioid use- overdose, by monitoring respiratory rate with mobile phones and trunk-mounted sensors (Nandakumar, Gollakota, and Sunshine 2019; Roth et al. 2021). Another study developed and validated an automated risk-modeling framework to predict opioid abstinence and medication adherence with machine-learning algorithms (Burgess-Hull et al. 2002). While these studies demonstrate the potential of ubiquitous sensors in the space of opioid use monitoring, they lack any pharmacokinetic information on how opioid concentration changes over time. Our knowledge of the biochemical behavior of opioids in the human body can help provide supplemental data about the timing and intensity of expected effects: harmonizing this data with real-time physiologic data can lead to improved precision of our predictions and smarter diagnostic tools.

Our pharmacokinetics-informed neural network is inspired by physics-informed neural networks (PINNs) which were first introduced in (Raissi, Perdikaris, and Karniadakis 2019) to solve partial differential equations (PDEs) of various physical phenomena in fluids (Zhu, Liu, and Yan 2021).

Traditional deep learning methods to solve PDEs are governed by a data-driven approach that does not use any underlying physical characteristics of the problem to approximate the partial differential equations. PINNs overcome this limitation by integrating available information about the equation into the loss function as a residual term (Jin et al. 2021; Pang, Lu, and Karniadakis 2019), thereby guiding the learning. In our work, while training the neural network with supervised learning to learn the pharmacodynamics (i.e., how opioid affects different physiological signals), we use the relative plasma drug concentration equation derived from opioid pharmacokinetics in the loss function. We demonstrate that the Pharmacokinetics-informed model can accurately predict Opioid use moments and can outperform traditional purely data-driven supervised models, which do not have any PK knowledge.

Data Collection

Study Protocol Overview: All study-related procedures were approved by the Institutional Review Board (IRB) of the UMass Chan Medical School. From a total of 42 unique participants, we collected 2,910 hours of wearable data in both inpatient and out-patient settings that contained a total of 227 oral opioid administration events. Eligible participants were hospital patients over 18 years old, receiving opioid analgesics for pain, who spoke English and could wear a wrist device on their non-dominant wrist. We excluded pregnant individuals, prisoners, or those unable to consent. During their hospital stay, participants wore the device continuously removing it only for clinical care or activities like showering. Participants used an event marker button on the device to record opioid administrations, which were cross-checked with Electronic Medical Records for accuracy and additional details.

Upon hospital discharge, participants had the option of continuing to wear the device for up to seven days (i.e., the outpatient portion of the study). Participants removed the device daily during sleep for charging. In addition to annotating all opioid administrations using the device’s event marker button, participants were instructed to self-report these events using either a paper or mobile app-based log. Information regarding outpatient opioid prescriptions (type of opioid, dosing instructions, dose per unit, quantity prescribed) was verified in the electronic medical records. The dose administered during each opioid use event was converted to morphine milligram equivalents (MME), a standardized metric to compare the relative potency of different types of opioids (Stone et al. 2018).

In our current work, we modeled opioid administrations in the inpatient and outpatient settings separately for two main reasons: 1) the distribution of physiological signals tends to vary because participants have more freedom to perform different activities in the outpatient setting than in the inpatient setting; and 2) the precision and accuracy of opioid administration event labeling in the outpatient setting are subject to recall bias and incomplete data, as participants were solely responsible for logging administrations.

Device: The device used for physiologic data collection was the Empatica E4 (Empatica, Milan, Italy). The E4 is a

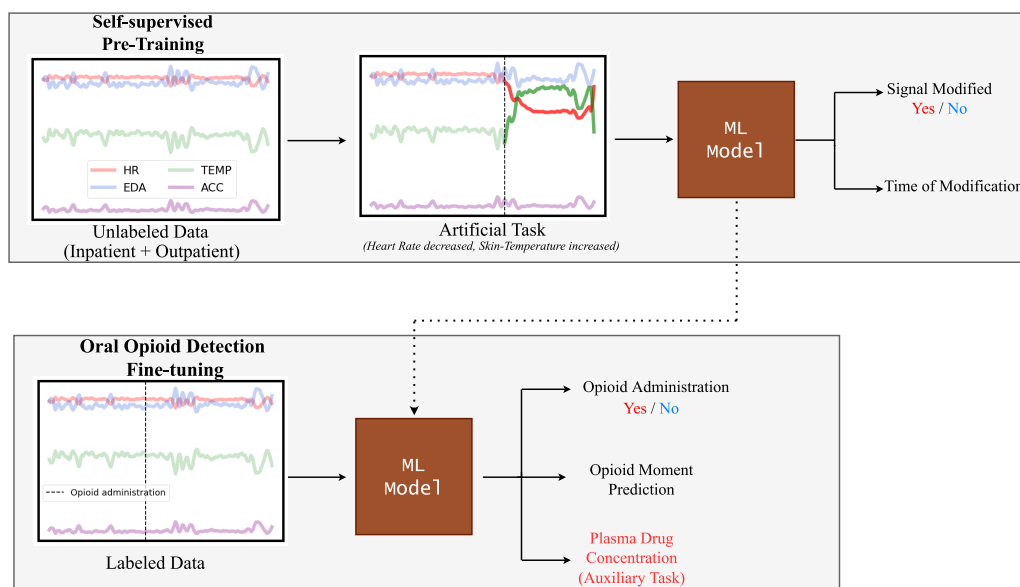


Figure 1: Illustration of our pharmacokinetics-informed approach for monitoring opioid administrations. We first train a machine learning (ML) model with a pretext task of detecting whether and where the signal is modified. The pre-trained model is then fine-tuned to detect opioid administration and predict the moment of administration in the time-window while also learning the plasma drug concentration over time, which is used as an auxiliary task.

research-grade wrist-worn device that continuously and passively captures various physiological signals. It is equipped with several sensors that measure instantaneous heart rate, skin temperature, electrodermal activity, and triaxial acceleration. Raw data are initially stored on the device’s onboard memory and then transferred to Empatica’s secure cloud-based server (Empatica Connect) for further analysis. The raw signals from the wearable were downsampled to the sampling frequency of 1Hz.

Before analyzing, we cleaned the raw data in a two-step process: 1) Periods where the participant was not likely wearing the device were identified and removed. Incompatible skin temperature ($\leq 20^{\circ}\text{C}$), lack of electrodermal activity (value of 0), and abnormally high heart rate (≈ 200 beats-per-minute) were used as indicators to identify these time periods. 2) Motion and noise artifacts were removed by passing all the signals through a fifth-order Butterworth low-pass filter.

Inpatient Data Acquisition: A total of 1,983 hours of E4 data were collected during this portion of the study. Seven participants were excluded from the analysis following data cleaning due to insufficient volume of physiologic data. Of the 35 participants included, 187 oral opioid administrations were recorded, all of which were oxycodone.

Outpatient Data Acquisition: A total of 927 hours of E4 data were collected from the 24 participants who took part in this portion of the study. Forty oral opioid administrations ($n=38$ oxycodone, $n=1$ morphine, $n=1$ tramadol) self-reported by participants were recorded.

Experiment and Results

Figure 2 shows the pipeline used in the current work for pharmacokinetics-informed learning. We at first partition the continuous sensor waveform/timeseries data into fixed-size time-windows using a sliding window mechanism where window length is 225 minutes and window slide is 20 minutes. We develop a model that can take a timeseries window as input and generate two types of output: (i) binary opioid administration/use detection inference (i.e., if the opioid was administered in the input time-window); and (ii) predict the exact moment of Opioid administration/use in the time-window. We train the model in two steps.

1. Self-supervised learning to teach the model how to detect sudden changes in different modalities in the wearable sensor data with a “pretext” task.
2. Pharmacokinetics-informed auxiliary task to teach our ML models Pharmacokinetics (PK) while simultaneously learning from the data.

Step 1: Self-Supervised Learning Framework

Self-supervision for representation learning in drug administration detection has not been explored previously. In our work, we showed the extent of its potential in detecting oral opioid administrations in both inpatient and outpatient settings. We first pre-trained our model using a “pretext” task based on domain expertise to extract useful latent representations from unlabeled sensor data. The design of the “pretext” task has been motivated by statistical findings of previously conducted work (Gullapalli et al. 2021; Carreiro et al. 2016) which demonstrated significant, immediate changes

Setting	Model	SSL	F1-score	Specificity	Sensitivity	NMAE	R^2	
Inpatient	BiLSTM	X	0.65	0.57	0.74	0.24	0.07	
		✓	0.74	0.67	0.82	0.18	0.05	
	LSTM-FCN	X	0.68	0.63	0.77	0.27	0.08	
		✓	0.74	0.66	0.83	0.20	0.05	
	ResNet	X	0.67	0.62	0.71	0.25	0.17	
		✓	0.76	0.71	0.75	0.19	0.21	
	Inception-Time	X	0.70	0.63	0.76	0.21	0.33	
		✓	0.75	0.70	0.74	0.15	0.45	
	CTA-TCN	X	0.73	0.65	0.88	0.15	0.49	
		✓	0.79	0.74	0.78	0.13	0.53	
	Outpatient	BiLSTM	X	0.48	0.35	0.61	0.37	-0.48
			✓	0.54	0.42	0.75	0.36	-0.43
LSTM-FCN		X	0.51	0.42	0.61	0.29	0.06	
		✓	0.57	0.55	0.61	0.19	0.14	
ResNet		X	0.65	0.62	0.69	0.27	0.13	
		✓	0.70	0.66	0.78	0.21	0.15	
Inception-Time		X	0.60	0.62	0.61	0.26	0.05	
		✓	0.69	0.64	0.67	0.17	0.43	
CTA-TCN		X	0.64	0.52	0.80	0.27	-0.64	
		✓	0.70	0.62	0.77	0.18	0.29	

Table 1: Modeling opioid administration with different time-series models for inpatient and outpatient settings with and without self-supervised learning (SSL). F1-score, Specificity, and Sensitivity metrics measure binary opioid use detection, while NMAE and R^2 metrics measure the moment of opioid administration.

in physiological signal time- and frequency-based features post-opioid administration. To replicate this, we randomly selected a moment (start time in minute) in the time-window. We then modified two physiological signals picked at random by either magnifying or reducing the signal with a randomly selected scalar (from a uniform distribution $U(0-0.3)$). Only a fraction of randomly chosen time-windows are modified in this manner.

In the pre-training step, the model has to predict if a signal has been modified in the input time-window or not, and if modified, predict the moment (time-minute) of modification. This process is similar to the downstream task of detecting if an opioid has been administered or not and if detected, predicting the opioid administration moment. We used all the data from inpatient and outpatient settings for the pre-training step. In figure 1, we show an overview of the pre-trained task. We fine-tuned the pre-trained model to detect opioid administration and predict the moment of administration (“opioid moment”) in the downstream setting. We used a hybrid loss function $Loss = \lambda L_{WCE} + (1 - \lambda)L_{KAPPA}$ to solve these tasks. A binary-weighted cross-entropy loss L_{WCE} was used in the hybrid loss to train the model to correctly detect whether the time series signals have been modified. We used the weighted kappa index L_{KAPPA} (Cohen 1968) between the actual and predicted start time of signal modification (in minute) in the time-window.

We trained the self-supervised framework to detect opioid administration and predict opioid moment using different machine learning models that are widely used for time-series modeling including bidirectional LSTM (BiLSTM) (Graves and Schmidhuber 2005), LSTM with fully convolutional network (LSTM-FCN) (Karim et al. 2017), ResNet (Wang, Yan, and Oates 2017), and InceptionTime (Ismail Fawaz

et al. 2020). Additionally, we also used Channel- Temporal Attention TCN (CTA-TCN) (Gullapalli et al. 2021) which was initially designed to detect intravenous (IV) opioid administrations in a hospital-based setting using physiological signals collected from E4 sensor data.

We evaluated the performance of opioid administration detection using weighted F1-score. We used a weighted F1-score instead of a binary F1-score to account for the class imbalance. To evaluate opioid moment prediction, we used normalized mean-absolute error (NMAE). For all analyses, we employed Leave-One-Subject-Out Cross-Validation (LOSOXV). As the data in the inpatient setting was collected in the presence of clinical personnel and consequently less noisy compared to the outpatient data, we present the results separately in the inpatient and outpatient settings. Table 1 highlights the main performance of different time-series models for opioid administration detection and opioid moment prediction, with and without self-supervised pre-training. Self-supervised learning improved both detection and moment prediction for all time-series models across both settings which clearly highlights a broader applicability of the technique. In the inpatient setting, the CTA-TCN model delivered optimal results, achieving an F1-score of 0.79 and an NMAE of 0.13 for moment prediction. In the absence of SSL, the metrics were an F1-score of 0.73 and an NMAE of 0.21, respectively. In the outpatient scenario, the CTA-TCN model again was best, with an F1-score of 0.70 and an NMAE of 0.18. Without the inclusion of SSL, these figures stood at an F1-score of 0.65 and an NMAE of 0.27, respectively. The improved F1-score and the decreased NMAE across all the models imply that self-supervised learning helped to regularize and improve the model.

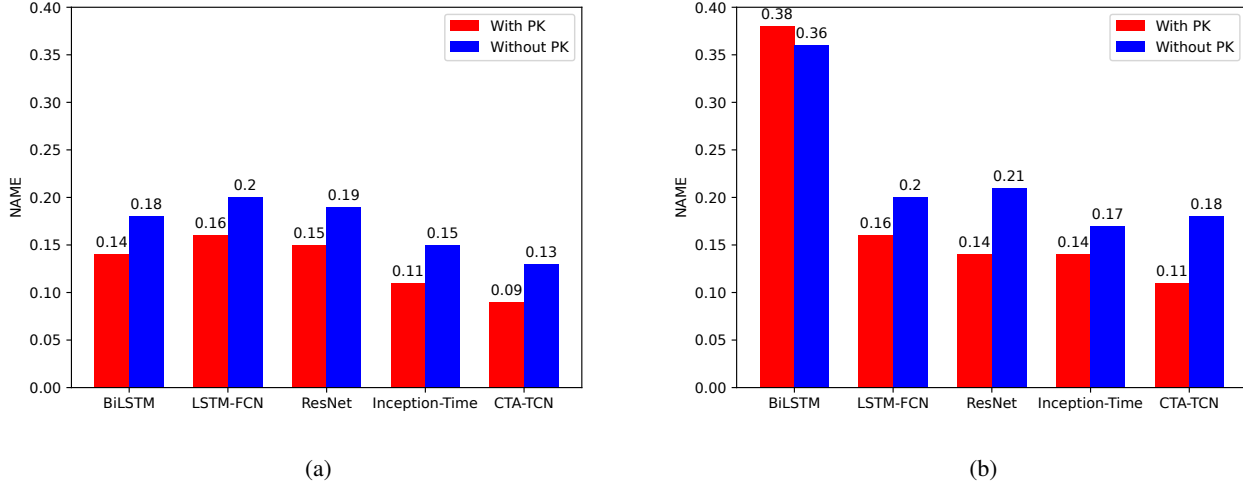


Figure 2: Barplot showing the impact of pharmacokinetics-informed learning on the performance of different Opioid use moment prediction model in both inpatient (left) and outpatient (right) settings. All models are pre-trained with the self-supervised task described in Step 1.

Step 2: Pharmacokinetics-Informed Auxiliary Task

In the previous subsection, we demonstrated how using self-supervised learning (SSL) improved the performance in opioid administration detection and opioid moment prediction across both inpatient and outpatient settings for various time-series models. However, all these predictions are purely data-driven by wearable data. To complement this, we used temporal information associated with the moment of opioid administration while training the model, in addition to the weighted kappa index loss, which only uses a one-hot encoding vector to represent this information.

To that end, we utilized the plasma drug concentration information that changes over time from the moment the drug is administered. Calculating the exact plasma drug concentration at any point in time requires drawing and analyzing a blood sample, which is not a practical solution for a longitudinal study. However, we can estimate a relative plasma concentration (RPC) as a function of time for opioids using a single-compartment model shown in equation 1. In this model, the drug is distributed instantaneously throughout the body, and absorption and elimination occur with a first-order reaction. While the exact plasma drug concentration over time also depends on numerous factors such as the bioavailability of the drug, salt factor, volume of distribution, and renal function, the relative plasma drug concentration over time can be described by equation 1.

$$D_t \propto \overbrace{D_0}^{\text{drug amount administered}} \times (e^{-K_e t} - e^{-K_a t}) \quad (1)$$

Where D_t = plasma drug concentration at time t , and K_e and K_a are the elimination and absorption rate constants that are dependent on the type of the opioid.

$$Loss = \lambda_1 L_{WCE} + \lambda_2 L_{KAPPA} + \lambda_3 L_{RMSE} \quad (2)$$

↑ Opioid moment prediction loss ↑ Auxiliary loss
↑ Opioid detection loss

To calculate the plasma drug concentration over time using the opioid pharmacokinetics equation 1, we used the elimination and absorption rate constants $K_e=0.17$, $K_a=4.19$ for orally ingested oxycodone ((MANDEMA et al. 1996)), as this represents the majority of opioid administrations in the dataset. The auxiliary task aims to supervise intermediate layers and create a robust representation of the input signals that can better predict the moment of opioid administrations. Throughout the paper, computed plasma concentration refers to computed RPC. The new hybrid loss function is now defined by equation 2. For the auxiliary task, we used a root mean square error (RMSE) between RPC over time derived from equation 1 and predicted plasma concentration over time. λ_1, λ_2 , and λ_3 are the hyperparameters used to weigh the importance of the loss terms. The only condition on these was $\lambda_1 + \lambda_2 + \lambda_3 = 1$. We evaluate the model’s performance only on opioid detection and opioid moment prediction, but not on the auxiliary task.

The opioid moment prediction results after using pharmacokinetics information in the auxiliary task are shown in figure 2. Utilizing pharmacokinetics-informed learning reduced NMAE for the majority of the models across both inpatient and outpatient settings. In the inpatient setting, we achieved the best NAME of 0.09 and when using SSL only, we get an NMAE of 0.13 ($\approx 4\%$ increase) and R^2 of 0.53 ($\approx 49\%$ decrease) compared to the model that used both SSL and auxiliary task. This improvement is much more prominent in the outpatient setting, with the best model CTA-TCN

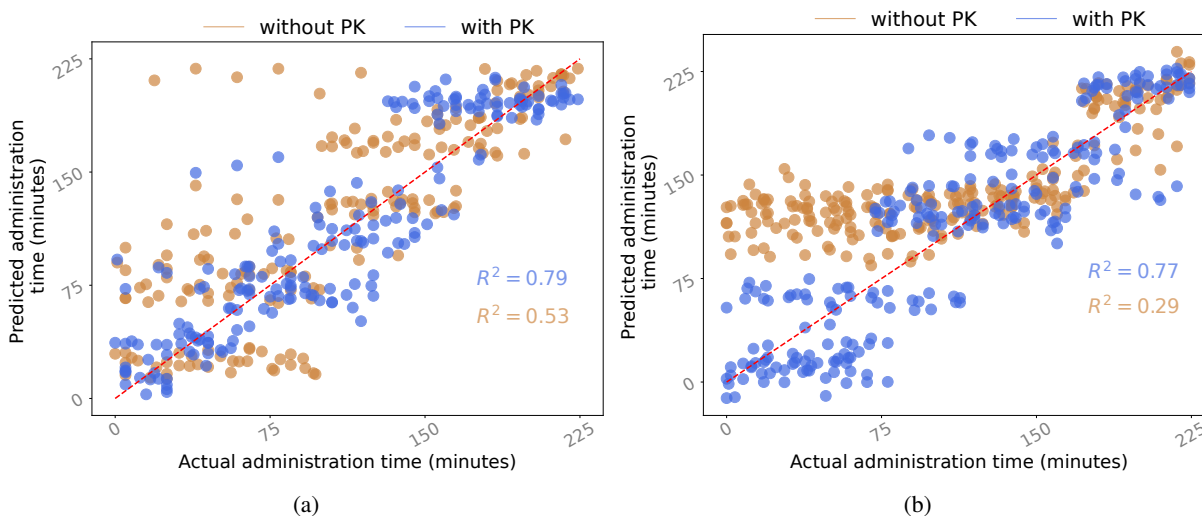


Figure 3: Scatter plot of moment predictions with the model which uses pharmacokinetics-informed learning (‘with PK’) and the model which does not (‘without PK’) in the: a) inpatient setting; and b) outpatient setting. Both ‘without PK’ and ‘with PK’ models have been pre-trained with the self-supervision task.

having a NMAE of 0.11 and R^2 of 0.77, a $\approx 165\%$ increase from just using SSL. Figure 3 shows the scatter plot between actual and predicted administration moment for both settings with and without using pharmacokinetics information to visualize the improvement in R^2 .

Conclusion, Limitations, and Future Work

In our present work, we trained machine learning models on extensive wearable data to detect oral opioid administrations, focusing on both the occurrence and timing of intake. Improving these models with self-supervised learning and opioid pharmacokinetic data, we found that self-supervised learning modestly improved performance, and incorporating pharmacokinetic data significantly increased the accuracy of predicting the timing of opioid use.

- 1. Interpreting results by medical professionals:** Medical professionals primarily rely on user self-reports to determine the frequency of opioid intake which can be biased and easily tampered with ((Vietri et al. 2014)). Passive measurement of opioid use using our proposed framework can assist medical professionals in tracking the subject’s frequency of opioid usage, allowing them to adjust prescriptions and dosages accordingly.
- 2. Opioid administration labels reliability and data quality:** The present data collection utilized a research-grade Emptaica device, which is impractical for long-term use in the general population due to its cost and aesthetics. Determining if similar results could be obtained on commercially available devices will be an important future step to ensure translation to the real world. Additionally, EHR confirmation of administration time was only available for inpatient (not outpatient) events. Outpatient events were confirmed using the patient reports only as the ground truth for ingestion, which may introduce additional noise.

- 3. Generalizing to other opioids and estimating plasma drug concentration with different equations:** The opioid oxycodone, examined in our study, exhibits first-order absorption and elimination kinetics. However, there is heterogeneity in opioid pharmacokinetics (PK). For instance, fentanyl adheres to a second-order compartment model, while morphine also follows first-order pharmacokinetics. The proposed PK-informed neural network training approach can generalize across different types of opioids following different PK equations and models.

Path to Deployment: The proposed Pharmacokinetics-Informed Neural Network can potentially be integrated into commercial smartwatches, as these devices are capable of collecting the necessary wearable signals for inference. However, for effective deployment, it is crucial to pre-identify the type of opioid commonly used by the user, to accurately incorporate it into the pharmacokinetics equation. Additionally, the model’s continuous measurements can result in false positives and negatives, which necessitates an effective management strategy that does not overwhelm the user. Furthermore, since users may take various medications that influence their physiology, our model requires further experimentation in different polysubstance use scenarios in the presence of different comorbidities prior to deployment.

Acknowledgments

This work is supported by the National Institutes of Health/National Institute on Drug Abuse under grant K23DA045242 (PI: Carreiro), National Science Foundation Smart and Connected Health program under grant 2124282 (PI: Rahman), Optum Labs (PI: Rahman), Google Research Scholar Award 2023 (PI: Rahman), and startup support from HDSI UCSD. We thank Dr. Tanzir Mortuza and Prof. Daniel Reker for their guidance and support.

References

- Burgess-Hull, A. J.; Brooks, C.; Epstein, D. H.; Gandhi, D.; and Oviedo, E. 2002. Using Machine Learning to Predict Treatment Adherence in Patients on Medication for Opioid Use Disorder. *Journal of Addiction Medicine*, 17(1): 28–34.
- Carreiro, S.; Wittbold, K.; Indic, P.; Fang, H.; Zhang, J.; and Boyer, E. W. 2016. Wearable biosensors to detect physiologic change during opioid use. *Journal of medical toxicology*, 12(3): 255–262.
- Chapman, B. P.; Gullapalli, B. T.; Rahman, T.; Smelson, D.; Boyer, E. W.; and Carreiro, S. 2022. Impact of individual and treatment characteristics on wearable sensor-based digital biomarkers of opioid use. *npj Digital Medicine*, 5(1): 1–12.
- Cohen, J. 1968. Weighted kappa: nominal scale agreement provision for scaled disagreement or partial credit. *Psychological bulletin*, 70(4): 213.
- Epstein, D. H.; Tyburski, M.; Kowalczyk, W. J.; Burgess-Hull, A. J.; Phillips, K. A.; Curtis, B. L.; and Preston, K. L. 2020. Prediction of stress and drug craving ninety minutes in the future with passively collected GPS data. *NPJ digital medicine*, 3(1): 26.
- Gibaldi, M.; Perrier, D.; et al. 1982. *Pharmacokinetics*, volume 15. M. Dekker New York.
- Graves, A.; and Schmidhuber, J. 2005. Framewise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural networks*, 18(5-6): 602–610.
- Gullapalli, B. T.; Carreiro, S.; Chapman, B. P.; Ganesan, D.; Sjoquist, J.; and Rahman, T. 2021. OpiTrack: A Wearable-based Clinical Opioid Use Tracker with Temporal Convolutional Attention Networks. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(3): 1–29.
- Hasin, D. S.; Fenton, M. C.; Beseler, C.; Park, J. Y.; and Wall, M. M. 2012. Analyses related to the development of DSM-5 criteria for substance use related disorders: 2. Proposed DSM-5 criteria for alcohol, cannabis, cocaine and heroin disorders in 663 substance abuse patients. *Drug and alcohol dependence*, 122(1-2): 28–37.
- Ismail Fawaz, H.; Lucas, B.; Forestier, G.; Pelletier, C.; Schmidt, D. F.; Weber, J.; Webb, G. I.; Idoumghar, L.; Muller, P.-A.; and Petitjean, F. 2020. Inceptiontime: Finding alexnet for time series classification. *Data Mining and Knowledge Discovery*, 34(6): 1936–1962.
- Jin, X.; Cai, S.; Li, H.; and Karniadakis, G. E. 2021. NSFnets (Navier-Stokes flow nets): Physics-informed neural networks for the incompressible Navier-Stokes equations. *Journal of Computational Physics*, 426: 109951.
- Karim, F.; Majumdar, S.; Darabi, H.; and Chen, S. 2017. LSTM fully convolutional networks for time series classification. *IEEE access*, 6: 1662–1669.
- Khatri, U. G.; and Perrone, J. 2020. Opioid use disorder and COVID-19: crashing of the crises. *Journal of Addiction Medicine*.
- Mahmud, M. S.; Fang, H.; Wang, H.; Carreiro, S.; and Boyer, E. 2018. Automatic detection of opioid intake using wearable biosensor. In *2018 International Conference on Computing, Networking and Communications (ICNC)*, 784–788. IEEE.
- MANDEMA, J. W.; KAIKO, R. F.; OSHLACK, B.; REDER, R. F.; and STANSKI, D. R. 1996. Characterization and validation of a pharmacokinetic model for controlled-release oxycodone. *British journal of clinical pharmacology*, 42(6): 747–756.
- Nandakumar, R.; Gollakota, S.; and Sunshine, J. E. 2019. Opioid overdose detection using smartphones. *Science translational medicine*, 11(474): eaau8914.
- Natarajan, A.; Parate, A.; Gaiser, E.; Angarita, G.; Malison, R.; Marlin, B.; and Ganesan, D. 2013. Detecting cocaine use with wearable electrocardiogram sensors. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, 123–132.
- Pang, G.; Lu, L.; and Karniadakis, G. E. 2019. fPINNs: Fractional physics-informed neural networks. *SIAM Journal on Scientific Computing*, 41(4): A2603–A2626.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378: 686–707.
- Roth, A. M.; Tran, N. K.; Cocchiario, B.; Mitchell, A. K.; Schwartz, D. G.; Hensel, D. J.; Ataiants, J.; Brenner, J.; Yahav, I.; and Lankenau, S. E. 2021. Wearable biosensors have the potential to monitor physiological changes associated with opioid overdose among people who use drugs: A proof-of-concept study in a real-world setting. *Drug and alcohol dependence*, 229: 109138.
- Senyurek, V. Y.; Imtiaz, M. H.; Belsare, P.; Tiffany, S.; and Sazonov, E. 2020. A CNN-LSTM neural network for recognition of puffing in smoking episodes using wearable sensors. *Biomedical Engineering Letters*, 10(2): 195–203.
- Stone, A. B.; Urman, R. D.; Kaye, A. D.; and Grant, M. C. 2018. Labeling morphine milligram equivalents on opioid packaging: a potential patient safety intervention. *Current pain and headache reports*, 22(7): 1–4.
- Vietri, J.; Joshi, A. V.; Barsdorf, A. I.; and Mardekian, J. 2014. Prescription opioid abuse and tampering in the United States: Results of a self-report survey. *Pain Medicine*, 15(12): 2064–2074.
- Wang, Z.; Yan, W.; and Oates, T. 2017. Time series classification from scratch with deep neural networks: A strong baseline. In *2017 International joint conference on neural networks (IJCNN)*, 1578–1585. IEEE.
- Zhu, Q.; Liu, Z.; and Yan, J. 2021. Machine learning for metal additive manufacturing: Predicting temperature and melt pool fluid dynamics using physics-informed neural networks. *Computational Mechanics*, 67(2): 619–635.