

# Using internet search trends to forecast short term drug overdose deaths: A case study on Connecticut

Sumit Mukherjee<sup>1</sup>, Nicholas Becker<sup>1</sup>, William Weeks<sup>2</sup>, Juan Lavista Ferres<sup>1</sup>

<sup>1</sup>Microsoft, AI for Good Research Lab

<sup>2</sup>Microsoft, Healthcare NExT

Redmond, USA

(summukhe,nibecker,wiweeks,jlavista)@microsoft.com

**Abstract**—In the United States, the opioid epidemic is a serious public health crisis which claimed over 130 lives per day in 2018, according to the CDC. While there are many efforts to design effective interventions to prevent drug related deaths, much of them are focused around better prescribing practices. A promising line of inquiry has focused on utilizing machine learning tools to predict addiction or overdose related hospital admission using prior health record information. However, these are strongly reliant on the private health record information of individuals. Here, we propose using publicly available historic death records along with publicly available internet search trends of drug related search terms to predict the number of overdose deaths in the upcoming week. Our model is able to predict both the number of, and spikes in drug overdose deaths with good accuracy compared to several baselines, demonstrating the utility of search data in forecasting overdose deaths. While we demonstrate this approach as a case study in the State of Connecticut, which collects and publishes overdose data, our findings could encourage other state governments to similarly invest in collection, publication and analysis of such data.

**Index Terms**—Drug overdose, search trends, temporal modeling.

## I. INTRODUCTION

Drug overdoses caused more than 70,000 deaths in the United States in 2017, with over 67 percent of those deaths being caused by opioids [1]. The numbers have been rising steadily since 1999 (when there were under 17,000 overdose related deaths) and the rate of death has doubled since 2010. Drug overdose is currently the leading cause of death in the United States for ages under 50 ahead of deaths due to firearms, car accidents and homicides. While a large factor in the growth of overdose deaths is the advent of powerful synthetic opioids, deaths due to other drugs such as cocaine and heroin have also been steadily rising. The burden of overdose deaths are augmented by ad-

diction; over 10 percent of Americans over 12 years of age were estimated have used illicit drugs in the last month (<https://www.cdc.gov/nchs/fastats/drug-use-illegal.htm>) and over 11 million Americans misused opioids in 2016 [1].

Despite these horrific statistics and the recent declaration of opioid abuse as a public health emergency, there is still a lack of understanding regarding the best interventions to reduce overdose deaths. Hence, there has been a lot of recent academic work focused on identifying different avenues of intervention and their utility [2], [3]. The interventions can be primarily classified into 3 broad categories: i) overdose and addiction prevention, ii) better treatment of chronic pain, iii) treatment of addiction. Machine learning methods could benefit each of these classes of interventions, although we are still in early days for such work.

Most of the current machine learning work on drug or substance abuse has focused on the prevention of overdose and addiction events. Several papers have used Electronic Health Records (EHR) and/or insurance claims information to identify opioid dependence diagnosis before it occurs [4]. A similar line of work has focused on classifying the severity of Emergency Room (ER) admission events [5]. The main drawbacks of these approaches are the non-public nature and local scope of the data sets required to train the models, which would likely prevent widespread adoption of such approaches. Another research team has used image and text data from Instagram to identify individuals at risk for substance abuse of various substance types [6]. While the use of social media information is intriguing, the approach seems to have demonstrated more success with alcohol than other substance types.

While these previous works focus primarily on prevention of overdose deaths through individual level interventions, to date no work has explored predicting overdose

related deaths at a population level in a geographic area. Given the known resource and personnel scarcity faced by many emergency departments [7], we anticipate such information would prove helpful to hospitals, public health officials, and local law enforcement for the purpose of more precise resource allocation.

To this end, we sought to develop a predictive model for short term drug overdose deaths using publicly available internet search trend data and overdose death data. We demonstrate that such an approach can predict near-future overdose deaths as well as spikes in overdose deaths. We further develop metrics to identify the search terms most important in predicting overdose deaths, which we use to generate novel hypotheses (see Figure 1 of overview of our approach). We demonstrate our results on multiple drugs (and all drugs) using publicly available overdose death information from the State of Connecticut (USA). While the current study only focuses on Connecticut, which collects and publishes such data, our framework is generalizable and we hope it will serve as a motivation for other states to make such data sets public. To further explore the potential power of such data sets, we also analyze the socio-demographic information from the death records and reveal how different demographic groups are differently affected by different drugs.

## II. RELATED WORKS

### A. Search trend based analysis in public health

While we are not aware of previously published work utilizing search trends to identify/forecast drug related trends, search trends have been used widely in various other public health domains [8]. This was first pioneered by Google in a pivotal study where researchers used Google search trends data to predict influenza outbreaks [9]. In this study the authors utilized searches correlated to the occurrence of influenza to build a simple linear model relating the volumes for those searches and estimates for influenza occurrence. Since then, search trend data has been utilized by many studies related to public health monitoring in various ways. One study showed an association between state level HIV infections and searches related to HIV at the state level [10]. Another paper examined the association of suicides with search terms related to 'suicide' between 2004–2009 [11]. A different line of work looks at how the search volumes relate to certain search query terms (acting as proxies for behavior) vary as a function of time of the year [12]. A more recent line of work has looked at forecasting events (such as flu and zika outbreaks) using prior occurrences of the events along with search trends [13], [14].

While we build on these approaches, our work is unique in the short term and local nature of our forecasting problem. We also incorporated historical data on overdose deaths and historical county level weather data to predict future overdose deaths. The use of weather data was motivated by literature on the effects of weather of mental well-being and happiness of individuals [15]. The weather data also serves as a 'control' to the search terms. A major difference of our work from previous search trend based forecasting works is the relatively stochastic nature of drug overdose deaths which complicates the problem of forecasting short term deaths and motivates looking at overdose death spikes instead.

## III. DATA SETS

Three distinct publicly available data sets were utilized in this study:

- *Overdose death information:* The government of Connecticut open data portal contains a de-identified list of all substance abuse related deaths between 2012–2018 (<https://data.ct.gov/Health-and-Human-Services/Accidental-Drug-Related-Deaths-2012-2018/rybz-nyjw>). Each death is time and location stamped with information about type of drugs found in the system of the deceased, along with demographic information. For the purpose of predictive modeling, the number of deaths in each day of the 2012–2018 date range was calculated for all drugs as well as for individual drugs.
- *Historic weather data:* Historic weather data is collected from the National Oceanic and Atmospheric Administration (NOAA) climate data online search tool [16]. We used the following city level data: i) maximum temperature, ii) minimum temperature, iii) average precipitation, iv) average snowfall. Because many data points were missing, we aggregated the information to the county level. This led to a total of 32 weather variables being used. While calculating averages, the missing values were ignored.
- *Google search trends data:* Google trends (<https://trends.google.com/trends>) is a service which tracks the variation in search volumes for different search terms over time with data from Google searches. The search terms used in this study were: i) drug names (of all drugs indicated as present at least one overdose death in Connecticut), ii) combined list of related search terms for all drug names, iii) search terms related to (and including) 'withdrawal'. A total of 105 search terms were

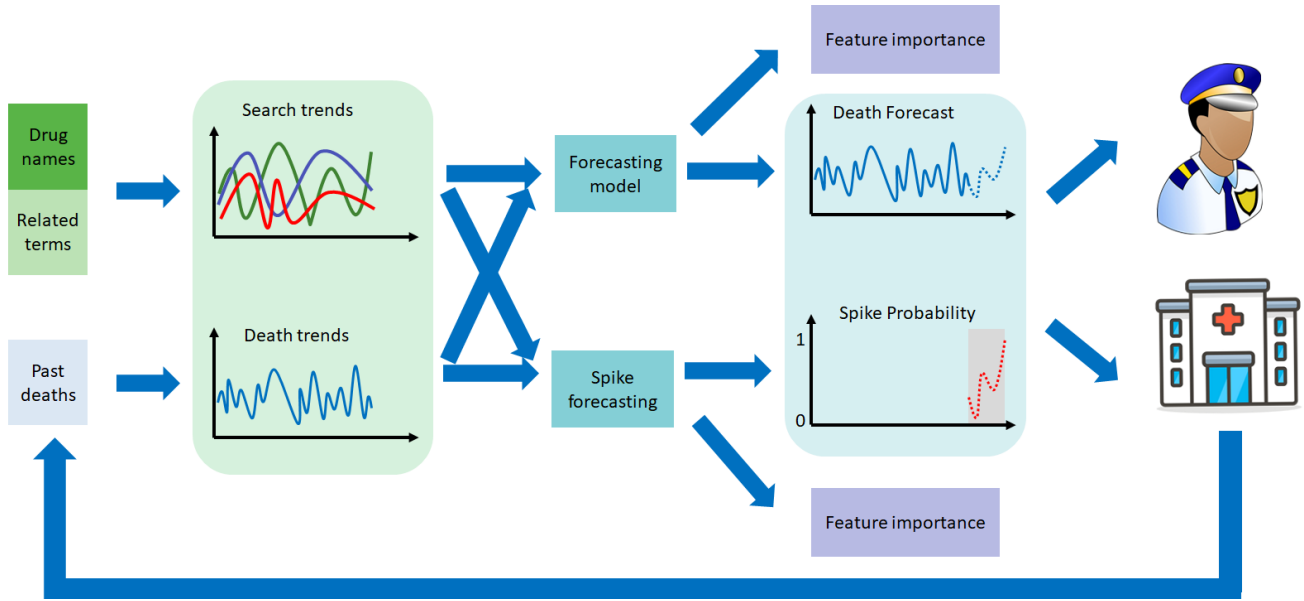


Fig. 1. Overview of our modeling approach. Drug names are used to obtain related search terms. The combination of these terms are used to obtain search trends which are combined with past death information to forecast number of deaths and probability of death spikes in a week ahead manner. We also obtain feature importance of each model.

used. The search trends were sub-set to the State of Connecticut, the data is collected every 24 hours for the time period between January 1st 2012 and December 31st 2018.

#### IV. METHODS

A key objective of this paper is to build a short term forecasting model using historical search trend, weather and overdose death data. We pose the problem in 2 separate ways: i) as a regression problem to predict the change in number of deaths in the upcoming week, ii) as a classification problem to predict whether or not there is a spike in number of deaths in the upcoming week.

##### A. Predictive modeling of drug overdose deaths

Due to sparseness of drug overdose events (particularly drug specific), it is extremely difficult to predict the exact number of deaths in any given day. However, a slightly more tractable problem is the prediction of sum of the overdose deaths in the upcoming week starting from any given day. However, as is evident from figure 2, such a time series is non-stationary due to the presence of long term trends. Hence, we modify the prediction problem to predict the change in the sum of overdose deaths in the upcoming week compared to the previous week using only historic data. Such a model is chosen keeping a practical deployment scenario in mind where we want to predict deaths in the following week

using data only available to us at the current time. For the purposes of this paper, we selected the forecasting window to be a week ahead following the majority of relapse prediction work which is measured in terms of weeks [17]. This can be mathematically formulated using a linear model as follows:

$$\Delta Y(t) = \sum_{i=1}^7 (\Delta Y(t-6-i)a_i + \sum_{j=1}^m \Delta U_j(t-6-i)b_{i,j}) + \epsilon_t$$

$$\text{where, } \Delta Y(t) = \sum_{i=t+1}^{t+7} Y(i) - \sum_{i=t-6}^t Y(i),$$

$$\Delta U_j(t) = \sum_{i=t+1}^{t+7} U_j(i) - \sum_{i=t-6}^t U_j(i),$$

$$\epsilon_t \sim WN(0, \sigma^2).$$

Here,  $Y(i)$  is the number of overdose deaths in day  $i$ ,  $U_j(i)$  is the value of the  $j$ th exogenous feature at time  $i$  and  $m$  equals the number of exogenous variables used in the modeling. Here, the exogenous features include both search terms as well as county level weather information.  $a_i$  and  $b_{i,j}$  are the model parameters to be estimated. The time series  $\{\Delta Y(i)\}$  was tested for stationarity using the Augmented Dickey-Fuller test [18] and the null hypothesis was rejected with a p-value  $\leq 0.05$  for all drugs. The same was repeated for the time series' corresponding to the search terms as well i.e.  $\{\Delta U_j(i)\}$ .

All variables were scaled to lie in the [0,1] range to help with identification of feature importance.

### B. Predictive modeling of drug overdose death spikes

While the previous modeling approach tries to predict the number of deaths, it can be argued that the raw numbers are not as important to the target users of such a model, for example government agencies and law enforcement. Instead predicting whether or not a spike in deaths will occur in the upcoming week, compared to the previous week, will be more useful. Such a model can be formulated using a logistic model as follows:

$$\log_e\left(\frac{p}{1-p}\right) = \sum_{i=1}^7 (\Delta Y(t-6-i)a_i) + \sum_{j=1}^m \Delta U_j(t-6-i)b_{i,j}$$

where,  $p = \Pr(\Delta Y(t) \geq \delta)$

Here  $\delta$  is the *spike threshold*, which is empirically selected to be one standard deviation above the mean in the observed distribution of  $\{\Delta Y(i)\}$ . This threshold was selected because of its simplicity and interpretability. As in the previous model, all variables were scaled to lie in the [0,1] range.

### C. Model training, validation and baselines

In a realistic deployment scenario, the model would be updated periodically with new  $Y(i)$  and  $U_j(i)$  to predict the outcome variables  $\Delta Y(i)$  and  $p$  respectively. The model training can either be done using all historical data or with a fixed rolling window. We use a  $L_1$ -norm penalty for feature selection in the the death prediction model, while a  $L_2$ -norm was used for the death spike prediction model. We use k-fold cross-validation to select the appropriate hyper-parameter value. The python scikit-learn package was used. For this work we used the latter approach with a window size of 1 year. We performed out-of-sample testing on the upcoming week [19] for both types of models and all reported error statistics are on out-of-sample tests.

For the task of predicting number of deaths, we use 3 baselines: i) Simple auto-regressive model without exogenous input (AR) , ii) no change model (which assumes no change in number from previous week), iii) auto-regressive model with random English words as search terms (same number of terms as full model). Furthermore, we tested an auto-regressive model with only search and not weather (AR+S) to explore the relative benefit provided using the search data alone. For the task of predicting spikes in deaths, we omitted the no change model.

### D. Identifying feature importance

To identify feature importance in linear models, it is common to rank by the absolute values of feature weights assuming the features are normalized. In our modeling approaches, since there are multiple parameters for each exogenous input and auto-regressive terms (due to lagged terms), it isn't enough to look at the absolute values of weights. We used 2 different approaches to quantify feature contribution: i) maximum absolute weight - which is the highest absolute weight corresponding to any given exogenous input or auto-regressive term i.e.  $\max_i(\{|a_i|\}_{i=1}^7)$  in the case of auto-regressive features and  $\max_i(\{|b_{i,j}|\}_{i=1}^7)$  in the case of exogenous inputs, ii) mean absolute weight - which is the average absolute weight corresponding to any given auto-regressive feature or exogenous input i.e.  $\frac{1}{7} \sum_{i=1}^7 |a_i|$  and  $\frac{1}{7} \sum_{i=1}^7 |b_{i,j}|$  respectively. The intuitive utility of the two feature importance rankings is slightly different. While the *maximum absolute weight* metric identifies if any of the lagged terms of the feature is important, the *mean absolute weight* identifies if a feature is important across all the lags considered in the model.

### E. Data set pre-processing

For the State of Connecticut, the substance abuse overdose death information is available for 7 years from 2012 through 2018. This information is time and location stamped and contains additional demographic information for each death. However, for the purposes of a short term predictive model, we only care about the number of deaths in a time period and ignore the demographic information associated with each death. Each death also mentions which of the drugs are found in the system, so deaths caused by any particular drug can be easily extracted by looking at which deaths listed the drug as present in the system. This leads to a time series corresponding to each drug as seen in Figure 2. However, since the number of deaths in each day is very noisy, we choose to predict sum for the upcoming week (as described in Methods), as summing over a window has the effect of reducing randomness.

### F. Identifying drug specific variation among categories

To identify whether the relative abundance of deaths due to certain drugs change among demographic categories, we first calculated the relative abundance of death due to each drug (count for a given drug divided by count for all of the drugs) for each year. We then performed a MANOVA [20] test to identify which drugs were significantly varying among the demographic categories

being considered. A p-value threshold of 0.05 was for statistical significance.

## V. RESULTS

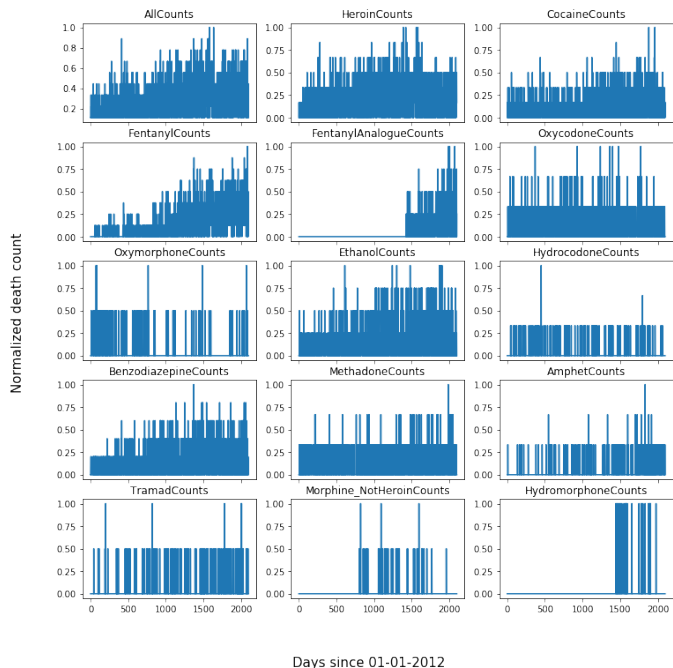


Fig. 2. Variation of daily substance abuse overdose death counts as a function of time for different substances.

### A. Search trends improve predictive modeling of drug overdose death and death spikes

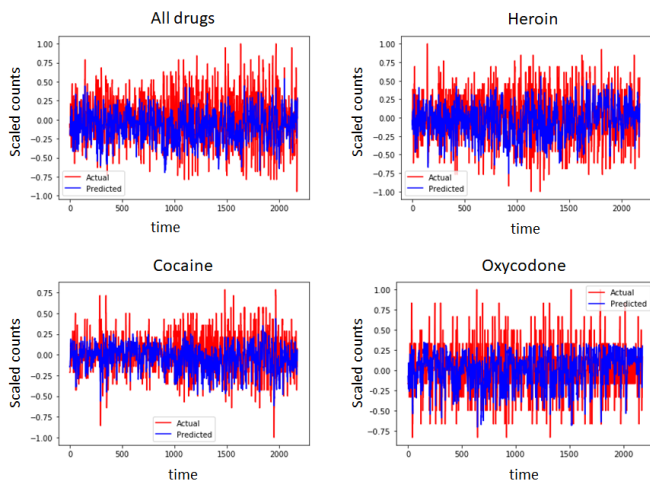


Fig. 3. True data and predicted drug overdose deaths for various drugs.

Since several of the drugs did not have many recorded deaths in the 7 year window, we chose 4 drugs (aside

from all drugs) to model, namely: i) heroin, ii) cocaine, iii) oxycodone and iv) fentanyl. Heroin and cocaine were chosen because they are known to be major causes of deaths, oxycodone was chosen because it is a commonly over-prescribed opioid, and fentanyl was chosen because it is a synthetic opioid and is fast becoming the most lethal abused drug. We trained the model for drug deaths using an increasing window training scheme. Under this scheme, the model was first trained on 1 year worth of data followed by testing on the next point in the time series i.e. the week immediately after (which wasn't used for training) and retraining with a warm start for every subsequent point. The mean absolute error (along with the standard error) of all the test points is then calculated to measure the overall accuracy of the model on the entire time series. We trained 4 separate models on the 4 drugs and 'all drugs' and compared them to 3 separate baselines mentioned in the 'Methods' section. We find that our model that uses search trends and weather, outperforms both baselines for all 5 drug types as seen in Figure 3 and Table I. Moreover, the model that uses search but not weather information performs almost as well for most drugs, although minor improvements are observed in some cases by utilizing weather data. This shows that most of the performance improvements are derived from the search terms and not weather data. Furthermore, seeing that using random English word search terms doesn't lead to the same improvements, shows that the improvements are linked to the relevance of the search terms. Since the testing is performed out of sample, this clearly illustrates that search volume information contains additional information that aids in the prediction of short term overdose deaths.

While the previous model predicts numbers of deaths, such information may not be the only relevant information to healthcare officials who can benefit from such a predictive model. Another type of information that might be valuable to various end users is whether or not there will be a spike in overdose deaths in the following week. Here we define a week to contain a spike if the number of deaths is over a certain threshold (corresponding to one standard deviation over the mean deaths in all weeks). We formulate and train the model (as described in 'Methods') and test it in a manner identical to the previous model. However, since spike prediction is a classification problem, we calculate the area under the ROC curve (AUC) instead of MAE. Similar to the death prediction models, we find that the full model that utilizes search and weather data improves upon the simple auto-regressive and random search baseline

Drug	Full	AR+Search	AR+Random	AR	No change
All drugs	<b>0.215</b> $\pm$ 0.004	<b>0.215</b> $\pm$ 0.004	0.223 $\pm$ 0.004	0.225 $\pm$ 0.004	0.252 $\pm$ 0.004
Heroin	<b>0.224</b> $\pm$ 0.004	0.225 $\pm$ 0.004	0.225 $\pm$ 0.004	0.228 $\pm$ 0.004	0.253 $\pm$ 0.004
Cocaine	<b>0.156</b> $\pm$ 0.003	<b>0.156</b> $\pm$ 0.003	0.157 $\pm$ 0.004	0.159 $\pm$ 0.003	0.177 $\pm$ 0.003
Oxycodone	<b>0.149</b> $\pm$ 0.003	<b>0.149</b> $\pm$ 0.003	0.215 $\pm$ 0.004	0.158 $\pm$ 0.003	0.234 $\pm$ 0.004
Fentanyl	<b>0.150</b> $\pm$ 0.003	<b>0.150</b> $\pm$ 0.003	0.152 $\pm$ 0.003	0.156 $\pm$ 0.003	0.167 $\pm$ 0.004

TABLE I

MEAN ABSOLUTE ERROR (MAE) OF VARIOUS DRUG DEATH PREDICTION MODELS FOR DIFFERENT DRUGS. STANDARD ERRORS ARE REPORTED IN THE SUBSCRIPT.

models for all drugs as seen in Figure 4 and Table II. Once again, we note that using relevant search data alone seems to provide much of the improvement over the auto-regressive model, although in some cases (such as for all drugs, cocaine and fentanyl), including weather data slightly improves the models.

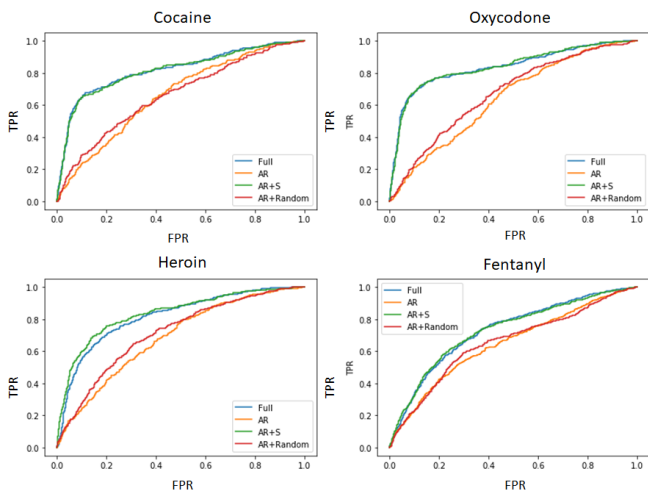


Fig. 4. ROC curves for overdose death spike prediction using search trends and weather compared various baseline models.

Drug	Full	AR + Search	AR + Rand.	AR
All drugs	<b>0.810</b>	0.805	0.699	0.717
Heroin	0.817	<b>0.835</b>	0.712	0.685
Cocaine	<b>0.819</b>	0.816	0.662	0.661
Oxycodone	<b>0.838</b>	<b>0.838</b>	0.672	0.641
Fentanyl	<b>0.738</b>	0.728	0.652	0.646

TABLE II

AUC OF VARIOUS OPIOID DEATH SPIKE PREDICTION MODELS FOR DIFFERENT DRUGS.

### B. Identifying important search queries

Having demonstrated the utility of search trends for predictive modeling, we look to understand the importance of various search terms. As mentioned previously,

identifying search term importance is not trivial for our model because each search term is used at multiple lag times. Hence, we used the 2 heuristic approaches described in 'Methods' to rank the features in the order of importance as seen in Figure 5. We focus on the death spike prediction models as those show more substantial improvement over the auto-regressive baselines. We notice that the most important search terms for all models correspond to drug name searches corresponding to major drugs such as heroin, fentanyl, oxycodone and cocaine. Moreover, several of the top search terms correspond to prescription pain medicine and medicine combinations such as acetaminophen, codeine, percocet, tramadol etc.. It is interesting to note that while several of the top search terms are similar between models for different drugs, there are also several search terms that are unique for each model. This could indicate a difference in the users of each drug, which is confirmed in the following sub-section. Interestingly, a non-intuitive search term that ranks highly for all models is the term 'pill identifier', which is a tool to identify what a drug is from unmarked pills [21]. While it is beyond the scope of this paper to explore the links between specific search terms and overdose deaths, we showed here that this approach could be useful as a hypothesis generation tool to further understand drug deaths.

### C. Identifying susceptible populations from overdose data

Aside from enabling predictive modeling, the socio-demographic information in the drug overdose death data set provides important information enabling the identification of groups that are more susceptible to particular drugs. In Figure 6, we visualize the relative abundance of overdose deaths for the 4 drugs we considered in our study, for several socio-demographic (and geographic) categories such as: i) age, ii) major races (which had at least one death per drug), iii) gender, and iv) county. We find no statistically significant variation

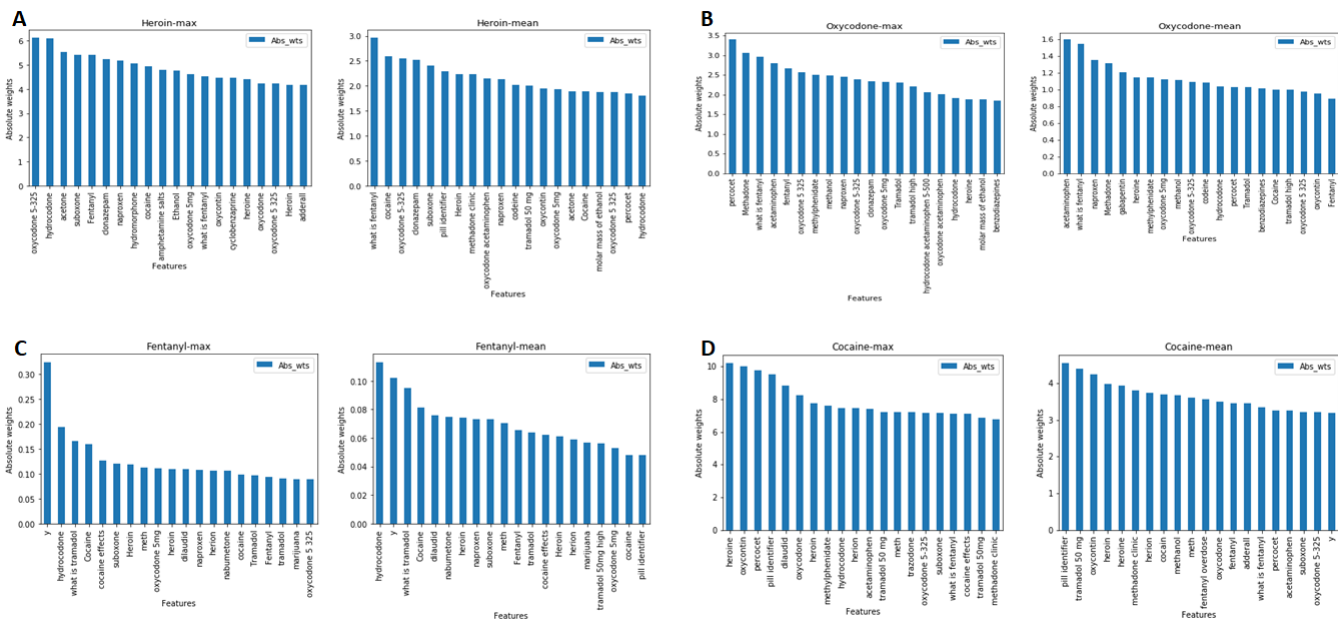


Fig. 5. Top 20 ranked features of the death spike prediction models for: a) heroin, b) oxycodone, c) fentanyl, d)cocaine.

in the case of gender or county ( $p \leq 0.05$ ), although we visually observe a higher variation for cocaine among the different counties. We find a statistically significant variation for the relative abundance of deaths due to oxycodone and upon visual inspection we find it to be higher among the 'White' demographic. We also find significant differences for 3 out of 4 drugs (cocaine, oxycodone and heroin) in the case of age. Notably, we find relative deaths due to oxycodone increase with age while relative deaths due to heroin decrease with age. While a more thorough analysis of drug specific variation of deaths is beyond the scope of this paper, we expect these results to demonstrate the utility of such data sets in identifying susceptible populations with high granularity.

## VI. CONCLUSION

We demonstrated the utility of search query data in predicting future drug overdose deaths. We utilized historical overdose death information in conjunction with search query volume and county level weather information to model the number of drug overdose deaths and spikes in overdose deaths in the upcoming week. We show that the incorporation of relevant search trend information into the predictive model improves the prediction performance for both tasks, for several common opioids and cocaine. We further showed that studying the feature ranking of the models offers important insights into which search terms are of predictive value. This in turn

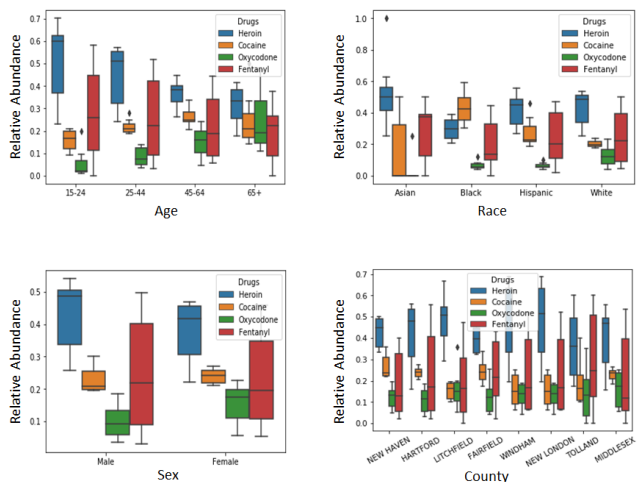


Fig. 6. Drug specific demographic trends of overdose deaths for the State of Connecticut.

may lend support to the generation of new hypothesis regarding overdose incidents in the future. One particular interesting hypothesis is that increased search of 'pill identifier' could indicate an ingress of illegal unmarked pills into the local market.

While this model is currently limited to the State of Connecticut where the data are available, it might also be applicable to other states and countries. Models like the one we developed have the potential to benefit law enforcement agencies, public health teams and health-

care institutions in being appropriately prepared for an upcoming spike in opioid deaths. Ideally, the models would inform interventions aimed at preventing overdose deaths.

A primary limitation of our is the lack of availability of similar data sets from other states and the limited granularity of the search data (currently only sub-setted to the state level). Future work will aim to reproduce the modeling using search data at a much greater geospatial resolution. Moreover, future work could include a collaboration between public health officials and internet search providers to pilot test the efficacy of early signal detection to drive resource allocation and prevention.

## REFERENCES

- [1] L. Scholl, P. Seth, M. Kariisa, N. Wilson, and G. Baldwin, "Drug and opioid-involved overdose deaths—united states, 2013–2017," *Morbidity and Mortality Weekly Report*, vol. 67, no. 5152, p. 1419, 2019.
- [2] J. H. Samet and S. G. Kertesz, "Suggested paths to fixing the opioid crisis: directions and misdirections," *JAMA network open*, vol. 1, no. 2, pp. e180218–e180218, 2018.
- [3] S. Wickramatilake, J. Zur, N. Mulvaney-Day, M. C. v. Klimo, E. Selmi, and H. Harwood, "How states are tackling the opioid crisis," *Public Health Reports*, vol. 132, no. 2, pp. 171–179, 2017.
- [4] R. J. Ellis, Z. Wang, N. Genes, and A. Ma'ayan, "Predicting opioid dependence from electronic health records with machine learning," *BioData mining*, vol. 12, no. 1, p. 3, 2019.
- [5] J. Badger, E. LaRose, J. Mayer, F. Bashiri, D. Page, and P. Peissig, "Machine learning for phenotyping opioid overdose events," *Journal of biomedical informatics*, vol. 94, p. 103185, 2019.
- [6] S. Hassanpour, N. Tomita, T. DeLise, B. Crosier, and L. A. Marsch, "Identifying substance use risk based on deep neural networks and instagram social media data," *Neuropsychopharmacology*, vol. 44, no. 3, p. 487, 2019.
- [7] S. Trzeciak and E. P. Rivers, "Emergency department overcrowding in the united states: an emerging threat to patient safety and public health," *Emergency medicine journal*, vol. 20, no. 5, pp. 402–405, 2003.
- [8] R. W. White, N. P. Tatonetti, N. H. Shah, R. B. Altman, and E. Horvitz, "Web-scale pharmacovigilance: listening to signals from the crowd," *Journal of the American Medical Informatics Association*, vol. 20, no. 3, pp. 404–408, 2013.
- [9] J. Ginsberg, M. H. Mohebbi, R. S. Patel, L. Brammer, M. S. Smolinski, and L. Brilliant, "Detecting influenza epidemics using search engine query data," *Nature*, vol. 457, no. 7232, p. 1012, 2009.
- [10] A. B. Jena, P. Karaca-Mandic, L. Weaver, and S. A. Seabury, "Predicting new diagnoses of hiv infection using internet search engine data," *Clinical infectious diseases*, vol. 56, no. 9, pp. 1352–1353, 2013.
- [11] A. C. Yang, S.-J. Tsai, N. E. Huang, and C.-K. Peng, "Association of internet search trends with suicide death in taipei city, taiwan, 2004–2009," *Journal of affective disorders*, vol. 132, no. 1-2, pp. 179–184, 2011.
- [12] L. J. Carr and S. I. Dunsiger, "Search query data to monitor interest in behavior change: application for public health," *PloS one*, vol. 7, no. 10, p. e48158, 2012.
- [13] Y. Teng, D. Bi, G. Xie, Y. Jin, Y. Huang, B. Lin, X. An, D. Feng, and Y. Tong, "Dynamic forecasting of zika epidemics using google trends," *PloS one*, vol. 12, no. 1, p. e0165085, 2017.
- [14] S. Yang, M. Santillana, J. S. Brownstein, J. Gray, S. Richardson, and S. Kou, "Using electronic health records and internet search information for accurate influenza forecasting," *BMC infectious diseases*, vol. 17, no. 1, p. 332, 2017.
- [15] Y. Tsutsui, "Weather and individual happiness," *Weather, Climate, and Society*, vol. 5, no. 1, pp. 70–82, 2013.
- [16] A. H. Young, K. R. Knapp, A. Inamdar, W. Hankins, and W. B. Rossow, "The international satellite cloud climatology project h-series climate data record product," *Earth System Science Data*, vol. 10, no. 1, pp. 583–593, 2018.
- [17] Z. Yang, L. H. Nguyen, and F. Jin, "Opioid relapse prediction with gan," in *2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pp. 560–567, IEEE, 2019.
- [18] J. H. Lopez, "The power of the adf test," *Economics Letters*, vol. 57, no. 1, pp. 5–10, 1997.
- [19] L. J. Tashman, "Out-of-sample tests of forecasting accuracy: an analysis and review," *International journal of forecasting*, vol. 16, no. 4, pp. 437–450, 2000.
- [20] R. G. O'Brien and M. K. Kaiser, "Manova method for analyzing repeated measures designs: an extensive primer," *Psychological bulletin*, vol. 97, no. 2, p. 316, 1985.
- [21] Y.-B. Lee, U. Park, and A. K. Jain, "Pill-id: Matching and retrieval of drug pill imprint images," in *2010 20th International Conference on Pattern Recognition*, pp. 2632–2635, IEEE, 2010.