Machine Learning for Drug Overdose Surveillance

Daniel B. Neill
Event and Pattern Detection Laboratory
Carnegie Mellon University
Pittsburgh, PA
neill@cs.cmu.edu

William Herlands
Event and Pattern Detection Laboratory
Carnegie Mellon University
Pittsburgh, PA
herlands@cmu.edu

1. INTRODUCTION

A recent epidemic of opioid overdoses has garnered national attention in the United States [13]. The CDC estimates that 47,055 overdose deaths occurred in the U.S. in 2014, 61% of which involved opioids (including heroin, pain relievers such as oxycodone, and synthetics) [9, 6]. These statistics motivate public health agencies to identify emerging trends in overdoses so that they can better target prevention and response measures. However, reliable identification of such trends requires us to solve several challenges. First, to accurately pinpoint the affected spatial area, we must account for the complex, correlated nature of spatio-temporal data, modeling these correlations and using them to distinguish between significant anomalous patterns and fluctuations due to correlated noise. Second, when a pattern of overdoses differentially affects some subpopulation (such as elderly males, or drug users who combine heroin with alcohol), it is important to identify and precisely target that subpopulation. To do so, we must scan over multiple dimensions of the data including demographics and behaviors as well as geographic areas. We expect that this will also increase power to detect emerging overdose clusters in their early stages, when case counts are low. Most importantly, we must consider how to integrate information across space, time, and subpopulation-level characteristics, since an event of interest may be spread across all of these features. Though individual anomaly detection methods are commonly used in practice [4, 12], simply considering the anomalousness of each individual point loses power to detect subtle or emerging anomalies.

2. METHODS

We consider two recently proposed extensions of the fast subset scan [5], which detects anomalous patterns by efficiently maximizing a log-likelihood ratio statistic over subsets of data points. These approaches, the Gaussian Process Subset Scan (GPSS) and Multidimensional Tensor Scan (MDTS), are described in more detail in [3] and [6] respectively. For each method, we frame the search as a log-likelihood ratio (LLR) comparison between a null hypothesis H_0, assuming no events of interest, and a set of alternative hypotheses $H_1(S)$, each representing the occurrence of an event in some subset S. GPSS monitors data $D = (x, y)$, where $x = \{x_1, \ldots, x_n\}$, $x_i \in \mathbb{R}^D$ are covariates, and $y = \{y_1, \ldots, y_n\}$, $y_i \in \mathbb{R}$ is a response variable. For ex-

ample, each x_i could represent the coordinates of a location in space and time, and the corresponding y_i could represent the number of overdose deaths for that location and time period. In MDTS, we again monitor data $D = (x, y)$, but now each x_i represents a set of discrete-valued characteristics. For example, each x_i could represent the characteristics of a given overdose victim or set of identical victims (gender = male, age group = 20-29, etc.) and the corresponding y_i would be the number of individuals with those characteristics (i.e., $y_i = 1$ for a single case). Both approaches follow the same general framework:

1. **Estimate the distribution of the observed responses y under the null hypothesis H_0.**

   GPSS assumes that points are drawn from a function with a Gaussian process prior, $y = f(x) + \epsilon$, where $f(x) \sim GP(m(x), k(x, x'))$ and $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2 I)$. Parameters of the Gaussian process are learned from the entire dataset D. For a given subset S, the y_i will follow a multivariate Gaussian distribution, and we condition on $D \setminus S$ to infer the posterior mean vector μ and covariance matrix Σ.

   Similarly, MDTS performs scalable tensor decomposition to model D as a sum of rank-one tensors. For a given subset S, MDTS uses the learned tensor decomposition to estimate the mean vector μ, and assumes that each y_i is drawn independently from a Poisson distribution with mean μ_i.

2. **Define a search space of subsets S to consider.**

   GPSS forms the local k-neighborhoods S_{ik} defined by each x_i and its $k - 1$ nearest neighbors according to some distance metric, then considers all $S \subseteq S_{ik}$. In other words, GPSS conducts an unconstrained search within each neighborhood to find the subset which maximizes the LLR. This provides GPSS with the flexibility to identify highly irregular shapes which still satisfy the spatial proximity constraint.

   MDTS considers the subspaces formed by Cartesian products of subsets of values for each attribute. For example, given attributes gender and age, one valid subspace would be "males and females in age groups 10-19 and 20-29", while the subset of "males aged 10-19 and females aged 20-29" would not be considered.

3. **Define the alternative hypothesis $H_1(S)$.**

   GPSS assumes that under $H_1(S)$, the values of y_i for $x_i \in S$ are drawn from a multivariate Gaussian with

a mean shift of β: $y \sim \mathcal{N}(\mu + \beta w, \Sigma)$, where w is a binary vector such that $w_i = 1\{x_i \in S\}$.

MDTS assumes that under $H_1(S)$, the values of y_i for $x_i \in S$ are increased by a multiplicative factor of $q > 1$, i.e., $y_i \sim \text{Poisson}(q\mu_i)$ for $x_i \in S$ and $y_i \sim \text{Poisson}(\mu_i)$ for $x_i \notin S$.

4. **Maximize the log-likelihood ratio over subsets.**

   For both methods, the LLR score of a given subset S is defined as $F(S) = \log \frac{Pr(Data \mid H_1(S))}{Pr(Data \mid H_0)}$. The subsets $S^* = \arg\max_S F(S)$ are identified as the most anomalous subsets, and randomization testing is performed to determine a threshold for statistical significance.

   To avoid the exponential time complexity of searching over all 2^k subsets for a neighborhood of size k, GPSS uses a novel approximation technique which extends [11] to iteratively compute conditionally optimal subsets. More details are provided in [3].

   Similarly, to avoid exhaustively searching over all 2^V subspaces, where V is the sum of the arities of all discrete-valued attributes in the data, MDTS uses an iterative conditional optimization technique which uses the linear-time subset scanning property [5] to efficiently and exactly compute the conditionally optimal subset of values for a given attribute, given the current subsets of values for all other attributes. This conditional optimization step is repeated until a local optimum is reached, and multiple restarts are used to approach the global optimum. More details are provided in [6].

GPSS and MDTS have multiple, complementary advantages. GPSS is useful for modeling and accounting for non-iid correlations in real-world systems, such as spatio-temporal urban data. GPs provide a natural means of learning covariance structure from data, extend to multiple dimensions, and seamlessly handle missing data. Given case data, MDTS has high power to detect and characterize emerging trends which may only affect a subset of the monitored population (e.g., specific ages, genders, neighborhoods, or users of particular drugs or combinations of drugs).

3. RESULTS

We now describe retrospective case studies applying our surveillance techniques to real-world overdose data from two different areas of the country. First, we use the Gaussian Process Subset Scan to detect and localize overdose clusters in aggregated spatio-temporal data from New York City. Second, we apply the Multidimensional Tensor Scan to identify subpopulation-level overdose trends in case data from western Pennsylvania, and explore the discovered clusters in collaboration with county public health officials.

3.1 Case study 1: aggregated count data

For our first case study [3], we analyzed monthly opioid overdose deaths in the New York City metropolitan area from 1999-2015 [1]. Data are provided at a county level and include Manhattan, Brooklyn, Queens, the Bronx, Nassau County, and Suffolk County. Data records are missing for some months in some counties. We compare GPSS against three competitive baseline algorithms, including GP

anomaly detection [4, 12], one-class SVM [10], and robust multivariate outlier detection using the Mahalanobis distance [7, 8]. We apply the GPSS and baseline approaches jointly to data across all locations and time steps, performing randomization testing at $\alpha = .05$ to identify significant clusters. GPSS identifies two statistically significant anomalous patterns in the data, shown in blue circles and red crosses in Figure 1 (left panel). The baseline anomaly detection methods (Figure 1, right panel) failed to discover a coherent anomalous pattern, instead selecting individual points across space and time.

The anomalies detected by GPSS correspond to important public health events. The blue circles at the end of 2015 indicate a surge in opioid deaths corresponding to a well known plague of fentanyl-related deaths in NYC [2]. The anomaly denoted by red crosses in 2006 is particularly interesting since it indicates a spike in opioid deaths immediately preceding the introduction of community training programs to administer a lifesaving naloxone drug. This may indicate a surge in fatalities that was cut short by making naloxone more widely available and educating communities in its use.

3.2 Case study 2: victim-level case data

For our second case study [6], we used the MDTS approach described above to analyze a publicly available dataset from the Allegheny County, PA medical examiner's office and to detect emerging overdose patterns and trends. The data consists of approximately 2000 fatal accidental drug overdoses between 2008 and 2015. For each overdose victim, we have date, location (zip code), age decile, gender, race, and the presence or absence of 27 commonly abused drugs in their system. The highest-scoring clusters discovered by MDTS were shared with Allegheny County's Department of Human Services and their feedback obtained.

One set of potentially relevant findings from our analysis involved fentanyl, a dangerous and potent opioid which has been a serious problem in western PA. In addition to identifying two well-known, large clusters of overdoses (14 deaths in Jan 2014 and 26 deaths in Mar-Apr 2015), MDTS was able to provide additional, valuable information about each cluster. For example, the first cluster was likely due to fentanyl-laced heroin, while the second was more likely due to fentanyl disguised as heroin (only 11 of the 26 victims had heroin in their system). Moreover, the second cluster was initially confined to the Pittsburgh suburb of McKeesport and the typical overdose demographic of white males ages 20-49, before spreading across the county. Our analysis demonstrated that prospective surveillance using MDTS would have identified the cluster as early as March 29th, enabling more targeted prevention efforts.

MDTS also discovered a previously unidentified, highly localized cluster of fentanyl-related overdoses affecting an unusual and underserved demographic (elderly black males near downtown Pittsburgh). This cluster occurred in Jan-Feb 2015, and may have been related to the larger cluster of fentanyl-related overdoses that occurred two months later. Finally, we identified multiple overdose clusters involving combinations of methadone, commonly used to treat heroin addiction, and the prescription drug Xanax between 2008 and 2012. We observed dramatic reductions in these clusters
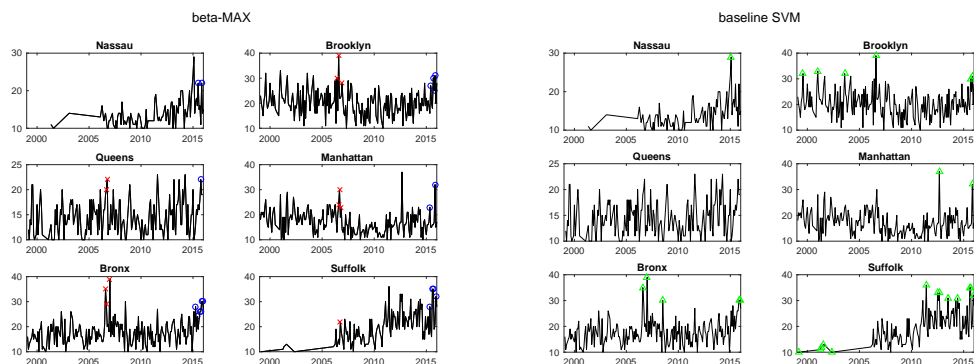
Figure 1: Monthly opioid overdose deaths in the NYC metropolitan area from 1999-2015. In the left panel, the two statistically significant anomalies detected by GPSS are depicted as red crosses and blue circles. The right panel shows significant anomalies detected by the one-class SVM.

after 2012, corresponding to the passage of the Methadone Death and Incident Review Act, which increased state oversight of methadone clinics and prescribing physicians.

4. CONCLUSIONS

Taken together, these two case studies suggest high potential utility for prospective surveillance of overdose data using subset scan methods. Detected clusters of overdoses which are geographically, demographically, and/or behaviorally similar can suggest details of the underlying process causing these deaths, and can help local public health agencies to target their prevention and response efforts to specific neighborhoods and subpopulations. These efforts can include outreach to increase the proportion of opioid addicts receiving Medication-Assisted Treatment such as methadone and suboxone, as well as patients' compliance with their treatment regimes. Additionally, increased supplies of the overdose treatment drug naloxone can be directed to potential overdose victims and their local communities. In addition to facilitating targeted and effective health interventions by prospective surveillance, these approaches can also be used on retrospective data to better understand the effects of drug legislation and other policy changes. Our ongoing work will integrate these two methods to better handle multidimensional, correlated space-time data, and apply them to new overdose-related data sources, including data from the state of Kansas's prescription drug monitoring program, in collaboration with public health partners.

5. ACKNOWLEDGMENTS

6. REFERENCES

[1] Centers for Disease Control and Prevention. WONDER Database. 2017.

[2] City of New York Office of the Mayor. HealingNYC: Preventing Overdoses, Saving Lives. 2017.

[3] W. Herlands, E. McFowland III, A. G. Wilson, and D. B. Neill. Gaussian process subset scanning for anomalous pattern detection in non-iid data. *Working paper*, 2017.

[4] K. Kowalska and L. Peel. Maritime anomaly detection using Gaussian process active learning. In *Proc. 15th Intl. Conf. on Information Fusion*, pages 1164–1171. IEEE, 2012.

[5] D. B. Neill. Fast subset scan for spatial pattern detection. *Journal of the Royal Statistical Society B*, 74(2):337–360, 2012.

[6] D. B. Neill. Multidimensional tensor scan for drug overdose surveillance. *Online Journal of Public Health Informatics*, 9(1):e20, 2017.

[7] P. J. Rousseeuw and A. M. Leroy. *Robust regression and outlier detection*, volume 589. Wiley, 2005.

[8] P. J. Rousseeuw and B. C. Van Zomeren. Unmasking multivariate outliers and leverage points. *J. Amer. Stat. Assoc.*, 85(411):633–639, 1990.

[9] R. Rudd, N. Aleshire, J. Zibbell, and R. Gladden. Increases in drug and opioid overdose deaths: United states, 2000-2014. *Morbidity and Mortality Weekly Report*, 64(40):1378–1382, 2016.

[10] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001.

[11] S. Speakman, S. Somanchi, E. McFowland III, and D. B. Neill. Penalized fast subset scanning. *J. Comput. Graph. Stat.*, 25(2):382–404, 2016.

[12] O. Stegle, S. V. Fallert, D. J. MacKay, and S. Brage. Gaussian process robust regression for noisy heart rate data. *IEEE Transactions on Biomedical Engineering*, 55(9):2143–2151, 2008.

[13] United States Department of Health and Human Services. The opioid epidemic: By the numbers. 2016.

[14] S. Yu, L.-C. Tranchevent, B. De Moor, and Y. Moreau. *Kernel-based data fusion for machine learning*. Springer, 2013.