Deep Learning & the Opioid Epidemic: Estimating Opioid-Related Mortality Risk in US Counties with Twitter Data

Zoe-Alanah Robert (*zrobert7@cs.stanford.edu*)
CS 230: Deep Learning, Stanford University

Predicting

Around 100 Americans die from an opioid-related overdose every day [1] According to the National Institute on Drug Abuse, an opioid overdose can be effectively reversed if the drug Narcan (naloxone) is given soon after overdose[2]. I built a deep model which can predict county-level opioid related mortality rates in 2016 using prescribing rates and mortality rates from the previous eight years. Such a model could be used to flag at-risk communities early on so that government prevention resources can be targeted towards these communities. In addition, the model could aid in redistribution of Narcan (which is often in short supply) to those who need it most. After finding promising results, I investigated the use of Twitter data (county-tagged tweets containing both medical and slang terms for opioid drugs) to predict future mortality rates in the user's county.

Data/ Features

A. **Historical Opioid Prescribing Rates** from CDC [3]: Opioid prescribing rate per 100 people per county in the US, per year (2008-2016)
B. **Opioid-Related Deaths** from CDC Wonder[4]: crude and age-adjusted mortality rates per county, per year (2008-2016); the data is further split into the Cause of Death Codes (CODs) related to Opioids as outlined by drugabuse.gov[5]
C. **Tweets Mentioning Opioids**[6]: a database of tweets that mention any of the official "street" or medical names for Opioids[7]

Historical Task
Input: 98-feature vector containing prescribing rates 2008-2016, 2015 population, and raw death counts per COD per year 2008-2015
Output: 2016 Opioid-related mortality rate per 1000 persons

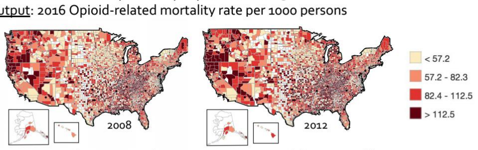

Figure 1: Opioid prescriptions per 100 people by us county[3]

Twitter Task
Input: a county-level Doc2Vec tweet embedding for all tweets mentioning Opioids in a single week - stripped of non-alphanumeric characters and links
Output: 2016 Opioid-related mortality rate per 1000 persons
As some counties contain fewer relevant tweets than others, all county-level Doc2Vec embeddings are padded to the max length, resulting in input vectors of 140300 features. It is also important to note that tweets are tagged with a (lat,lon) which is referenced against the official FIPS[8] county center (lat,lon) and an average county radius to obtain county labels.

'Them percs put me in my feelings every time'

'Taskforce Completes Successful Opioid Bust Near Tribal Reservations in Arizona - Native News Online <link>'

'@kasie When I was 1st prescribed Oxy for pain, I was told it was not addictive. For the next 6 years I was hopeless…

'where da percs at'

'Order quality Oxycodone, Percocet, Dilaudid, Codeine, xanax, viagra, Nembutal, Ritalin, Valium, Roxicodone and many… <link>'

Figure 2: Example tweets

Models

This project involves two distinct regression prediction tasks. For both, I explored the following elements with 10-fold cross validation:

- learning rate (exponential search)
- #epochs
- batch size
- MSE loss vs R^2 loss
- effects of gradient clipping
- dropout & dropout hyperparameters
- input normalization
- Adam vs RMSProp

Historical Task: Each of NNx (NN1, NN2, and NN3) contain x dense layers with relu activation. NN4 structure is dense layer + dropout + dense layer + dropout + dense layer + relu activation.

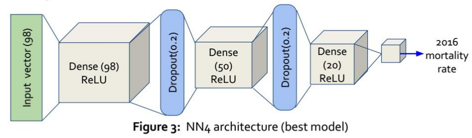

Figure 3: NN4 architecture (best model)

Twitter Task: Each of NNx_twiter contain x Dense layers with relu activation and MSE loss. NNx_twiiter_r2 uses R^2 loss instead. And NN1_twitter_r2dropout is similar to NN1_twitter_r2 but with additional dropout layers after the dense layers as shown below.

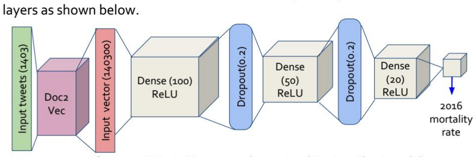

Figure 3: NN3_twitter_mse_dropout architecture (best model)

Results

Historical	Model Architectures	Train R^2	Test R^2
	Baseline linear*	0.37	0.27
	NN1	0.08	0.03
	NN2	0.19	0.01
	NN3	0.39	0.20
	NN4	0.57	0.37
Twitter	**Baseline linear ***	0.53	-0.58
	NN1_twitter	-0.06	-0.24
	NN3_twitter	0.62	-0.40
	NN1_twitter_r2	-0.04	-0.13
	NN3_twitter_r2	0.69	-0.46
	NN1_twitter_r2dropout	-0.02	-0.07
	NN3_twitter_r2dropout	0.58	0.08
	NN3_twitter_mse_dropout	0.61	0.12

Table 1: Results for all tasks and all models, avg of 10-fold (train/dev/test = 461/86/61)
*For the Baseline linear, four separate linear models were run (linear regression, ElasticNet, Lasso, Ridge) and the results from the best performing are displayed here.

Discussion

Although there are over 3000 counties in the US, due to limited data availability, this set was reduced to 608 after preprocessing. Given the small data set and the relatively small input feature vector size for both tasks, a batch size of the full length of X was chosen. All training, dev, and test set input data was normalized. In all models, gradient clipping proved effective in reducing the effects of gradient "cliffs." Hyperparameter tuning resulted in a top learning rate of 0.001 for both learning tasks. It is important to note that the MSE values for these top performing models reached values as low as 0.01 on test data for both tasks.

Historical Task
The NNx results in the linear baseline models for the Historical Task suggest that increasing the number of dense layers improves training results, but that the model does not generalize well. With the addition of the Dropout layers in NN4, we see a big improvement in the model's performance on test data. The limited improvement between NN4 and the best linear baseline model suggest that the problem of 2016 mortality rate prediction from previous years' mortality rates and opioid prescribing rates is a relatively trivial task.

Twitter Task
Comparing the results from different model architectures, we see that the Twitter task is much more difficult to learn. The addition of dense layers improves training metrics but worsens test metrics, suggesting the need to reduce model complexity. The use of 1-R^2 loss over MSE improves both training and test set metrics ,likely due to the relatively small numerical value of mortality rates. Finally, with many dense layers and the addition of dropout layers, we achieve the best test set performance.

Future

Given additional time and resources, further improvements could be made in the Historical Task with the addition of various publicly available datasets (drug addiction facility admissions, poverty rates, median income, historical opioid price data, self reported use & misuse, etc). I would explore these along with different means of representing missing data elements other than zero initializations. For the Twitter task, a major hurdle in this project has been the lack of financial resources to access the Premium Twitter API - the Twitter data collected for this project was limited to the Standard free Search API, which only contains tweets from the last seven days. Ideally, I would use tweets from 2015 to predict 2016, for example. Additionally, I believe it would be useful to explore tweet classification beyond the mention of an opioid related term. (i.e. How can we filter out news articles about the Opioid crisis, song lyrics, and other less-relevant tweets?)

References

1. "Opioid Overdose," *Centers for Disease Control and Prevention*, 30-Aug-2017. [Online]. Available: https://www.cdc.gov/drugoverdose/epidemic/index.html. [Accessed: 07-Jun-2018].
2. National Institute on Drug Abuse, "Opioid Overdose Reversal with Naloxone (Narcan, Evzio)," *NIDA*, 04-Apr-2018. [Online]. Available: https://www.drugabuse.gov/related-topics/opioid-overdose-reversal-naloxone-narcan-evzio. [Accessed: 07-Jun-2018].
3. "Opioid Overdose," Centers for Disease Control and Prevention, 31-Jul-2017. [Online]. Available: https://www.cdc.gov/drugoverdose/maps/rxcounty2008.html. [Accessed: 07-Jun-2018].
4. "CDC WONDER", Wonder.cdc.gov, 2018. [Online]. Available: https://wonder.cdc.gov/. [Accessed: 07- Jun- 2018].
5. "Opioid Summaries by State", Drugabuse.gov, 2018. [Online]. Available: https://www.drugabuse.gov/drugs-abuse/opioids/opioid-summaries-by-state. [Accessed: 07- Jun- 2018].
6. "Premium search", Developer.twitter.com, 2018. [Online]. Available: https://developer.twitter.com/en/docs/tweets/search/overview/premium. [Accessed: 07- Jun- 2018].
7. Drugfreeworld.org, 2018. [Online]. Available: https://www.drugfreeworld.org/drugfacts/prescription/opioids-and-morphine-derivatives.html. [Accessed: 07- Jun- 2018].
8. G. Branch, "2010 FIPS Code Files for Counties - Geography - U.S. Census Bureau", Census.gov, 2018. [Online]. Available: https://www.census.gov/geo/reference/codes/cou.html. [Accessed: 07- Jun- 2018].