Deep Learning for Predicting Hospital Admissions due to Chronic Disease Exacerbation during Periods of Bad Air Quality

Govind Chada
Stanford University
gchada@stanford.edu

Ekin Tiu
Stanford University
ekintiu@stanford.edu

Kabir Jolly
Stanford University
kjolly@stanford.edu

David Kim*
Stanford University
davidak@stanford.edu

Abstract

Given the recent increases seen in wildfires, especially in areas such as the Western United States, prolonged periods of poor air quality have become more prevalent in nearby residential areas. Consequently, certain populations are exposed to dangerous levels of PM 2.5, posing a threat to medically at-risk individuals. Respiratory conditions such as chronic obstructive pulmonary disease (COPD) and heart conditions such as coronary artery disease (CAD), as well as neurological conditions such as Parkinson's and Alzheimer's, have been shown to worsen as a result. This study focuses on leveraging longitudinal deep learning models to predict if a patient will be admitted to a hospital due to these conditions during a period of poor air quality. This process takes place in two steps: first, embeddings are created from patient cohorts using a Word2Vec model and second, these embeddings are fed into an LSTM model to create the final predictions. Using this methodology, the model obtained AUROC scores of 0.93 for COPD, 0.94 for CAD, 0.97 for Parkinson's, and 0.96 for Alzheimer's. This supports the notion that the model is able to exhibit strong predictive performance across these four conditions and has the potential to greatly aid intervention efforts in order to mitigate the consequences of chronic diseases during periods of bad air quality.

1 Introduction

1.1 Problem Statement

This study investigates which patients will have Emergency Department/hospital admissions during periods of bad air quality due to chronic disease (e.g. COPD, CAD) exacerbations. Furthermore, it focuses on determining which features account for this risk. The information was given to us by Dr. David Kim regarding the clinical relevance of this question: the reason this question is interesting is because the basic epidemiology is still poorly understood. In effect, it is unknown how bad these levels of smoke are for people, and for whom in particular. Therefore, combining a descriptive and predictive approach could have a significant impact. This problem is also especially topical considering the increasing prevalence of wildfires, as seen in Northern California, leading to periods of bad air quality.

1.2 Novel Contributions

There have been no prior applications of longitudinal models to patient's hospital visits to predict if they will be hospitalized during periods of bad air quality due to wildfires. The goal of this work is to

*Dr. David Kim is not enrolled in the course.

34th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada.

associate historical air quality data with EHR data to predict who will suffer severe health effects during these periods of bad air quality. A longitudinal deep learning model will be used to predict who is likely to suffer during these periods, so as to target preventive interventions.

1.3 Summary of Results

The longitudinal model is able to predict COPD with an AUROC of 0.93, CAD with an AUROC of 0.94, Parkinson's with an AUROC of 0.97, and Alzheimer's with an AUROC of 0.96; the model also had sensitivities greater than 0.89 for all 4 diseases that were investigated. These results show that a longitudinal model is able to accurately predict who is likely to be hospitalized due to chronic diseases during periods of bad air quality, allowing for possible preventative interventions in the future.

2 Related Work

While there is no existing work directly addressing the problem we are investigating, we review related work in the following areas: (1) longitudinal deep learning models applied to the field of healthcare, (2) the effects of wildfires on poor air quality, and (3) studies of the effects of bad air quality on certain conditions.

Leveraging machine learning and deep learning models to analyze longitudinal healthcare data has been a topic of interest in recent years. More specifically to our work, recurrent neural networks (RNNs) and long short-term memory networks (LSTMs) have been shown to effectively analyze health records for a variety of tasks. Notable work in this space includes using EHR data to predict cardiovascular events (Zhao et al., 2019), patient readmission (Ashfaq et al., 2019), and enhancement of disease coding (Rashidian et al., 2019). Increasing model efficiency and accuracy is still a product of experimentation in all of these studies. Zhao et al. have attempted both machine learning and deep learning based approaches and found that longitudinal models perform with the overall best AUROC scores when predicting the 10-year risk for cardiovascular diseases in the sample dataset. Ashfaq et al. further extend the work of Zhao et al., focusing on not only the reappearance of the heart condition, but hospital readmission as a result of heart failure for the patient. In a similar manner, we will be analyzing COPD and CAD, but linking occurrence of these diseases with periods of bad air quality. The work by Rashidian et al. relates closely to our methodology as we also plan on working towards reducing the inherent variability found within diagnosis coding by testing out different architectures and hyperparameters.

The second component of our related work is looking at the connection between wildfire events and periods of bad air quality. Related work from Phuleria et al. examine the effects of wildfires in Southern California on air quality levels and how they result in fine particulate matter (PM 2.5) concentrations that have adverse health effects (2003). Similarly, Viswanathan et al. study the effects of San Diego wildfires on air quality. The findings from these two papers explore not only the Air Quality Index (AQI) concentrations, but also respiratory condition counts and a county-wise computational analysis in locations relevant to our target region (California). As we work towards exploring a potential connection between wildfire events and patient condition recurrence, this related work helps create the partial link between wildfires and poor AQI levels. Much of our early data setup work revolved around finding extended periods of poor air quality resulting from wildfire events to properly select patients who visited the hospital during this period of time due to chronic diseases.

Risk of health conditions resulting from bad air quality is the last aspect to our work. A 2015 study by Haikerwal et al. describes the effects of PM2.5 on cardiovascular health. This is directly connected to the recent work conducted by Shi et al., which looks at PM2.5 as well and its effect on neurological disorders (2020). Related work also includes studying the effects of air pollution on the risk of death resulting from COVID-19, a topic that is especially relevant in the current research landscape (Pozzer et al., 2020). As for a qualitative analysis of our results, as well as understanding the potential impacts of our work, research done by Jiang et al. in 2016 and Schikowski et al. in 2014 clarifies the connection between air pollution levels and its effect on COPD, one of the target conditions in our work.

This work uses an established technical method, LSTMs, in a novel application area. From a review of the literature discussed above, it is evident that wildfires create periods of bad air quality, which in

turn causes increased prevalence of certain medical conditions. The goal is to capture this connection using deep learning on EHR data, as we seek to employ LSTMs to predict condition occurrence during these periods of unhealthy air quality.

3 Data

Two datasets were used for this research study.

The first dataset, referred to as the "EHR Dataset" is a private 5.3GB csv containing emergency department/hospital admissions information from electronic health records. The data was collected from October 2015 - December 2018 in California county hospitals. This data is securely provided to our research group by our mentor, David Kim.

The second dataset, referred to as the "AQI Dataset" is a public air quality index dataset from the EPA. This dataset consists of air quality information from 2015-2018 recorded on a daily basis and organized by county. A subset of this dataset was selected, containing information about California counties, and is dependent on PM2.5-based AQI measurements, as PM2.5 is the primary pollutant in wildfire smoke. The processed dataset has 25,644 data points.

3.1 Features of EHR Dataset

The EHR Dataset consists of the following features. Each feature is supplemented with a brief description. The specific features we are paying attention to include:

Constants

- Type = ED visit or hospital admission
- Sex = male/female/unknown
- Race = patient-declared race
- Language = primary language of patient.
- Payer = primary payer for that visit (Medicare, Medicaid, private, etc)
- Pt county = county codes of the patient
- Fac county = county of facility/hospital

Variables

- Age = age in years at time of visit/admission
- Dispo = patient's disposition from that visit, e.g., home, admitted, transferred, died, etc.
- Dx10 prin, Dx10 1-24 = up to 25 ICD-10 diagnosis codes associated with that visit
- Proc prin, Proc 1-20, Proc10 prin, Proc10 1-20 = Procedure codes.

3.2 Features of AQI Dataset

The AQI Dataset consists of the following features. Each feature is supplemented with a brief description. The features below were used for pre-processing and cohort creation:

- State Name = name of the state in which the measurements were taken
- County Name = name of the county in which the measurements were taken
- County Code = FIPS county codes
- **Date = Y/M/D date string**
- **AQI = air quality index at time and place of measurement**
- **Category = category of AQI: Good, Moderate, Unhealthy for Sensitive Groups, Unhealthy, Very Unhealthy, and Hazardous**
- Defining Parameter = parameter used to determine measurement; only entries with PM2.5 as their defining parameter will be used in this study to focus on the impact of wildfire smoke pollution

3.3 Sample of AQI Data

State Name	county Name	County Code	Date	AQI	Category	Defining Parameter
California	Alameda	1	2015-01-01	61	Moderate	PM2.5
California	Alameda	1	2015-01-02	76	Moderate	PM2.5
California	Alameda	1	2015-01-03	96	Moderate	PM2.5
California	Alameda	1	2015-01-04	99	Moderate	PM2.5
California	Alameda	1	2015-01-05	91	Moderate	PM2.5
...	...	...	...	...	...	...
California	Yolo	113	2018-12-27	33	Good	Ozone
California	Yolo	113	2018-12-28	32	Good	Ozone
California	Yolo	113	2018-12-29	45	Good	PM2.5
California	Yolo	113	2018-12-30	75	Moderate	PM2.5
California	Yolo	113	2018-12-31	31	Good	Ozone

Figure 1: Sample rows of AQI Dataset.

3.4 Data Usage

We use the above features in our model in order to predict whether each patient will have an ED visit or hospital admission for certain diseases during a pre-selected period of bad air quality.

With regards to the AQI dataset, we have pre-selected a period of poor air quality (PM2.5 > 100) for each California county represented in the EHR dataset. The time period for each county will be taken into account when training the model.

We plan to analyze the admission of patients who were diagnosed particularly with respiratory and cardiovascular diseases. These include ICD-10 codes J00-J99 (respiratory) and ICD-10AM, I00 to I99 (cardiovascular that have been linked to concentrations of PM2.5 in the atmosphere).

3.5 Relevant Statistics

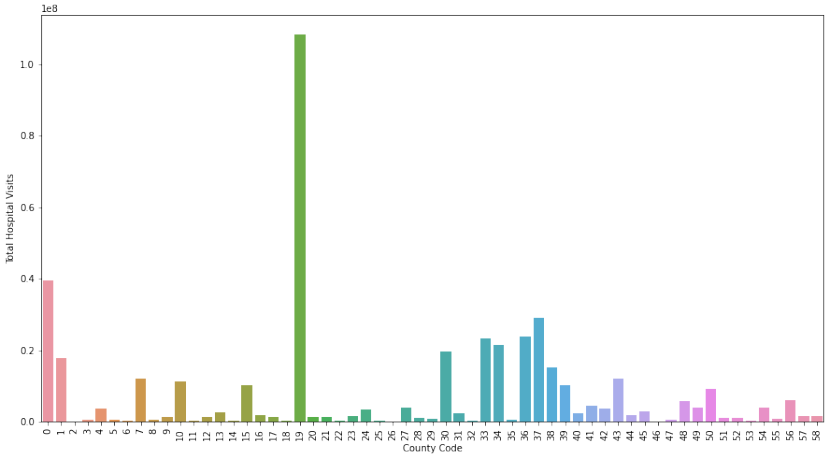

Figure 2: Countplot of number of visits per California county.

Figure 2 aided in the preliminary stages of our project by helping us choose sample counties that had a significant number of COPD and CAD occurrences for data exploration.

Some relevant statistics from our patient dataset include:

- There are a total of 27,977,932 patient visits in our general dataset.

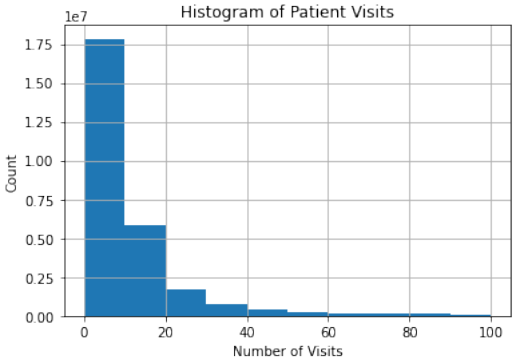

Figure 3: Histogram of EHR Visit Counts.

- There are 2,602,731 patients who have >= 30 visits. This is relevant in determining the maximum number of timesteps (number of visits prior to predefined period of bad air quality) we want to allow for each patient. Our group decided that < 30 visits was appropriate.

4 Approach

Our methods for cohort selection for this project are as follows:

We began by defining a time period containing measurements indicative of bad air quality (AQI > 100 with defining parameter PM2.5), as defined by the EPA (US Environmental Protection Agency, 2019). We selected a unique time period for each county that is represented by the population of patients in the EHR Dataset.

We then selected a cohort of patients who were diagnosed with the relevant diseases (COPD, CAD, Alzheimers, Parkinsons) during our predetermined periods of bad air quality. We create a dataset for each disease by filtering by the following ICD-10 Codes: J44 (COPD), I25 (CAD), G20 (Parkinson's), G30 (Alzheimer's). These patients are our 'positive cases'.

We also sampled 'negative cases' by randomly selecting patients who were not diagnosed with relevant diseases during our pre-defined time periods. The number of negative samples is equal to the number of positive samples to create a balanced dataset.

We pre-processed this dataset by applying Word2Vec on our features in order to obtain embedded 100-vectors for each feature. This was then provided as input to our baseline and longitudinal models. Please see the following sections for more details.

4.1 Word2Vec Implementation

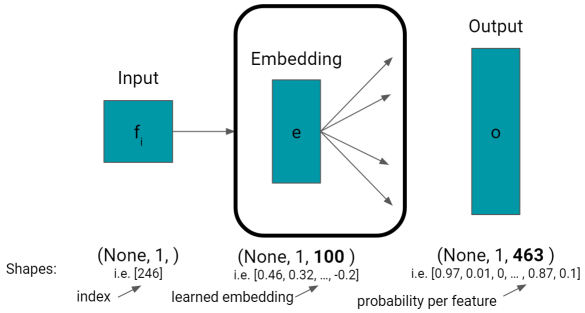

Figure 4: Visual representation of the Word2Vec Model trained.

We created a 100-vector for each feature (i.e. each diagnosis, each procedure; more concretely ICD10 Codes and CPT Codes) with at least 15 occurrences in our dataset.

We created our Word2Vec training set by treating all of a patient's medical information as a 'document', the medical trajectory (sequence of codes) as 'sentences' and codes for each feature as 'words'.

Each training sample is a pair of features of a patient. For instance, if a patient has diagnosis code D70 (not a real code) and a procedure code P20, a training sample will be (D70, P20). We do the same for each pair. We care about the order of pairs as well.

Our implementation process is detailed below:

1. We obtain our dataset via our cohort selection detailed above. This will result in a dataset consisting of many samples where each sample is a patient trajectory. A patient trajectory is a set of n vectors, each representing a visit of the patient before the pre-defined period of bad air quality. (n is defined as the number of preceding visits <= NUM_VISITS_CAP hyperparameter). We then performed an 80-10-10 train-test-val split.
2. Using all of the patient trajectories from the train samples (including all preceding visits), we created pairs for Word2Vec using the following steps:
3. For each patient visit, we created a pair between all of the features per patient visit. In other words, we obtained all possible permutations between a random subsample of patient visits from our dataset. We considered both orderings (i.e. (a, b) and (b, a)).
4. We trained a Word2Vec implementation (similar to Skip-Gram with minor modifications) with a random sample of pairs obtained from the preceding scheme. This model uses an Input Layer, one Embedding Layer, and one Output Layer. We used the outputs of our Embedding layer as our embeddings. The Embedding layer will have EMBEDDING_LENGTH neurons, where EMBEDDING_LENGTH is a hyperparameter that represents the size of our embedding vectors for each feature.
5. We implemented the model on all of our patient's features. We then averaged across features to obtain one embedding vector for each patient visit.

4.2 Longitudinal Model

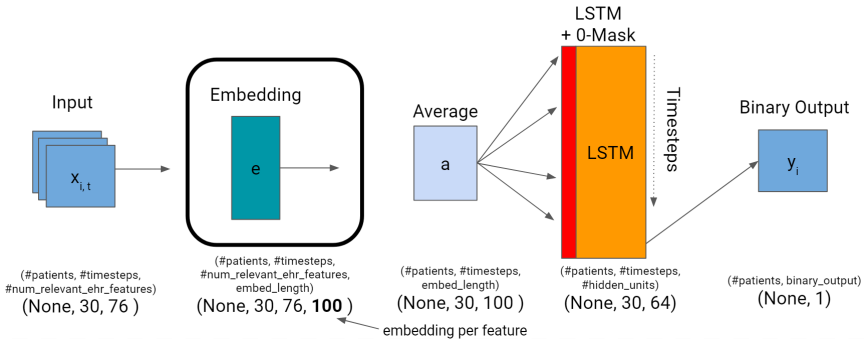

Figure 5: Visual representation of the LSTM model trained.

With the aforementioned feature inputs and a defined output, we train a longitudinal model (LSTM) on the preceding visits of each patient in the cohort, and predict whether each patient will have an ED visit or hospital admission during the period of bad air quality.

The longitudinal model takes as input a dataset with shape (num_patients, num_timesteps, num_relevant_ehr_features). The model output and labels shape will be (num_patients, 1), as the model outputs a binary result. The foundation of this code is based on the LSTM subsection of Assignment 2 of BIODS220.

Lastly, we evaluate the model performance with the methods elaborated upon in the subsequent section, iterate, and gauge if the model performs better than random chance. The model's performance

is also compared to a linear baseline model. We then use interpretability/feature importance techniques to determine what characteristics of patients predict their incurring additional visits during periods of bad air quality.

Our research is attempting to answer a rather novel problem, so there are no related implementations to our knowledge.

4.3 Baselines

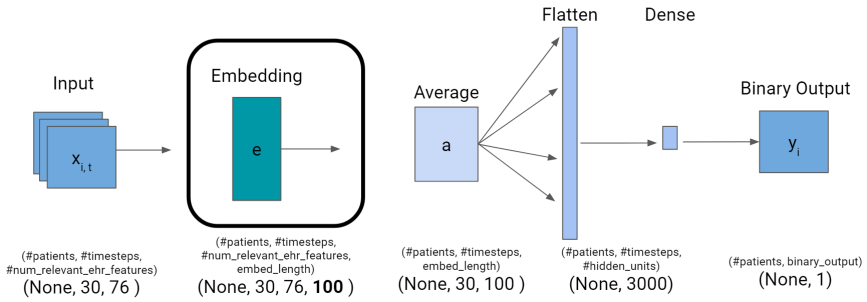

Figure 6: Visual representation of the baseline model trained.

Since there is no established baseline from previous work, we created a linear baseline model in order to evaluate the efficacy of our longitudinal model approach. This linear baseline model accepts the same input as the longitudinal model, and outputs a binary result. This baseline model first flattens the input dataset in the time steps and embedded feature dimensions in order to feed the data into a single neuron Dense layer, which is intended to act as a linear regressor. A sigmoid function is applied in the Dense layer in order to map the result between 0 and 1. Both the baseline model and the longitudinal model are evaluated with the same metrics.

5 Experiments

We look at the following results to evaluate the outcomes of our experimentation. We start by monitoring overall model performance using the standard loss and accuracy metrics for training and validation datasets. We also analyze the receiver operating characteristic (ROC) curve for a comparison between model predictive performance and random chance. In addition to the AUROC metric, we also obtain sensitivity and specificity values by optimizing Youden's index. After tuning hyperparameters and experimenting with various model architectures, we record final results on a separate test dataset for all of the discussed metrics. The ROC evaluation process is built upon code from Assignment 2 of BIODS220.

5.1 Word2Vec

In order to create embeddings to be passed into our predictive models, we trained a Skip-Gram model with context pairs created by taking permutations of features within patient visits.

The model was trained for 2 epochs with the RMSProp optimizer and the binary crossentropy loss function, achieving a maximum accuracy of 0.9981 on the train dataset and 0.9978 on the validation dataset.

5.2 Baseline Dense Model

The baseline model was trained on the embeddings created by the Word2Vec model for COPD, CAD, Parkinson's and Alzheimer's. The baseline model was able to perform reasonably well as shown in Table 1, but was outperformed by the LSTM as will be discussed in section 5.4.

Table 1: Baseline Performance Metrics by Disease

Disease	AUROC	Sensitivity	Specificity	Accuracy
COPD	0.93	0.87	0.86	0.86
CAD	0.92	0.89	0.82	0.85
Parkinson's	0.96	0.86	0.95	0.89
Alzheimer's	0.96	0.92	0.90	0.90

5.3 LSTM

The LSTM was also trained on the embeddings created for COPD, CAD, Parkinson's, and Alzheimer's. The longitudinal LSTM model both surpasses random chance, as well as outperforms the baseline model, showing that our approach is promising for predicting chronic disease exacerbation. The models were each trained for 20 epochs, with best results attained using the RMSProp optimizer and the binary crossentropy loss function. Model evaluation was conducted using test data that was not used for training or validation. 80% of the data was used for training, 10% was used for validation, and 10% was used for testing. Figure 7 shows the resulting ROC curves when evaluating the trained model on the test set.

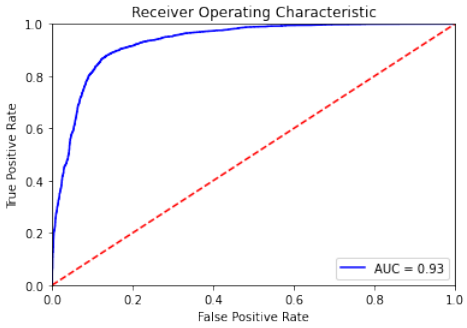

(a) ROC Curve from LSTM model on COPD

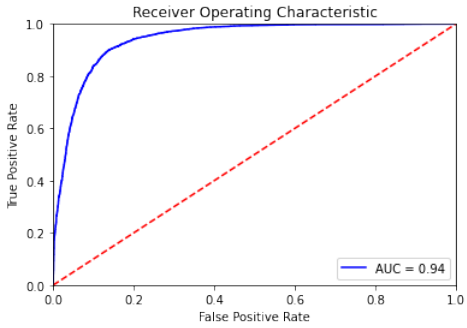

(b) ROC Curve from LSTM model on CAD

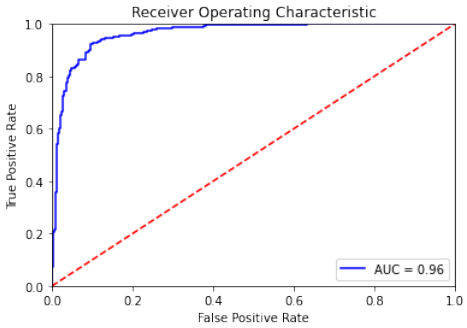

(c) ROC Curve from LSTM model on Alzheimer's

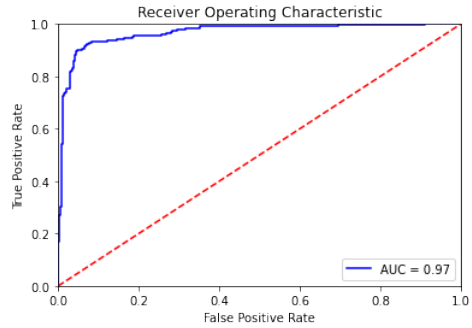

(d) ROC Curve from LSTM model on Parkinson's

Figure 7: Quantitative evaluation of the LSTM model using ROC curves and AUROC scores.

Table 2 reflects the results of the highest performing models upon tuning learning rates and hidden units for each of the condition-specific LSTM models. The RMSProp optimizer was chosen for the final results due to its increase in performance over the Adam optimizer. The learning rate that was found to be the best was 0.001, a value that was able to determine the patterns within the dataset well without overfitting or underfitting. LSTM hidden units were also tuned, and a value of 64 was found to perform the best for the chosen conditions.

Table 2: LSTM Performance Metrics by Disease

Disease	AUROC	Sensitivity	Specificity	Accuracy
COPD	0.93	0.89	0.86	0.88
CAD	0.94	0.90	0.86	0.88
Parkinson's	0.97	0.90	0.95	0.92
Alzheimer's	0.96	0.93	0.90	0.89

5.4 Analysis

The LSTM was able to predict COPD with an AUROC of 0.93 and accuracy of 0.88, compared to an AUROC of 0.93 and accuracy of 0.86 for the baseline model. Additionally, using the LSTM improved the sensitivity from 0.87 to 0.89. This demonstrates that a longitudinal model is better at finding positive cases, probably due to its advantage when using temporal data like hospital visits over time.

Similar to the trend in comparisons seen in COPD results, LSTM performance on CAD had an AUROC of 0.94 and an accuracy of 0.88. This is an improvement over the baseline performance of 0.92 AUROC and 0.85 accuracy value. The biggest difference was in specificity, where the LSTM had a significantly higher true negative rate of 0.86 compared to the baseline's 0.82.

We believe that the results for COPD and CAD show that a longitudinal model is able to take advantage of the temporal relationship between patient visits, which allows for higher predictive ability. This demonstrates how the progression of a patient over time can be a very strong predictor of future diagnoses when paired with a longitudinal model.

For Parkinson's, the LSTM exhibited a stronger predictive power for true positives, resulting in a sensitivity of 0.90 which was 4% higher than that of the baseline model. Furthermore, the LSTM demonstrated a 3% improvement in accuracy over the baseline model. The remainder of the results for Parkinson's, as well as all of the results for Alzheimer's, were very similar between the baseline model and LSTM. The results indicate that the Dense architecture was just as effective at analyzing the majority of neurological cases as the LSTM model.

While the longitudinal model does show improvement for Parkinson's when compared to the baseline, the performance of the 2 models were very similar for Alzheimer's. We believe that this is because the temporal relationship between patient visits matters less when looking at neurological diseases as opposed to respiratory and cardiovascular conditions.

Overall, strong performance metrics were seen for AUROC, sensitivity, specificity, and binary accuracy metrics across the 4 target diseases examined in this study. In most cases, the LSTM model outperformed the baseline Dense model, and was otherwise at par with the baseline. Both architectures significantly outperformed random chance. Evidently, the longitudinal model is capable of accurately predicting the occurrence of these conditions during periods of bad air quality with low false positive and false negative rates.

6 Conclusion

6.1 Key Results and Learnings

Our longitudinal model shows promising results for the prediction of all 4 chronic diseases that we investigated. The model predicted COPD with an AUROC of 0.93, CAD with an AUROC of 0.94, Parkinson's with an AUROC of 0.97, and Alzheimer's with an AUROC of 0.96. Additionally, the model had a sensitivity of greater than 0.89 for all 4 diseases, showing that cases of the diseases are rarely missed. These results demonstrate the model's ability to exhibit strong predictive performance for these chronic diseases. Furthermore, these results show that there is potential to use longitudinal models to target preventative interventions by accurately predicting who will be hospitalized with these chronic diseases during periods of bad air quality.

6.2 Future Directions

There are multiple approaches that can be used in the future to improve our models' performance and interpretability.

The first approach is modifying our methodology for obtaining context pairs. Currently, we only take permutations of features within singular patient visits. This approach can extend to permutations across visits temporally. This will allow the model to learn temporal feature relationships, thereby improving the embedding of features with respect to understanding patient trajectory over time. This also makes certain permutations within visits redundant, thereby decreasing the computational complexity of the pairs matching algorithm.

The second approach is to augment the Skip-Gram model with techniques such as negative sampling. This approach has been demonstrated to increase the performance of Word2Vec models while simultaneously decreasing computational complexity. Thus, this seems to be a logical next step for our project.

Another approach is to add an attention mechanism as part of the pipeline, which would allow for a detailed analysis of which features are most predictive of hospital admission during periods of bad air quality. Another future step would be to perform a sanity check on the the outputs of our Word2Vec model by using t-SNE to visualize the 100-vector to ensure that the embedding captures the empiric relatedness of certain classes of events via clustering in space.

Another step that can be taken in the future is analyzing which cases the model is not performing well on to see if there is a systematic failure on certain groups of patients.

Contributions

Paper Contributions

K.J. wrote the Abstract, Related Work, and Experiments. E.T wrote the Approach, Data, and Future Directions. G.C. wrote the Introduction, Conclusion, Analysis, and Data. G.C., K.J., and E.T. contributed to final revisions for the entire paper.

Project Contributions

G.C. wrote the code for defining the periods of bad air quality for each county and used this data to create cohorts for each disease that was investigated. E.T. developed the deep learning models (Word2Vec, Baseline, LSTM), and evaluation metric code and packaged them into scripts for easy usage and experimentation. K.J. worked on writing the code for creating context pairs and tuning hyperparameters to create the final models for each disease. D.K. provided our research group with the EHR dataset for patients from 2015-2018 and high-level guidance.

References

[1] Zhao, J., Feng, Q., Wu, P. et al. Learning from Longitudinal Data in Electronic Health Record and Genetic Data to Improve Cardiovascular Event Prediction. *Oxford University Press*, 717 (2019). https://doi.org/10.1038/s41598-018-36745-x

[2] Ashfaq, A., Sant'Anna, A., Lingman, M., Nowaczyk, S. (2019). Readmission prediction using deep learning on electronic health records. *Journal of Biomedical Informatics*, 97. doi:https://doi.org/10.1016/j.jbi.2019.103256

[3] Rashidian, S., Hajagos, J., Moffitt, R. A., Wang, F., Noel, K. M., Gupta, R. R., Tharakan, M. A., Saltz, J. H., Saltz, M. M. (2019). Deep Learning on Electronic Health Records to Improve Disease Coding Accuracy. *AMIA Joint Summits on Translational Science*, 2019, 620–629.

[4] Phuleria, H. C., Fine, P. M., Zhu, Y., and Sioutas, C. (2005), Air quality impacts of the October 2003 Southern California wildfires, *J. Geophys. Res.*, 110, D07S20, doi:10.1029/2004JD004626.

[5] Viswanathan, S., Eria, L. , Diunugala N., Johnson, J., McClean, C. (2006). An Analysis of Effects of San Diego Wildfire on Ambient Air Quality, *Journal of the Air Waste Management Association*, 56:1, 56-67, doi: 10.1080/10473289.2006.10464439

[6] Haikerwal, A., Akram, M., Monaco, A. D., Smith, K., Sim, M. R., Meyer, M., . . . Dennekamp, M. (2015). Impact of Fine Particulate Matter (PM 2.5) Exposure During Wildfires on Cardiovascular Health Outcomes. *Journal of the American Heart Association*, 4(7). doi:10.1161/jaha.114.001653

[7] Shi, L., Wu, X., Yazdi, M. D., Braun, D., Awad, Y. A., Wei, Y., . . . Zanobetti, A. (2020). Long-term effects of PM2.5 on neurological disorders in the American Medicare population: A longitudinal cohort study. *The Lancet Planetary Health*. doi:10.1016/s2542-5196(20)30227-8

[8] Pozzer, A., Dominici, F., Haines, A., Witt, C., Münzel, T., Lelieveld, J. (2020). Regional and global contributions of air pollution to risk of death from COVID-19, Cardiovascular Research, https://doi.org/10.1093/cvr/cvaa288

[9] Jiang, X. Q., Mei, X. D., Feng, D. (2016). Air pollution and chronic airway diseases: what should people know and do?. *Journal of thoracic disease*, 8(1), E31–E40. https://doi.org/10.3978/j.issn.2072-1439.2015.11.50

[10] Schikowski, T., Mills, I., Anderson, R., Cohen, A., Hansell, A., Kauffmann, F., . . . Kunzli, N. (2014). Ambient air pollution: A cause of COPD? *European Respiratory Journal*, 43, 250-263. doi:10.1183/09031936.00100112

[11] US Environmental Protection Agency (2019) What are the Air Quality Standards for PM? Retrieved from https://www3.epa.gov/region1/airquality/pm-aq-standards.html