Wildfire Forecasting

Bryant McArthur, Sam Goldrup, Max Nielsen, Jonah Sundrud

March 2023

Abstract

The purpose of this paper is to find ways to predict the amount and intensity of wildfires that will happen in a given time period across a certain region in the United States. Using models such as ARIMA, HMM, and Facebook's "Prophet" model, we attempt to model both the amount of new fires started and the size of these fires across various regions and states. Our hypothesis is that we will be able to make some fairly accurate predictions about fire risk in the near future with our mathematical models.

1 Problem Statement and Motivation

Wildfires are some of the most destructive and unpredictable natural disasters in the United States. In 2022 alone, 68,988 different wildfires burned around 7.6 million acres in the U.S. [Kat23]. Resources to fight wildfires can be relatively scarce, and the amount of land that firefighting agencies must protect is tremendous. Predictions about when certain areas will be especially prone to fire danger in the future are extremely valuable because authorities can concentrate the necessary resources in the region ahead of time. Current models that evaluate fire risk only make predictions out to a few days in advance using observations of factors like temperature, precipitation, etc. Our goal is to make more long-term predictions that will aid in the efficient allocation of resources and manpower.

2 Data

Our data is the Fire Program Analysis fire-occurrence database (FPA FOD) downloadable on kaggle.com. This dataset contains the records of over 1.88 million US fires between the years of 1992 and 2016. The data was collected by federal, state, and local fire authorities to help the national Fire Program Analysis (FPA) system.

The database originally contained several features for every fire occurrence including: the start day of year, year, fire size (acres burnt), fire size class, state, exact coordinate location, and many others. This database had several fires per day, so in order to create a time-series dataset we needed a single data point at each regular time interval. To deal with this complexity we feature engineered our rows to count either the total number of fires, acreage burnt, or the fire class of the worst fire per day.

We created several time-series corresponding to different geographical areas. We kept the whole continental US as our first time-series, then split by the geographical areas: South, Southwest, Midwest, West, and Northeast. We also consider the case of forecasting wildfires for a given state, such as Utah.

3 Methods

3.1 Classical Decomposition

The first method that we made use of is the Classical Decomposition $Z_t = T_t + S_t + R_t$. We used the statsmodel package to decompose our observed data into the trend, seasonal, and residual components. By taking out the seasonal component we are able to make our time-series more covariance stationary to improve our accuracy of parameter fitting and model prediction.

3.2 ARMA Model

We fit ARMA models on the residuals of our decomposed **daily** and **monthly** time-series data for each geographical region. We found the best parameters for our ARMA model to be 2 for p and 1 for q by a simple grid-search testing p and q values from 1-4 to give us an AMRA(2,1) model.

We trained the models from January 1, 2000 to July 1, 2014. We dropped the data from before the year 2000 because the data seemed inconsistent and we are unsure if all the fires were accurately reported. Next, we used our fit ARMA model to forecast a full year into the future from July 2, 2014 to July 1, 2015. Afterward we added back in the trend and seasonal components to get a more realistic forecast.

3.3 HMM i

We implement a Hidden Markov Model to predict the most severe fire on a given day. Wildfire severity classifications are denoted by a letter A through G, which corresponds to the amount of acreage burned.

We engineer wildfire data to get the most severe wildfire on a day. Letters A through G are mapped to $1, 2, 3, 4, 5, 6, 7$ If there is no wildfire document, a 0 is assigned.

Despite an interest in investigating Wildfire patterns in Utah, there is much greater variation in (worst) wildfire severity in California during the data collection period. Using California data we conduct the following experiments:

1. For each month in 2013, use an HMM trained on the previous month to make predictions.

2. For each wildfire season in the years 2009-2015, generate predictions from an HMM trained on the previous year's series. According to the `Western Fire Chiefs Association`, the wildfire season in California is May-October [Ass22].

For each prediction period (month/year) we take the first observation z_0 and choose $\mathbf{x_0}$ via $\underset{x_0}{argmax} P(z_0|x_0)$ (or $\underset{x_0}{argmax}\ B_{z_0,x_0}$). This departs from the standard from sampling from the initial state distribution π. However, given considerable heterogeneity across sampling periods (especially months, as the results will show), this method gives us a better initial x_0 to start the prediction sequence $\mathbf{z}$. In order to choose the dimension of the latent space, we implement `this AIC Calculation`: $AIC = -2\ell(z|\Theta) + 2p$, where $p = m^2 + nm - 1$, and m, n are the dimension of the latent space and observation space, respectively [Wit19].

3.4 HMM ii

We also implemented a Continuous State Gaussian Hidden Markov Model in order to predict the amount of fires in a given region during a future time period. We are able to treat our space as continuous if we consider a big enough region because there will almost always be some fire activity at any time of the year. We decided that it would be best to first classically decompose the data so that we could train on the residuals and then later add the seasonal and trend components back into our predictions. Before we could make our predictions, we had to estimate the model parameters and decide how many components to use in our hidden states. hmmlearn's "fit" method uses the Kalman filter coupled with expectation-maximization to find the most likely parameters for the system given a predesignated number of components to use. We iterated through models with 5-10 components and kept the model with the highest score (as measured by the log probability under the model). We then made daily and monthly predictions of the hidden states for the training data and used the last hidden state coupled with the transition matrix given by the expectation-maximization step to make forecasts (assuming zero noise) one and two years into the future. For each estimate, we would find the most likely hidden state of the next day/month (as given by our transition matrix), and draw a sample from the distribution of the corresponding observed state. Finally, we added the trend and seasonality components back into to our forecasts and plotted the results against the true values.

3.5 Prophet

The prophet algorithm is easy to understand and implement. As opposed to the ARMA, and HMM model the prophet model creates its own decomposition of the trend and seasonality effects to fit its parameters. For the trend, rather than using a moving average model, it may fit a number of different trends for linear, piece-wise and nonlinear growth based on the training data. The model is also able to account for multi-period seasonality. It does this by relying on a Fourier series to provide a malleable model for various seasonal effects hidden in the state space [Rob19].

4 Results

4.1 Classical Decomposition

Below are included a few examples of decompositions featuring the whole U.S. and the Midwest. The first is a decomposition of the U.S. daily time series, and the second is the U.S. monthly time series, while the third is the Midwest daily time series.

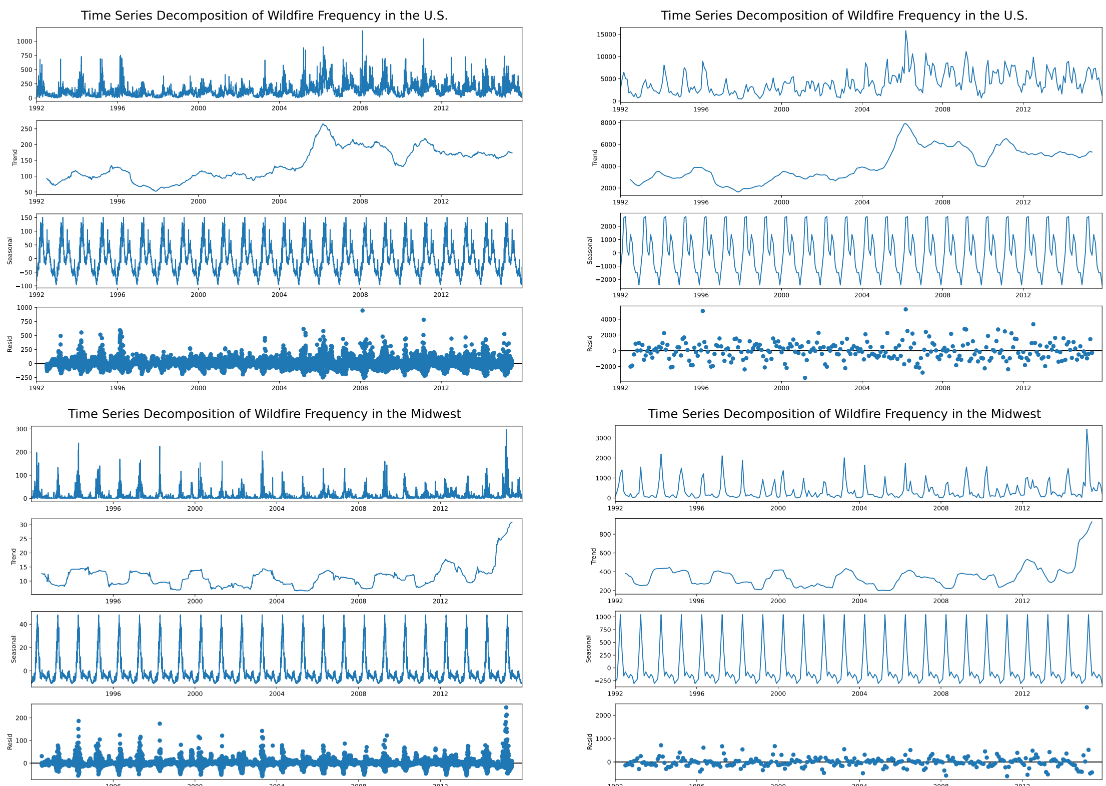

Figure 1: Time Series Decompositions of wildfires. The top two subfigures are decompositions of the wildfires across the entire U.S. and the bottom two are decompositions of the wildfires in the MidWest. The two subfigures on the left are decompositions of the daily time series, while the two on the right are of the monthly time series.

The results of these decompositions are fairly intuitive. As one would expect, we see a strong seasonal cycle in the occurrence of wildfires. The warmer months every year see many more fires than the cooler months. We can also infer from the trends that the total number of wildfires in the United States has been on the rise for the last 20 years. This makes our research all the more relevant. As you can see the residuals of the monthly time series appear to be less covariance stationary, and hence will hopefully perform better in an ARIMA model.

4.2 ARMA Model

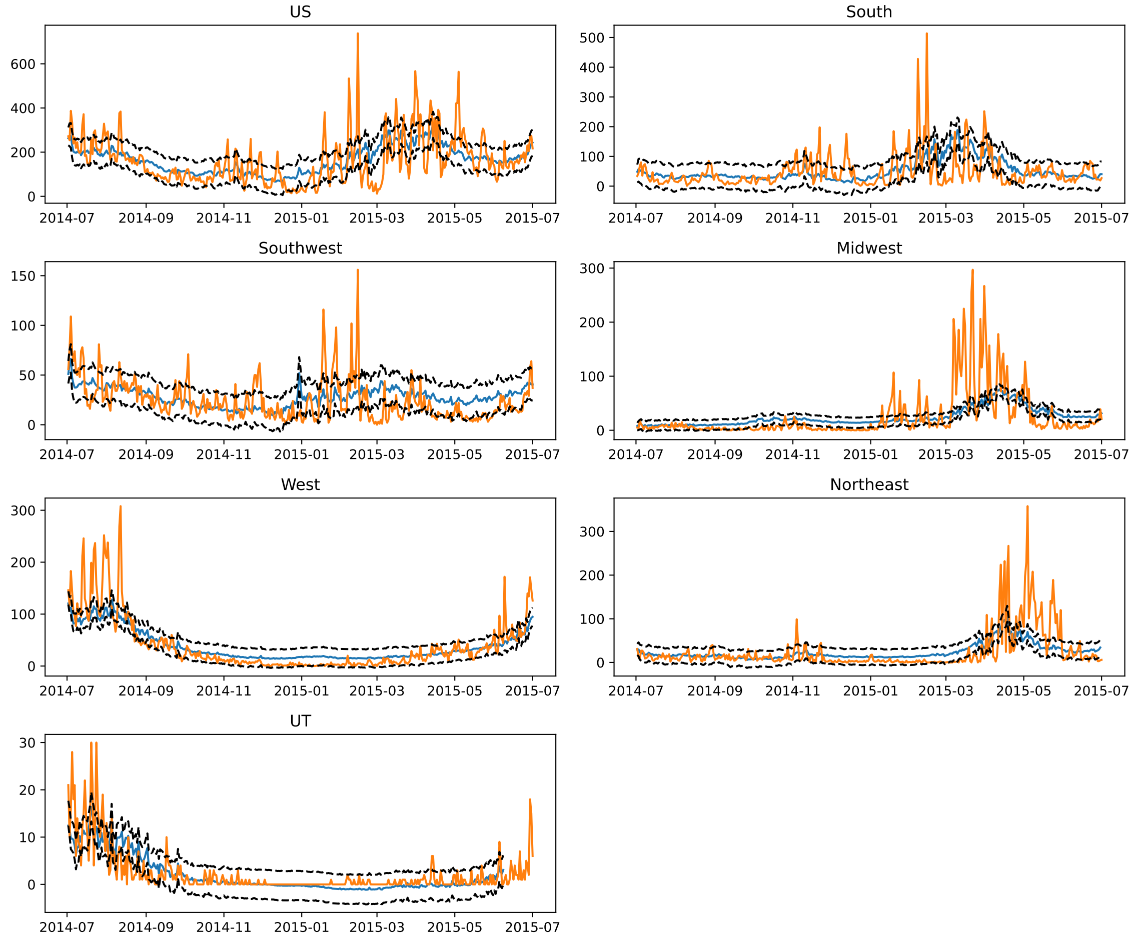

Figure 2: ARMA(2,1) daily forecasts with confidence intervals. Forecasts are in blue, the actual data is in orange, and the confidence intervals are the black dotted lines. This is standard for all following ARMA and Prophet models.

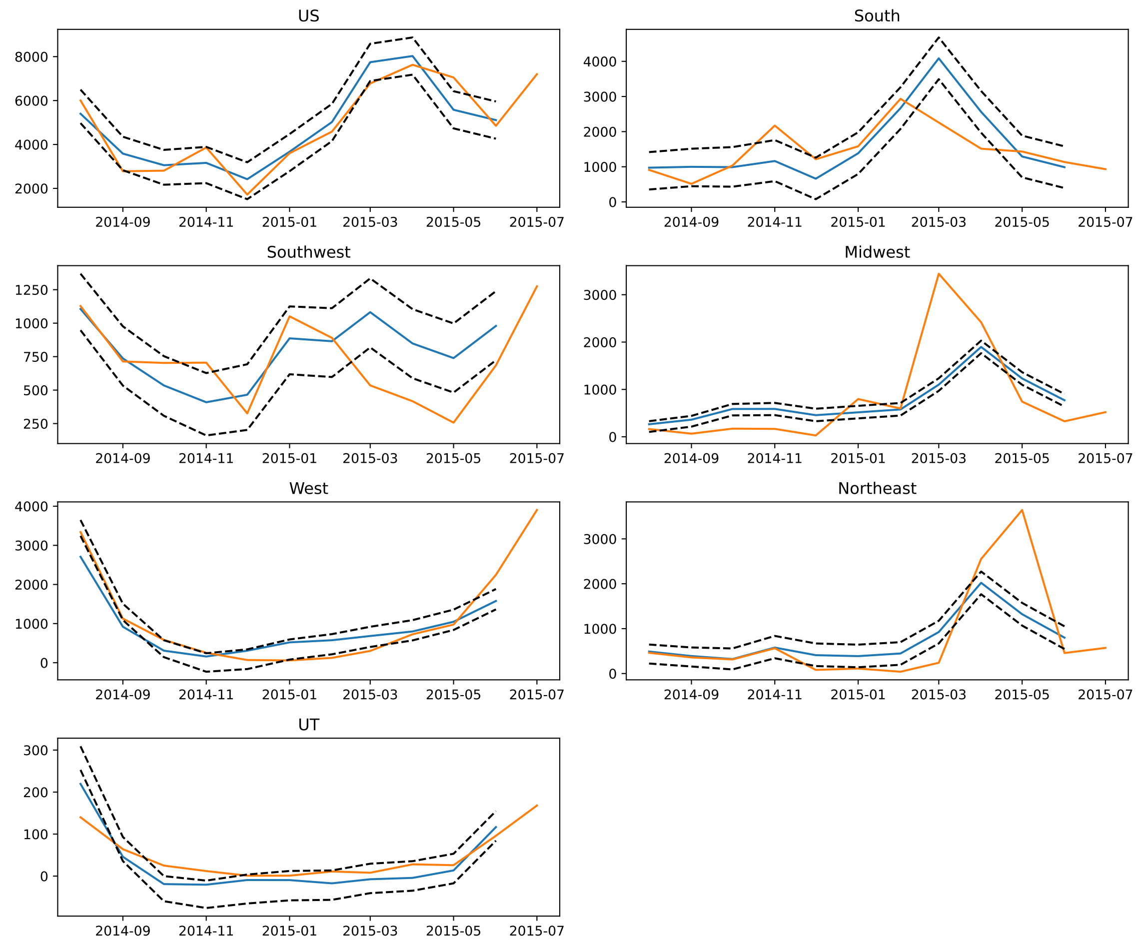

Figure 3: ARMA(2,1) monthly forecasts with confidence intervals

As we can see the ARMA model predictions seems to follow the actual forest fires fairly well for most regions, however it is not able to account for the outliers especially in the daily case. The nature of the monthly time-series removes most outliers and is sometimes able to predict the number of wildfires pretty well, especially on the entire U.S. One thing to note is that our forecast for the summer of 2015 in the Midwest is much lower than the actual number of fires. If you look back at the original trend of the Midwest fires from 2000-2016 you'll notice that the summer of 2015 was abnormally ablaze in the Midwest compared to every previous year that our model was trained on.

4.3 HMM i

Here are results for the last four months of 2013. In the title of each plot, there is a percentage denoting the percent accuracy improvement over the naive method (predicting the most common category from the previous month/season for all time steps). There is one month where the HMM is an accuracy improvement.

The orange denotes the true outcome, the shading of the green dots corresponds to the probably of observing each category. The most probably categories are underscore by a magenta bar as the prediction.

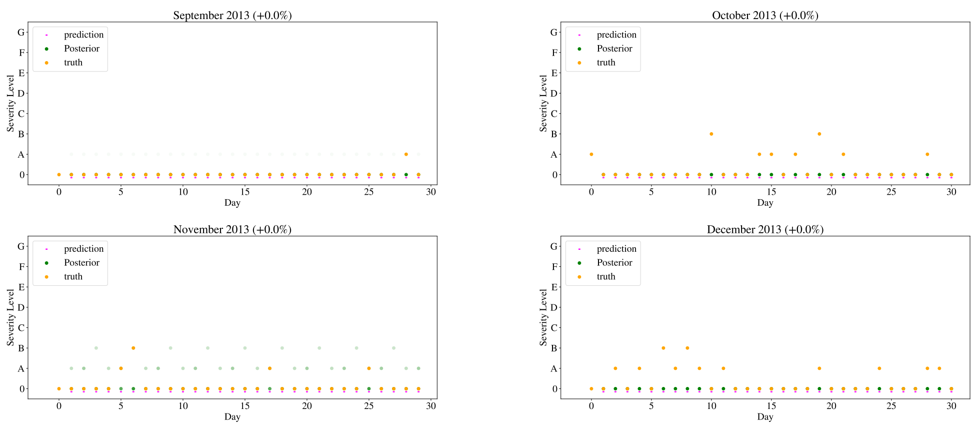

Figure 4: Predicting the Worst Wildfires throughout the months of September-December

Below are the results for wildfire seasons in California. The HMM appears to be worse than the naive prediction.

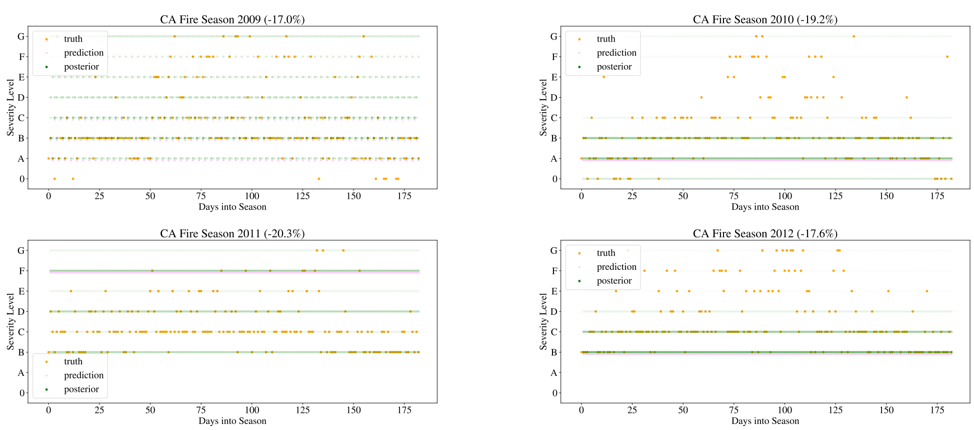

Figure 5: Predicting the worst wildfires through the fire seasons of 2009-2012

4.4 HMM ii

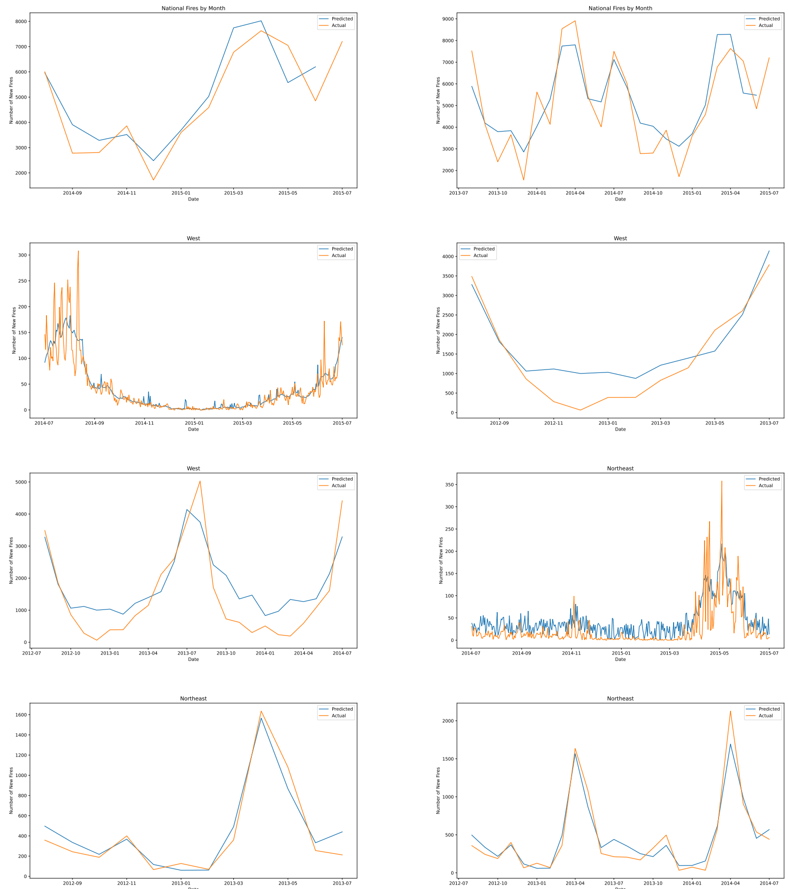

Figure 6: One year monthly and daily predictions of fire occurrences using a continuous state HMM for various regions.

The model performs fairly well over both one and two year time periods, but especially with the daily predictions there can be some pretty significant variance from one prediction to the next—it is nearly impossible to get a smooth curve. The model tends to struggle to predict values at the extremes of the range (i.e. very low or very high values). This is most likely due to the fact that the probability of transitioning to these states in our stochastic process is very low.

4.5 Prophet

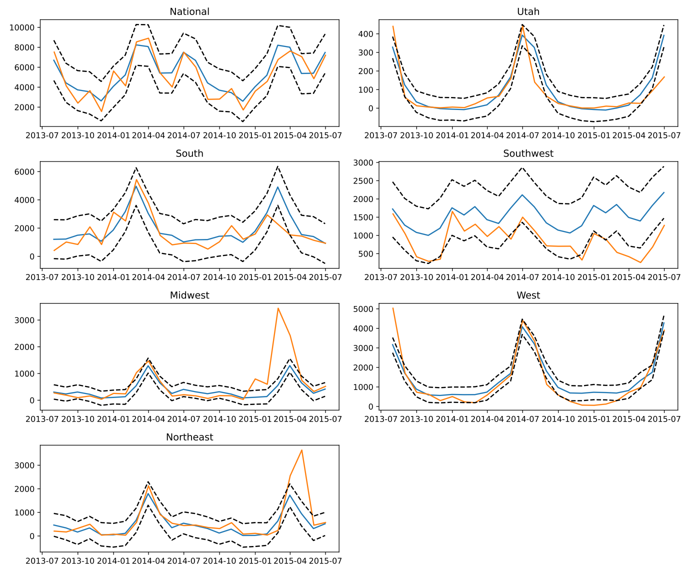

Figure 6: Monthly two year prediction of fire occurrences using Prophet model

As we can see the Prophet model worked better on the monthly data than the daily (see appendix for daily Prophet predictions) and was able to pretty well accurately forecast the next 2 years out slightly for most regions, and do so slightly better than the ARMA model. Again, we notice that the summer of 2015 was extreme for the Midwest and the prophet model was not able to predict that.

5 Final Analysis

It seems that the discrete time HMM has better predictions on the most severe wildfire when we do monthly training and prediction. Contrast this with a six-month long fire season. It is unreasonable to expect that one Markov chain can accurately simulate a latent state evolution throughout an entire period. There is a good chance that the Markov chain is not temporally homogeneous.

It is clear from the results that predicting wildfires by month is much more effective than by day. There is just too much noise in the day-to- day data. One day there can be zero new fires, and the next day there may be fifty. It is difficult to effectively train models when there is so much variation in the training data. We should also consider the fact that because we are training a continuous state HMM, we would definitely prefer to have very few data points with value zero. The monthly data does a better job at satisfying this requirement.

It is difficult to say which model worked best to predict future amounts of wildfires. All do a relatively good job at predicting the overall evolution of wildfire occurrences but suffer the same difficulties with regard to predicting extreme values. That being said, the Prophet model and continous state HMM do seem to handle the extreme values a little better than ARIMA. It should be kept in mind that our goal is not necessarily to make one-hundred percent accurate predictions about exact future fire numbers, but rather to predict when certain areas will be particularly vulnerable. Even though our estimations often tend to be on the conservative side, they still give very valuable information about future wildfire danger.

6 Ethical Considerations

We do not see many ethical concerns for this data or analysis. The wildfire data is public and general information that has nothing to do with any individual's privacy.

There may be some elements to consider with the interpretation of this paper. Every reader must be aware that our forecasts are not completely reliable, and we are not suggesting governments should or should not allocate funds towards fire departments based solely off our findings.

7 Conclusion

In conclusion, we achieved surprisingly accurate results of forecasting wildfires one or even 2 years out in the future across geographical regions. We may not be able to tell you exactly where and when the next 1000 wildfires will be, but we hope that our results are still useful for environmental and governmental agencies to get an idea of how the next fire season or two is shaping up to be. If authorities at least have an idea of which regions will be hit the worst, then as described before, they can allocate the proper resources ahead of time to the most ideal locations. We hope that there is much room for improvement in predicting the nature of future fire seasons, but there are so many critical factors like precipitation, temperature, and wind that are currently nearly-impossible to forecast long-term. In light of such difficulty, we are pleased with how our results turned out and think that our work demonstrates the potential of the described models for predicting wildfire occurrence.

References

[Ass22] Western Fire Chiefs Association. "California Fire Season: In-Depth Guide". In: (2022).

[Kat23] Laura Hanson Katie Hoover. "Wildfire Statistics". In: (2023).

[Rob19] Winston Robson. "The Math of Prophet". In: (2019).

[Wit19] Vitali Witowski. "AIC and BIC Value for a Discrete Time Hidden Markov Model". In: (2019).

8 Appendix

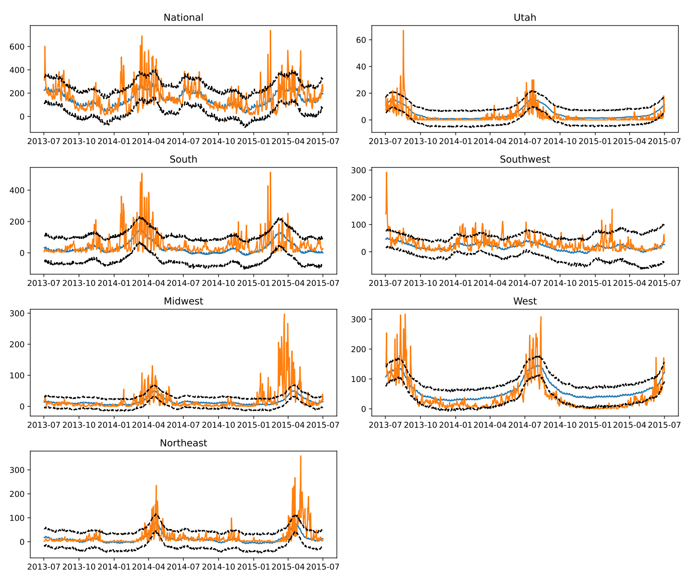

Figure 7: Daily two year prediction of fire occurrences using Prophet model