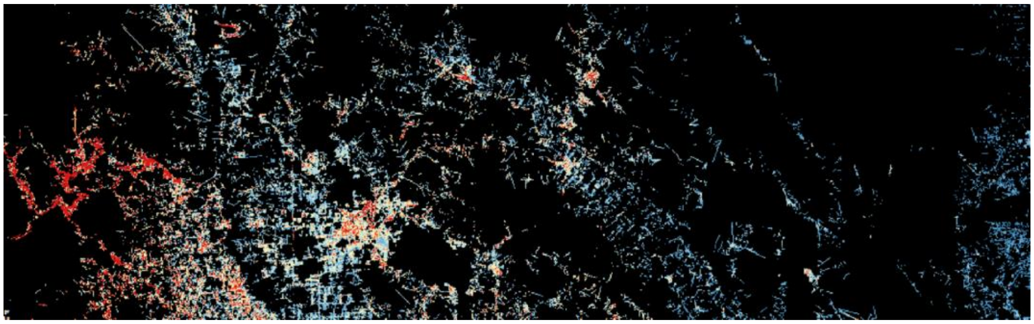

Dx Risk 2021 circuit prioritization: Update and deliverables

Convergence Data Analytics, Salo Sciences, Presence Product Group

Maxent approach to calculating annual veg-caused ignition, outage & wiredown probabilities

Introduction

This document was prepared by the Dx Risk team to support PG&E's upcoming 2021 work plan. We computed a series of event probability scores for each conductor, which were grouped by Circuit ID to predict annual event rates on a per-circuit basis. These event types were vegetation-caused ignitions, outages & wiredowns.

Commented [OY1]: Why are you calling probability scores, rather than just probability?

Commented [OY2]: Why were they grouped? Is the probability at a conductor level useful as well?

Commented [OY3]: Why don't you utilize the relationship between the three event types in estimating probability of ignition? If you are estimating the probability of ignition, outage and wiredown independently, I would think that you could utilize information on the likelihood of ignitions given the outage or wiredown, to better estimate the locations that didn't have ignitions before, given that the ignition probability is really small in each conductor level.

Also don't you utilize any information on the conductor types? How would you capture a significant reduction in the probability of ignition for covered conductor vs bare conductor for example? What about the changes in ignition probability for conductors that the system hardening work is done recently?

Failure probability scores were modeled using the Maxent approach developed in Phase 1 / Milestone 1. We combined spatially-explicit event data (locations where ignitions, outages & wiredowns occurred) with statewide covariates, which included environmental data (weather, vegetation, topography) and distribution grid properties (e.g. conductor density, jacket type). The model assigns event probability scores to each grid location based on the environmental and grid conditions at that location. Event probabilities for any one asset are typically low as these events are rare; 343 veg-caused ignitions were recorded in 2019 from a network of over 4.5 million conductors. When aggregated across a circuit, the failure probabilities sum to the number of events predicted to occur over a year on that circuit.

Interpreting outputs & performance metrics

Models for ignitions, outages and wiredowns were independently trained and evaluated using AUC, a metric of separability. An AUC score of 0.8 can be interpreted to mean the model is 80% confident it can distinguish between locations where events are and are not likely to occur based on continuous probability scores. We also calculated event 'exposure' scores for each conductor. This was an effort to discriminate, in a binary sense, conductors where events are likely to occur from where they are not. Using an omission rate threshold of 5% we created a 95% confidence area that identifies areas where events could occur. We calculated recall scores to evaluate whether this threshold was well-calibrated; a recall score of 0.96 can be interpreted to mean that 96% of event locations occur within this 'exposure' threshold. Regarding the drivers of model performance, covariates describing the density of fall-in trees, the height of nearby trees, and the number of conductors in an area were consistently the top drivers of model performance.

One key challenge in interpreting this work, and in performing it, is reconciling differences between the gridded predictions and the linear conductor data. We calculated event probabilities using a 100 meter grid, but a grid cell may include multiple conductors. We report normalized probability scores based on the number of conductors per grid cell to avoid double-counting, but this difference manifests elsewhere: while only 40% of the grid *area* may be 'exposed'. 80% of grid *assets* may be exposed when there are multiple conductors in high probability areas.

Notes on modeling details:

- The ignitions and wiredown data were filtered by date from 2015-01-01 to 2018-12-31.
- We report the total counts separately for 2015-2018 and 2018-2019 because it appears there were changes in reporting requirements starting some time in 2018 that led to increases in reported ignitions and wiredowns.
- There are separate ignitions databases for 2015-2018 and 2018-2019, and only data from the former were included in model training.
- Ultimately these counts were used to compute **τ**, the Maxent scaling parameter. This corresponds to the average likelihood of observing an event across the whole system, computed as the number of events in a year divided by the area where these events could be observed. Setting this was as much art as science. We erred on the side of using a higher **τ** value using the higher reported event frequencies of the past few years on the assumption that these events still occurred in past years but simply weren't reported.
- Independent model testing has not yet been performed. All reported performance metrics are training metrics.
- In the metrics below, percent contribution scores compute how model performance (AUC) changes as a function of randomly altering fitted feature coefficients. Permutation scores are calculated by randomly altering the values of the underlying covariates, not their fitted coefficients, and comparing changes in model performance. Each of these effects are rescaled from 0-100 to estimate relative importance.
- More precise details can be found in the software documentation provided by Phillips, S. (2006) A brief tutorial on Maxent. AT&T Research. Available at: http://www.cs.princeton.edu/~schapire/maxent/tutorial/tutorial.doc:

Percent contribution: *"While the Maxent model is being trained, it keeps track of which environmental variables are contributing to fitting the model. Each step of the Maxent algorithm increases the gain of the model by modifying the coefficient for a single feature; the program assigns the increase in the gain to the environmental variable(s) that the feature depends on. Converting to percentages at the end of the training process, we get the percent contribution."*

"The percent contribution values are only heuristically defined: they depend on the particular path that the Maxent code uses to get to the optimal solution, and a different algorithm could get to the same solution via a different path, resulting in different percent contribution values. In addition, when there are highly correlated environmental variables, the percent contributions should be interpreted with caution."

Permutation importance: *"...for each environmental variable in turn, the values of that variable on training presence and background data are randomly permuted. The model is reevaluated on the permuted data, and the resulting drop in training AUC is shown in the table, normalized to percentages."*

Commented [OY4]: What about 2019?

Commented [OY5]: I thought the change started from 1/1/2018, not sometime in 2018.

Commented [OY6]: 2018-2019 data includes more ignitions since we started to record ignitions that do not meet CPUC reportable ignitions. However, if you look at a subset of the CPUC reportable ignitions from 2018-2019, the reporting criteria should be consistent with previous years.

Commented [OY7]: Why exclude 2019? Is this due to the impact of PSPS?

"The permutation importance measure depends only on the final Maxent model, not the path used to obtain it. The contribution for each variable is determined by randomly permuting the values of that variable among the training points (both presence and background) and measuring the resulting decrease in training AUC. A large decrease indicates that the model depends heavily on that variable. Values are normalized to give percentages."

Vegetation-caused ignitions

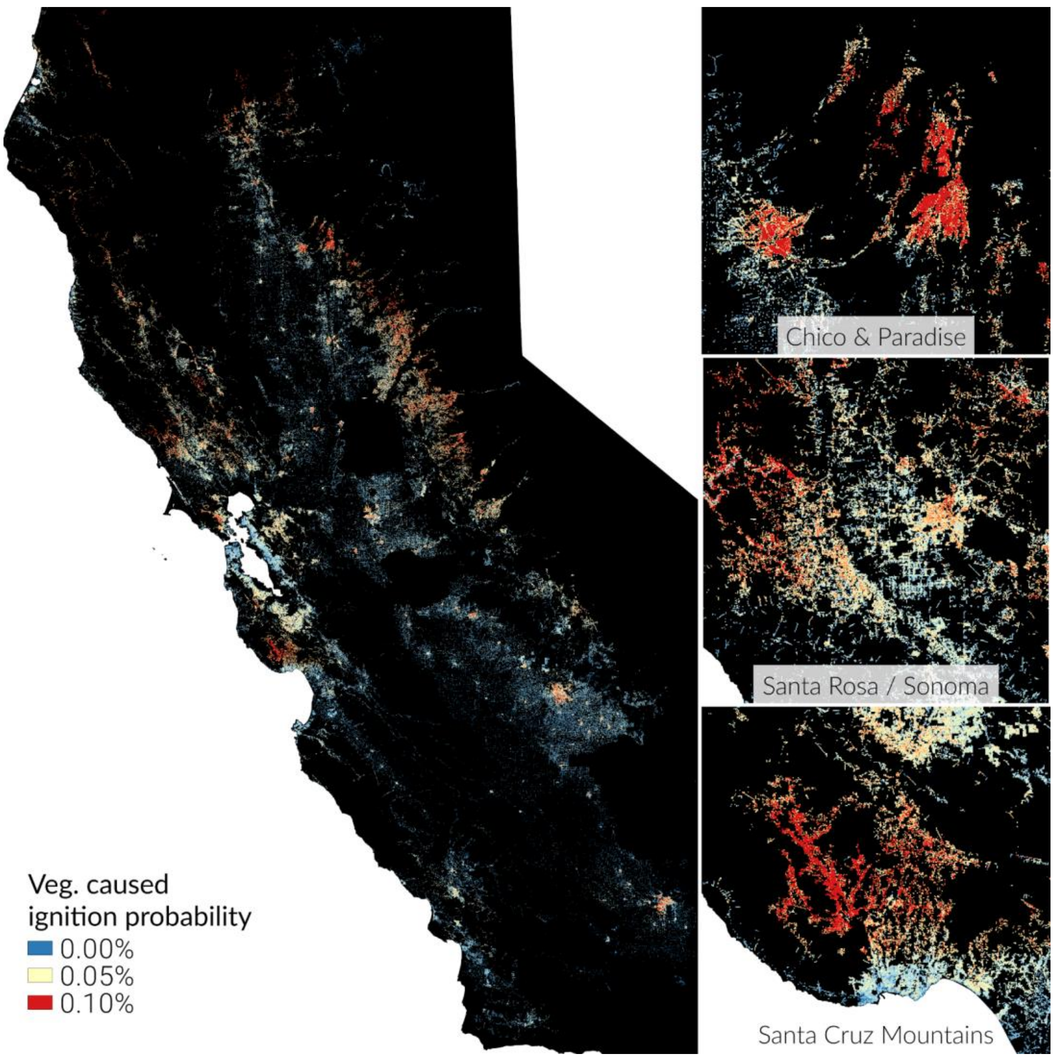

Summary statistics

- Predicted annual ignitions (average): **348.3**
- Observed annual ignitions (2018-2019): 343
- Observed annual ignitions (2015-2018): 119
- Grid area exposed to ignitions: 42.9%
- Grid assets exposed to ignitions: 79.1%
- Top 10 feeders (name, predicted count):
  - APPLE HILL 2102 3.25
  - PLACERVILLE 2106 2.82
  - BRUNSWICK 1106 2.61
  - WEST POINT 1102 2.32
  - PINE GROVE 1102 2.10
  - CAMP EVERS 2106 1.99
  - SHINGLE SPRINGS 2109 1.96
  - MOLINO 1102 1.95
  - NARROWS 2105 1.74
  - HOOPA 1101 1.73

Model performance

- AUC - 0.73
- Recall - 0.98
- Precision - 0.65

Event data

- 461 ignition locations
- Time frame: 2015-2018
- Filtered by: 'Contact from Object' == 'Vegetation'
- Removed points with spurious lat/lon
  - Outside continental US
  - Greater than 400m away from any conductor

Variable importance rankings

Variable	Percent contribution	Permutation importance
tree-fall-in	32	9.1
tree-height-max	25.2	21.9
conductor-count	8.7	13.3
vapor-pressure-deficit-avg	8.1	15.6
hftd	5.4	2.1
precipitation-avg	4.8	10.3
impervious	3.7	10.9
specific-humidity-avg	3	1.2
canopy-stress	3	1.9
temperature-avg	2.7	2.5
1000-hour-fuels-avg	1.1	1.6
tree-height-avg	0.9	4.1
100-hour-fuels-avg	0.6	0.9
local-topography	0.5	1.5
energy-release-avg	0.3	2.2
burn-index-avg	0	0.7

Commented [OY8]: This seems to be the total # of ignitions over two years, not annual.
Thus it would be 171.5/yr

However, if we only look at CPUC reportable ignitions, 236/2 = 118/yr over 2018-2019. This should be consistent with the dataset before 2018.

Commented [OY9]: Need definition for grid area and grid assets

Commented [OY10]: Why not just use CPUC reportable ignitions only so that the model can be validated with 2019 data.

Commented [OY11]: How many points are these? How about using the mapping with the outage database with ignition for these.

Commented [OY12]: How do you account for different miles of circuits.

Vegetation-caused outages

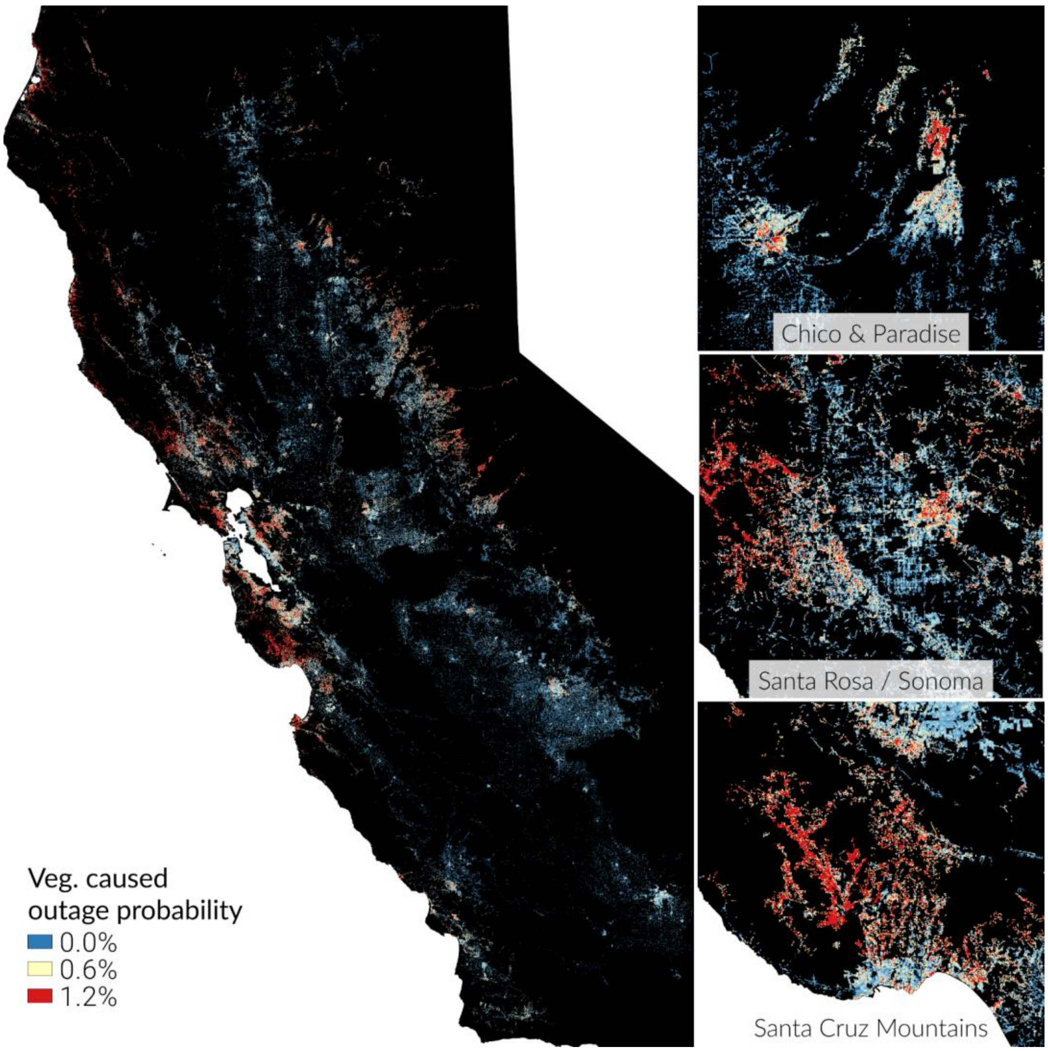

Summary statistics

- Predicted annual outages (average): **3,499.4**
- Observed annual outages (2015-2018): 3,471
- Grid area exposed to veg-caused outages: 36.7%
- Grid assets exposed to veg-caused outages: 80.0%
- Top 10 feeders (name, predicted count):
  - CAMP EVERS 2105 36.79
  - MOLINO 1102 36.19
  - CAMP EVERS 2106 34.78
  - MONTE RIO 1113 29.67
  - MONTE RIO 1111 25.35
  - BIG BASIN 1101 24.99
  - FORT BRAGG A 1102 23.31
  - EL DORADO PH 2101 22.82
  - PINE GROVE 1102 22.16
  - ARCATA 1122 22.12

Variable importance rankings

Variable	Percent contribution	Permutation importance
tree-fall-in	47.7	27.7
tree-height-max	25.3	43.1
conductor-count	13.9	16.4
specific-humidity-avg	4.3	4.7
vapor-pressure-deficit-avg	3.5	3.7
temperature-avg	3.2	0.9
precipitation-avg	0.8	1.4
operating-voltage	0.7	1.6
local-topography	0.3	0.3
tree-height-avg	0.2	0.2
canopy-stress	0.1	0
jacket-type	0	0

Model performance

- AUC - 0.82
- Recall - 0.96
- Precision - 0.80

Event data

- 12,885 outage locations
- Time frame: 2015-2018
- Records from the Veg. Outages database
- Removed points with spurious lat/lon
  - Outside continental US
  - Greater than 400m away from any conductor

Vegetation-caused wiredowns

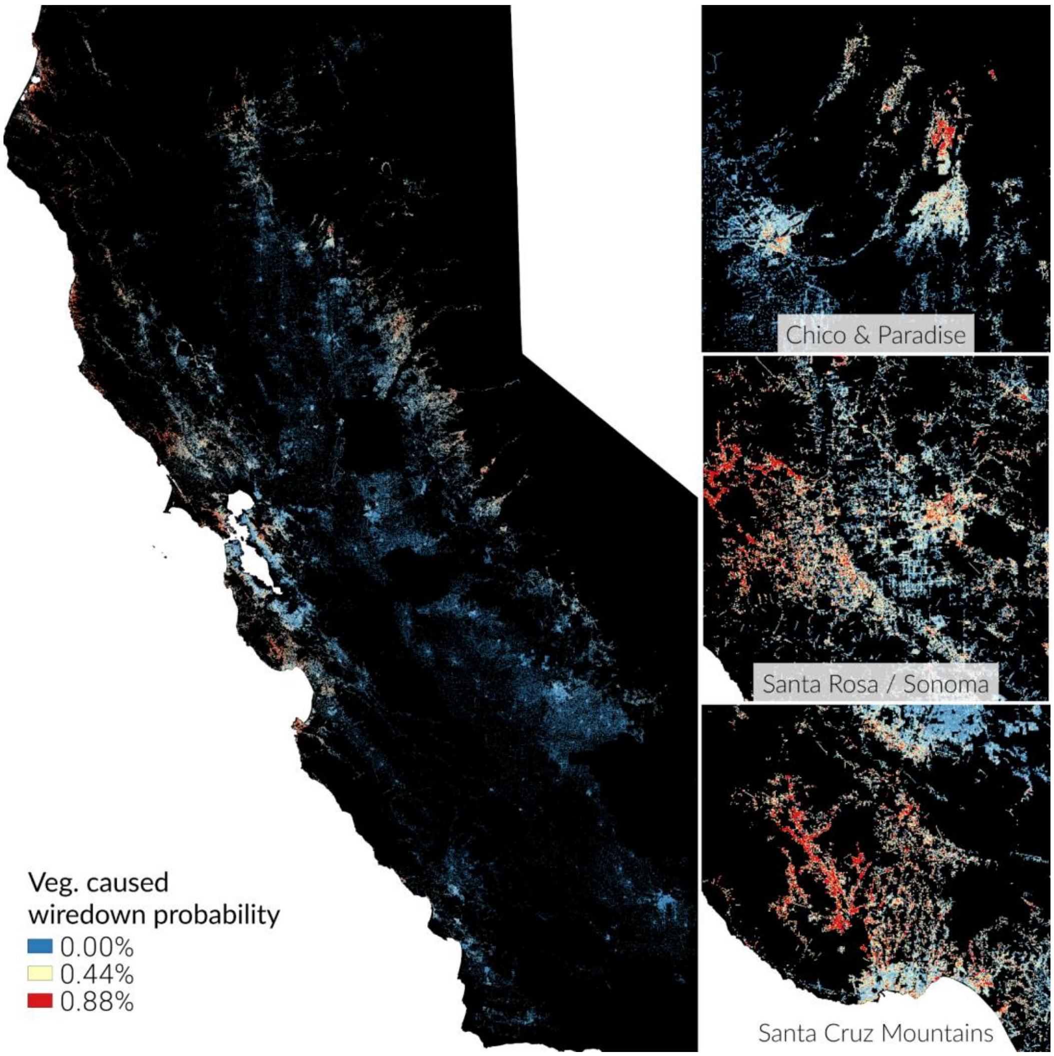

Summary statistics

- Predicted annual wiresdown (average): **2,721.9**
- Observed annual wiresdown (2018-2019): 2,718
- Observed annual wiresdown (2015-2018): 690
- Grid area exposed to wiresdown: 36.7%
- Grid assets exposed to wiresdown: 73.2%
- Top 10 feeders (name, predicted count):
    - MOLINO 1102 27.28
    - FORT BRAGG A 1102 25.45
    - BIG RIVER 1101 22.38
    - MONTE RIO 1113 22.27
    - CAMP EVERS 2105 21.58
    - ARCATA 1122 20.86
    - MONTE RIO 1111 20.75
    - BRUNSWICK 1106 20.06
    - CAMP EVERS 2106 19.62
    - APPLE HILL 2102 19.59

Model performance

- AUC - 0.82
- Recall - 0.96
- Precision - 0.73

Event data

- 1,844 wiredown locations
- Time frame: 2015-2018
- Filtered by 'Basic Cause' ∈
    - 'Vegetation'
    - 'Environmental/Extern',
    - 'Environmental/External'
- Removed points with spurious lat/lon
    - Outside continental US
    - Greater than 400m away from any conductor

Variable importance rankings

Variable	Percent contribution	Permutation importance
tree-fall-in	46.5	28.6
vapor-pressure-deficit-avg	15.6	16.6
tree-height-max	15.2	19.8
temperature-avg	6.9	4.4
precipitation-avg	5.2	15.1
conductor-count	5.2	5.8
operating-voltage	3	2.5
specific-humidity-avg	0.9	1.3
local-topography	0.7	0.8
tree-height-avg	0.5	4.9
canopy-stress	0.2	0.2
jacket-type	0	0

CSV data

A tally of expected failure counts produced with this approach has been shared along with this document.

column	description
lineMiles	the sum of conductor line miles on the circuit
conductorCount	the number of conductors on the circuit
nPrimary	the number of primary conductors on the circuit
nSecondary	the number of secondary conductors on the circuit
predictedAnnualIgnitions	the sum of expected ignitions across all conductors in a circuit. These were calculated on a raster grid.
predictedAnnualWiresdown	the sum of expected wiredown counts across all conductors in a circuit
predictedAnnualOutages	the sum of expected veg-caused outage counts across all conductors in a circuit
predictedAnnualUnknownOutages	the sum of expected unknown-cause outage counts across all conductors in a circuit
pctExposedIgnitions	the percent of all conductors flagged as 'exposed' to veg-driven ignitions – based on the omission rate calculation (used in recall score assessment)
pctExposedWiresDown	the percent of all conductors flagged as 'exposed' to veg-driven wiredowns
pctExposedOutages	the percent of all conductors flagged as 'exposed' to veg-driven outages

Predicted vs. observed counts

With a model that assigns an annual probability of failure (from 0 to 1) to each small set of grid components, the expectation value of the annual count of failures on a feeder can be computed as the sum of the expected count of failures across all of its components. A simple example with three components that can only fail once, with a 50% chance of failure each, would have an expected failure count of 0.5 each and produce and expected count of 0.5 + 0.5 + 0.5 = 1.5 failures.

The above calculation is a simplification of what we are actually doing. Recall that we are not modeling assets directly. We are modeling grid pixels that still have outage rates are so sparse in both space and time that no individual "grid pixel" has a large probability of annual failure. It is the nature of probabilities that an event with a tiny probability of happening over the course of a year, which is itself the integral of smaller probabilities on finer timescales, could happen 0, 1, 2, 3 or more times, but with a small annual rates, the average rate of failures over each feeders pixels might be expected to be 1 in 100 years or more. If you use a Poisson distribution to model such rare events, the difference between 1*p(y=1) and sum(p(y>=i) * i) for 1 to infinity is something like 0.0001, or 1 in 10,000. In other words we have taken a small shortcut by equating E[fail] for one year with p(fail). We might casually say that we are taking the sum of probabilities across all of a feeder's components, but we are actually implicitly converting to E[fail] for each pixel in a feeder and summing them. It is true that if we were modeling a less sparse occurrence, we could not take this shortcut.

What the ME model in particular is doing with each failure is recording the covariate values at that time/location and adding them to the distributions of values for observed failures. To the model it is irrelevant where they came from. So if you have many failures in the exact same spot, that signal boosts the probabilities of all pixels with similar values, with no particular affinity for the original location. So in the hypothetical case where a single asset repeatedly fails, say n times, that "count" will be spread across environmentally similar pixels, resulting in an expected increase in E[fail] of n. The expected failure counts for this work were calculated in this manner and can be compared to historical rates of outages of different types.

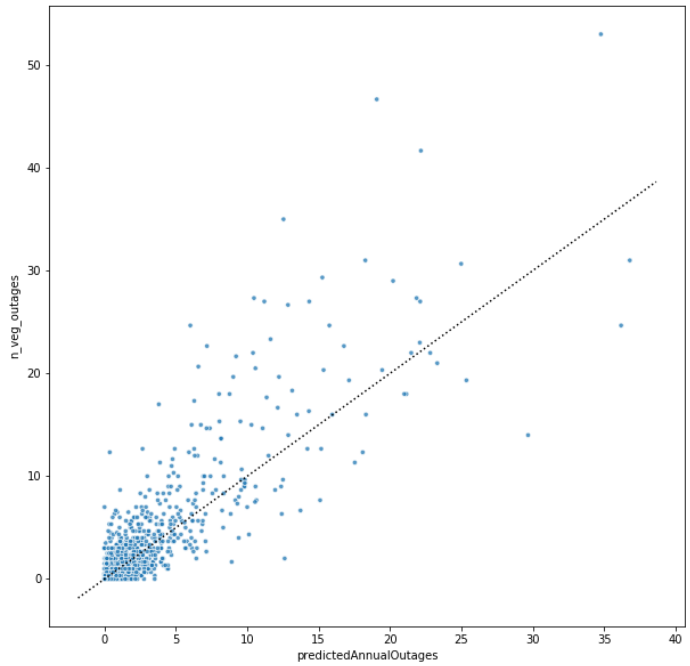

The figure above plots a point for every feeder with the model prediction of the expected failure counts on the x-axis and the 2015-2018 mean count of outages with vegetation listed as their cause on the y-axis. The dotted line runs through the points where the two numbers are equal (perfect prediction). Feeders above the line suffer more outages than predicted and feeders below the line suffer fewer than predicted. Feeders that fail more often than the model predicts are likely to have exacerbating factors beyond the ones used to train the model influencing their outcomes.

Statistical properties and time-series modeling (progress to date)

There are just over 3100 feeders in the GIS data set of primary conductors studied for this work. The properties of 1.4 million individual "conductors" – each with their own row in the ED GIS system, were mapped (with asset count, length, phase count, and coastal indicators preserved) to 1.3 million 10m square geo-spatial "pixels" to allow their augmentation via 10m scale raster data from tree presence, height, and fall-in geometry geo tifs. Note that a separate pdf document on the sources and development of those tree data products has been included with this deliverable. The 10m "pixels" of data were aggregated to 100m scale (770,000 of those) to be augmented by a much wider variety of data, including population density, elevation, HFTD designations, and land cover classifications as well as the outputs from the MaxEnt modeling, again accurately preserving information on asset count, line length, etc. Next those 100m "pixels" were aggregated to 1000m scale (82,000 of those), to be augmented with time series geo-spatial weather conditions from GRIDMET and RTMA. Finally, the 1000m "pixels" can be aggregated together into a single representation of all asset, spatial, and weather characteristics per feeder (3100 feeders with several years of weather data each).

The panel of figures below illustrates the location of the feeders and illustrates some of their characteristics.

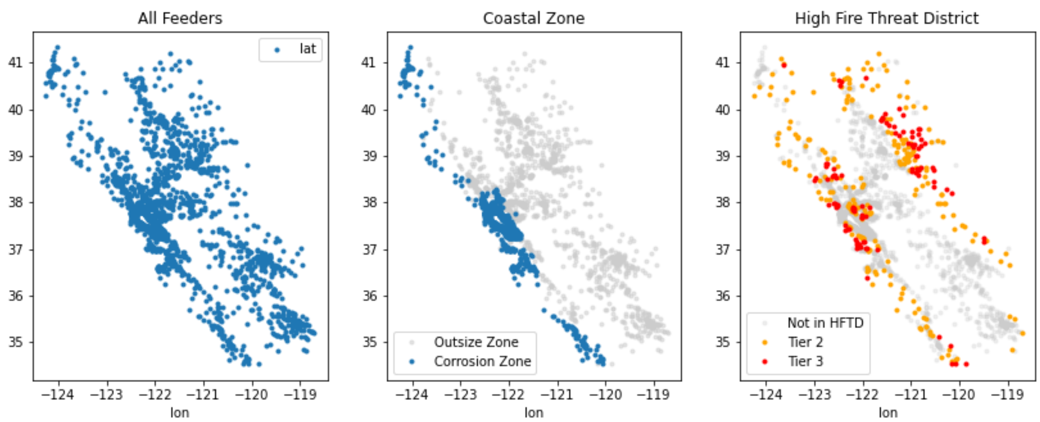

The next panel highlights the locations of outages of different types being modeled as a part of this project.

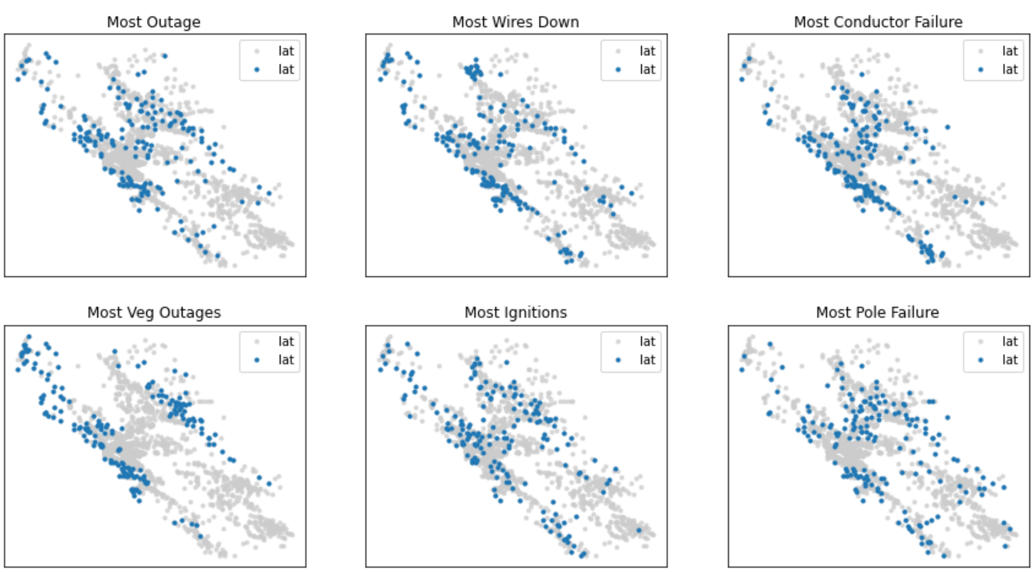

The panel of figures below presents pairwise comparisons of feeder characteristics of 100 "problem feeders" and 100 randomly selected feeders. Problem feeders are those that have: (1) experienced either ignitions or wire downs, and (2) are in the top 200 for most outages. Problem circuits tend to be:

- Longer overall, but not dramatically so
- Proportionally more 2- than 3-phase circuit miles
- Greater # of trees w/in falling distance
- Taller trees nearby

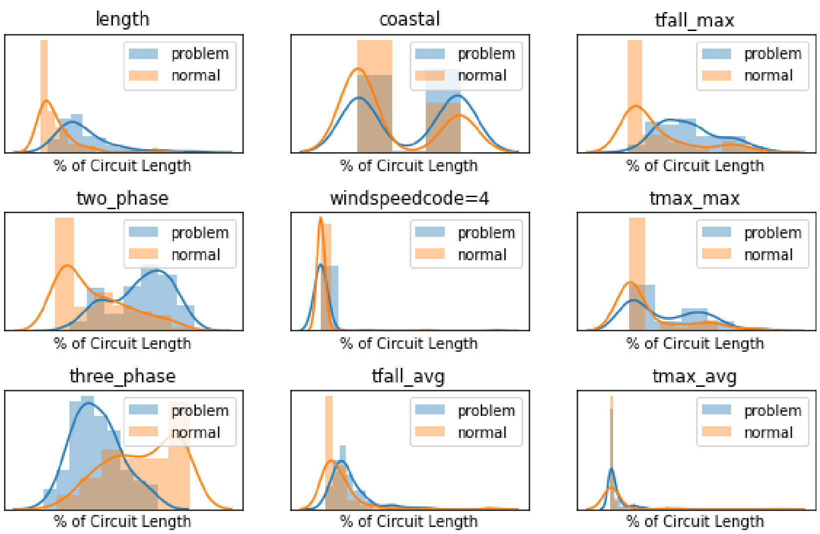

The figure below provides an analysis of the “weather signals” indicator values most often associated with failures of different types, ranked in order of importance. It is notable that pole and conductor failures are associated with both high heat and winter storms, clearly via different failure mechanisms. This importance ranking also underscores the relationship between PSPS days and ignitions – an indication that such days are well targeted.

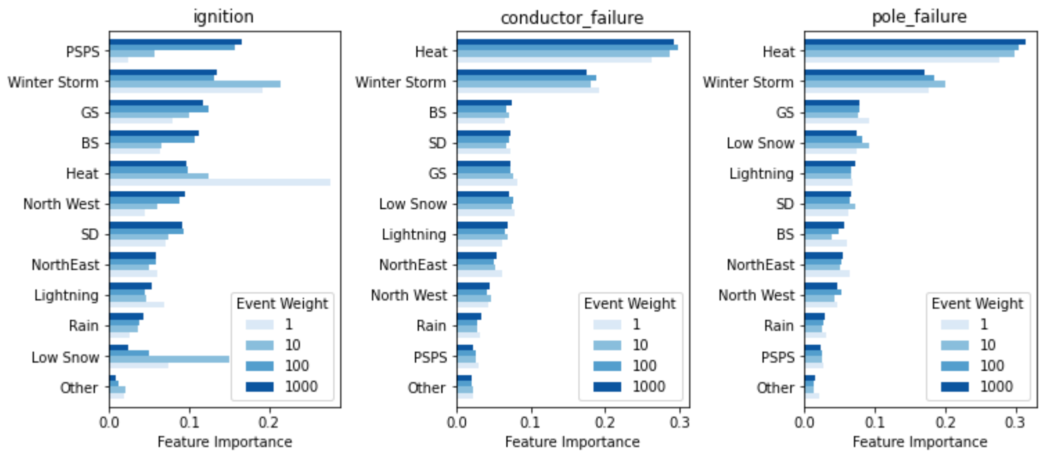

This next figure illustrates how the timeseries data on counts of outages of different types and "Storm Days" in grey line up in time. This figure illustrates that the occurrence of storms in general and the failures associated with them peak during winter months, with storm days producing a wide variety of correlated failures. However, the number of ignitions peak during summer and early fall when the rates of outages in general are at or near their background levels. It is likely that the same failure modes that produce ignitions during the summer are occurring all year round, but fortunately without the enabling conditions for fires to grow large enough to be observed and reportable.

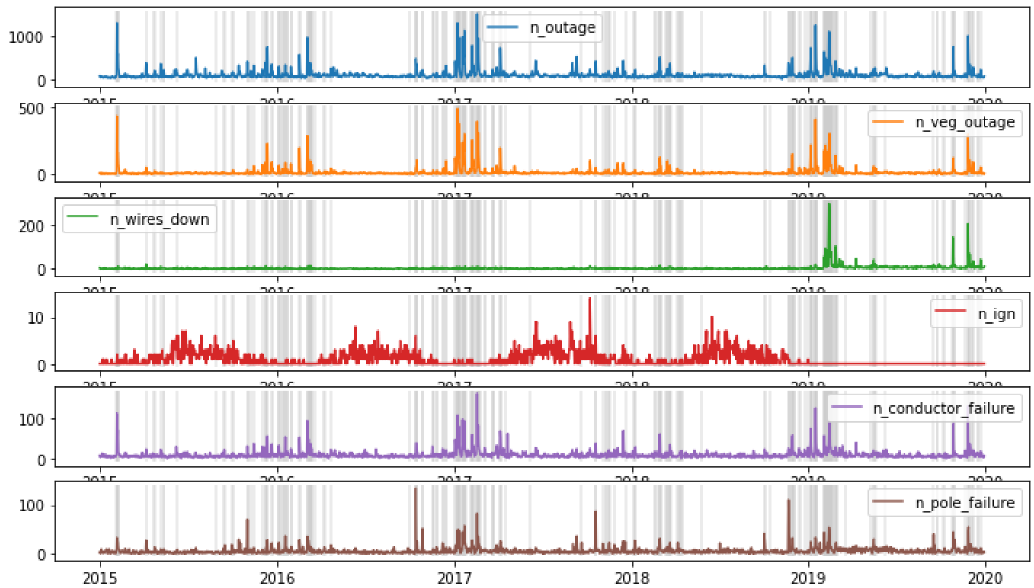

The figure above summarizes the time series of outages across the whole system. The figure below summarizes the count of outages per week (lighter colors represent more outages) for each feeder. Each row corresponds to a separate feeder_id, sorted by latitude (although the low latitudes are on top) and column correspond with weekly data spanning 2016-2019. Horizontal banding is due to chronic failures of the same or nearby feeder(s). Vertical banding is due to events in time that cause widespread outages (e.g. storms). This is the complete data set our modeling need to explain.

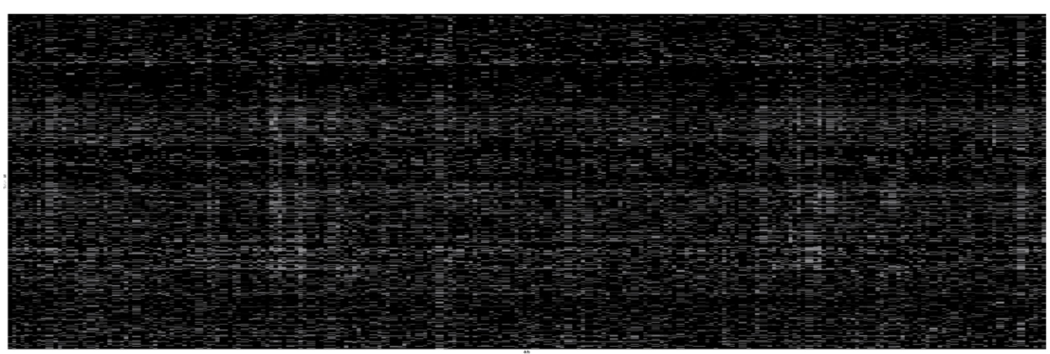

The figure below presents our RTMA weekly maximum gust speed metric with the same logic to its formatting. Lighter colors are higher wind speeds. Each row presents the gust speeds experienced by a single feeder and each column covers a week in time. Vertical bands correspond to wind events that spanned across the system.

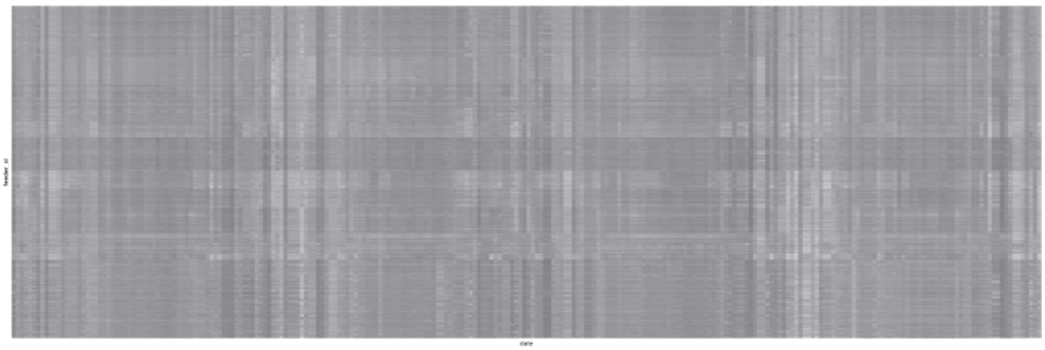

Finally, the figure below here presents the GRIDMET fm100 standardized metric of 100 hr fuel moisture (a typical fire modeling input related to how likely fuels are to burn if ignited). Darker patches of lower fuel moisture represent periods of increased fire risk. Lighter vertical bands correspond with wet storm events that moisten fuels.

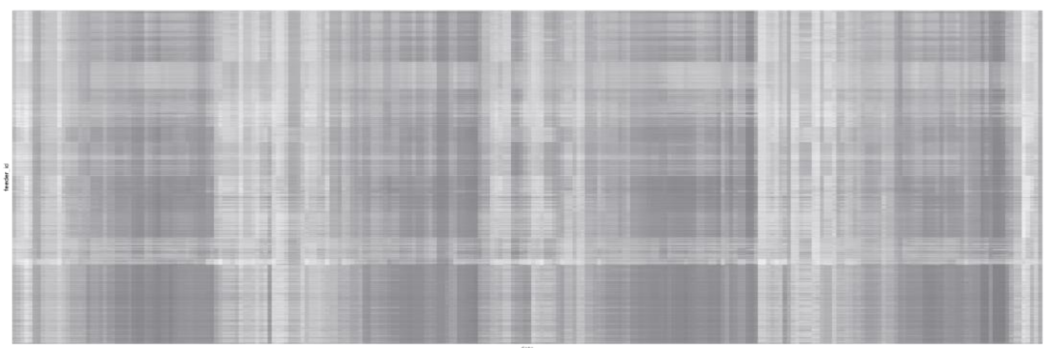

We assembled the above figures by selecting their underlying data out of our data processing pipeline to give a sense of the patterns we are trying to fit and the data requirements of the models we are running. You can also build some intuition for the drivers of outages and their spatial and temporal patterns.

Current modeling effort

Using the above data, along with the a few additional weather covariates and feeder attributes derived from their conductor attributes and locations, like length, whether they are coastal, number of phases, the estimated count of fall-in trees, the p(fail) results from the Maxent modeling, etc., we are currently training regression models that predict outage counts per-feeder-week. Those results can be aggregated into annual per-feeder outage rate estimates that are highly relevant to the circuit prioritization process. We are not yet ready to publish "official" estimates using this method, but we include some performance metrics below to ensure that we are more fully representing our status and current modeling capabilities. The models are under rapid development, so this provides a preview for where our work is headed in the near term. We estimate that such outputs will improve upon the results of Maxent alone and will also be capable of predicting at finer timescales and with some natural hooks for explaining why we are predicting what we predict and how we would expect the process of mitigation in the field to impact the modeled results.

The first figure below is a normalized histogram (area = 1), with smoothed density line, for the absolute value of the annual feeder outage count errors (aka the number of outages we over or under-estimated by) of a model run only accounting for feeder attributes and spatially derived data (no weather data and no feeder indicators). Years are from 2016-2019, so each feeder contributes 4 data points. The mean is 2.79 and the median is 2.0, with 6.3 at the 90th percentile.

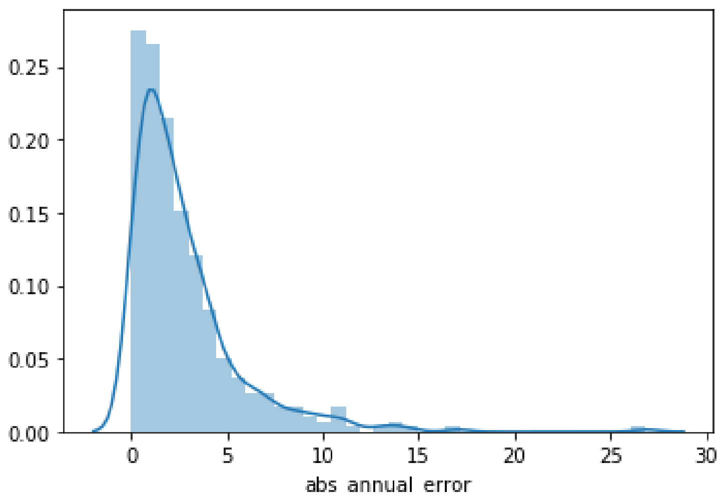

Another view of model performance is presented below as a scatter plot of actual vs. predicted outage counts per feeder (again with 4 contributions from each feeder over 4 years). This plot shows a strong correlation between predicted and actual outage rates, with an R^2 of 0.85. Feeders above the diagonal line fail more often than predicted. Feeders below fail less often – the reasons for either will likely be of interest to the group of people prioritizing preventative activity by circuit.

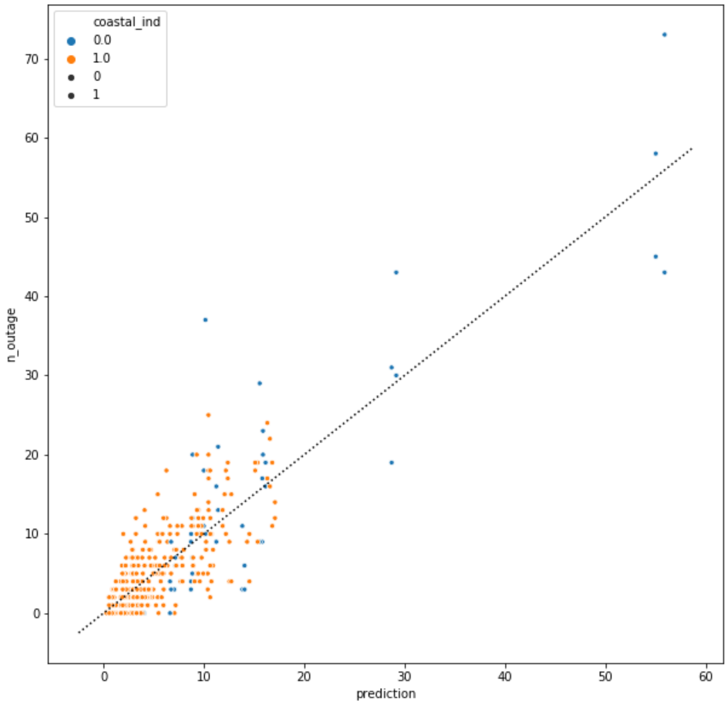
coastal_ind
0.0
1.0
0
1
n_outage
prediction
0
10
20
30
40
50
60
70