

Dx-Risk Milestone 3 Model Specification

A Multi-Layered Model for Characterizing Ignition Probability and Mitigation Potential

Introduction

This report provides a specification of the Dx-Risk models as of the conclusion of Phase 1 (P1) of work. This document satisfies part of the Milestone 3 (M3) documentation delivery requirements for the project by providing a high level motivation for and summary of all modeling efforts, including the technical details of our model specifications. This document does not cover the specific data, implementation, and modeling choices made to satisfy the requirements of Milestone 3 via circuit-level outage prediction or the results of those efforts. Each of those efforts is covered in its own document. The table below summarizes the rest of these M3 (or earlier) deliverable documents and where they can be found. All are referenced in context by this document.

M1 Maxent methods summary: still the best description of how we approached MaxEnt modeling of grid events.	Milestone 1 analysis documentation
M3 Maxent model applied to the question of circuit prioritization: summary of results for ignitions, wires down, vegetation caused outages, and outages of unknown type.	https://docs.google.com/document/d/1nZxT-ta1g4sh5AiiM6Swa6RCnJ-dF-PfDamcu8KrdY/edit
M3 Arrival Process model: development and application to circuit prioritization.	DxRisk P1 M3: arrival process model
M3 Event Classification model: given an outage, what are the probabilities the outage leads to a wire down, or ignition, conditional on asset and environmental conditions.	DxRisk P1 M3: event classification model
M3 software and data pipelines: documentation on the details of data preparation and model execution sufficient to retrace our steps and reproduce our results.	DxRisk P1 M3: software and data pipelines

We begin by framing the problem and revisiting modeling objectives. Notably, these are: (1) to characterize ignition probabilities and (2) to provide a basis for quantifying risk reduction associated with candidate mitigation measures.

We go on to discuss the statistical properties of the system and requirements of the model that make this a problem challenging to solve. Our discussion is informed by the data, a survey of the statistical methods available, and by our assessment of what opportunities there are to develop novel tools for advancing the state of the art. We summarize solutions to these problems that we have found and implemented to date, as well as those that can be used to extend and improve upon this work in the future.

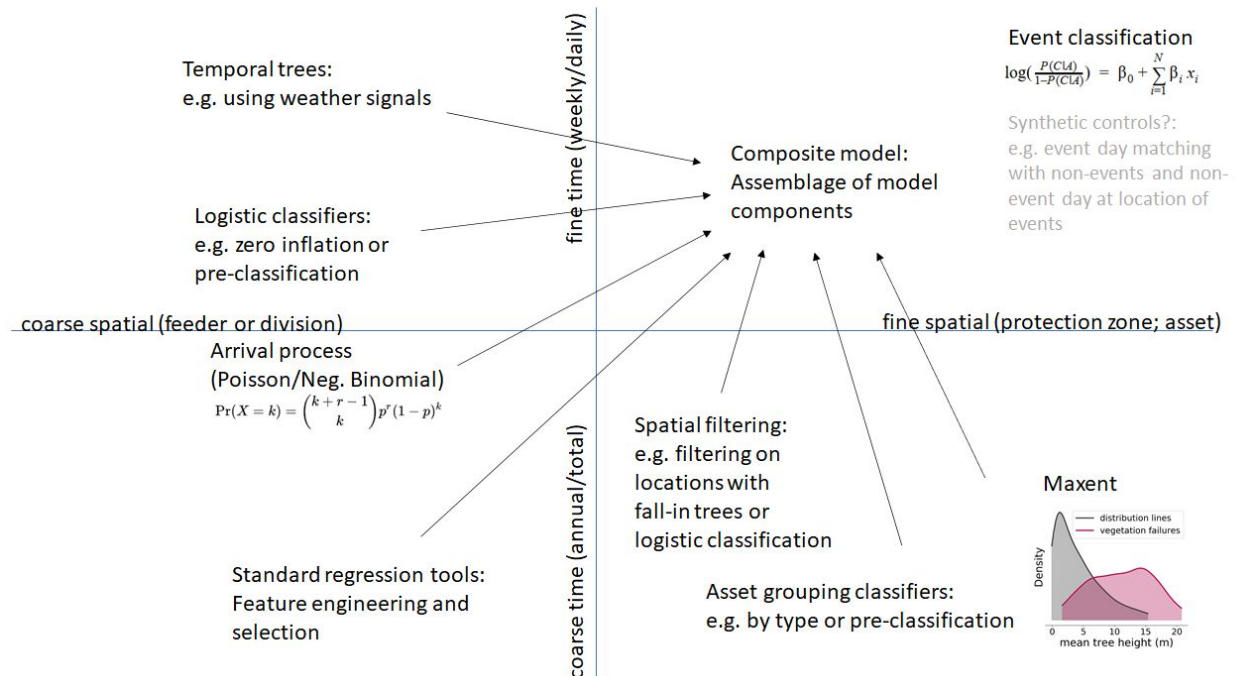
We go on to describe the structure of the models we have developed thus far. We have specified and fit multiple model types because each has different strengths in dealing with the spatial distribution, timing, and types of outage events.

“Modeling” vs. Models

The core deliverables of Phase 1 includes the specific models described in this document, but also the data infrastructure and modeling tools required to build them. The data infrastructure and tools are expected to deliver more value over time than the specific models used to provide the final quantitative results of the phase. This is because there is no such thing as a one-size-fits-all “best” model – different strategies process data in different ways, allowing them to be more or less responsive to specific patterns in underlying data. The primary tradeoffs in modeling outages and ignitions relate to the degree of spatial and temporal precision of the results and the ability to express occurrence probabilities and arrival rates in a manner that is conditional on asset and environmental characteristics.

The figure below places various modeling strategies along two axes that span spatial scale from coarse to fine (x-axis) and time scale from coarse to fine (y-axis). It can be verified that the arrival process model operates with moderately coarse spatial and temporal scales, while the Maxent model operates at fine spatial scales but coarse time scales.

The event classification model is depicted on the upper right of the axes: it can be thought of as filtering both time and space to just the moments and locations of outages as a prerequisite of its operation. The event classification model is tuned to provide estimates of downstream event outcomes (wire down, ignition, etc.) that are conditional on asset and environmental characteristics, where some of those conditions can be modified through maintenance, repairs, and replacements, to reduce risk.



The Phase 1 modeling deliverable consists of three models, which can be integrated as components of more complicated models, each tuned to emphasize a different aspect of the outage/ignition prediction/classification problem (broadly termed “event prediction”), which can be summarized as **where**, **when**, and **what** types of events occur:

- Where (maximum entropy):** To answer the question of *where* events occur, we have built a model that draws on maximum entropy models pioneered in the modeling of ecological ranges of species. This model is trained on outage or ignition locations and spatial environmental and grid asset data. The data can draw from a specific time period, but the model itself is dedicated to spatial, not temporal, patterns. The Maxent model provides relative scores or probabilities for the expected occurrence of the type of event it was trained on and offers the most finely spatially resolved event predictions, thus the focus on “where”. [Documented here](#) and applied to circuit prioritization [here](#).
- When (arrival process):** To answer the question of *when* events occur, we have trained models that frame the occurrence of events as arrival processes. Specifically, we are using models that estimate the probabilities associated with potential counts of events over time using Poisson-family distributions (Generalized Linear Models with Poisson and Negative Binomial distributions). Given historical rates of event occurrences (for example event counts per-feeder per year) and associated asset and environmental covariates (for example weather data, meteorological indicators, circuit length, fall-in tree count estimates, and circuit physical characteristics) these models assign probabilities to counts of events per-feeder per-year, given a complete set of covariates for the prediction period. [Documented here](#), with results applied to circuit prioritization.
- What type (event classification):** At the heart of scenario based modeling of risk mitigation is the ability to assign probabilities that an outage event will also be an ignition

event (or some related estimate of downstream outcomes). The event classification model takes data on outage, wire down, and ignition events to derive the probability that a given outage will lead to an ignition (or wire down), given the attributes of the outage, where outage attributes include asset characteristics and environmental conditions at the time and location of each outage. The event classification model is based on a regularized fit using logistic regression and is [documented here](#), with separate sets of results by event type and equipment type.

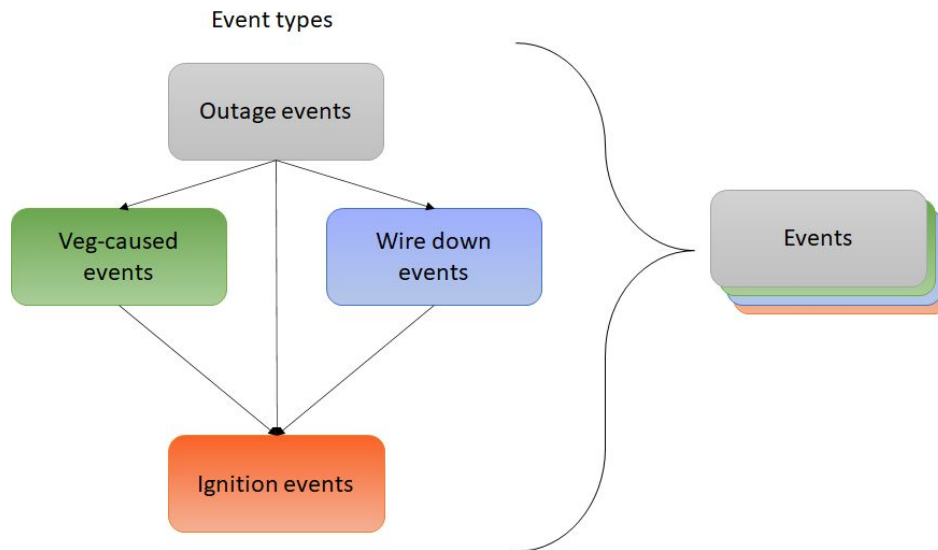
This document provides high level motivation and structure for each model, points to existing documentation on usage and outcomes, and provides context on how each of these models can be fit as a component in a larger model. We describe how these models can be interpreted to estimate ignition probabilities, and to characterize risk reduction associated with mitigation actions.

Problem Statement

The objective of this work is to specify a model for characterizing the probability of an ignition, given what is known about the environment and the condition of the grid. These probabilities are to be coupled with fire spread models to describe the consequences of an ignition should it occur. The primary application for the model will be to examine how different grid hardening measures can mitigate wildfire risks that are inherent to operating the grid.

In light of these requirements, relevant models must provide several capabilities:

- Data-driven estimates describing the probability of an ignition;
- Temporally and spatially explicit estimates of when and where ignitions are likely to occur;
- Characterization of causal factors that give rise to ignition;
- Model structure suitable for inferring how mitigation measures reduce the probability that an ignition will occur.



The models described here constitute the current state of the art in probabilistic risk modeling for wildfire mitigation. The current report provides a mathematical framework for quantifying these risks. It is well documented in the literature, however, that accurate parameterization of rare events such as wildfire ignition is a notoriously challenging problem to solve. Here, we summarize why the characteristics of the problem and the requirements of such models make this problem challenging to solve.

What Makes This a Hard Problem

Sparse Data

The tools that are most used to fit and evaluate machine learning models are predicated on the assumption that the data are sufficient to train a model. These assumptions may break down when data are inaccurate or when observations are sparse. If observations are too sparse to capture the full range of outcomes that could occur, it is possible to arrive at an interpretation of the data that is blind to these outcomes.

Probabilistic methods for fitting models can address many of these issues by calling into question the validity or representativeness of the data, and by exploring the possibility that the “optimal” model fit is merely a construct of the data that are present. If data are suspect, alternative interpretations of the system may be incorporated in the form of priors. Specifying priors can allow subject matter expertise to be internalized to explore the merits and limitations of data, and build on insights that may not be evidenced in the historical data.

The data science libraries most widely used and most thoroughly vetted (Scipy and Statsmodels) rely on frequentist methods and do not support probabilistic fitting methods. The current work describes a modeling structure that can be adapted to more probabilistic

implementations. In the interest of furthering our understanding of the functional form and features of the model, however, work to date has focused on implementations that build on more standard Python libraries.

Zero Inflation

Zero inflation describes a property of certain random variables where it is not always possible for a particular outcome to occur. For example, daily pole fire data for every pole on the system would be considered a “zero inflated” dataset. Most poles on the system never catch fire, and those that do may spend decades in operation before they ignite. The data would be overwhelmingly dominated by “zeros” to denote poles that *did not* catch fire on a particular day.

Problems where there are substantially more observations of non-events than there are observations of events suggests that most of the data describe conditions that could be completely irrelevant to characterizing the sensitivities that are of interest. A model trained on zero-inflated data can yield parameter estimates more heavily influenced by conditions of little interest than by those conditions that are indeed problematic, thereby masking sensitivities to conditions that pose the most severe risks.

Zero inflated models address this challenge by breaking the problem in two. The process is described by an interaction between two random variables—one characterizing when and where events *do not* occur, and the other describing the number of events that could occur under amenable conditions. We elaborate on this approach below in our description of the Arrival Process.

Characterizing Probabilities

Characterizing probabilities is an inherently statistical undertaking. Elevated risk offers no guarantee that a particular outcome will occur, and there are undoubtedly instances in the data where ignitions *did not* occur despite relatively higher risk.

The current work describes a mathematical basis for estimating the probability that a particular outcome will occur, and documents methods for training, tuning, and testing the model using event data from the past. However, it is worthwhile to note that the true probability distributions that give rise to the data can never be known.

Inferring Causality

Ignitions result from complex interactions between weather, ecosystems, and grid assets. These dynamics are not perfectly understood and often are monitored only indirectly. The number of splices on a particular span, for example, does not tell us about the condition that those splices are in.

While we may find correlations to suggest that certain attributes of the system contribute to elevated risk, it may not be possible to infer causal relationships. This poses a challenge to quantifying risk reduction associated with mitigation measures.

The current work describes an approach that uses expert judgement to inject assumptions about how mitigation measures will alter correlations observed in the data. To develop robust statistical methods for doing “causal inference” typically requires randomized controlled trials or other managed interventions rather than the “natural experiment” of past outcomes this work is based on.

Data Sources and Synthesis

- What decisions were made
- What differs between Milestone 2 documentation

Methods

Here, we describe the underpinnings of our model. We formulate a three-stage model comprised of:

1. The MaxEnt model
2. The Arrival Process model
3. The Event Classification model

Below we elaborate on the documentation provided for each of the three model components to show how the pieces fit together, and how the outputs can be synthesized to generate conditional failure probabilities.

This multi-component structure provides a number of distinct advantages over alternative formulations. First, by training the Arrival Process Model on outage data (rather than ignitions), the model can capitalize on a deep time history of outage data which dates back to 2007. Thus only the Event Classification Model requires ignitions data, which only dates back to 2015. Second, we build on insights gained from the Maxent and Arrival Process models to inform our understanding of when and where events are (and are not) likely to occur. Leveraging the outputs from those models mitigates issues with zero inflation in the Event Classification model.

The MaxEnt Model

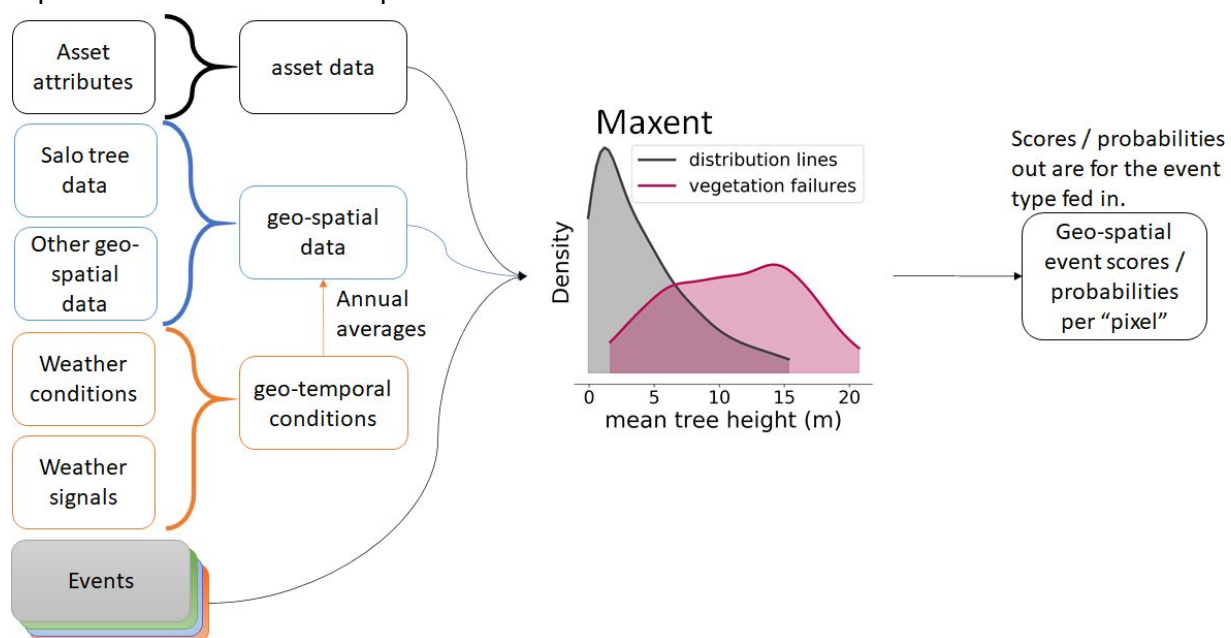
The principle of maximum entropy (MaxEnt) refers to the expectation that all things being equal, systems will tend to be found with the macroscopic properties that can arise from the greatest number of underlying micro-scale configurations - a system's entropy is closely related to the number of micro-scale configurations. Just as some models can be estimated by selecting parameters that offer the maximum likelihood solution, MaxEnt models are solved by tuning

their parameters to maximize entropy. In essence, MaxEnt applies a mathematical analog to Ockham's razor: the least unique solution is the most likely one.

MaxEnt models have been successfully applied in ecology to the problem of estimating a species' range, given a set of locations where members of that species have been observed and the corresponding environmental conditions at those locations and all candidate locations for the range. In that context, the model assigns a score to every location that captures how similar the conditions at that location are to the locations where the species was observed. The correspondence between MaxEnt applied to species observations and ranges and outage/ignition locations and at-risk locations is fairly obvious. We have applied MaxEnt methods to event occurrences and their proximate asset and environmental conditions contrasted with the background conditions everywhere else along the distribution grid to identify the locations most likely to experience similar events in the future.

The MaxEnt model has been [documented in detail](#) and its application demonstrated to the Circuit Prioritization problem is [available here](#). The MaxEnt model internalizes information about long time-scale trends and geographic features to draw conclusions about where different types of events are more (or less) likely to occur. Outputs from the MaxEnt model are incorporated into downstream models as a feature capturing spatial variability in risk exposure.

Note: the software package used to train MaxEnt models is called Maxent (lowercase "e"), and the use of a lowercase "e" in referring to the model should invoke the specific software implementation we have adopted.



Arrival Process Model

The Arrival Process model and its application to Circuit Prioritization is [documented here](#). Here, we describe the mathematical underpinnings of the model, and discuss how the specification could change to support other applications.

An “Arrival Process” is a component of queuing systems that characterizes the random process by which new arrivals enter the queue, thus accumulating over time. The canonical example in queuing literature describes computational jobs waiting to be processed by a server. Another common (and more intuitive) example describes the number of people in line waiting to be served by tellers at the bank. The defining characteristic is that arrival processes follow a count distribution--such as a Poisson or Negative Binomial distribution--where realizations take on non-negative integer values. Arrival processes are parameterized by an “arrival rate” (or alternatively by the inverse, the “interarrival time”).

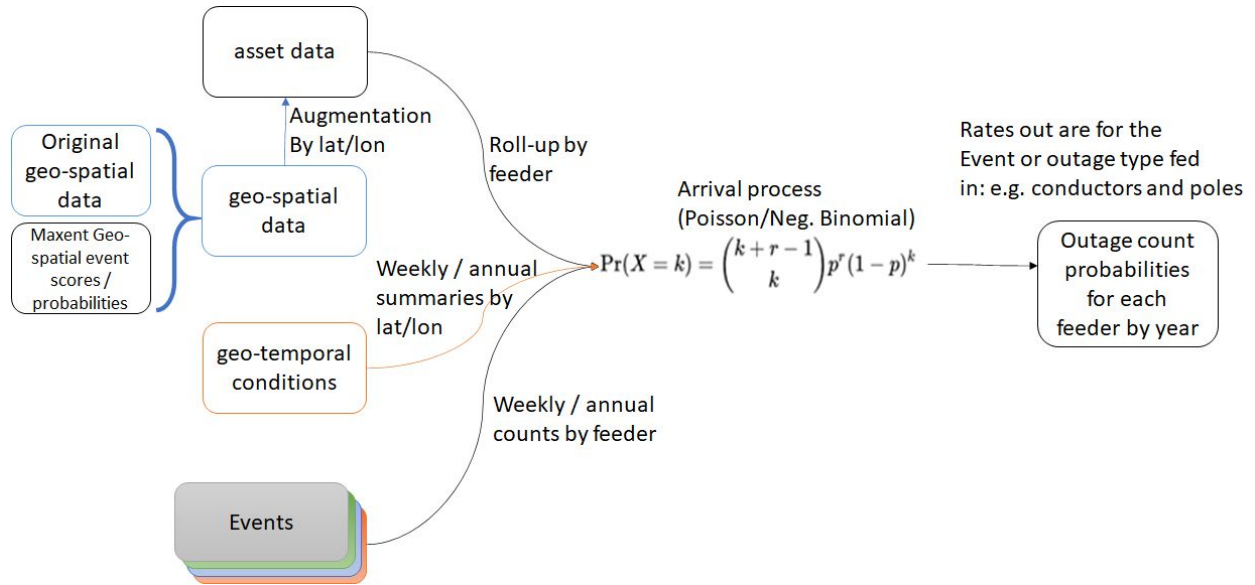
Spatio-Temporal Variability

Outages are an “inhomogeneous” arrival process with an arrival rate that varies in both space and time. A key challenge in characterizing the rate is that it cannot be observed directly, but is a derived property that must be inferred from arrivals, whose arrival rate is not constant. Further, if we attempt to characterize the rate at a very granular scale, sparse observations will cause us to overestimate the rate in areas where outages have happened before, and underestimate the rate in areas where they have not. The outage record must be sufficient to represent the underlying data-generating process; the amount of data needed depends on the complexity of that underlying process.

Aggregation of arrival data must weigh data sufficiency against the insights needed to support a particular decision-making process. The defining factors that make two outages “similar” or “different” may change, for example, depending on what questions are being asked of the model. For circuit prioritization, arrival rates were calculated in terms of the number of outages on each feeder per year. Different possibilities could include groupings by asset type, cause, weather signals, or downstream risks. We reiterate, however, that sparse data will set a lower limit on how granular the model can be.

Formulation

The diagram below illustrates the data flow of the arrival process model, which is trained on asset data, aggregated to grid-topology-defined sub-groupings (feeders in this case), environmental conditions found at the locations of the assets (summarized for each sub-group), and weather conditions found at the locations of the assets (summarized spatially for each sub-group and temporally into annual extracts of weather conditions).



We use a library of generalized linear models suitable for modeling count processes. A generalized linear model (or GLM) builds on techniques used in linear regression, but allows a broader range of probability distributions to be specified. Selected from among several candidate count distributions, we have adopted a Negative Binomial distribution (specifically, the NB2 form, where the variance in counts over time is equal to the mean plus a coefficient α times the square of the mean, where α can be derived from the results of a Poisson fit to the same data). See [Greene \(2007\)](#) for details on the NB2 formulation and [statsmodels](#) for the implementation.

To provide intuition about the formulation, we begin with a linear regression model given by:

$$\hat{y} = \beta_0 + \sum_{i=1}^N \beta_i x_i$$

Here, the response variable Y is simply a linear combination of the features X , where each feature x_i is given some weight β_i . This formulation assumes that both X and Y are normally distributed, and the predictions \hat{y} are simply the expected value $E[Y]$. GLMs use similar concepts but allow the response variable to take on any arbitrary distribution.

In the context of GLMs, the function relating the parameters of Y (e.g., the mean) to the inputs x_i is sometimes called the “link” function. The link function may take on any functional form that makes sense given the nature of the data. In the context of count models, for example, it is conventional to use an exponential link function describing the arrival rate λ (or arrivals per unit time) given as follows:

$$\lambda = \exp \left[\beta_0 + \sum_{i=1}^N \beta_i x_i \right]$$

The use of the exponential link function ensures that the arrival rate is non-negative. For a Poisson arrival process, the response variable N is the count process, and the probability of observing a particular value n is given by:

$$P[N = n] = \frac{(\lambda t)^n}{n!} e^{-\lambda t}$$

One attribute of the Poisson distribution that can limit its ability to characterize certain count processes is that the variance is equal to the mean. When the variance is greater than the mean, there is said to be “over-dispersion” in the data. The Negative Binomial distribution provides an additional parameter α (the “coefficient of dispersion”) that allows the variance to differ from the mean.

Choosing a Count Distribution

We use the Cameron-Trivedi dispersion test to check for over-dispersion and to select between the Poisson and Negative Binomial distributions. To perform the test we compute α , which is given as follows:

$$\frac{(n_i - \lambda_i)^2 - n_i}{\lambda_i} = \alpha \cdot \lambda_i + 0$$

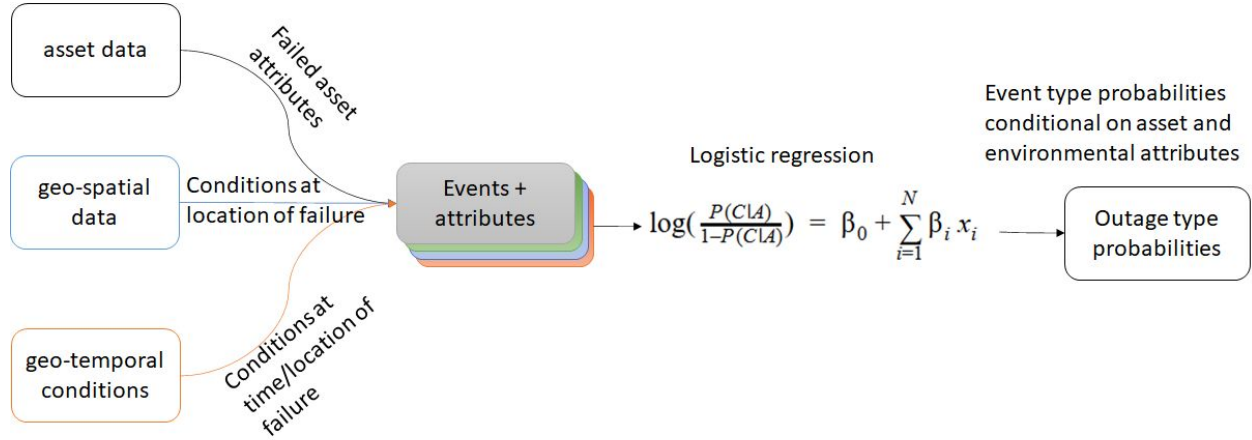
where n_i is the number of outages during time interval i and λ_i is our estimate of the rate.

We set the left hand side of the equation above equal to some random variable Y . Given the relationship $Y = \alpha \cdot \lambda + \varepsilon$, we use OLS to estimate α . If the confidence interval on α is inclusive of zero, the test concludes that the Poisson distribution is a good fit. When α is positive there is said to be overdispersion, and the Negative Binomial distribution is used instead.

Event Classification Model


Details about the development and implementation of the event classification model are [documented here](#). This model internalizes asset and environmental context to characterize how real-time dynamics change the probability that an outage will escalate into an event that poses a greater ignition risk.

The figure below illustrates the data sources and their flow through the event classification model.



Model Structure

We formulate the event model as an acyclic graph where each node represents an event (e.g., an outage, a wire-down, or an ignition), and each edge indicates a probabilistic relationship between two events (e.g., between an outage and an ignition). In other words, given the knowledge that an outage occurred, the graph structure informs us of the possibility that an ignition could also occur.



The nodes specified here and the directionality of the arrows describe a simple representation of causal relationships between events informed by our own judgement and interpretation of the data. One could imagine that changing the graphical structure underlying the model could shed light on relationships between different types of events of unique interest to a particular use case or decision-making process.

Mathematical Formulation

Conditional probabilities are estimated using logic regression. The logistic function is given by:

$$\log\left(\frac{P(C|A)}{1-P(C|A)}\right) = \beta_0 + \sum_{i=1}^N \beta_i x_i$$

For ease of notation, the left hand side gives the “odds ratio” $\frac{p}{1-p}$ (one could also imagine solving for p). The probability, $P(C|A)$ describes the conditional probability of Event C , given the knowledge that Event A has occurred. The features X describe characteristics of the grid and of the environment at the time of the outage.

In words, the probability of an ignition is conditioned on the knowledge that an outage occurred. Training these probabilities on the subset of the data where outages (or other upstream events) were observed can mitigate issues with zero inflation. We partition the outage data by equipment type to internalize differences in sensitivity to specific ambient conditions (e.g., wind

speed, or proximity to the coast) changes the regularization parameter and features selected. This partitioning also allows us to identify the types of outages where data are sparse.

In the results presented here, we model ignitions and wire-downs as distinct events. If we consider the various sequences of events that could give rise to an ignition, the probability of observing an ignition $P(C)$ becomes:

$$P(C) = P(C|A) + P(C|B)P(B|A)$$

Work to date has focused on probabilistically classifying outages into downstream events $P(C|A)$ and $P(B|A)$. Future work will examine different sequences of events that lead to the same outcome, as well as more complex graph structures.

The coefficients β_i are solved for using maximum likelihood estimation with L1 regularization. The effect of the regularization is to add a penalty to model complexity that is capable of forcing parameter estimates to zero, thereby eliminating features that lack explanatory power.

Model Structures, Complexity, and Mitigation

Explicitly specifying links where conditional probabilities between two events can enable analysis of how changing these probabilities could alter the risk of downstream events. Certain mitigation measures could remove a particular link from the graph; for example, undergrounding would eliminate the risk of wire-down events. Other mitigation measures may change the underlying conditional probabilities; for example, advanced protective equipment could reduce the likelihood that a wire-down will lead to an ignition. Where the data are too sparse to characterize these probabilities with confidence, beliefs about how mitigation measures could improve performance can be informed by expert judgement.