

# Predictive Models of Maintenance Needs for Power Distribution Wood Poles Using Machine Learning—A Conceptual Case Study



Alexandre Cesa, Carla Boehl and Kecheng Shen

**Abstract** The maintenance of a large network of power distribution wood poles is a challenging activity, in part due to the extensive geographical area in which the network is distributed. It is not uncommon that an operator has to maintain an infrastructure with hundreds of thousands of wood poles installed over an area bigger than many countries. Behind the challenge, there are various factors that drive the maintenance and renewal requirements of wood poles including their age, wood species, insect activity, soil composition, weather conditions and many others. As these factors highly interact with each other, it is not immediately obvious or simple to decide which poles to prioritise for the maintenance/renewal in lieu of others. Therefore, an improved process capability of decision making in this regard will have direct positive implications on the ownership cost, safety and level of service such as availability of the network. This paper discusses and illustrates how machine learning techniques could support decision making processes by predicting wood poles maintenance needs, i.e. by concomitantly assessing the multiple factors that drive the pole failures/integrity loss, through the use of diverse and disparate data sources such as geographical area features and scanned hand-written records of historical maintenance.

## 1 Introduction

The effective and efficient maintenance of a wood pole power distribution network is a complex topic. A few key challenges include the large geographical area where the wood poles are installed and the different variables (also called features

---

A. Cesa (✉) · C. Boehl · K. Shen  
Curtin University, Perth, Australia  
e-mail: [19107084@student.curtin.edu.au](mailto:19107084@student.curtin.edu.au)

C. Boehl  
e-mail: [carla.boehl@curtin.edu.au](mailto:carla.boehl@curtin.edu.au)

K. Shen  
e-mail: [kechengshen@gmail.com](mailto:kechengshen@gmail.com)

throughout the paper) that drive their economical and safe life. It is not uncommon a power distribution operator having under its management hundreds of thousands of wood poles—each one potentially requiring frequent inspection to assess its condition and maintenance needs.

A report issued by the Western Australia General Auditor in 2013 [1] highlights the topic's complexity and its relevance. It was found the practices that Western Power (operator) had followed to maintain its 600+ thousand wood poles across the state were not adequate, with direct adverse implications to the safety and level of service of the network. In this particular case, decades of underinvestment means an expenditure in excess of AU\$1 Billion, above and beyond normally expected maintenance expenditures, would be required within the next decade to re-instate the safety of the network. In all, the report suggests that Western Power's Asset Management System was not adequate [1].

This paper discusses some conceptual and practical aspects of the analytical techniques that can support a better decision making process concerning the wood poles problem setting. It explores how machine learning methods could be used to handle large amounts of disparate data. In particular, the K-Nearest Neighbour method is used to demonstrate how geographical features of a large footprint wood poles network can be attained. In addition, an Artificial Neural Network method is used to attain features from hand written historical inspection sheets scans. Finally, the concept of how the disparate data sets could be analysed together and used to drive better informed decisions for maintaining this particular asset type are discussed.

## 2 Background and Motivation

The expenditure figures the report [1] suggests are fundamentally based on future maintenance requirements given an average wood pole expected life and the age distribution of the poles' population installed on this particular network. This "maximum age threshold" approach suggested by the report is likely to improve this particular poor situation observed but it would not be much cost and risk effective. In regards to the cost, it will over-maintain poles that do not yet require maintenance or replacement. It will also under-maintain poles with a less-than average life—which presents a safety risk to the community. For these reasons, it can be argued a more robust method of wood poles maintenance management is necessary.

More comprehensive approaches than solely using the age threshold criteria for wood poles replacement already exist. References [2–4] are examples, where documented standards and decision trees support determining maintenance needs of individual wood poles. Such documents are the result of many years of experience making it possible to achieve an acceptable level of asset integrity and performance. Often, these guidelines are a collection of technical considerations, practical experience and heuristics shortcuts [3].

Life of power distribution wood poles is affected by at least 34 different variables including soil type, weather condition, wood species, pole load and many others [5–7]. Mapping many of the influences that individual variables have on wood poles life could be easily achievable through life distribution such as Weibull and other simple modelling techniques like linear regression. These models would typically illustrate e.g. “wood pole life distributions by type of wood species” or “wood pole life distributions by geographical areas”. One shortly realizes that these variables are hardly independent from each other. For instance, life by wood species is affected by fungi rot, which in turns depends on soil type and a “catalyst” like rain to grow. Fungi rot would also be affected by which chemical treatments the wood pole might have received and how long ago it happened [6, 7]. In all, this very complex structure of variables’ dependency is what ultimately drive the lives of individual wood poles. This leads to the realization that single-variable or more generally low dimensionality models (i.e. models with few variables) as so far described could only be effective to some very limited extent.

In all, it can be said an ideal maintenance decision model would be one capable of predicting with a great deal of accuracy the life characteristics of each individual pole. The achievement of such granular predictive capability would enable one to truly optimise this asset system’s performance, cost and risk profiles as a whole. There are different reasons why simple maintenance models as before discussed are suggested and ultimately adopted. These include the complex multivariate data analysis and the computational capability required to undertake it, the difficulty in retrieving and consolidating historical data stored in different formats and others. As time elapses and information technologies evolve, these challenges will become easier to deal with.

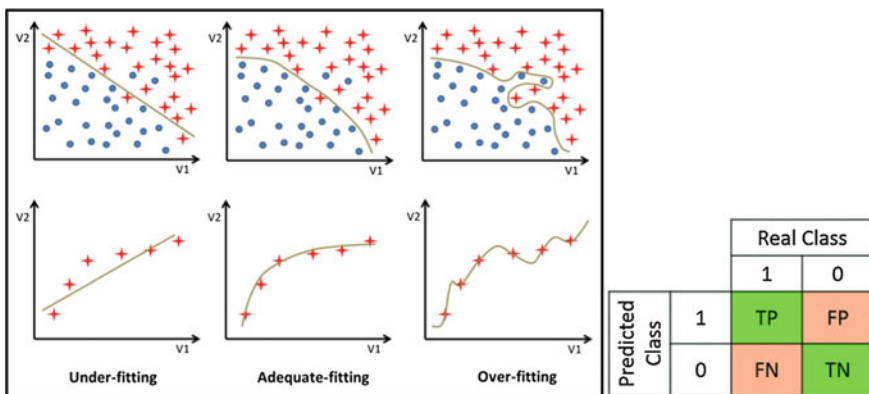
### 3 Literature Review

Machine learning (ML, also known as Statistical Learning) refers to a large set of tools for modelling and understanding complex data structures. With the ever increasing number of “big data” problems, ML has become a very hot field in many scientific areas (as well as other disciplines). Broadly speaking, supervised ML involves building a statistical model learnt from the data itself for predicting, or estimating, an output based on one or more inputs [8]. In other words, ML models differ from other modelling techniques that depend on prior knowledge of first principles and cause and effects. This means one can successfully build complex and yet accurate ML models without domain knowledge of the topic (though domain knowledge sometimes helps to increase model assertiveness).

There are different ML techniques; a few popular ones include Linear and Logistic Regression, Support Vector Machines (SVM), K-Nearest Neighbour (KNN) and Artificial Neural Networks (ANN) [9, 10]. Regardless of the technique selected, the idea of supervised learning involves using a dataset with known inputs and outputs to “train” a selected model hypothesis. By training the model we mean

iteratively conducting model parameterisation with the objective of minimising the model’s output error: i.e. given e.g. inputs “A”, “B” and “C” the model should return output “X” (and not “Y”, for instance). It is quite easy for a multi-parameter model hypothesis to deliver a very low or zero-error output. Although this might sound promising, such a model would be of limited value to predict outputs from inputs not yet seen, meaning the model does not generalize well. Such a problem is related to the bias-variance trade-off and is illustrated in Fig. 1 by three model hypothesis fits to the same dataset: an under-fitting case (left, high bias), a reasonable model fit case (middle) and an over-fitting case (right, high variance). It can be observed in the under-fitting case that a linear boundary separation is not adequate as it is too simplistic. Alternatively, a highly “wiggly” model fits the training data well but would not likely generalize well to new unseen data. An adequate model hypothesis would be one in which the fit to the dataset minimizes both bias and variance simultaneously, i.e. the reducible error. As illustrated in both classification and regression examples in Fig. 1, in general there will be cases in which the model prediction (regardless of how good it is) would still yield what is known as an irreducible error. In the classification setting, this irreducible error means that in some instances, the model will miss-classify the correct output. This is easily illustrated through a Confusion Matrix as shown in Fig. 1 (right) for a binary classifier case.

There are different techniques to address the issue of model generalization. A very basic approach is to divide the dataset in “training dataset” (e.g. 70% of total size) and “validation/test” dataset (e.g. 30% of total size). This enables one to parameterise the model using both training and validation datasets and later, experimenting by means of the test dataset, the accuracy of the responses predicted by the model. This basic process is conducted iteratively, until the point when the model generalises sufficiently well to meet the user’s expected predictive capability.



**Fig. 1** Left: Different model hypothesis illustrating classification (top) and regression settings. Right: Confusion matrix—TP (true positive), TN (true negative), FP (false positive or type I error) and FN (false negative or type II error)

Additional iterative approaches include feature engineering, using different models (SVM, ANN, etc.), optimization techniques, model hyper-parameters and others. Resources [8, 11, 12] are gentle, self-contained and comprehensive introductions to the ML topic for the interested reader.

## 4 Methodology

For the purpose of model development, maintenance requirements of individual poles could be classified as Do nothing, Repair and Replace. A Repair typically involves doing (e.g.) some form of reinforcement of it the at the cost of \$1,000 per pole while a Replacement costs approximately \$9,000 [1]. Given the safety risks involved and large capital required to sustain level of service of this asset system, more accurate predictions of individual poles is obviously highly desirable.

Two practical uses of Machine Learning techniques are demonstrated. The first one illustrates how to attain geographical features from colour coded maps. A second method illustrates how to attain hand-written features from scanned inspection sheets. Finally, the overarching architecture concept of how these disparate models are collated for later individual poles' classification is illustrated and some relevant aspects of the task are discussed.

### 4.1 *Attaining Geographical Features*

Figure 2 illustrates how geographical features of Average Temperature, Average Rainfall, Soil Type and Termite Hazard were attained using Matlab (Ref. [13] contains a web-link to the code). As a first step, a network footprint "mask" is placed over the different feature maps in a process called image registration. Later, the KNN technique is used to discern the different colour levels of these maps, which are translated into different classes (e.g. three classes for the Soil type: Poor Drainage = 0, Average Drainage = 1 and Good Drainage = 2). Finally, feature maps are overlaid on each other and an arbitrary point is selected to show the different variables encoding and their actual physical meaning.

It is not surprising that the maintenance needs observed in particular areas may tend to cluster towards particular outcomes. However, outcomes might still vary widely within a small area and the "geographical map" approach can no longer be extended to successfully predict each individual pole's maintenance needs. Within the small area considered, obviously there will be poles with different ages, which received different levels of service over the years and other idiosyncrasies.

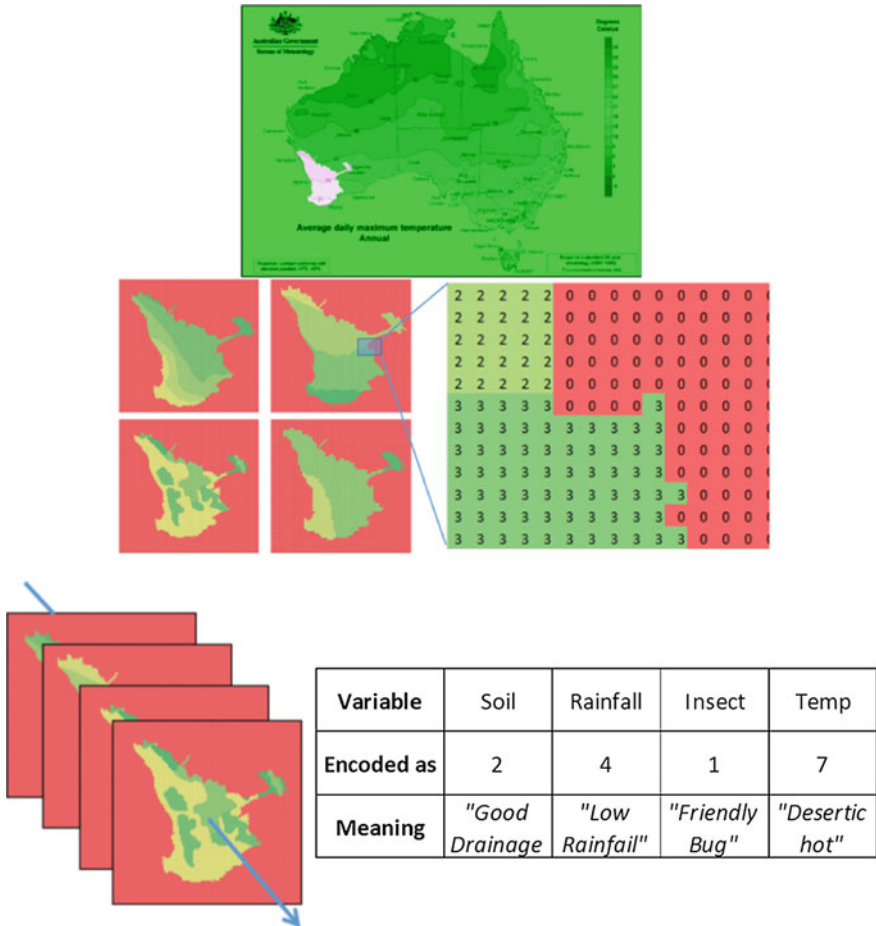
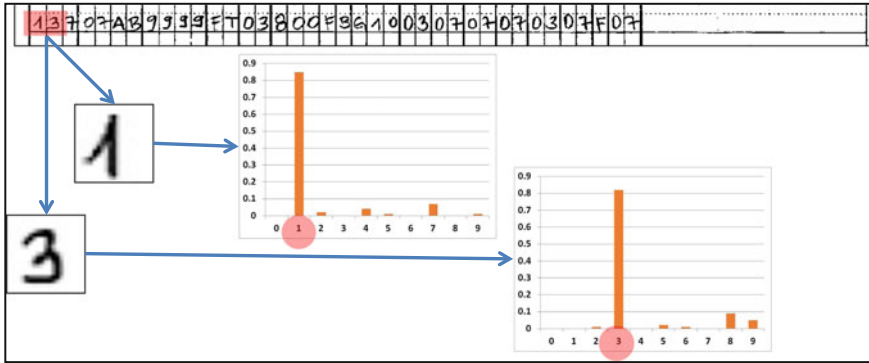


Fig. 2 Western power’s footprint image registration in a particular geo-feature map; four feature maps with close-up detail, maps lined up for illustrating features of a particular location

### 4.2 Attaining Hand Written Features

Individual poles’ inspection history is a data source that contains a wealth of information, likely recorded in different formats over the past decades: paper based and perhaps in electronic format more recently. Ignoring historical paper based inspection records would be naïve, so a method to retrieve this data is paramount. For this particular situation, the capability to recognize hand-written digits and letters would of great value. Yan LeCunn pioneers this task utilising the “MNIST dataset”—which is a collection of 70 thousand hand-written numbers from zero to nine [14]. The MNIST dataset is a popular dataset used by the ML community to benchmark algorithm performance in regards to the image recognition task. State of



**Fig. 3** Written digits scan from (e.g.) a wood pole inspection sheet and ANN correct probabilistic outputs for two example digits, “1” and “3” data encodings

the art algorithms currently reach 99.7% overall accuracy in correctly classifying hand written digits [12]—which in practice is as good as humans can be in the task of classifying these hand-written digits.

Figure 3 shows the bottom scan of a typical wood pole inspection sheet (see e.g. [4] as an example), with hand written records (encodings) regarding e.g. different treatments of the wood pole and their associated dates in between other featured records. The bottom of Fig. 3 shows two digits classification, “one” and “three” outputs from the classifier developed in Matlab [13]. It can be seen the ML ANN algorithm used performs well in the classification task, achieving approximately 95% overall accuracy for the MNIST benchmark.

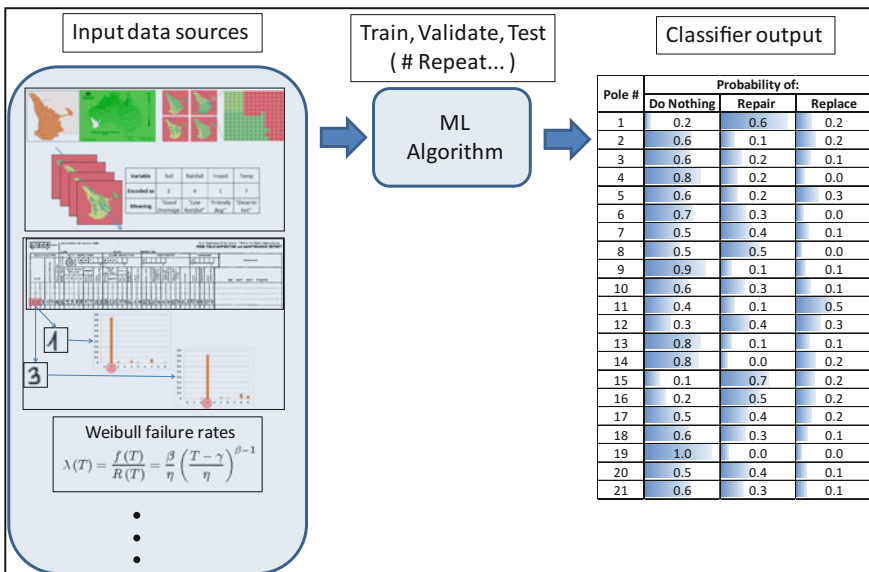
### 4.3 Final Classifier—Conceptual Architecture

The grouping of geo-features and individual poles historical inspection data records, as demonstrated, would very likely enable one achieving a sufficiently good predictive model of wood poles future maintenance needs. This is in part due to the fact the available dataset is relatively large and there are lots of examples to train a final classifier to properly classify the “Do nothing”, “Repair” and “Replace” classes. In the particular case of Western Power network, there are more than 600 thousand wood poles with an average age of 45 years and, likely, millions of inspections have been completed over time. In the ML community, there is the popular saying that more data beats complex algorithms [15], which truly well applies in this case.

Nevertheless, it is relevant to note that the “Do nothing”, “Repair” and “Replace” classes are not well balanced, meaning the classes proportions are not similar (i.e. each class does not comprise of ~33% of the dataset). Over the life of a pole, there will be far more events of “Do nothing” than “Repair” and similarly, more “Repair” events than “Replace” ones. Although the training data might be abundant indeed, it is often difficult to achieve good classification performance for all the classes on a heavily unbalanced dataset [16, 17]. ML algorithms trained with

unbalanced data will tend to misclassify the less representative classes, which is of concern if these are “Replace” cases that had been misclassified as (e.g.) “Do nothing”.

One might rightfully attempt to improve the final model classifier assertiveness by (e.g.) increasing the number of explanatory variables or through feature engineering—which typically requires some domain knowledge. For instance, the training dataset that will be used to train the final classifier could be augmented by having other explanatory input variables not yet discussed. Potential candidates could be (e.g.) Weibull distribution Failure Rates regarding wood species and/or geographical locations. It can be argued that the inputs to the Weibull distributions suggested are somewhat reflected by the existing variables already being used—like the geo-feature maps. Such a case will lead to the problem of Collinearity, meaning two or more variables being considered are highly correlated and hence, not much additional information could be achieved by having these additional variables. The Collinearity issue can be addressed by performing dimensionality reduction through different techniques, like e.g. Principal Component Analysis—PCA [8]. The “correct” number of explanatory variables is dependent on the user expectations of accuracy and it is a trial process that trades-off different desirable aspects. In general, having too many (too few) explanatory variables increase (decrease) computational power requirements and might increase (decrease) the model assertiveness. The final classifier dataflow structure is conceptually illustrated in Fig. 4. It can be seen that outputs from previous ML models are used as inputs to another ML layer. Also, the required iterative nature of the process is mentioned and typical classifier probabilistic outputs are shown.



**Fig. 4** Schematic dataflow of final classifier. Different inputs, iterative ML classifier algorithm process, classifier outputs for individual poles



In the wood pole classification problem, it is obvious the adverse risk of misclassification, in particular not identifying the poles that require Repair or Replacement i.e. False Negatives (Error Type II) cases. It is known that the failure to identify wood poles that require immediate or more frequent attention had resulted in catastrophe in the past [18]. Such a risk can be minimized by calibrating the decision boundary threshold between the three classification classes as shown in Fig. 4—which inevitably would lead to more False Positive (Type I Error) outcomes and increased operational costs. The fine-tuning of the classification model output should be an iterative process driven by the risk appetite the operator and/or stakeholders are prepared to accept in line with their existing Asset Management Policies, Strategies and Plans.

Generally, ML models need to be continuously “re-trained” as more data becomes available over the time. This is the case because it is reasonable to assume the model’s explanatory variables might not be Identically Distributed—meaning that their behaviour (as well as the whole asset system behaviour) changes as time elapses. For instance, the quality of chemical treatments on wood poles conducted over the past decade is better than those conducted prior to this period or the inspections during the 1980s were not as comprehensive as after the 2000s. These changes over the time obviously affect the asset system performance and therefore, the predictive model needs to evolve accordingly.

There are many other nuances to the topic that are not addressed here but it suffices to state that the development of successful Machine Learning applications requires a substantial amount of “black art” that is hard to find in textbooks [15]. In all, it can be said one has to trial different approaches to achieve adequate model’s performance as each case has its own peculiarities.

## 5 Conclusion

This paper offers an overview of how the topic of machine learning could support maintenance decisions for the case of a wood poles power distribution network asset system. It demonstrates two applications of ML methods to extract information from disparate raw datasets and conceptually, how these could be later used as inputs for a final predictive classifier. It also discusses in general the importance of considering multiple variables concomitantly, which is indeed paramount to many problems involving physical asset systems. The principles herewith discussed would benefit the decision making processes for various other asset systems too and potential developments in this regard are many. In the particular context of Engineering Asset Management, this includes e.g. predicting process plant upsets by continuously monitoring and processing process historian databases.

## References

1. Western Australia Auditor General's Report (2013) Western power's management of its wood pole assets. Report 17, Nov 2013
2. Ausgrid, Network Standard NW000-S0098 (2015) NS145 Pole Inspection and treatment Procedures, Sept 2015
3. ActewAGL (2014) Asset specific plan. Poles, May 2014
4. USA Bureau of reclamations (1992) Wood poles maintenance. Facilities instructions, standards & Techniques, vol 4–6, Aug 1992
5. Cesa A, Shen K (2015) Use of multi-folds asset operational performance data to prioritise business improvement actions: challenges and opportunities. Presentation delivered at the WA Engineers Australia, 13 Aug 2015
6. Payhak S (2011) Leakage current in wooden structures used for power distribution. PhD thesis, Engineering School, RMIT University, Mar 2011
7. Rahman A (2003) Modelling inground decay of wood poles for optimal maintenance decisions. Queensland University of Technology, Master of Engineering Thesis, Jan 2003
8. James G, Witten D, Hastie T, Tibshirani R (2014) An introduction to statistical learning: with applications in R. Springer Publishing Company, Incorporated
9. Liu B (2007) Web data mining: exploring hyperlinks, contents, and usage data. Springer Science & Business Media
10. Friedman J, Hastie T, Tibshirani R (2001) The elements of statistical learning, vol 1. Springer series in statistics, Springer, Berlin
11. Ng A (2015) Machine learning from Stanford University by Coursera. <https://www.coursera.org/learn/machine-learning>
12. Nielsen MA (2015) Neural networks and deep learning. Determination Press
13. Cesa A (2017) Matlab code repository for the ML cases presented on this paper. <https://github.com/alecesa/Digit-recognition-ANN>
14. Le Cun Y, Boser B, Denker JS, Howard RE, Habbard W, Jackel LD, Henderson D (1990) Handwritten digit recognition with a back-propagation network. In: Touretzky DS (ed) Advances in neural information processing systems 2. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp 396–404
15. Domingos P (2012) A few useful things to know about machine learning. Commun ACM 55 (10):78–87. <http://dx.doi.org/10.1145/2347736.2347755>
16. He H, Garcia EA (2009) Learning from imbalanced data. IEEE Trans Knowl Data Eng 21 (9):1263–1284. <https://doi.org/10.1109/tkde.2008.239>, <http://dx.doi.org/10.1109/TKDE.2008.239>
17. Seiffert C, Khoshgoftaar TM, Van Hulse J et al (2007) Mining data with rare events: a case study. In: IEEE international conference on tools with artificial intelligence, Patras, Greece, 2007, vol 2, pp 132–139
18. Electrical Incident Report (2009) Power line fault and bush fire near cape Naturaliste lighthouse Dunsborough. Western Power, Perth, Western Australia, 7th Feb 2009. [https://www.commerce.wa.gov.au/sites/default/files/atoms/files/dunsborough\\_fire\\_report\\_feb\\_2009.pdf](https://www.commerce.wa.gov.au/sites/default/files/atoms/files/dunsborough_fire_report_feb_2009.pdf)