

Data Mining in the Electric Power Industry

¹Manuel Mejia-Lavalle ²Guillermo Rodríguez O.

Instituto de Investigaciones Eléctricas, Gerencia de Sistemas Informáticos
Reforma 113, 62490 Cuernavaca, Morelos, México

¹mlavalle@iie.org.mx, ²gro@iie.org.mx

Paper received on 30/07/08, accepted on 10/09/08

Abstract. We present our experiences in three Mexican electric power industry domains where we applied diverse data mining techniques. The first domain is about electric generator diagnosis. The second one is related to flashover forecasting in high-voltage insulators. The third case is about obtaining expert knowledge, applying data mining techniques to hydroelectric and thermoelectric utilities databases. We outline successful and bad practices, and comment about possible solutions or future work that we think it have to be done to maximizing the usefulness of the data mining approach.

1 Introduction

Data mining has been employed with success in various fields and in many real world problems [1]. Data mining is applied to huge volumes of historical data mainly with the expectation of finding knowledge, or in other words, when it is sought to determine trends or behavior patterns that permit improve the current procedures of marketing, production, operation, maintenance, or others. In summary data mining, or more widely expressing, knowledge discovery, is the nontrivial extraction of implicit, previously unknown, and potentially useful information from data [2]. Some of the traditionally used computer techniques to accomplish data mining are: neural networks, induction of decision trees, decision rules and case-based methods.

In this paper we present our experiences in several Mexican power electric industry domains where we applied diverse data mining techniques. The first domain is about electric generator diagnosis using expert systems plus a novel neural network paradigm. The second one is related to flashover forecasting in high-voltage insulators, where we proposed several tools to approach this problem. The third case is about obtaining expert knowledge, applying and comparing well known data mining techniques to hydroelectric and thermoelectric utilities databases. We outline successful and bad practices, and comment about possible solutions or future work that we think it have to be done to maximizing the usefulness of the data mining approach.

2 Electric Generator Diagnosis

In this Section we present an Electric Generator Failure Diagnosis (EGFD) system. The EGFD system combines two artificial intelligence approaches: expert

systems and artificial neural networks. With our expert system shell we capture the human expertise. With our neural net paradigm, we obtain knowledge from data. For instance, human experts on electric generation failures know that:

- a) Internal partial discharges can occur within the ground wall insulation at delaminations or areas where the bonding material is missing or incompletely cured.
- b) Such discharge activity is particularly common in older insulation systems such as mica folium and asphalt-mica.
- c) The main characteristic of this mechanism is that the positive and negative partial discharge activity is about equal.

Then a production rule to detect this condition becomes as follows:

RULE 3

IF: - Insulation system is made of <mica folium> or <asphalt-mica>.

and - [Positive PD at 50 mV] * 1.1 >= [Negative PD at 50 mV].

and - [Positive PD at 50 mV] * 0.9 <= [Negative PD at 50 mV].

and - [Positive PD at 200 mV] * 1.1 >= [Negative PD at 200 mV].

and - [Positive PD at 200 mV] * 0.9 <= [Negative PD at 200 mV].

THEN:

- Bonding material is missing or incompletely cured. Certainty 7. and- Exe (PHAFII).

A number is assigned to each rule: '3' in this example. Then, the keyword 'IF' indicates the beginning of the list of conditions, premises or antecedents of the rule. The first condition is true, if and only if, the user answer to the question: 'Insulation system is made of?' is 'mica folium' or 'asphalt-mica'. The second condition is true, if and only if, the variable [positive PD at 50 mV] multiply by 1.1 is greater or equal to the variable [Negative PD at 50 mV]. The same applies to the rest of the conditions. If one of these conditions happens to be false, because the user answer is different than expected, or because some variable value do not match the required condition, then the rule is false, and the inference machine of the expert system searches for another rule.

On the other hand, if all the conditions of a rule are true, then the rule is true and its conclusion is 'fired': 'Bonding material is missing or incompletely cured'. The word 'Certainty' at the end of a rule means 'the degree of certainty' or belief that the human expert has on the rule and it ranges from 0 to 10, where 10 means that the expert is absolutely certain of what the rule states.

With this production rule, the expert system can identify, with 70% certainty or reliability, that 'Bonding material is missing or incompletely cured' if the 'insulation system is made of mica folium or asphalt-mica', as stated in first condition, and if the 'positive partial discharge activity' is similar (within ten percent) to the 'negative partial discharge activity' at 'pulse magnitudes' of 50 mV and 200 mV, as stated in the rest conditions.

Then, our neural net paradigm is called using the command *Exe(PHAFII)*, where PHAF II is the module that handles the neural net. Algorithmic details of PHAF II are in [3]. We used the neural net to take advantage of the enormous amount of information currently available in many electric generator databases. Data from the

partial discharge graphs are normalized within the range [0,1] and then fed to the PHAF II neural net. The neural net, previously trained with normalized data from graphs which are typical patterns of abnormal situations, performs the recognition of the fed graph and computes the percentage of similarity using three criteria:

- a) 10,000 base, where the graphs are compared using a lineal scale from 0 to 10,000 of frequency units. With this criterion, the differences or likenesses of the graphs have the same weight at high and low frequencies.
- b) Logarithmic base, where the graphs are compared using a logarithmic scale from 0 to 10,000 of frequency units. With this criterion, the differences or likenesses of the graphs are adjusted with more weight given to the differences in the low frequencies (0 to 100) and less weight to the differences at the high frequencies (100 to 10,000).
- c) Central base, where the graphs are compared using as a reference the pattern graph.

With the mean (average) of these three criteria, we obtain a final certainty factor. This factor indicates the similarity of the fed graph and the pattern graph. If the final factor is greater than 70% the system displays the screen shown in Fig. 1.

From Fig. 1, it is observed that the system displays the graph being recognized, the diagnosis, and eight certainty factors. Four of these correspond to a 'Global' analysis (GCF), where the certainty factor was computed as the mean of the likenesses or differences at all the points of the graphs. The other four certainty factors, called 'Local' (LCF), are obtained from the same point on the graphs where there exists the greatest distance between the graphs (the test graph and the pattern graph).

We are planning, as future work, to incorporate more human and data knowledge to this system. To make more efficient this phase we will investigate about automatic elicitation tools.

3 Flashover Forecasting

To approach the flashover on high-voltage insulators forecasting, we developed and integrated four data mining tools that combine the ID3 algorithm [4] and the nearest neighbor case-based reasoning method [5]. The first tool uses data mining to build a classification or decision tree from historic data, the second generates production rules, the third operates the decision tree as an expert system, and the last, makes tests with unseen cases to evaluate forecasting accuracy. The results were compared against other classic machine learning tools like C4.5, FOIL, CN2 and OC1, and we obtain similar or better solutions.

To perform the experiments, data from an N-120P high-voltage insulator were registered during 21 days (504 examples of meteorological and surface resistance values). The attributes used were: hour of the day, wind direction, wind velocity, temperature, precipitation, dew temperature, barometric pressure, relative humidity and absolute humidity. The class attribute is the *surface resistance*, a variable correlated with the flashover phenomenon.




Fig.1 EGFD final display

Three classes were assigned to the surface resistance: "low" (for values between 1,013 and 3,489 kilo ohms), "medium" (for the interval 3,490/ 7,992) and "high" (for the interval 7,993/ 11,482).

With these tools, we could successfully determine:

- The relationship among the environmental variables and the surface resistance class.
- The variables that have more impact on the phenomenon, and the variables that hardly affect the surface resistance behavior.
- The circumstances that cause the surface resistance to be low.
- The surface resistance value 24 hours ahead, with an accuracy of 83%.

Table 1 shows comparison results. In the future, we will work in the simplification and reduction of the production rule quantity generated for the proposed tool, to facilitate the interpretation of the discovered knowledge. We will work too with more data to try to forecast with more anticipation the flashover event.

Table 1. Results comparison.

TOOL	Proposed	C4.5	FOIL	CN2	OC1
TEST-Acc	82.9	82.9	62.9	82.4	80.1
TRAIN-Acc	99.6	85.3	96.8	92.0	81.5
# Rules	496	9	61	52	NA
Default	Unknown	Low	NA	High	NA
Time(secs)	3.68	143	10.7	74	643

where:

TEST-Acc = Accuracy to forecast unknown cases (%).

TRAIN-Acc = Accuracy to forecast the same examples used in learning stage (%).

Rules = Number of generated rules.

Default = Default class (if no rule applies).

Time (secs) = Seconds required to generate the results.

NA = Characteristic not available by the tool.

4 Electric Utilities Data Mining

This case is about extracting knowledge from data. Some well-known data mining tools (C4.5, CN2, FOIL and PEBLS) were applied and evaluated for the task to obtain expert knowledge. This was done on a real power generation database with thermoelectric and hydroelectric Mexican electric utilities information over eight years of historic data. We evaluated accuracy, knowledge amount reduction and processing time. We analyzed the expert system rules (extracted knowledge) and we propose an architecture of an integrated knowledge discovery system for this electric power generation database [6].

For this research, personnel of the Performance Control and Informatics Unit of Federal Commission of Electricity in Mexico (CFE) selected the data. One table was built with 32 variables and 1,110 records corresponding to thermo and hydroelectric information for the years 1988 to 1995.

The 32 variables include the following: power plant identifier, date, plate and effective capacity; unavailability by type of failure; outage equivalent hours due to decrements, number of outages and outage hours due to failure and due to routine and corrective maintenance and other causes; fuel kilocalories; net and gross generation; permanent workers used in maintenance, in operation, and in other activities; additional workers used in maintenance, in operation, and other activities; equivalent substitution workers in maintenance, in operation, and in other activities; total personnel positions; accidents that cause lost of time; accidents in transit; days lost due to accidents; days lost due accidents in transit; sum of disabilities in percent; and various expenses.

With this data set, a supervised data mining knowledge extraction was outlined. We used the variable Power Plant Factor (PF), as the "class" or focus of attention for the experiment. To calculate the PF the following formula was used:

$$(1) \quad PF = \frac{\text{Gross Generation}}{\text{HP} * \text{Net Generation}} * 100$$

where HP (hours per period) is equal to 8,760 hours (365 days by 24 hours). It was found that only one rule describes the knowledge for 'excellent' plant factor for hydroelectric utilities, with 85.7% certainty:

IF:
 Unavailability due to failure (%) <= 9.375 and
 Unavailability due to maintenance (%) <= 0.520 and
 Unavailability due to other causes (%) <= 32.560 and
 Permanent Workers (Rest) > 822
 THEN: The Plant Factor is Excellent [certainty 85.7%]

Only one rule describes the knowledge for 'excellent' plant factor for thermoelectric utilities, with 92.3% certainty:

IF:
 Effective Capacity (MW) <= 298 and
 Gross Generation (GWH) > 1,721
 THEN: The Plant Factor is Excellent [certainty 92.3%]

We found that the variables that most affect the hydroelectric plant factor turned out to be: Unavailability due to other causes (63%), Gross generation (59%), and Effective capacity (48%). For thermoelectric utilities, the variables were: Gross generation (87%), Effective capacity (83%), and Unavailability due to failure (48%). A summary of the results is shown in Table 2. From the experience obtained in the development of the experiments described, the need of having a system to facilitate the process of knowledge discovery using data mining algorithms and the exploration of various alternatives that improve the quality of the extracted knowledge (expert system rules) was evident.

Table 2. Comparison of the number of Errors*

Plant Factor	C4.5	C4.5*	CN2	FOIL	FOIL*
Very-Low	4	7	8	1	1
Low	37	83	125	3	14
Regular	31	54	79	3	23
Average	23	67	87	10	7
Good	28	42	61	1	9

Very-Good	12	29	39	6	7
Excellent	5	22	23	0	0
Total	140\13.5%	304\29.2%	422\40.6%	24\2.3%	61\5.8%

Errors* = Number of cases misclassified using unseen data

Commentaries: FOIL has the better classification efficiency, followed by C4.5

C4.5 = results of the 'composite rule set'

C4.5* = results of the 'trial 0'

FOIL* = using similar attributes grouping

EXECUTION TIMES:

C4.5 = more than 30 mins.

CN2 = 10 mins.

FOIL = 128.1 secs.

FOIL* = 189.6 secs.

So we proposed the creation of the following data mining modules:

- User Interface: allows the user to have an integrated environment, which shows the user a screen from which he can choose different options to accomplish the data mining and to obtain the results.
- Pre-Processing: this module handles different options to prepare the information of the database before the application of the mining algorithm. This module allows, among other things, the addition or deletion of columns and rows, clustering (using several methods like ChiMerge, 1R, Chi2, etc.) of continuously valued variables to group them in (a few) labeled classes, feature selection methods and to automatically prepare the data to the format required by the mining tool.
- Mining tools: the user selects from among several data mining tools, the one to be applied to the preprocessed data. Usually, it is necessary to try different algorithms because there does not exist a perfect tool, but rather, depending on the data, some algorithms perform better than others.
- Post-Processing: through this module, the user may request the conversion of the extracted knowledge by the mining tool in a representation that it will be easier for him to understand; again, it does not exist "the best" representation of knowledge, since it depends on the user preferences. Some knowledge representations are: production rules, decision trees, graphics (OLAP), characteristic tables (prime relation tables and feature tables), Horn clauses, and prototypes.

These ideas were proposed by us before tools like Weka, Orange, Elvira, and others, arrived to the data mining community. However, still nowadays several issues related with the proposed modules are open for research, like data quality tools (profiling, cleansing, etc.), knowledge representation and visualization tools, etc.

5 Conclusions and Future Work

We have presented three data mining applications in the Mexican power industry, and the way that we approached each one of them. From the experimentations presented, we think that our proposed methods represent promising alternatives, compared to other methods, because of its acceptable performance. At Table 3 we resume our experiences: we outline advantages, drawbacks and possible solutions that we think it have to be done in the near future to maximizing the usefulness of the data mining techniques that we used in our works.

In our opinion, a great variety of Mexican power industry applications are still waiting to be tackled with data mining techniques, but we need develop more and sophisticated tools to accomplish the challenges. Some future work includes problems with real and very large power system databases such as the national power generation performance database, the national transmission energy control databases, the de-regulated energy market database, and the Mexican electric energy distribution database. Also, we need apply statistical tests to observe if the differences in accuracies, processing time or another parameter of the proposed methods are really significant.

Table 3. Our experiences and recommendations.

Approach used	Advantage	Drawback	Possible solution
Expert System	Representation of human-expert knowledge in a natural way.	Complex elicitation process.	Develop more sophisticated and computer aided elicitation tools.
Neural Network	Captures knowledge from numeric data.	It needs manual tuning. Discovered knowledge is in a black box.	Develop tools for dynamical tuning and to extract knowledge from neural inter-connections.
Induction Tree	Captures and shows knowledge from nominal data in an explicit way.	It needs previous data discretization. Obtained results are not very precise.	Develop tools for automatic and efficient data discretization. Improve output thru post-processing-visualization tools.
Data Mining	Discovers and shows hidden knowledge from data.	It needs an integration of the pre-processing, processing and post-processing phases.	Construct a integrated system with: data quality process, final user easy of interpret knowledge representation and visualization tools.

References

1. The 24th Annual International Conference on Machine Learning (ICML-2007) June 20-24, Oregon State University, USA, 2007.
2. Piatetsky-Shapiro, G. et al, Knowledge Discovery in Databases: An Overview, In Knowledge Discovery in Databases, Piatetsky-Shapiro, G. eds., Cambridge, MA, AAAI/MIT, 1991, pp 1-27.
3. Mejía, M., Rodríguez, G., A New Neural Network Paradigm for Power Systems Applications, Proceedings of the IASTED International Conference on Power Systems and Engineering, Vancouver, Canada 1992, pp. 41-48.
4. Quinlan, J., Discovering Rules by Induction from Large Collections of Examples, Expert Systems in the Micro-Electronic Age, Michie, D., (ed), Edinburgo, Escocia, Edinburgh University Press, 1979.
5. Mejía, M., Rodríguez, G., Montoya, G., Knowledge discovery in high-voltage insulators data, Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, Proceedings of the Tenth International Conference, Atlanta, Georgia, USA, June 1997, pp. 223-230.
6. Mejía, M., Rodríguez, G., Obtaining expert systems rules using data mining tools from a power generation database, Expert Systems with Applications, J.Liebowitz (ed), 14(1/2) Pergamon, 1998, pp. 37-42.