Machine-Learning Approaches to Power-System Security Assessment

Louis Wehenkel, University of Liège

A framework uses machine learning and other automatic-learning methods to assess power-system security. The framework exploits simulation models in parallel to screen diverse simulation scenarios of a system, yielding a large database. Using data-mining techniques, the framework extracts synthetic information about the simulated system's main features from this database.

Security assessment is a major concern in planning and operating electric power systems. It consists of evaluating the power system's ability to face various contingencies, and proposing ways to counter its main weaknesses when necessary. Contingencies may be external or internal events (for instance, faults subsequent to lightning versus operator-initiated switching sequences) and may consist of small/slow or large/fast disturbances (for example, random behavior of the demand pattern versus generator or line tripping).

Usually, numerical (for example, time-domain) simulation of the corresponding scenario assesses the effect of a contingency on a power system in a given state. However, the nonlinear nature of the physical phenomena and the growing complexity of real-life power systems make security assessment difficult. For example, monitoring a power system every day calls for fast analysis, sensitivity analysis to identify the salient parameters driving the phenomena, and suggestions on how to act on the system so as to increase its level of security. On the other hand, increasing economic and environmental pressure make the conflicting aspects of security and economy even more challenging. To meet these challenges, we need methods different from the standard time-domain simulation approaches.

This article describes ongoing research and development of machine learning and other automatic-learning techniques and their adaptation to the specific needs of power-system security assessment. In particular, I describe a framework that integrates several of these techniques so that users can extract relevant information tailored to their decision-making needs. Among the many other potential applications of automatic learning in power systems, security assessment is probably the most needed and versatile.[1]

The machine-learning framework

Figure 1 shows the framework for applying machine-learning methods to security assessment. Random-sampling techniques screen all relevant situations in a given context, and existing numerical-simulation tools are exploited—in parallel, if necessary—to derive detailed security information. Machine-learning methods, the heart of the framework, extract and synthesize relevant information and reformulate it in a suitable way for decision making. This involves transforming the database of case-by-case numerical simulations into a power-system security knowledge base. As Figure 1 illustrates, the framework integrates a large variety of automatic-learning methods in a data-mining toolbox, according to the type of information that these methods exploit or produce. The final step involves using the extracted synthetic information (decision trees, rules, statistical or neural network approximators) either in real-time, for fast

0885-9000/97/$10.00 © 1997 IEEE

and effective decision making, or in the offline study environment, so as to gain new physical insight and derive better system- or operation-planning strategies.

How will this automatic-learning-based framework complement classical system-theory-oriented methods (relying on analytical power-system models such as numerical simulation) for security assessment? We can expect important contributions along three dimensions: computational efficiency, interpretability, and management of uncertainties.

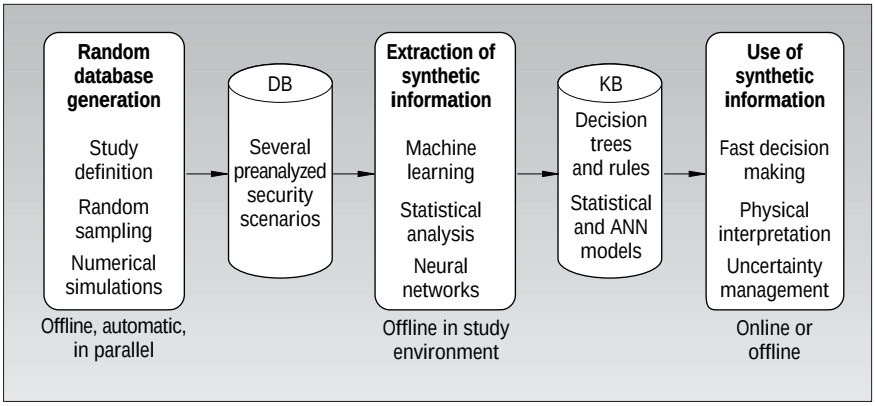

Figure 1. Machine-learning framework for security assessment.

Computational efficiency. By using synthetic information extracted by automatic learning rather than using analytical methods, the framework enables much faster real-time decision making. Moreover, regarding data requirements, analytical methods require a full description of the system model; however, the framework lets users tailor approximate models constructed through automatic learning, letting them exploit only the significant input parameters. Computational efficiency was actually the motivation of Tom Dy Liacco, when he first envisioned (in the late 1960s) the use of automatic learning (at that time, statistical pattern recognition) for real-time security assessment. Even today, despite the significant increase in computing power in the last 25 years, this remains a strong motivation.

However, the synthetic information extracted by automatic-learning methods may itself be complementary to and generally more powerful than that provided in a case-by-case fashion by existing analytical methods. In particular, power-system engineers are providing much more attention nowadays to interpretability and management of uncertainties.

Interpretability. The use of automatic learning to provide physical insight into nonlinear system behavior was first proposed by Yoh-Han Pao, Tom Dy Liacco, and Isil Bozma in the mid-1980s.[2] In the meantime, others have shown that machine learning is indeed an effective way to generate reliable and interpretable security rules from very large bodies of simulated examples,[3] even for complex, large-scale power systems. The extracted rules express explicitly problem-specific properties, as a human expert might do. Engineers in charge of security studies can more easily appraise, criticize, and eventually adopt these rules. This means the framework can also help maintain and enhance utility expertise. The machine-learning framework's flexibility lets users tailor the resulting information to analysis, sensitivity analysis, and control applications.

Management of uncertainties. The need to devise a rational way to make decisions whenever there are major uncertainties about the power-system state is becoming increasingly more apparent. Today, for example, operators are often sorely missing guidance in the context of unusual system states reached after major disturbances, where reliable real-time information is generally lacking. Tomorrow, technological and economic changes will probably lead to a higher and physically more irrational distribution of decision making, and thus to more uncertainties in routine operation and planning activities. On the one hand, new devices, such as Facts (flexible alternating-current transmission systems), may cause stronger interactions among remote components of very large interconnections. On the other hand, increased competition among economic actors may further reduce their willingness to share information on their respective subsystems, despite the stronger physical interactions. Such circumstances will create an urgent need for approaches that can manage uncertainties, such as the above framework based on automatic learning.

Applying the framework to security assessment

Despite repetitive attempts, there are still no large-scale industrial applications of the machine-learning framework to power-system security assessment. The main reason is that, until recently, the existing automatic-learning methods were not powerful enough, and the amount of possible security studies was limited by available simulation hardware and software.

Today, however, all the required conditions are met. Present-day computer networks along with fast simulation tools help generate many detailed studies. At the same time, researchers have recently made much progress in applying automatic-learning methods to large-scale power systems. Hence, automatic information-synthesis tools to help engineers compare and interpret extensive elementary results and extract and appraise useful synthetic information are strongly needed and, at the same time, technically feasible.

Therefore, while my colleagues and I expect additional progress in learning methods and application methodologies, we believe that some important electric power companies—for example, in North America or Europe—will soon start using this approach more or less routinely for security studies.

Aspects of automatic learning

Now we introduce classes of potentially useful automatic-learning methods for synthesizing security-assessment information (see the "Aspects of power-system security problems" sidebar). We first give a definition of the generic *supervised* learning problem and introduce three important classes of algorithms for this problem. Then we comment on the use of *unsupervised* learning methods.

Supervised learning problem. The generic problem of supervised learning from examples can be formulated as follows:

> Given a learning set of examples of associated input/output pairs, derive a general model for the underlying input/output relationship, which may be used to explain the observed pairs or predict output values for any new unseen input.

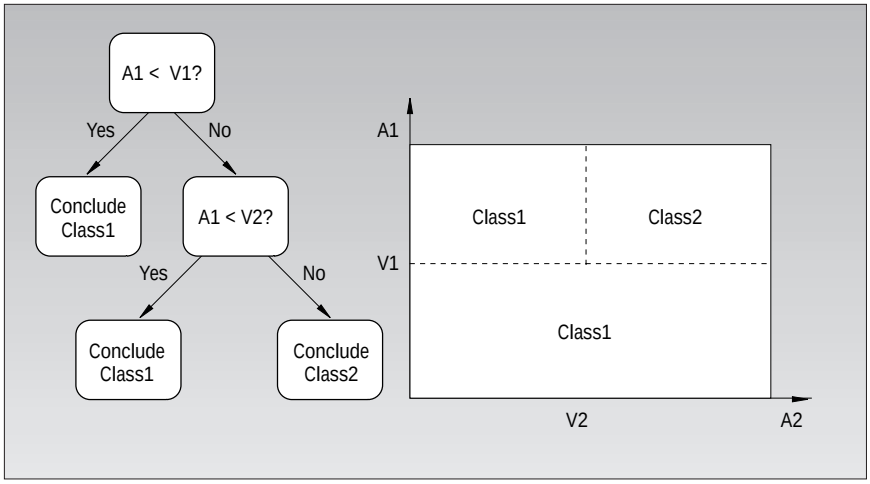

Figure 2. Hypothetical decision tree and its corresponding input-space decomposition.

For security assessment, an example corresponds to a given operating situation. The input attributes would (hopefully) be relevant parameters describing this situation's electrical state and topology, and the output could be information concerning its security in the form of either a discrete classification (for instance, secure, marginal, or insecure) or a numerical value derived from security margins.

The solution of this learning problem has several subtasks, which we now discuss.

Representation. This involves choosing appropriate input attributes to represent the power-system state, defining the output security information, and choosing a class of models suitable to represent input/output relations.

The representation problem is left to the engineer, who must compromise between using very elementary standard operating parameters and sophisticated compound features.

Feature selection. This subtask aims at reducing the input space's dimensionality by dismissing attributes that do not carry useful information to predict the considered security information. This lets us exploit the local nature of many security problems.

Model selection. Learning typically identifies, in the predefined class of models, the one that best fits the learning states. This generally requires choosing model structure and parameters, using an appropriate search technique.

The distinction between feature selection and model selection is somewhat arbitrary, and some of the methods actually solve these two problems simultaneously rather than successively.

Interpretation and validation. These are very important for understanding the physical meaning of the synthesized model and determining its range of validity. They involve testing the model on a set of unseen test examples and comparing its information with prior expertise about the security problem. From the interpretation-and-validation point of view, some supervised learning methods provide information that is difficult to interpret, while others provide explicit and very transparent models that are easy to compare with prior knowledge.

Model use. This subtask involves applying the model to predict the security of new situations on the basis of the values assumed by the input parameters and, if necessary, inverting the model to provide information on how to modify input parameters to achieve a security-enhancement goal. In using the model for fast decision making, we notice speed variations of several orders of magnitude between various techniques, but most of the methods are sufficiently fast in the context of power-system security analysis for control centers.

Supervised learning methods. Now we consider only nonparametric automatic-learning methods.[4] Parametric methods may be useful in some particular circumstances, but they are not powerful enough to treat the wide variety of practical security problems. We discuss three classes of methods providing three complementary types of information. Although we have selected them from three different paradigms (machine learning, neural nets, and pattern recognition), we insist on the type of information provided rather than on the paradigm itself.

Symbolic knowledge via machine learning. Machine learning is the subfield of AI concerned with the design of automatic procedures that can learn from examples. *Concept learning from examples* denotes the process of deriving a logical description of the necessary and sufficient conditions corresponding to a class of objects—that is, a rule in some given representation language. A major concern is finding adequate compromises between rule complexity and data fit, so as to avoid overfitting and to privilege interpretability.

Top-down induction of decision trees, popularized by J. Ross Quinlan,[5] is one of the most successful classes of such methods. Figure 2 shows a hypothetical binary decision tree: to infer the output information corresponding to given input attribute values, we traverse the tree, starting at the top node, and sequentially apply the dichotomous tests encountered to select the appropriate successor. When a terminal node is reached, the output information stored there is retrieved.

Tdidt approaches decision-tree learning in a divide-and-conquer fashion, progressively building up a decision tree, starting with the top node and ending with the terminal nodes. At each step, the algorithm considers a tip (or pending) node of the growing tree and decides whether it will be a terminal node or should be further developed. To develop a node, this method first identifies an appropriate attribute, along with a dichotomy on its values. The method then splits the subset of the learning examples corresponding to the node into two subsets corresponding to the current node's successors, according to this dichotomy. The method gives the terminal nodes appropriate information on the output values derived from learning examples—for example, the majority class label or probabilities—or expected value and standard deviation of numerical output information.

The right part of Figure 2 shows how the decision tree decomposes its input space into nonoverlapping subregions. The number of such regions should ideally be as small as possible; at the same time, the states contained by each region should belong to the same class. Thus, to build good decision trees, an algorithm must rely on appropriate *optimal-splitting* and *stop-splitting* rules.

Optimal splitting has to do with selecting a dichotomy at a test node so as to provide a maximum amount of information on the output value (that is, separate states of different classes), whereas stop splitting must identify situations where further splitting would either be useless or lead to performance degradation because of overfitting.

Aspects of power-system security problems

This sidebar is a guided tour on power-system security, for the unfamiliar reader. We focus on security problems involving *large* disturbances corresponding to nonlinear system behavior. Although such disturbances are generally very unlikely to happen, their potential consequences can be extremely important and may lead to complete system blackouts, freezing the economic activity of a whole country for many hours.

Classifying operating states

Tom Dy Liacco has defined the different operating modes of a power system. Figure A shows a more detailed description of the Dy Liacco state diagram.

Preventive security assessment raises the question of whether a system in its normal state can withstand every plausible disturbance. If it cannot, preventive control would involve moving this system state into a secure operating region. Because predicting future disturbances is difficult, preventive security assessment aims essentially at balancing the reduction of the probability of losing integrity with the economic cost of operation.

Emergency state detection aims at assessing whether the system is in the process of losing integrity, following an actual disturbance inception. This is a more deterministic evolution, where response time is critical and economic considerations become temporarily secondary. Emergency control aims at taking fast last-resort actions to avoid partial or complete service interruption.

When both preventive and emergency controls have failed to bring system parameters back within their inequality constraints, automatic local protective devices preserve power-system components operating under unacceptable conditions from undergoing irrevocable damages. This leads to further disturbances, which may result in system splitting and partial or complete blackouts. Consequently, the system enters the restorative mode, where the operator must minimize the amount of undelivered energy by resynchronizing lost generation as soon as possible and picking up the disconnected load, in the order of priority. (The main text of this article focuses only on preventive and emergency aspects.)

Physical classification of security problems

Various security problems are distinguished according to the time scales of the dynamics, the characteristic symptoms (low voltage, large angular deviations, and so on), and the control means (reactive power, switching, and so forth) to alleviate problems.

Transient stability. This concerns the ability of a power system's generators to recover synchronous operation following the electromechanical oscillations caused by a large disturbance. In this context, dynamic performance is a matter of seconds and is mainly affected by switching operations and fast power controls (such as fast valving or high-voltage DC converters), and voltage support is affected by the automatic voltage regulators of synchronous generators and static VAr compensators (SVCs). To determine the degree of stability, we evaluate a fault's critical clearing time, which is the maximum time it may take to clear the fault without the system losing its ability to maintain synchronism.

Voltage security. The fastest voltage instabilities, characterized by sudden voltage collapses, may develop at the same or even higher speeds than loss of synchronism. More common is the midterm voltage instability, which corresponds to a typical time frame of one to five minutes. In this case, voltage collapse is mainly driven by automatic transformer on-load tap changers trying to restore voltage nearby the loads. There is a third, even slower time frame, corresponding to the so-called *long-term* voltage instability, which involves a gradual buildup in load demand. This interacts with classical static security and is well within the scope of operator intervention.

Although a voltage collapse may result in wide-spread degradation of the voltage profile and subsequent loss of synchronism, it is normally initiated by a local deficiency in reactive power reserves or a reduced reactive-power transmission capability into a given load area. A load power margin, the maximum additional amount of power that may transfer safely from the generation to a given load area, may evaluate the distance to voltage insecurity.

Static security. This concerns essentially thermal overload problems of generation-transmission system components, where phenomena span over significantly longer periods of time. For example, line overloads may be tolerated for 30 to 60 minutes under favorable weather conditions.

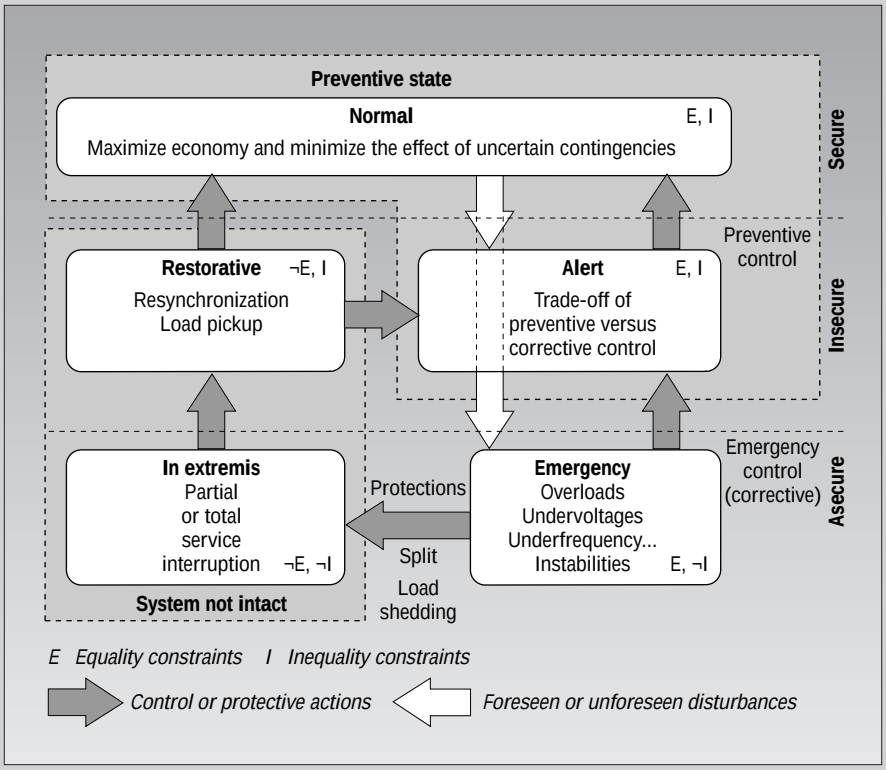

Figure A. Operating states and transitions. Adapted from Lester Fink and Kjell Carlsen.[1]

Practical application domains

Table A shows the practical study contexts or environments in security-assessment applications. The first column identifies the study context; the second specifies how long in advance

Table A. Security assessment environments (adapted from Wehenkel and Pavella[2]).

Environment	Time scale	Problems	Operator	Expert
System planning	1–10 years	Generation Transmission Protection	No	Yes
Operation planning	1 week–1 year	Maintenance Unit commitment Protection settings	No	Yes
Online operation	1 hour–1 day	Preventive mode Security assessment	Yes	Partly
Real-time* monitoring	Seconds–minutes	Emergency control Protective actions	No**	No
Training	Days–months	Improve operator skill	Yes	No

*Here we distinguish real-time monitoring, which considers dynamic situations following a disturbance inception, from mere online monitoring, which considers static predisturbance situations.
**Except for static security corrective control.

(with respect to real time) studies may be carried out; the third indicates the type of subproblems generally considered in a given environment; the fourth indicates whether an operator is involved in decision making; and the last column indicates whether an expert in the field of power-system security is available.

In the first three contexts, we rely mostly on the intervention of human experts exploiting the numerical simulation tools. In real-time monitoring and emergency control, the reduced time available calls for more automatic procedures.

System planning. Multitudinous system configurations must be screened for several load patterns, with many contingencies possible for each. An order of magnitude of 100,000 different scenarios per study would be realistic for a medium-sized system. Although enough time might be available to run so many security simulations, there is still room for improved data-analysis methods to exploit the results more effectively in order to identify structural system weaknesses and to provide guidelines in order to improve reliability.

Operation planning. As Table A suggests, operation planning concerns a broad range of problems, including maintenance scheduling (one year to one month ahead), design of operating strategies for usual and abnormal situations, and setting of protection delays and thresholds. The number of combinations of situations that must be considered for maintenance scheduling is also generally very large, and automatic-learning approaches would be equally valuable to better use the available information and exploit the system more economically.

Similarly, for a closer real-time determination of operating security criteria, machine learning is particularly well-adapted. It would let engineers screen more systematically representative samples of situations, to identify critical operating parameters and determine their security-limit tables needed for online operation. This would involve automating and enhancing the manual approaches that many utilities use.

Online operation. In the context of this framework, we would exploit online the security knowledge bases set up offline (for instance, in operation planning). Appropriate strategies are necessary to update this information when major changes happen in the system. For example, several weeks in advance, we could design routine security criteria for a forecast range of topologies, load levels, and generation schedules. Then, closer to real time (perhaps a day or a few hours before), we might refresh these criteria to handle previously unexpected situations. To be compatible with the way operators usually appraise their system, it is particularly important for the synthetic information extracted by automatic learning to be as simple as possible to interpret.

Real-time monitoring. Here, the purpose is to design criteria to more or less automatically trigger emergency-control actions, so as to prevent a disturbed system state from evolving toward a blackout. An important aspect in designing the security criteria is the use of appropriate models to reflect the disturbed power-system behavior. (Depending on the context, we use the term "model" either to denote the physical-power system model or the synthetic information extracted by automatic learning.) Furthermore, along with minimal data requirements and ultrahigh speed, readily available system measurements as inputs to the derived emergency-control rules is often an operational constraint.

Researchers have extensively studied decision trees in the context of various security-assessment problems.[6] A main asset lies in the explicit and logical representation of the induced classification rules and the resulting unique explanatory capability. In particular, the method provides systematic correlation analyses among different attributes and identifies the most discriminating attributes at each tree node. From a computational viewpoint, it is efficient at the learning stage as well as at the prediction stage.

Two generalizations of decision trees useful for security assessment are

- *regression trees*, which infer information about a numerical output variable, and
- *fuzzy trees*, which use fuzzy logic instead of standard logic to represent output information smoothly.

Both approaches let us infer information about security margins, as do the techniques discussed below.[7,8]

Smooth nonlinear approximations via artificial neural networks. The field of artificial neural networks has grown, since the early work on perceptrons, to an important and productive research field. In this article, I focus on only multilayer perceptrons.[9]

The single-layer perceptron is basically a simple linear threshold unit with an error-correcting learning algorithm. This perceptron can represent a linear boundary in its input space. Its limited representation capabilities have motivated the consideration of more complex models composed of multiple interconnected layers of perceptrons. Figure 3 illustrates the classical feed-forward MLP, compared with a single perceptron. The first, or input, layer in the MLP corresponds to the attribute values; the last, or output, layer cor-

Training. During operator training, the security criteria derived in either of the preceding contexts might be usefully exploited as guidelines, provided they're presented in an intelligible way. In addition, these models might be used internally in a training-simulator's software, to set up particular scenarios that present particular insecurity modes.

Analytical tools

Many numerical methods are available for security assessment in the different time frames mentioned. We call them *analytical* tools, because they exploit analytical power-system models, in contrast to *synthetic* tools, which are extracted by automatic-learning techniques. Some of them are based on general-purpose power-system dynamic simulation packages and have a very broad scope. Others are based on simplified models or approaches representing only those features that are relevant for the particular study. The latter methods might be significantly more efficient, although at the expense of being restricted to some particular physical phenomena or some particular (types of) power systems. We briefly discuss the analytical tools, because they provide the raw input data exploited by the automatic-learning methods to synthesize the high-level security information.

Transient stability. There are two main classes of analytical tools for transient stability assessment: a time-domain (or step-by-step) simulation approach and the direct methods, which are based on the second Lyapunov method.

Time-domain simulation. The general power-system dynamic model contains mixed algebraic and differential equations, strongly nonlinear, and typically involving a few thousand discrete or continuous time-state variables. To assess transient stability, the time-domain approach involves simulating the system's during- and post-fault behavior for a given disturbance, and observing its electromechanical angular and voltage swings during a few seconds. Practical criteria vary from one utility to another, but an unacceptable performance would generally imply too large or undamped angular deviations (for example, pole slips) or excessively large variations of voltage or frequency.

To obtain stability margins, we must run repetitive simulations for various prefault operating states or various assumptions concerning the delays of protection devices. Although this approach is still considered CPU-intensive, within the last three years the time required for a typical power-system simulation with high-order models has shrunk from one hour to a few minutes.

Direct Lyapunov-type methods. These methods identify when the system leaves its stability domain without further integration of the system trajectory. By not simulating the post-fault trajectory, they reduce the simulated time period to a fraction of a second instead of the several seconds taken by time-domain methods. Some of them can thus provide a rich stability assessment (margins, sensitivities, and mode of instability) within a fraction of the time required for a single time-domain simulation. A major drawback is their difficulty in accurately exploiting models of generators and control loops as well as nonlinear or dynamic loads. However, since the first multimachine direct methods of the late 1960s, researchers have made much progress in incorporating more realistic models.

Voltage stability and security. Tools for voltage-security assessment range from static load-flow calculations to full short-term or midterm time-domain simulations. Because of the rather recent emergence of voltage-security problems, modeling practices have not yet reached maturity, compared with those used in transient-stability studies. In particular, one intrinsic difficulty of analyzing voltage collapses is that overall system behavior depends strongly on the load behavior, for which good models are generally missing.

Short-term or mid-term dynamic simulations. Because voltage collapses may involve time constants ranging from a fraction of a second to a few minutes, a variable step-size numerical integration method with stiff-system simulation capability is preferable for the sake of efficiency and accuracy, in contrast to transient stability, where fixed step-sized methods have been widely used.

Simplified simulations. Because many voltage-security problems are essentially driven by automatic on-load tap-changer mechanisms, it is possible to sometimes neglect the faster interactions among load and generation dynamics. Equilibrium equations then replace the differential equations corresponding to the faster phenomena, and we model only the slower dynamics. With the intrinsic limitation of neglecting problems caused by the fast dynamics, this kind of approach can drastically reduce computing times.

Post-contingency load flow. A further simplification involves totally neglecting the dynamics and using only purely static postcontingency load-flow calculations. Typically, this lets us compute maximal loading limits on the basis of successive computations or even direct optimization.

References

1. L.H. Fink and K. Carlsen, "Operating under Stress and Strain," *IEEE Spectrum*, Vol 15, No. 3, Mar. 1978, pp. 48–53.
2. L. Wehenkel and M. Pavella, "Advances in Decision Trees Applied to Power System Security Assessment," *Proc. IEE Int'l Conf. Advances in Power System Control, Operation and Management*, Inst. Electrical Engineers, Hong Kong, 1993, pp. 47–53.

responds to the desired security classification or margin information. Intermediate layers let the network arbitrarily approximate complex input/output mappings, provided the network's topology and weights are chosen appropriately.

The discovery of the back-propagation algorithm has been central to the success of MLPs. This algorithm lets us efficiently and locally compute the gradient of the network's output error with respect to its weights and thresholds. We may exploit this algorithm iteratively to adjust the weights so as to reduce the total mean-square output error for learning examples. In recent years, researchers have made much progress in improving the efficiency of optimization techniques for the learning procedures of MLPs, but MLPs are still very slow at the learning stage, which may prevent extensive experimentation for database sizes typical of security assessment of realistic power systems.

As with decision trees, an interesting property of MLPs is their ability to achieve feature extraction and learning in a single step: the weights connecting the input layer with the first hidden layer essentially project the input vector in particular directions, realizing a linear transformation of the input space, which, in subsequent layers, approximates outputs. However, one of the difficulties with MLPs comes from the very high number of weights and thresholds related in a nonlinear fashion, which significantly limits any insight into the relationship learned—the input/output model corresponding to the MLP after the MLP's weights have been adapted by back-propagation. All in all, MLPs offer a function-approximation approach that is flexible and easy to apply, but difficult to interpret.

Many similar methods exist, such as radial-basis functions and projection-pursuit

regression techniques. They offer the possibility of translating the case-by-case information provided in the learning sets into an approximate but closed-form numerical model. The latter might be useful for quickly assessing unseen situations and directly computing sensitivities.

Memory-based reasoning via statistical pattern recognition. The previous two approaches essentially compress detailed information about individual simulation results into general, global security characterizations. We can provide additional information, however, in a case-by-case fashion, by matching an unseen (for example, real-time) situation with similar situations found in the database.[10] We do this by defining generalized distances so as to evaluate similarities among power-system situations, along with appropriate fast database-search algorithms.

One such technique is the well-known *K nearest neighbors* (K-NN) method, which can complete decision trees and multilayer perceptrons. The method classifies a state into the majority class among its *K* nearest neighbors in the learning set. This method's main characteristics are high simplicity yet sensitivity to the type of distances used. In particular, to be practical, ad hoc algorithms are needed to choose the distances on the basis of the learning set. Although in the past this generally involved exploiting a few sophisticated ad hoc input features manually selected on the basis of engineering judgment, today the emphasis is more on the research of automatic distance design methods exploiting the learning states.

Clustering and unsupervised learning. In contrast to supervised learning, where the objective is clearly defined in terms of modeling the underlying correlations between some input variables and some particular output variables, unsupervised learning methods are not oriented toward a particular prediction task. Rather, they try to identify existing underlying relationships among a set of objects characterized by a set of variables or among a set of variables used to characterize a set of objects.

Thus, one of the purposes of clustering is to identify homogeneous groups of similar objects, to represent many objects by a few representative prototypes. Graphical, 2D scatter plots may help in analyzing the data and identifying clusters. Another application of the same techniques is to identify similarities (and redundancies) among the different attributes used to characterize objects. In the context of power-system security, both applications may be useful as complementary data-analysis and preprocessing tools.

Researchers have proposed unsupervised learning methods under the three umbrellas given above to classify supervised learning methods, termed *cluster analysis* in the statistics literature, *conceptual clustering* in the machine-learning community, and *self-organizing maps* or *vector quantization* in the neural net community.

Applying automatic learning to power-system security

Here, we describe a hypothetical application of the automatic learning-based framework to a hypothetical security problem. Then we provide a short overview of some real-life applications to large-scale security problems.

A hypothetical illustration of the framework. The machine-learning framework is flexible enough to be applied to a large variety of security assessment problems, ranging from system planning to the design of special protection schemes. In this section, I will describe its application to voltage security assessment, which is one of the areas where we can expect systematic, real-life applications in the near future.

A security problem. Let us imagine that in our hypothetical power system voltage security is limited in some reactive power-weak area. Let us also suppose that this security problem was discovered in a preliminary screening security study, where possibly constraining disturbances were also identified.

A practical problem would be to characterize security regions with respect to these disturbances, so as to provide operators with preventive security-assessment criteria and effective preventive control to alleviate potential insecurities, such as optimal rescheduling of available reactive power resources.

Another, different problem would be the design of emergency-state indicators to be applied in case of a disturbance. Ideally, these indicators would be highly anticipative and reliable, and would provide information on appropriate emergency-control means such as on-load tap-changer blocking and load shedding.

How can we generate a database? To provide a representative sample of voltage-security scenarios for the above problems, we would first talk with planning and operation-planning engineers and system operators to gather information about known system weaknesses and operating practices.

From this information, we would design database-building software to generate randomized samples representative of normal operating conditions, including a sufficient number of unusual situations, deemed relevant for security characterization. In particular, with respect to real-life operating statistics, this sample would typically be biased toward the insecure regions of the state space.

According to that sampling procedure, we would generate an initial database, typically comprising several thousand states. The security of each state would be preanalyzed with respect to the studied disturbances. For example, we could compute post-contingency load-power margins for real large-scale power-system models on existing computer networks within some hours of response time, by using efficient simulation software and exploiting trivial parallelism. In addition to this information, we could predetermine appropriate preventive or emergency-control information for the insecure states, and secure economic-generation dispatch for the secure ones.

Furthermore, we would compute a certain number of attributes, which we would propose as input variables to formulate security criteria. In the preventive-mode security-assessment problem, these attributes would typically be *contingency-independent* prefault operating parameters, such as voltages, reactive power generation and compensation reserves, power flows, and topology indicators. For the emergency-state-detection problem, we would rather use raw system measurements (for example, voltage magnitudes, power flows, transformer ratios, or breaker status) of the intermediate state that immediately follows the disturbance. In contrast to the preventive-mode attributes, the emergency-state attributes would depend on the disturbance and on the short-term system modeling, in addition to the prefault operating state. Care must be taken to appropriately take into account uncertainties about this information by adding random-noise terms wherever necessary (load distribution and sensitivity to voltage, external systems, measurement noise, and delays).

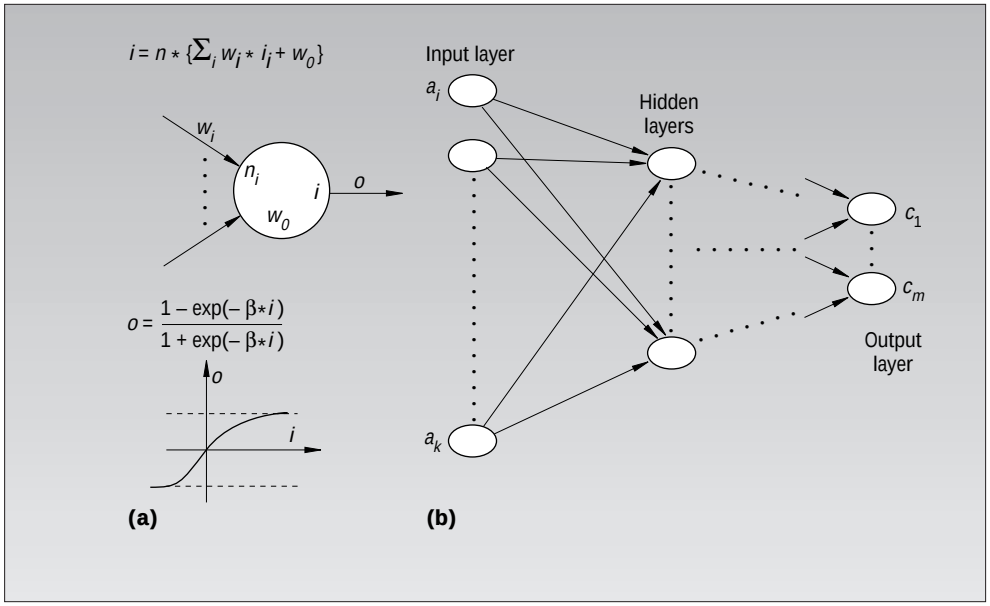

Figure 3. Comparison of single- and multilayer perceptron models: (a) single-layer perceptron; (b) feed-forward multilayer perceptron.

Unsupervised learning for data preprocessing. In practical security problems, many different attributes provide equivalent information because of the very strong physical correlations among a power system's geographically close components. Thus, clustering methods can help define a few representative attributes from a larger number of elementary variables.

For example, let us consider the case of voltage magnitudes. We can easily compute correlation coefficients among any pair of bus voltages on the basis of the database statistical sample. A clustering algorithm searching for a reduced number of voltage *coherent* regions can then use them as similarity measures. For each region, we'd use an equivalent (for example, mean) voltage as an attribute instead of individual bus voltages, and we would reduce the computational burden of the subsequent supervised learning of security criteria yet improve robustness and interpretability. For example, we can exploit 2D Kohonen feature maps to visualize the relationships among voltage regions and easily compare them with the geographic location of busbars in the power system.

In addition to the above feature-extraction application, researchers have proposed more conventional clustering techniques that identify groups of similar power-system operating states. One possible purpose is to partition a very large database into smaller subsets for which the security-assessment problem could be easier to solve. Another interesting application would be to condense the full database into a reduced number of representative prototypes, thereby decreasing the number of required security simulations and shortening the associated computation delays.

Supervised learning of security criteria. Given a database of examples, with security margins determined for several contingencies and several candidate attributes computed, supervised learning would derive appropriate security criteria. First, though, we would partition the database into disjoint learning and test samples. We'd use the learning sample to build the synthetic security criteria. We'd use the test states to assess the security criteria's reliability, by comparing the security information predicted by them with the "real" information determined by the simulation. Along with the unseen test states generated automatically with the learning states, a test sample that is representative of actual operating statistics from historical online records should be collected.

We must first define security classes by setting appropriate thresholds on the security margin. Then, the decision-tree building includes

- the automatic identification of the subset of attributes among the candidate ones relevant for the prediction of the security class (say ten to twenty among one or two hundred), and
- the definition of appropriate threshold values for these attributes so as to provide an approximate model of the voltage-security region of the area of the power system studied.

In addition to a global DT covering all disturbances simultaneously, single-contingency DTs can also be constructed to provide more specific information and additional insight. Furthermore, we can construct DTs for various security-margin threshold values, to discriminate between marginally secure and very secure situations. Depending on whether normal predisturbance or only after-disturbance attribute values are used, we can use the DTs in either a preventive or an emergency-wise approach. If there are too many nondetections of insecure states, then before rebuilding a tree we can increase the threshold value used to define the secure class in terms of the security margin. If there are too many false alarms, we should use additional candidate attributes or learning states.

The DTs provide a simplified view of security in terms of a discrete model relating a few security classes and thresholds on attribute values. We might also wish to provide a continuous security margin—at least, in the neighborhood of the threshold values used to define security classes. As I've mentioned, one of the strong points of the MLP is its nonlinear modeling capability. On the other hand, the decision tree identifies the attributes in strong correlation with the security class. Thus, in a hybrid approach, we might use the latter attributes as input variables to an MLP model, and a normalized security margin as output information.

In practice, we might need to proceed by trial and error to determine an appropriate number of hidden neurons and topology for the MLP structure. Once its structure and weights have been adapted on the basis of the learning states, the MLP provides a closed-form and differentiable security approximator, which we can use for fast margin prediction for any seen or unseen state and to compute margin sensitivities to attribute values. Practical experiments with various security problems show that this leads to richer and more reliable security-assessment information.

With the previous two approaches, we essentially compressed detailed information about individual simulation results into general global-security characterizations. This let us provide the required physical understanding, thanks to the data-analysis component of decision trees and attribute-clustering tech-

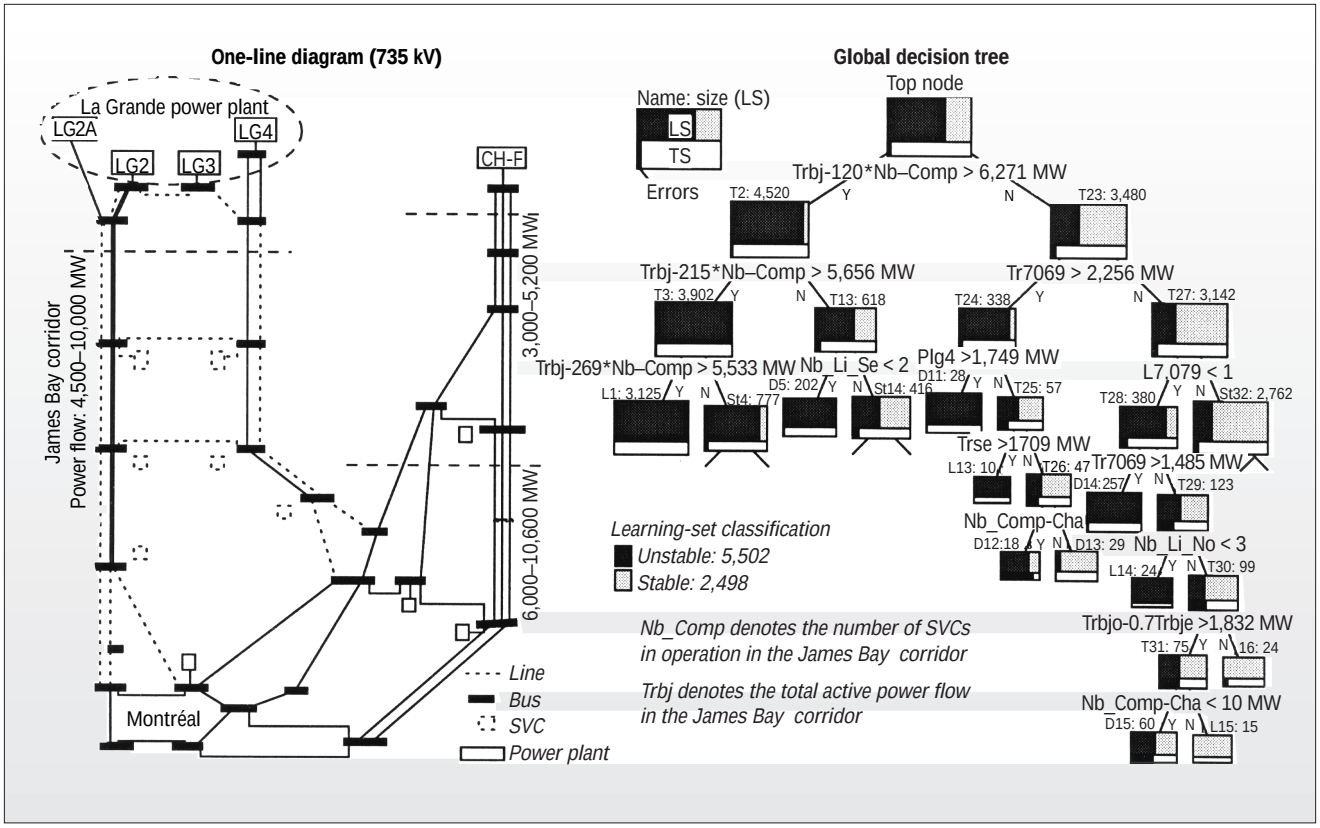

Figure 4. Transient stability assessment of the Hydro-Québec system.

niques. In addition, we can efficiently use the derived models for online security analysis.

In this latter context, we can get further information via memory-based reasoning, exploiting appropriate distances to find the most similar preanalyzed situations to the real-time state. Once identified, we can use these situations in many ways. For example, their distance to the current state would provide a measure of confidence of the security information provided by any model derived from the database (DT or MLP). If the latter were too large, we would then conclude that, for the current state, no reliable security information may be derived from the database. If, on the contrary, the nearest neighbors were sufficiently close to the current state, then we could extrapolate various kinds of detailed and specific security information from these states to the current situation and show them to the operator, along with a detailed contingency analysis and preventive or emergency controls.

Overview of some real-life applications. Below, we provide more specific information about feasibility studies of the automatic-learning approach, made for actual power-system security problems.

Transient stability. Together with Electricité de France (EDF), we initiated a first large-scale feasibility study in early 1990, for preventive transient stability assessment of an important generation plant within the large-scale extra-high-voltage (EHV) system of EDF.[11] (EHV is 225 kV to 400 kV; HV (high voltage) is 63 kV to 90 kV; MV (medium voltage) is less than or equal to 20 kV) More recently, a detailed study was carried out on the Hydro-Québec system.

The 735-kV system of Hydro-Québec is illustrated in Figure 4. Its normal operating condition is considered secure if it withstands any permanent single-phase-to-ground fault, followed by line tripping, fast reclosure, and subsequent permanent tripping. This system is mainly constrained by its transient stability limits, caused by the very large power flows and long transmission distances.

More specifically, in our investigations, we considered only faults occurring within the James Bay transmission corridor in the Western part of the system. With respect to such faults, the stability is mainly influenced by the power flows and topology within the same corridor. A manual approach had previously developed transient stability limits; operation-planning engineers determined offline, on the basis of carefully chosen simulation scenarios, approximate limit tables relating the system topology and power flows to a stable/unstable classification. Hydro-Québec implemented the limit tables on the real-time computer via an ad hoc database tool called Limsel. The purpose of our investigation was to evaluate the automatic-learning approach's ability to provide a more systematic and potentially more efficient methodology to derive these operating guidelines.

We generated a database of 12,500 normal operating states via random sampling and chained load-flow computations; it comprises more than 300 different combinations of up to six line outages, about 700 different combinations of reactive voltage-support equipment in operation, and a wide variety of power-flow distributions. The dashed lines in Figure 4 show the variable-topology part of the 735-kV system. For each state, we obtained the corresponding classification, stable/unstable, from Limsel, running on the backup online computer. The results were 3,939 stable states and 8,561 unstable ones, among which 393 were marginally unstable and 8,168 were fairly unstable.

To describe the operating states and characterize their stability, we computed the following types of candidate attributes: active power flows through important lines and cut sets in the James Bay corridor; total active power generated in the four La Grande (LG) power plants and various combinations; sev-

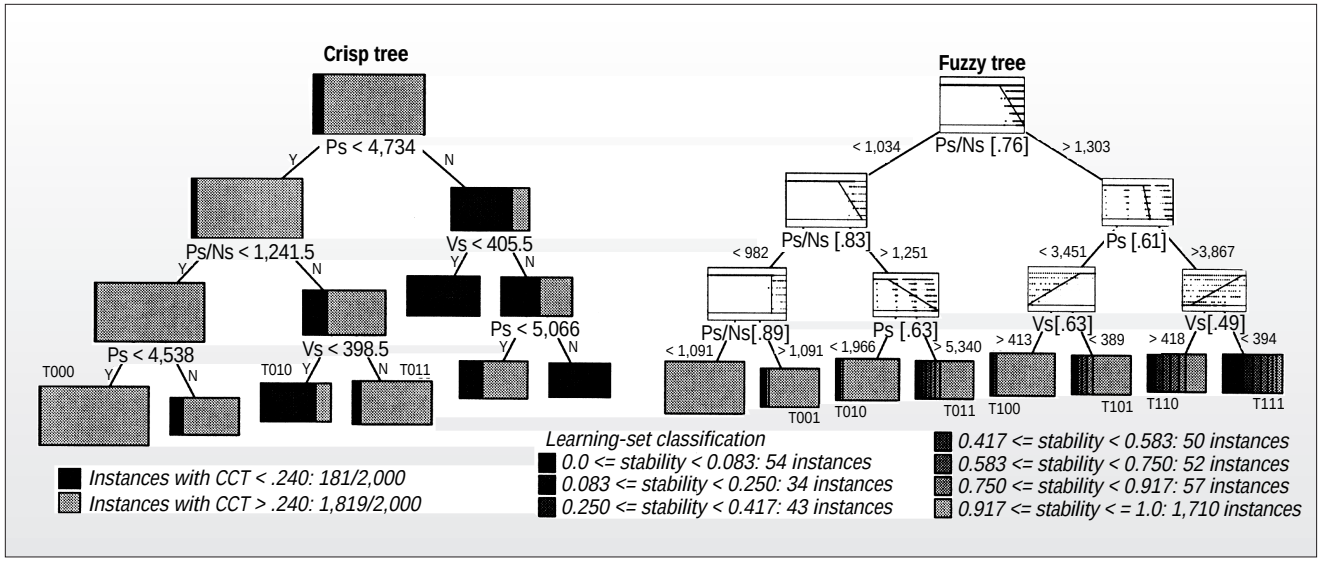

Figure 5. Crisp decision trees and fuzzy decision trees.

eral SVCs in operation within the six substations in the James Bay corridor; and logical indicators (in/out) for important lines. We determined this set, composed of 67 candidate attributes, with the help of an expert in charge of transient-stability studies at Hydro-Québec. From previous studies, we already knew that, along with the topological information and the total number of SVCs, the total power flow through the corridor would be an important attribute.

We built the tree that is partially represented in the right-hand part of Figure 4 on the basis of the database's first 10,000 states (8,000 to grow the DT and 2,000 to determine its optimal degree of pruning) and 87 candidate attributes—including, in addition to the above 67 ones, four linear-combination attributes and some other combined ones. Figure 4 shows its most important parts near the top node. The notation used for a typical node is also represented at the top left-hand side of the tree: each node is represented by a box, the upper part of which corresponds to the proportions of stable and unstable learning states relative to this node. Test nodes are identified by the label T*i* or S*ti*, the latter corresponding to subtrees that have not been drawn on the picture. Terminal nodes are identified by a label L*i* for leafs and D*i* for dead ends. A leaf is a terminal node that contains a strong enough majority of learning states of a single class (the algorithm expressed this in terms of an entropy measure), whereas a dead end is a node that corresponds to a subtree pruned to avoid overfitting. The label indicates the type of a node, and the node's number of learning states is indicated next to it.

All in all, the tree comprises 57 test nodes and 58 terminal ones. This tree has identified among the candidate attributes the 24 most relevant ones. Among others, at several test nodes (including the top node), the algorithm has selected a linear combination of the total power flow Trbj in the James Bay corridor and the number of SVCs in operation, `Nb_Comp`, which thus confirms prior knowledge. Thus, the threshold values of Trbj are functions of `Nb_Comp`. For example, if `Nb_Comp` = 12, then the leftmost terminal node L1 in Figure 4 corresponds to a limit value of max{6,271 + (12 ∗ 120), 5,656 + (12 ∗ 215), 5,533 + (12 ∗ 269)} = 8,761 MW. Above this value, the tree declares a state unconditionally unstable for at least one line fault in the corridor.

To evaluate the tree's generalization capability, we tested it with an independent test set comprising the 2,500 states of the database not used for growing or pruning the tree, yielding an overall error rate of 4.3%. (The lower part of each node box in Figure 4 depicts the proportion of erroneous classifications of each subtree's test states. Of the 1,622 fairly unstable states, the tree classified only 30 as stable, yielding 1.85% dangerous errors. On the other hand, the tree classified 23 marginally unstable states as stable, leading to small nondetection errors. There were also 52 false alarms—that is, stable test states that the tree classified as unstable.

To improve accuracy, we exploited the same database further by building a multilayer perceptron (with a single hidden layer of 20 neurons) on the basis of the same 10,000 learning states, leading to a reduced test-set error rate of 2.4%. Computational requirements (in CPU time), determined on a Sun Sparc10 workstation, are

- about 1 week to generate the database generation;
- 1 hour to build the decision tree, and 1 second for testing; and
- 60 hours to learn the MLP weights, and 10 seconds for testing.

Researchers have also investigated decomposing the database into various topology classes and obtaining simpler and more interpretable trees.[12,13]

Fuzzy decision trees. To illustrate the potential of fuzzy decision trees in the context of security assessment, let us consider a simplified problem derived from the transient stability study carried out on the EDF system.[11] Here, we measure a fault's degree of stability, using its critical clearing time. Thus, we define stability classes for crisp trees, using thresholds on the CCT.

The left side in Figure 5 gives a partial view of such a crisp decision tree; its right side shows a corresponding fuzzy tree. The former was built for a classification threshold of 0.240 seconds. The fuzzy tree was built on the basis of a fuzzy classification: the stability degree of a state varies continuously from 0 to 1 as its CCT increases from 0.215 to 0.265 seconds. The crisp tree uses attribute thresholds to propagate a state either to right or left successors, and the fuzzy tree uses transition regions defined by two thresholds. Outside the transition region, a state propagates only to one successor, but inside it go both successors, with a weight varying progressively as a function of the attribute value and the thresholds.

The fuzzy-tree growing algorithm is very similar to the crisp Tdidt method.[8] It auto-

matically determines the transition regions at each test node, recursively partitioning the learning set, though in a fuzzy way. We designed a pruning technique similar to those used for crisp trees, to keep the tree complexity minimal. In the above example, this algorithm let us significantly improve accuracy by reducing classification-error rates from 3.3% to 1.3%. At the same time, the algorithm provided more refined information about the system's stability.

Thus, fuzzy trees can express continuously varying degrees of security in a very natural and effective way, as with smooth regression techniques. At the same time, fuzzy trees provide easily interpretable information, as do symbolic machine-learning techniques. Some research is still needed to improve the computational performances of the fuzzy-tree growing and pruning algorithms, and to further validate them on different test problems. However, this is a very promising technique—particularly in the context of security assessment, where the output information often varies continuously with input attributes.

Voltage security. We carried out a second, rather extensive feasibility study for voltage security, on a test problem concerning the Brittany region of the EDF system. We considered both preventive security assessment and emergency-state detection.[7] The left side of Figure 6 shows the one-line diagram of the related part of the EDF system. Its subregions correspond to voltage-coherent load areas, determined with respect to the behavior of HV voltage magnitudes just after the loss of a generator in Plant 1. These regions were automatically determined in a preliminary study by unsupervised learning, using a Kohonen feature map.

The independent variables used during the random sampling of the predisturbance states concerned the following: topology (single or double line or transformer outages); regional load level, unit commitment, and generation dispatch; reactive support by synchronous condensers; and gas turbines. To account for uncertainties, we randomized the following quantities: secondary voltage-control set points, individual HV load-distribution and power factors, MV shunt compensation, and voltage sensitivities of active and reactive load powers.

The sampling drew a total of 13,513 random variants, yielding 5,000 predisturbance states. (The remaining 8,513 variants led to power-flow computation divergence or nonconvergence.) For each state, we computed about 200 attributes, corresponding to key variables such as topological indicators, important EHV power flows, 400-kV voltages, numbers of units in operation in power plants, total load demand, reactive shunt compensation reserves in the study region, and reactive generation reserves.

All in all, this broad study considered 26 different contingencies, corresponding to a synchronous condenser, a generator or line tripping, and busbar faults.[7] The difference between pre- and postdisturbance load-power margins in the Brittany region determined a disturbance's severity. Thus, besides computing the predisturbance margin, we also computed the corresponding 26 post-disturbance margins for each operating state, yielding a total of 135,000 load-power margin computations. Overall, the database generation required about one month of CPU time on a Sun Sparc10 workstation.

We built several tens of multilayer perceptrons and even more decision or regression trees, for different disturbances and both preventive security assessment and emergency-state detection. In addition, we also tried out various nearest-neighbor classifiers. To illustrate, let's look at the regression tree depicted on the right-hand side in Figure 6, built to estimate the severity of the loss of Circuit 1 of an important 400-kV line (see the one-line diagram in the figure). A box represents each node of the tree. The box graphically represents the contingency severity's distribution of values in this node's learning set, along with its sample mean value and standard deviation, and the number of its learning states. At the top node, N = 2,775 corresponds to the total number of learning states used to build the tree.

The total predisturbance reactive reserve available in Plant 2 is automatically selected as the best test attribute at the top node, with a threshold of 191 MVAr. The learning set splits into two subsets, corresponding to 1,219 and 1,556 states. This reduces the variance from $106^2 = 11{,}236$ at the top node to a mean value of $(1{,}219/2{,}775)67^2 + (1{,}556/2{,}775)116^2 = 9{,}517$ at its successors.

Proceeding to both successors, we see that the selected test consists of checking whether Circuit 2 is in operation, which lets us further reduce the overall variance to a mean value of $(1{,}146/2{,}775)21^2 + (73/2{,}775)80^2 + (1{,}464/2{,}775)38^2 + (92/2{,}775)155^2 = 1{,}817$. Thus, the regression tree explains $100 * (1 - (1{,}817/11{,}236)) = 84\%$ of the severity's variance.

Once we've constructed the tree, we can use it to estimate an unknown state's contingency severity. The algorithm directs the state from the top node to the appropriate successor according to the state's reactive reserve, and then to a terminal node according to the status of Circuit 2. There, the mean severity of the corresponding learning states is an estimate of the severity.

This very simple tree accurately estimates the disturbance's severity. Admittedly, we might improve it by further developing some of its terminal nodes, using other attributes carrying complementary information. However, when we apply this simple model to a representative independent test sample, the difference between its estimate and the actual precomputed severity yields an overall mean error of –0.5 MW and standard deviation of 43.6 MW. This is, indeed, almost negligible compared to the study region's overall load level, which varies between 5,000 MW and 7,700 MW.

Probabilistic global dynamic-security assessment. In the summer of 1995, another long-term research collaboration started, through the initiative of Electricité de France, with the objective of developing a probabilistic method to globally evaluate power-system failure modes. This included assessing their probability, their actual consequences, and their prevention. The approach proceeded in the following way:

- Set up a detailed probabilistic model of the possible causes of insecurity: multiple disturbances, bad coordination or misuse of protective devices, or over-optimistic preventive-security strategies caused by uncertainties in modeling parameters.
- Sample representative combinations of these causes, and do extensive simulations to determine the effect on the power system.
- Analyze the database of dynamic simulation results to identify a posteriori the system's main weaknesses.
- Evaluate the most effective countermeasures (for example, in the form of new system protections), and validate them through a cost/benefit analysis on the scenarios stored in the database.

The ongoing research project has reached

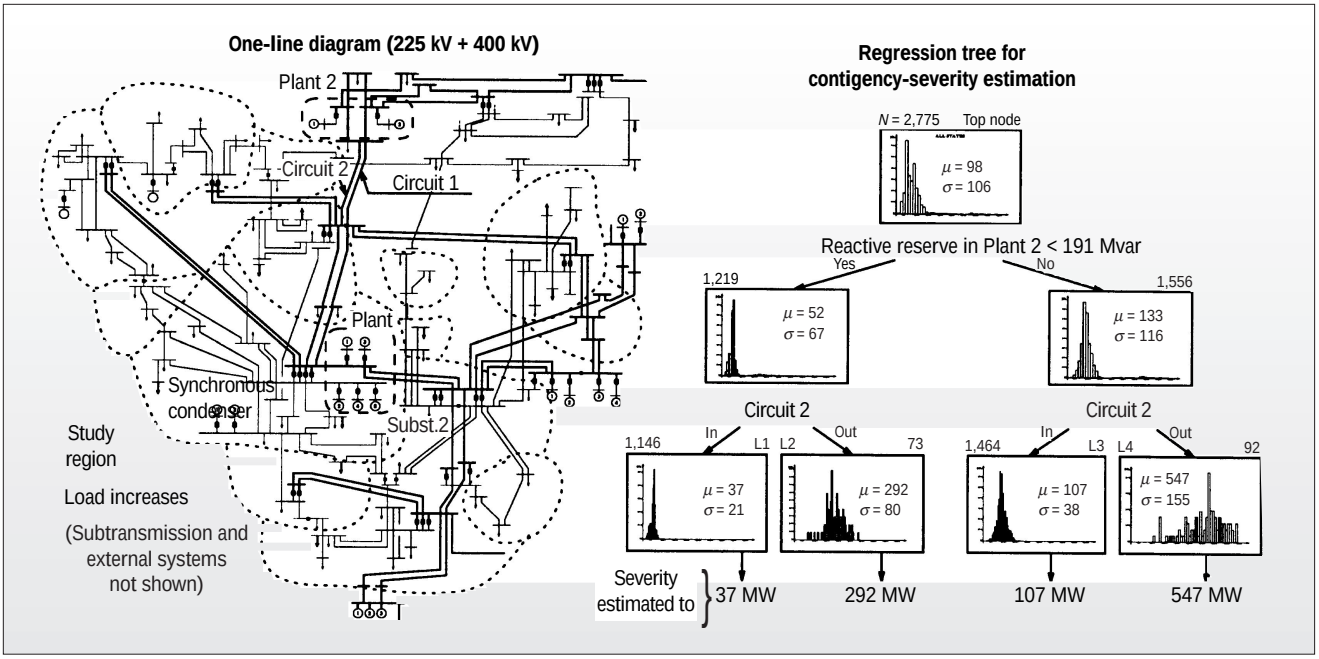

Figure 6. Voltage-security assessment of the Brittany system.

the following status: We have identified main plausible causes of security problems in the system, along with their relative probability distributions for random sampling (study definition). We have designed a global, dynamic simulation model to simulate fast (transient) phenomena as well as slower dynamics (up to 40 minutes); this model (11,000 state variables) of the French system comprises

- an EHV (400 kV and 225 kV) transmission-system model with 1,550 branches and 1,150 buses, as well as line overload and busbar low-voltage protections;
- 196 machines with detailed generator models and secondary voltage control, over-and-under frequency, and over-and-under voltage protections, as well as local loss-of-synchronism line-tripping relays and a coordinated defense plan; and
- voltage-sensitive load models, along with 300 (EHV/HV/MV) automatic under-load tap-changing transformers, an under-voltage tap-changer blocking device, and under-frequency load-shedding.

We use a variable time-step integration method (Eurostag software), which requires about 10,000 steps per scenario. Thus, each scenario simulation requires about eight hours on a high-end workstation and generates about 500 Mbytes of raw output data. From the latter, we've extracted a subset of about 500 relevant system variables (EHV voltages and power flows, main machine variables, regional load behavior, and so on), yielding only 1 Mbyte per scenario of compressed information stored in the database. We have developed a parallel database-generation module (random sampling and numerical simulations in parallel on a cluster of workstations). We have simulated 1,500 scenarios (about 1,500 Mbytes of data), and we've conducted some preliminary analyses, which confirm the software's good performance.

To take full advantage of this type of database, we need to meet some technical challenges. In particular, in contrast to the issues discussed previously, where scalar attributes represented information about operating points (a system snapshot in a given state, we need to represent and manipulate temporal information—that is, attributes varying with time along the system-dynamic trajectory. Thus, database sizes are multiplied by a factor larger than 100, and the data-mining software must be scaled up to handle such volumes. The second aspect relates to the automatic-learning algorithms themselves, which we must enhance to properly cope with the problem's temporal nature. In the context of supervised learning, this is still an open problem.

To conclude, the above approach is cause-driven, which is in strong contrast with the usual deterministic practice. In that practice, the overall security problem is a priori decomposed into subproblems corresponding to the different expected possible consequences. The decomposition generally corresponds to different phenomena, such as transient, midterm, and long-term dynamic instabilities, or different weak geographical areas. The subproblems are essentially studied independently of one another. Therefore, in this approach, a problem is considered only when it is already known to be a "true" problem, which presupposes that the main weaknesses of the system are known a priori. Hence, it is no wonder that the history of power-system blackouts includes many unexpected problems.

By more systematically exploiting available computing power and developing more sophisticated automatic-learning algorithms, this new approach can hopefully provide earlier warnings when a new problem arises. We also hope that it will allow a more objective arbitration and a better coordination of countermeasures to competing problems.

WE BELIEVE OUR APPROACH may have many applications in engineering and other complex, large scale, nonlinear systems. Simply stated, it exploits appropriate simulation models in parallel to screen a diversity of simulation scenarios of a system, yielding a large database of detailed infor-

mation. Then we apply data-mining techniques to these scenarios to extract synthetic information about the simulated system's main features, from various complementary viewpoints.

Nevertheless, for automatic-learning methods to be successful, we need a human expert to help derive security information. For example, to guide the security studies, we must exploit his prior expertise and let him criticize, assimilate, or accept the new information. Therefore, we need to provide the results in a form compatible with his own way of thinking. In the general class of automatic-learning approaches, machine learning is the only one that can meet this requirement; therefore, it is a key element in the data-mining toolbox.

However, machine learning, as well as other learning methods, can produce interesting security information only when it exploits representative databases. The initial investment required to obtain these databases is very important for each new security problem, but the subsequent database generations take full advantage of the previous ones. To further enhance the approach, we are developing powerful parallel simulation environments to transparently allocate simulations on virtual machines composed of several elementary workstations, available through local- or wide-area networks.

After 10 years of research, we conclude that automatic-learning methods can indeed provide interesting security information for various physical problems and practical contexts. Actually, in their way of approaching problems, they are quite similar to existing practices in power-system security studies, where limits are derived from simulations—albeit manually. But automatic-learning approaches are more systematic and easier to handle and master—in short, they are more reliable and more powerful.

These possibilities open up new ways for power-system engineers to respond to the challenge of planning and operating future power systems with an acceptable level of security despite the growing levels of complexity and uncertainty and the increasing economical and environmental pressures.

Acknowledgments

Several collaborations with Electricité de France and Hydro-Québec have supported our research, and we are pleased to acknowledge the valuable discussions with and suggestions by the engineers of the R&D division of Electricité de France and the Operations Department of Hydro-Québec.

References

1. L. Wehenkel, *Automatic Learning Techniques in Power Systems*, Kluwer Academic Publishers, Boston, to be published in 1997.
2. Y.H. Pao, T.E. Dy Liacco, and I. Bozma, "Acquiring a Qualitative Understanding of System Behavior through AI Inductive Inference," *Proc. IFAC Symp. Electric Energy Systems*, 1985, pp. 35–41.
3. L. Wehenkel and M. Pavella, "Decision Tree Approach to Power System Security Assessment," *Int'l J. Electrical Power and Energy Systems*, Vol. 15, No. 1, Feb. 1993, pp. 13–36.
4. S.M. Weiss and C.A. Kulikowski, *Computer Systems That Learn*, Morgan Kaufmann, San Francisco, 1991.
5. J.R. Quinlan, "Learning Efficient Classification Procedures and Their Application to Chess Endgames," in *Machine Learning: An Artificial Intelligence Approach*, R.S. Michalski, J. Carbonell, and T. Mitchell, eds., Morgan Kaufmann, 1983, pp. 463–482.
6. L. Wehenkel and M. Pavella, "Advances in Decision Trees Applied to Power System Security Assessment," *Proc. IEE Int'l Conf. Advances in Power System Control, Operation and Management*, Inst. of Electrical Engineers, Hong Kong, 1993, pp. 47–53.
7. L. Wehenkel, "Contingency Severity Assessment for Voltage Security Using Nonparametric Regression Techniques," *IEEE Trans. Power Systems*, Vol. PWRS-11, No. 1, Feb. 1996, pp. 101–111.
8. X. Boyen and L. Wehenkel, "Fuzzy Decision Tree Induction for Power System Security Assessment," *Proc. Second IFAC Symp. Control of Power Plants and Power Systems*, Inst. de Investigaciones Electricas, Mexico City, Dec. 1995, pp. 151–156.
9. S. Haykin, *Neural Networks: A Comprehensive Foundation*, IEEE Press, Piscataway, N.J., 1994.
10. R.O. Duda and P.E. Hart, *Pattern Classification and Scene Analysis*, John Wiley & Sons, New York, 1973.
11. L. Wehenkel et. al, "Decision Tree Based Transient Stability Method: A Case Study," *IEEE Trans. Power Systems*, Vol. PWRS-9, No. 1, Feb. 1994, pp. 459–469.
12. L. Wehenkel et. al, "Automatic Learning Approaches for Online Transient Stability Preventive Control of the Hydro-Québec System—Part I: Decision Tree Approaches," *Proc. Second IFAC Symp. Control of Power Plants and Power Systems*, 1985, pp. 231–236.
13. L. Wehenkel, I. Houben, and M. Pavella, "Automatic Learning Approaches for Online Transient Stability Preventive Control of the Hydro-Québec System—Part II: A Toolbox Combining Decision Trees with Neural Nets and Nearest Neighbor Classifiers Optimized by Genetic Algorithms," *Proc. Second IFAC Symp. Control of Power Plants and Power Systems*, 1995, pp. 237–242.

AUTOMATIC-LEARNING APPROACHES ARE MORE SYSTEMATIC AND EASIER TO HANDLE AND MASTER—IN SHORT, THEY ARE MORE RELIABLE AND MORE POWERFUL.

Louis Wehenkel is a research associate of the Belgian National Fund for Scientific Research at the University of Liège, Belgium, where he also teaches a course on applied automatic learning. His research interests include developing AI and probabilistic methodologies for electric power-system planning, operation, and control. He received an electrical (electronics) engineering degree, a PhD in electrical engineering, and the Agrégation de l'Enseignement Supérieur, all from the University of Liège, Belgium. He is a member of the IEEE. Readers can contact Wehenkel at the Dept. of Electrical Eng., Univ. of Liège, Institut Montefiore, Sart-Tilman B28, Liège B-4000, Belgium; lwh@montefiore.ulg.ac.be; http://www.montefiore.ulg.ac.be/~lwh/.