

Electric Power System Anomaly Detection Using Neural Networks

Marco Martinelli¹, Enrico Tronci¹, Giovanni Dipoppa², and Claudio Balducelli²

¹ Dip. di Informatica Università di Roma “La Sapienza”
Via Salaria 113, 00198 Roma, Italy
marco.martinelli@quipo.it tronci@dsi.uniroma1.it

² ENEA - Centro Ricerche Casaccia
Via Anguillarese 301, 00060 Roma, Italy
{giovanni.dipoppa, c.balducelli}@casaccia.enea.it

Abstract. The aim of this work is to propose an approach to monitor and protect Electric Power System by learning normal system behaviour at substations level, and raising an alarm signal when an abnormal status is detected; the problem is addressed by the use of autoassociative neural networks, reading substation measures. Experimental results show that, through the proposed approach, neural networks can be used to learn parameters underlying system behaviour, and their output processed to detecting anomalies due to hijacking of measures, changes in the power network topology (i.e. transmission lines breaking) and unexpected power demand trend.

1 Introduction

Monitoring and protecting *Large Complex Critical Infrastructures* (LCCIs) is becoming more and more important, as the growth of structures interdependencies, and their increasing complexity make them vulnerable to failures or to deliberate attacks.

Our goal is to detect anomalies in the dynamics (i.e. evolution over time) of the measure vectors coming from the substations of an Electric Power System. In this paper, a neural network based approach for novelty detection is presented, on the same lines proposed by Thompson et al. [1], but in a different setting. The use of autoassociative neural networks is aimed at learning *normal behaviour* of a LCCIs subcomponents, for a low level, distributed monitoring approach: dangerous attack or accidental fault within the system would probably bring significant deviations at this level, thus causing *novelty* detection.

Research, implementation and test have been developed in the operative framework of the *Safeguard*¹ project.

¹ *Safeguard* is a European project investigating new ways of protecting Large Complex Critical Infrastructures, developed by ENEA, QMUL, AIA, LiU, and Swisscom. For further details refer to [5] .

2 Basics

2.1 Electric Power Systems

An *Electric Power System* (EPS) can be seen as a set of nodes, called *substations*, connected each other by transmission lines. Each substation, usually monitored by a *Remote Terminal Unit* (RTU), is composed by several components, each playing a specific role in the power generation/consuming process.

Electric power is generated by generators, distributed through transmission lines, consumed by loads, which demand may usually vary hourly, weekly and monthly.

2.2 Artificial Neural Networks

An *Artificial Neural Network* (ANN) is built out from simple, non-intelligent units (neurons) which are connected together, becoming able to perform complex signal processing.

In the learning phase, an ANN is presented with input data set and is trained to fire out the desired values at output layer. The training algorithm iteratively modify weights on connections through which signals are transmitted, in order to minimize gap between network output and desired one.

The Autoencoder Model - An *Autoassociative Neural Network Encoder* (or simply *autoencoder*) has two primary features:

- *Autoassociative Feature*: the network is trained to reproduce at output layer same values presented as input. For this reason input and output layer have the same size (i.e. the same number of neural units).
- *Bottleneck Layer*: at least one of the hidden layers of the network must be smaller than input and output.

The architecture selected in this work consists of an input layer, 3 hidden layers, and an output layer (see Fig. 1).

The three hidden layers shape a “feature detection” architecture in which the *bottleneck layer* plays the key role in the identity mapping, as it forces the network to develop a compact representation of the training data that better models the underlying system parameters.

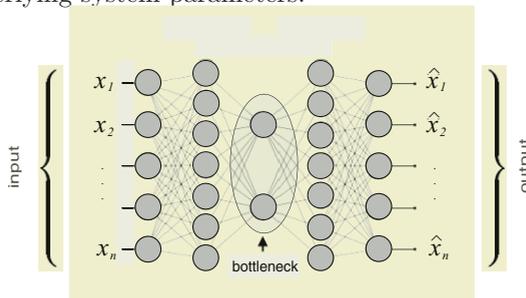


Fig. 1. Sample of autoassociative neural network encoder

3 Problem Definition

The aim of this work is to build a system able to perform strict on-line monitoring on substations belonging to an electric power network, reading measures by RTUs, and able to raise an alarm signal in case of anomaly detection. One of the major difficulties in LCCIs monitoring is due to the non-linear nature of its behaviour; the problem become even harder when a large amount of non predictable abnormal states can arise, due to local or generalized faults.

Numerical methods are usually time and resources consuming and could be not proper for an on-line measure monitoring with a small sampling time. Presently, numerical estimation algorithms are often used to rebuild the state of the power system in case of missing and/or corrupted data: however, state estimator approach does not address the problem of giving a normal/abnormal state assessment, and in some cases could tend to hide traces of an ongoing attack or of other anomalies. Moreover state estimators efficiency and accuracy depend on the size of the network, and the estimation of state is often based on prior knowledge about substations sensor reliability.

Electric system peculiarities and problem specific features suggested the use of neural networks as objects able to deal with continuous values coming from physical fields, with good performances, suitable for on-line data processing, and able to implicitly learn data underlying aspects featuring the system behaviour.

4 Power Net Monitoring

In our approach, a specific autoencoder is deployed on each substation (i.e. through a software-agent platform). Due to peculiarities of each substation in terms of components, geographic location, role in power generation and/or distribution, it's necessary to have a different, specific training session for each deployed neural network.

4.1 Substation Measures as Autoencoder Input

Some data preprocessing is needed in order to feed the autoencoder with substation measures: it can be useful to make a selection among available measures in order to reduce input and output layer size, thus saving time in training phase.

In order to obtain some learning of data variation over time, measures have been composed in a r -slot sliding window, so that each autoencoder is feeded with a data vector containing current as well as past measures. The gap between the slots should be wide enough to have a sensible delta of variation for measured values. Since it turns out that we only need the signal first derivative, we use $r = 2$ in our sliding window. It can be noticed that granularity of substation monitoring depends on closeness of the sampling rate by which measures are retrieved from the electric field and stored in the buffer. Let's call $\mathbf{v}(t)$ the substation measure vector at the time t , T the sampling rate and w the time gap between the sliding window slots; then, at step k , a slot is filled with $\mathbf{v}(kT)$

while the other slot is filled with $v(kT-w)$. How can be noticed, there is no dependence between parameters T and w , as past measures are buffered and ready to be processed at the right time (see Fig. 2).

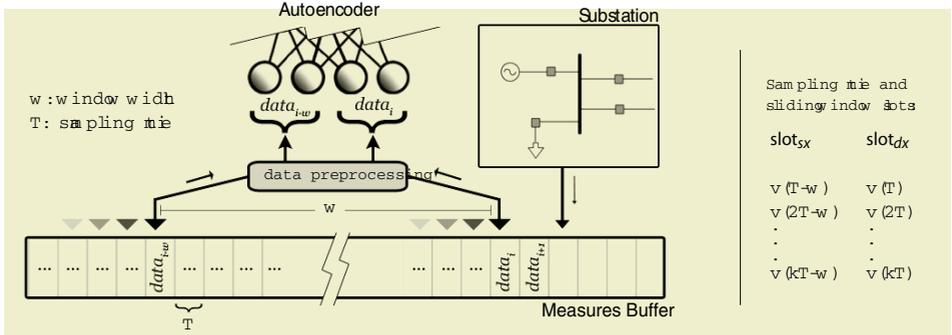


Fig. 2. The sliding window technique: measures coming from power network are preprocessed and stored in a buffer, so that two measures vectors at once are read by the neural network

4.2 Autoencoder Output and Novelty Assessment

In the proposed approach, a novelty assessment is given out measuring absolute gap vector $y(t)$ between input data set and network output, and then performing some post-processing on it (see Fig. 3).

A properly trained autoencoder is able to successfully reproduce *normal* data sets: it will be now showed how autoencoder output is processed to have a measure of anomaly level.

Let $v(t)$ and $\hat{v}(t)$ be respectively the input and output data vector for the autoencoder at the time t , both composed by n measures. First of all, let's consider the vector $y(t)$, where each component is calculated as $y_i(t) = |v_i(t) - \hat{v}_i(t)|$. Let $m(t)$ be the mean value of the components of $y(t)$ at the time t , that is $m(t) = \frac{1}{n} \sum_{i=1}^n y_i(t)$. The value of $m(t)$ itself could be used to have a measure of anomaly level detected by the autoencoder, but it's opportune to introduce

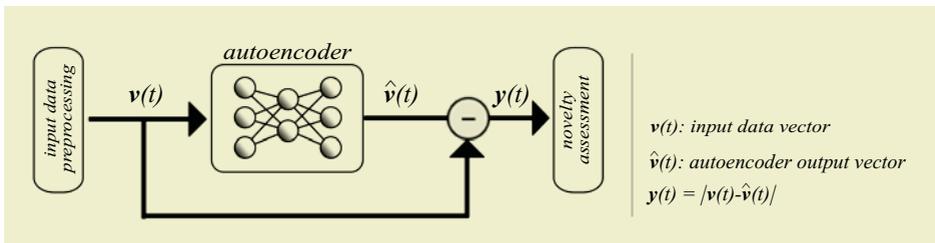


Fig. 3. Comparison between input data vector and autoencoder output vector. Vector $y(t)$ is the absolute difference between autoencoder input and output vector

some smoothing, averaging this value on a sliding window. Thus, being w the width of the sliding window, a new value is introduced as:

$$z(t) = \frac{1}{w} \sum_{k=0}^{w-1} m(t-k)$$

As some measures in the autoencoder output vector are more sensible to anomalies than others, it has been also used the average of the absolute deviation of the measures from their mean. Also in this case this value is observed over a sliding window, so that the final calculated value at the time t is:

$$\gamma(t) = \frac{1}{w} \sum_{k=0}^{w-1} \mu(t-k)$$

where $\mu(t)$ is the average absolute deviation of the measures $y_i(t)$ from their mean at time t , that is $\mu(t) = \frac{1}{n} \sum_{i=1}^n |y_i(t) - m(t)|$.

Values $z(t)$ and $\gamma(t)$ can be used to take a decision about novelty detection. Analyzing these two values during system normal behaviour, two threshold values are defined: let's call them τ_z and τ_γ , an alarm signal is raised if $z(t) > \tau_z$ or $\gamma(t) > \tau_\gamma$.

5 Experimental Results

Experimental tests have been conducted implementing the electric network model IEEE RTS-96 [6] in an electric power network simulator.

5.1 Training the Autoencoder

Training session has been carried out by a 72 hours data set, consisting of 432 training patterns. Using *backpropagation* training algorithm a root mean squared error (RMSE) equal to 0.015 can be reached in about 6000 epochs. Experimental results have shown that a similar value for RMSE permits the autoencoder to generalize well, recognizing as normal data which have small variation respect to training set.

5.2 Testing the Autoencoder

The following experimental tests are aimed to verify if it is possible to discriminate values of $z(t)$ and $\gamma(t)$ generated in case of normal data processing from the ones coming from processing data containing anomalies.

Normal Behaviour Data - Data set used to test network during *normal* behaviour consisted of vectors $v(t)$ obtained simulating a load demand added with zero average uniformly distributed random perturbation. For robustness, calling $l(t)$ the nominal load demand value at time t , at each step the simulation is

executed with a load demand value $\hat{l}(t) = l(t) * (1 + \alpha)$, where α is randomly generated in a suitable interval (i.e. $\alpha \in [-0.01, 0.01]$ for a 1% random perturbation).

Several test sets have been generated with the above criteria: the trained neural network was able to generalize well and successfully reproduced substation behaviour. As stable (small) values for $z(t)$ and $\gamma(t)$ were obtained, it has been easy to choose thresholds τ_z and τ_γ to discriminate *normal* values of $z(t)$ and $\gamma(t)$ from *abnormal* ones.

Next step is to prove that using the autoencoder and the chosen thresholds it is possible to raise an alarm signal in case of anomaly.

Novelty Detection - One of the critical point of this work is the lackness of a priori knowledge about system behaviour in case of anomaly. To test the autoencoder on non-normal values, the following kinds of editing has been made on the data sets:

1. introduction of random noise on each measure vector;
2. changing the curve shape of load demand over time;
3. changing the electric network topology or components status.

The first approach is aimed to verify if relationships among data have been embedded in the autoencoder connection weights. The original measures vector produced by simulation has been perturbed varying each component value by a certain percentage. Being $\mathbf{v}(t)$ the measure vector at the time t , the value of each component has been recalculated as $v_i(t) = v_i(t) * (1 + \beta_i)$ where β_i is a zero average uniformly distributed random noise.

Results of this kind of test are particularly significant: values of $z(t)$ and $\gamma(t)$ given by the autoencoder output postprocessing are very different from normal ones when the random noise introduced is just 2% (i.e. $\beta \in [-0.02, 0.02]$). Plots of anomalous values and thresholds are shown in Fig. 4 (middle plot).

Several simulation have been performed scheduling different demand curves: the right side of plot in Fig. 4 shows values of $z(t)$ and $\gamma(t)$ becoming greater than established thresholds where demand trend is different by the learned one.

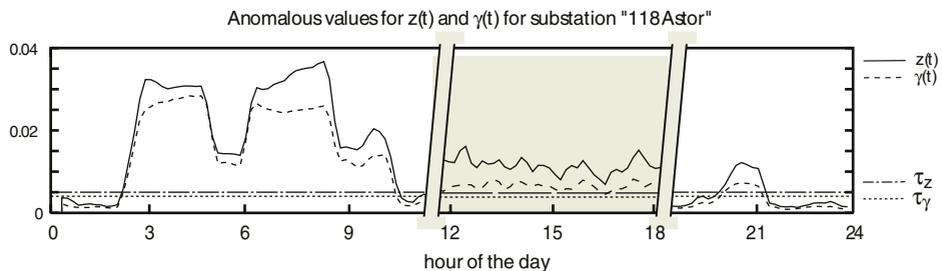


Fig. 4. Plot for $z(t)$ (solid) and $\gamma(t)$ (dashed) in case of anomalous state: generators fault in a neighbour substation (on the left), data hijacking with 2% of noise (center) and unexpected demand trend (on the right); threshold lines are also plotted

With the third data set we want to investigate if the autoencoder is able to detect variations in network topology (i.e. due to a transmission line breaking or a power generator fault). Left side of plot in Fig. 4 shows peak values in correspondence of power generators (simulated) fault in a substation belonging to the same area of the one monitored by the autoencoder. Values for $z(t)$ and $\gamma(t)$ in case of system normal behaviour can be seen in the left plot, for $t < 2$, and in the right plot, for $t > 22$.

6 Conclusions

As shown, using autoassociative neural networks it is possible to build a module to monitor electric power system substations: after a training on data concerning components normal activity, the autoencoder became able to map the system behaviour. Through the processed values $z(t)$ and $\gamma(t)$, the proposed approach successfully detects anomalies in the measures due to sensor failures or intentional data hijacking, network topology changes (i.e. components breaking or transmission lines interruption) or unexpected power demand trend.

References

1. B. Thompson, R. Marks, J. Choi, M. A. El-Sharkawi, M. Huang and C. Bunje, *Implicit Learning in Autoencoder Novelty Assessment*, International Joint Conference on Neural Networks, 2002 IEEE World Congress on Computational Intelligence, Hawaii, May 12-17, 2002
2. S. Haykin, *Neural Networks: A comprehensive Foundation (2nd Edition)*, Prentice Hall, 1998.
3. M. Markou, S. Singh, *Novelty Detection: A Review - part 2: Neural network based approaches*, Signal Processing, vol. 83, 2003.
4. T. M. Mitchell, *Machine Learning*, McGraw-Hill International Editions, 1997.
5. The *Safeguard Project* website: www.ist-safeguard.org
6. Reliability Test System Task Force of the application of probability methods subcommittee, *The IEEE Reliability Test System - 1996*, IEEE Transaction on Power Systems, Vol 14, No. 3, August 1999.