

The Challenge of Non-Technical Loss Detection Using Artificial Intelligence: A Survey

Patrick Glauner¹, Jorge Augusto Meira¹, Petko Valtchev¹², Radu State¹, Franck Bettinger³

¹ *Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg,
4 Rue Alphonse Weicker,
Luxembourg, 2721, Luxembourg*

E-mail: {patrick.glauner, jorge.meira, radu.state}@uni.lu

² *University of Quebec in Montreal,
PO Box 8888, Station Centre-ville,
Montreal, H3C 3P8, Canada*

E-mail: valtchev.petko@uqam.ca

³ *CHOICE Technologies Holding Sàrl,
2-4 Rue Eugene Ruppert,
Luxembourg, 2453, Luxembourg*

E-mail: franck.bettinger@choiceholding.com

Received 29 December 2016

Accepted 25 February 2017

Abstract

Detection of non-technical losses (NTL) which include electricity theft, faulty meters or billing errors has attracted increasing attention from researchers in electrical engineering and computer science. NTLs cause significant harm to the economy, as in some countries they may range up to 40% of the total electricity distributed. The predominant research direction is employing artificial intelligence to predict whether a customer causes NTL. This paper first provides an overview of how NTLs are defined and their impact on economies, which include loss of revenue and profit of electricity providers and decrease of the stability and reliability of electrical power grids. It then surveys the state-of-the-art research efforts in a up-to-date and comprehensive review of algorithms, features and data sets used. It finally identifies the key scientific and engineering challenges in NTL detection and suggests how they could be addressed in the future.

Keywords: Covariate shift, electricity theft, expert systems, machine learning, non-technical losses, stochastic processes.

1. Introduction

Our modern society and daily activities strongly depend on the availability of electricity. Electrical power grids allow to distribute and deliver electricity from generation infrastructure such as power plants or solar cells to customers such as residences or fac-

tories. One frequently appearing problem are losses in power grids, which can be classified into two categories: technical and non-technical losses.

Technical losses occur mostly due to power dissipation. This is naturally caused by internal electrical resistance and the affected components include generators, transformers and transmission lines.

The complementary non-technical losses (NTL) are primarily caused by electricity theft. In most countries, NTLs account for the predominant part of the overall losses as discussed in Ref. 1. Therefore, it is most beneficial to first reduce NTLs before reducing technical losses as proposed in Ref. 2. In particular, NTLs include, but are not limited to, the following causes reported in Refs. 3 and 4:

- Meter tampering in order to record lower consumptions
- Bypassing meters by rigging lines from the power source
- Arranged false meter readings by bribing meter readers
- Faulty or broken meters
- Un-metered supply
- Technical and human errors in meter readings, data processing and billing

NTLs cause significant harm to economies, including loss of revenue and profit of electricity providers, decrease of the stability and reliability of electrical power grids and extra use of limited natural resources which in turn increases pollution. For example, in India, NTLs are estimated at US\$ 4.5 billion in Ref. 5. NTLs are simultaneously reported in Refs. 6 and 7 to range up to 40% of the total electricity distributed in countries such as Brazil, India, Malaysia or Lebanon. They are also of relevance in developed countries, for example estimates of NTLs in the UK and US that range from US\$ 1-6 billion are reported in Refs. 1 and 8.

We want to highlight that only few works on NTL detection have been reported in the literature in the last three to four years. Given that NTL detection is an active field in industrial R&D, it is to our surprise that academic research in this field has dropped in the last few years.

From an electrical engineering perspective, one method to detect losses is to calculate the energy balance reported in Ref. 9, which requires topological information of the network. In emerging economies, which are of particular interest due to their high NTL proportion, this is not realistic for the following reasons: (i) network topology undergoes continuous

changes in order to satisfy the rapidly growing demand of electricity, (ii) infrastructure may break and lead to wrong energy balance calculations and (iii) it requires transformers, feeders and connected meters to be read at the same time.

A more flexible and adaptable approach is to employ artificial intelligence (AI), which is well covered in Ref. 10. AI allows to analyze customer profiles, their data and known irregular behavior. This allows to trigger possible inspections of customers that have abnormal electricity consumption patterns. Technicians then carry out inspections, which allow them to remove possible manipulations or malfunctions of the power infrastructure. Furthermore, the fraudulent customers can be charged for the additional electricity consumed. However, carrying out inspections is costly, as it requires physical presence of technicians.

NTL detection methods reported in the literature fall into two categories: expert systems and machine learning. Expert systems incorporate hand-crafted rules for decision making. In contrast, machine learning gives computers the ability to learn from examples without being explicitly programmed. Historically, NTL detection systems were based on domain-specific rules. However, over the years, the field of machine learning has become the predominant research direction of NTL detection. To date, there is no authoritative survey that compares the various approaches of NTL detection methods reported in the literature. We are also not aware of any existing survey that discusses the shortcomings of the state of the art. In order to advance in NTL detection, the main contributions of this survey are the following:

- We provide a detailed review and critique of state-of-the-art NTL detection research employing AI methods in Section 2.
- We identify the unsolved key challenges of this field in Section 3.
- We describe in detail the proposed methods to solve the most relevant challenges in the future in Section 4.
- We put these challenges in the context of AI research as a whole as they are of relevance to many other learning and anomaly detection problems.

2. The State of the Art

NTL detection can be treated as a special case of fraud detection, for which general surveys are provided in Refs. 11 and 12. Both highlight expert systems and machine learning as key methods to detect fraudulent behavior in applications such as credit card fraud, computer intrusion and telecommunications fraud. This section is focused on an overview of the existing AI methods for detecting NTLs. Existing short surveys of the past efforts in this field, such as in Refs. 3, 13, 14 and 15 only provide a narrow comparison of the entire range of relevant publications. The novelty of this survey is to not only review and compare a wide range of results reported in the literature, but to also derive the unsolved challenges of NTL detection.

2.1. Features

In this subsection, we summarize and group the features reported in the literature.

2.1.1. Monthly consumption

Many works on NTL detection use traditional meters, which are read monthly or annually by meter readers. Based on this data, average consumption features are used in Refs. 1, 7, 16, 17 and 18. In those works, the feature computation used can be summarized as follows: For M customers $\{0, 1, \dots, M-1\}$ over the last N months $\{0, 1, \dots, N-1\}$, a feature matrix F is computed, in which element $F_{m,d}$ is a daily average kWh consumption feature during that month:

$$x_d^{(m)} = \frac{L_d^{(m)}}{R_d^{(m)} - R_{d-1}^{(m)}}, \quad (1)$$

where for customer m , $L_d^{(m)}$ is the kWh consumption increase between the meter reading to date $R_d^{(m)}$ and the previous one $R_{d-1}^{(m)}$. $R_d^{(m)} - R_{d-1}^{(m)}$ is the number of days between both meter readings of customer m .

The previous 24 monthly meter readings are used in Refs. 19 and 20. The features computed are the monthly consumption before the inspection, the consumption in the same month in the year before

the consumption in the past three months and the customer's consumption over the past 24 months. Using the previous six monthly meter readings, the following features are derived in Ref. 21: average consumption, maximum consumption, standard deviation, number of inspections and the average consumption of the residential area. The average consumption is also used as a feature in Refs. 22 and 23.

2.1.2. Smart meter consumption

With the increasing availability of smart meter devices, consumption of electric energy in short intervals can be recorded. Consumption features of intervals of 15 minutes are used in Refs. 24 and 25, whereas intervals of 30 minutes are used in Refs. 26 and 27.

The $4 \times 24 = 96$ measurements of Ref. 25 are encoded to a 32-dimensional space in Refs. 6 and 28. Each measurement is 0 or positive. Next, it is then mapped to 0 or 1, respectively. Last, the 32 features are computed. A feature is the weighted sum of three subsequent values, in which the first value is multiplied by 4, the second by 2 and the third by 1.

The maximum consumption in any 15-minute period is used as a feature in Refs. 29–31 and 32. The load factor is computed by dividing the demand contracted by the maximum consumption.

Features from the consumption time series called shape factors are derived from the consumption time series including the impact of lunch times, nights and weekends in Ref. 33.

2.1.3. Master data

Master data represents customer reference data such as name or address, which typically changes infrequently. The work in Ref. 22 uses the following features from the master data for classification: location (city and neighborhood), business class (e.g. residential or industrial), activity type (e.g. residence or drugstore), voltage (110V or 200V), number of phases (1, 2 or 3) and meter type. The demand contracted, which is the number of kW of continuous availability requested from the energy company and

the total demand in kW of installed equipment of the customer are used in Refs. 30–32. In addition, information about the power transformer to which the customer is connected to is used in Ref. 29. The town or customer in which the customer is located, the type of voltage (low, median or high), the electricity tariff, the contracted power as well as the number of phases (1 or 3) are used in Ref. 23. Related master data features are used in Ref. 33, including the type of customer, location, voltage level, type of climate (rainy or hot), weather conditions and type of day.

2.1.4. Credit worthiness

The works in Refs. 1, 17 and 18 use the credit worthiness ranking (CWR) of each customer as a feature. It is computed from the electricity provider's billing system and reflects if a customer delays or avoids payments of bills. CWR ranges from 0 to 5 where 5 represents the maximum score. It reflects different information about a customer such as payment performance, income and prosperity of the neighborhood in a single feature.

2.2. Expert systems and fuzzy systems

An ensemble pre-filters the customers to select irregular and regular customers in Ref. 19. These customers are then used for training as they represent well the two different classes. This is done because of noise in the inspection labels. In the classification step, a neuro-fuzzy hierarchical system is used. In this setting, a neural network is used to optimize the fuzzy membership parameters, which is a different approach to the stochastic gradient descent method used in Ref. 16. A precision of 0.512 and an accuracy of 0.682 on the test set are obtained.

Profiles of 80K low-voltage and 6K high-voltage customers in Malaysia having meter readings every 30 minutes over a period of 30 days are used in Ref. 26 for electricity theft and abnormality detection. A test recall of 0.55 is reported. This work is related to features of Ref. 7, however, it uses entirely fuzzy logic incorporating human expert knowledge for detection.

The work in Ref. 1 is combined with a fuzzy logic expert system postprocessing the output of the SVM in Ref. 7 for ~100K customers. The motivation of that work is to integrate human expert knowledge into the decision making process in order to identify fraudulent behavior. A test recall of 0.72 is reported.

Five features of customers' consumption of the previous six months are derived in Ref. 21: average consumption, maximum consumption, standard deviation, number of inspections and the average consumption of the residential area. These features are then used in a fuzzy c-means clustering algorithm to group the customers into c classes. Subsequently, the fuzzy membership values are used to classify customers into NTL and non-NTL using the Euclidean distance measure. On the test set, an average precision (called average assertiveness) of 0.745 is reported.

2.3. Neural networks

Neural networks are loosely inspired by how the human brain works and allow to learn complex hypotheses from data. They are well described for example in Ref. 34. Extreme learning machines (ELM) are one-hidden layer neural networks in which the weights from the inputs to the hidden layer are randomly set and never updated. Only the weights from the hidden to output layer are learned. The ELM algorithm is applied to NTL detection in meter readings of 30 minutes in Ref. 35, for which a test accuracy of 0.5461 is reported.

An ensemble of five neural networks (NN) is trained on samples of a data set containing ~20K customers in Ref. 20. Each neural network uses features calculated from the consumption time series plus customer-specific pre-computed attributes. A precision of 0.626 and an accuracy of 0.686 are obtained on the test set.

A self-organizing map (SOM) is a type of unsupervised neural network training algorithm that is used for clustering. SOMs are applied to weekly customer data of 2K customers consisting of meter readings of 15 minutes in Ref. 24. This allows to cluster customers' behavior into fraud or non-fraud. Inspections are only carried out if cer-

tain hand-crafted criteria are satisfied including how well a week fits into a cluster and if no contractual changes of the customer have taken place. A test accuracy of 0.9267, a test precision of 0.8526, and test recall of 0.9779 are reported.

A data set of ~22K customers is used in Ref. 22 for training a neural network. It uses the average consumption of the previous 12 months and other customer features such as location, type of customer, voltage and whether there are meter reading notes during that period. On the test set, an accuracy of 0.8717, a precision of 0.6503 and a recall of 0.2947 are reported.

2.4. Support vector machines

The Support Vector Machines (SVM) introduced in Ref. 36 is a state-of-the-art classification algorithm that is less prone to overfitting. Electricity customer consumption data of less than 400 highly imbalanced out of ~260K customers in Kuala Lumpur, Malaysia are used in Ref. 17. Each customer has 25 monthly meter readings in the period from June 2006 to June 2008. From these meter readings, daily average consumption features per month are computed. Those features are then normalized and used for training in a SVM with a Gaussian kernel. In addition, credit worthiness ranking (CWR) is used. It is computed from the electricity provider's billing system and reflects if a customer delays or avoids payments of bills. CWR ranges from 0 to 5 where 5 represents the maximum score. It was observed that CWR is a significant indicator of whether customers commit electricity theft. For this setting, a recall of 0.53 is achieved on the test set. A related setting is used in Ref. 1, where a test accuracy of 0.86 and a test recall of 0.77 are reported on a different data set.

SVMs are also applied to 1,350 Indian customer profiles in Ref. 25. They are split into 135 different daily average consumption patterns, each having 10 customers. For each customer, meters are read every 15 minutes. A test accuracy of 0.984 is reported. This work is extended in Ref. 28 by encoding the $4 \times 24 = 96$ -dimensional input in a lower dimension indicating possible irregularities. This encoding technique results in a simpler model that is faster to train while not losing the expressiveness of

the data and results in a test accuracy of 0.92.

Consumption profiles of 5K Brazilian industrial customer profiles are analyzed in Ref. 29. Each customer profile contains 10 features including the demand billed, maximum demand, installed power, etc. In this setting, a SVM slightly outperforms K-nearest neighbors (KNN) and a neural network, for which test accuracies of 0.9628, 0.9620 and 0.9448, respectively, are reported.

The work of Ref. 28 is extended in Ref. 6 by introducing high performance computing algorithms in order to enhance the performance of the previously developed algorithms. This faster model has a test accuracy of 0.89.

A data set of ~700K Brazilian customers, ~31M monthly meter readings from January 2011 to January 2015 and ~400K inspection data is used in Ref. 16. It employs an industrial Boolean expert system, its fuzzified extension and optimizes the fuzzy system parameters using stochastic gradient descent described in Ref. 37 to that data set. This fuzzy system outperforms the Boolean system. Inspired by Ref. 17, a SVM using daily average consumption features of the last 12 months performs better than the expert systems, too. The three algorithms are compared to each other on samples of varying fraud proportion containing ~100K customers. It uses the area under the (receiver operating characteristic) curve (AUC), which is discussed in Section 3.1. For a NTL proportion of 5%, it reports AUC test scores of 0.465, 0.55 and 0.55 for the Boolean system, optimized fuzzy system and SVM, respectively. For a NTL proportion of 20%, it reports AUC test scores of 0.475, 0.545 and 0.55 for the Boolean system, optimized fuzzy system and SVM, respectively.

2.5. Genetic algorithms

The work in Refs. 1 and 17 is extended by using a genetic SVM for 1,171 customers in Ref. 18. It uses a genetic algorithm in order to globally optimize the hyperparameters of the SVM. Each chromosome contains the Lagrangian multipliers ($\alpha_1, \dots, \alpha_i$), regularization factor C and Gaussian kernel parameter γ . This model achieves a test recall of 0.62.

A data set of ~1.1M customers is used in Ref. 38. The paper identifies the much smaller class

of inspected customers as the main challenge in NTL detection. It then proposes stratified sampling in order to increase the number of inspections and to minimize the statistical variance between them. The stratified sampling procedure is defined as a non-linear restricted optimization problem of minimizing the overall energy loss due to electricity theft. This minimization problem is solved using two methods: (1) genetic algorithm and (2) simulated annealing. The first approach outperforms the other one. Only the reduced variance is reported, which is not comparable to the other research and therefore left out of this survey.

2.6. *Rough sets*

Rough sets give lower and upper approximations of an original conventional or crisp set. The first application of rough set analysis applied to NTL detection is described in Ref. 39 on 40K customers, but lacks details on the attributes used per customer, for which a test accuracy of 0.2 is achieved. Rough set analysis is also applied to NTL detection in Ref. 23 on features related to Ref. 22. This supervised learning technique allows to approximate concepts that describe fraud and regular use. A test accuracy of 0.9322 is reported.

2.7. *Other methods*

Different feature selection techniques for customer master data and consumption data are assessed in Ref. 33. Those methods include complete search, best-first search, genetic search and greedy search algorithms for the master data. Other features called shape factors are derived from the consumption time series including the impact of lunch times, nights and weekends on the consumption. These features are used in K-means for clustering similar consumption time series. In the classification step, a decision tree is used to predict whether a customer causes NTLs or not. An overall test accuracy of 0.9997 is reported.

Optimum path forests (OPF), a graph-based classifier, is used in Ref. 31. It builds a graph in the feature space and uses so-called “prototypes” or training samples. Those become roots of their optimum-path tree node. Each graph node is classi-

fied based on its most strongly connected prototype. This approach is fundamentally different to most other learning algorithms such as SVMs or neural networks which learn hyperplanes. Optimum path forests do not learn parameters, thus making training faster, but predicting slower compared to parametric methods. They are used in Ref. 30 for 736 customers and achieved a test accuracy of 0.9021, outperforming SMVs with Gaussian and linear kernels and a neural network which achieved test accuracies of 0.8893, 0.4540 and 0.5301, respectively. Related results and differences between these classifiers are also reported in Ref. 32.

A different method is to estimate NTLs by subtracting an estimate of the technical losses from the overall losses reported in Ref. 27. It models the resistance of the infrastructure in a temperature-dependent model using regression which approximates the technical losses. It applies the model to a data set of 30 customers for which the consumption was recorded for six days with meter readings every 30 minutes for theft levels of 1, 2, 3, 4, 6, 8 and 10%. The respective test recalls in linear circuits are 0.2211, 0.7789, 0.9789, 1, 1, 1 and 1, respectively.

2.8. *Summary*

A summary and comparison of models, data sets and performance measures of selected work discussed in this review is reported in Table 1. The most commonly used models comprise Boolean and fuzzy expert systems, SVMs and neural networks. In addition, genetic methods, OPF and regression methods are used. Data set sizes have a wide range from 30 up to 700K customers. However, the largest data set of 1.1M customers in Ref. 38 is not included in the table because only the variance is reduced and no other performance measure is provided. Accuracy and recall are the most popular performance measures in the literature, ranging from 0.45 to 0.99 and from 0.29 to 1, respectively. Only very few publications report the precision of their models, ranging from 0.51 to 0.85. The AUC is only reported in one publication. The challenges of finding representative performance measures and how to compare individual contributions are discussed in Sections 3.1 and 3.6, respectively.

Table 1. Summary of models, data sets and performance measures (two-decimal precision).

Ref.	Model	#Customers	Accuracy	Precision	Recall	AUC	NTL/theft proportion
1	SVM (Gauss)	< 400	0.86	-	0.77	-	-
7	SVM + fuzzy	100K	-	-	0.72	-	-
16	Bool rules	700K	-	-	-	0.47	5%
16	Fuzzy rules	700K	-	-	-	0.55	5%
16	SVM (linear)	700K	-	-	-	0.55	5%
16	Bool rules	700K	-	-	-	0.48	20%
16	Fuzzy rules	700K	-	-	-	0.55	20%
16	SVM (linear)	700K	-	-	-	0.55	20%
17	SVM	< 400	-	-	0.53	-	-
18	Genetic SVM	1,171	-	-	0.62	-	-
19	Neuro-fuzzy	20K	0.68	0.51	-	-	-
22	NN	22K	0.87	0.65	0.29	-	-
23	Rough sets	N/A	0.93	-	-	-	-
24	SOM	2K	0.93	0.85	0.98	-	-
25	SVM (Gauss)	1,350	0.98	-	-	-	-
27	Regression	30	-	-	0.22	-	1%
27	Regression	30	-	-	0.78	-	2%
27	Regression	30	-	-	0.98	-	3%
27	Regression	30	-	-	1	-	4-10%
29	SVM	5K	0.96	-	-	-	-
29	KNN	5K	0.96	-	-	-	-
29	NN	5K	0.94	-	-	-	-
30	OPF	736	0.90	-	-	-	-
30	SVM (Gauss)	736	0.89	-	-	-	-
30	SVM (linear)	736	0.45	-	-	-	-
30	NN	736	0.53	-	-	-	-
33	Decision tree	N/A	0.99	-	-	-	-

3. Challenges

The research reviewed in the previous section indicates multiple open challenges. These challenges do not apply to single contributions, rather they spread across different ones. In this section, we discuss these challenges, which must be addressed in order to advance in NTL detection. Concretely, we discuss common topics that have not yet received the necessary attention in previous research and put them in the context of AI research as a whole.

3.1. Class imbalance and evaluation metric

Imbalanced classes appear frequently in machine learning, which also affects the choice of evaluation metrics as discussed in Refs. 40 and 41. Most NTL detection research do not address this property. Therefore, in many cases, high accuracies or high recalls are reported, such as in Refs. 17, 22, 23, 31 and 38. The following examples demonstrate why those performance measures are not suitable for NTL detection in imbalanced data sets: for a test set containing 1K customers of which 999 have regular use, (1) a classifier always predicting non-NTL has an accuracy of 99.9%, whereas in contrast, (2) a classifier always predicting NTL has a recall of 100%. While the classifier of the first example has a very high accuracy and intuitively seems to perform very well, it will never predict any NTL. In contrast, the classifier of the second example will find all NTL, but triggers many costly and unnecessary physical inspections by inspecting all customers. This topic is addressed rarely in NTL literature, such as in Refs. 20 and 42, and these contributions do not use a proper single measure of performance of a classifier when applied to an imbalanced data set.

3.2. Feature description

Generally, hand-crafting features from raw data is a long-standing issue in machine learning having significant impact on the performance of a classifier, as discussed in Ref. 43. Different feature description methods have been reviewed in the previous section. They fall into two main categories: features computed from the consumption profile of cus-

tomers, which are from monthly meter readings, for example in Refs. 1, 7, 16–22 and 23, or smart meter readings, for example in Refs. 6, 24–32, and 33, and features from the customer master data in Refs. 22, 23, 29–32 and 33. The features computed from the time series are very different for monthly meter readings and smart meter readings. The results of those works are not easily interchangeable. While electricity providers continuously upgrade their infrastructure to smart metering, there will be many remaining traditional meters. In particular, this applies to emerging countries.

There are only few works on assessing the statistical usefulness of features for NTL detection, such as in Ref. 44. Almost all works on NTL detection define features and subsequently report improved models that were mostly found experimentally without having a strong theoretical foundation.

3.3. Data quality

In the preliminary work of Ref. 16, we noticed that the inspection result labels in the training set are not always correct and that some fraudsters may be labelled as non-fraudulent. The reasons for this may include bribing, blackmailing or threatening of the technician performing the inspection. Also, the fraud may be done too well and is therefore not observable by technicians. Another reason may be incorrect processing of the data. It must be noted that the latter reason may, however, also label non-fraudulent behavior as fraudulent. Handling noise is a common challenge in machine learning. In supervised machine learning settings, most existing methods address handling noise in the input data. There are different regularization methods such as L_1 or L_2 regularization discussed in Ref. 45 or learning of invariances allowing learning algorithms to better handle noise in the input data discussed in Refs. 46 and 47. However, handling noise in the training labels is less commonly addressed in the machine learning literature. Most NTL detection research use supervised methods. This shortcoming of the training data and potential wrong labels in particular are only rarely reported in the literature, such as in Ref. 19, which uses an ensemble to pre-filter the training data.

3.4. Covariate shift

Covariate shift refers to the problem of training data (i.e. the set of inspection results) and production data (i.e. the set of customers to generate inspections for) having different distributions. This fact leads to unreliable NTL predictors when learning from this training data. Historically, covariate shift has been a long-standing issue in statistics, as surveyed in Ref. 48. For example, The Literary Digest sent out 10M questionnaires in order to predict the outcome of the 1936 US Presidential election. They received 2.4M returns. Nonetheless, the predicted result proved to be wrong. The reason for this was that they used car registrations and phone directories to compile a list of recipients. In that time, the households that had a phone or a car represented a biased sample of the overall population. In contrast, George Gallup only interviewed 3K handpicked people, which were an unbiased sample of the population. As a consequence, Gallup could predict the outcome of the election very well.

For about the last fifteen years, the Big Data paradigm followed in machine learning has been to gather more data rather than improving models. Hence, one may assume that having simply more customer and inspection data would help to detect NTL more accurately. However, in many cases, the data may be biased as depicted in Fig. 1.



Fig. 1. Example of spatial bias: The large city is close to the sea, whereas the small city is located in the interior of the country. The weather in the small city undergoes stronger changes during the year. The subsequent change of electricity consumption during the year triggers many inspections. As a consequence, most inspections are carried out in the small city. Therefore, the sample of customers inspected does not represent the overall population of customers.

One reason is, for example, that electricity suppliers previously focused on certain neighborhoods

for inspections. Concretely, the customers inspected are a sample of the overall population of customers. In this example, there is a spatial bias. Hence, the inspections do not represent the overall population of customers. As a consequence, when learning from the inspection results, a bias is learned, making predictions less reliable. Aside from spatial covariate shift, there may be other types of covariate shift in the data, such as the meter type, connection type, etc.

To the best of our knowledge, the issue of covariate change has not been addressed in the literature on NTL detection. However, in many cases it may lead to unreliable NTL detection models. Therefore, we consider it important to derive methods for quantifying and reducing the covariate shift in data sets relevant to NTL detection. This will allow to build more reliable NTL detection models.

3.5. Scalability

The number of customers used throughout the research reviewed significantly varies. For example, Refs. 17 and 27 only use less than a few hundred customers in the training. A SVM with a Gaussian kernel is used in Ref. 17. In that setting, training is only feasible in a realistic amount of time for up to a couple of tens of thousands of customers in current implementations as discussed in Ref. 49. A regression model using the Moore-Penrose pseudoinverse introduced in Ref. 50 is used in Ref. 27. This model is also only able to scale to up to a couple of tens of thousands of customers. Neural networks are trained on up to a couple of tens of thousands of customers in Refs. 20 and 22. The training methods used in prior work usually do not scale to significantly larger customer data sets. Larger data sets using up to hundreds of thousands or millions of customers are used in Refs. 16 and 38 using a SVM with linear kernel or genetic algorithms, respectively. An important property of NTL detection methods is that their computational time must scale to large data sets of hundreds of thousands or millions of customers. Most works reported in the literature do not satisfy this requirement.

3.6. Comparison of different methods

Comparing the different methods reviewed in this paper is challenging because they are tested on different data sets, as summarized in Table 1. In many cases, the description of the data lacks fundamental properties such as the number of meter readings per customer, NTL proportion, etc. In order to increase the reliability of a comparison, joint efforts of different research groups are necessary. These efforts need to address the benchmarking and comparability of NTL detection systems based on a comprehensive freely available data set.

4. Suggested Methodology

We have reviewed state-of-the-art research in machine learning and identified the following suggested methodology for solving the main research challenges in NTL detection:

4.1. Handling class imbalance and evaluation metric

How can we handle the imbalance of classes and assess the outcome of classifications using accurate metrics?

Anomaly detection problems are particularly imbalanced, meaning that there are much more training examples of the regular class compared to the anomaly class. Most works on NTL detection do not reflect the imbalance and simply report accuracies or recalls, for example in Refs. 17, 22, 23, 31 and 38. This is also depicted in Table 1. For NTL detection, the goal is to reduce the false positive rate (FPR) to decrease the number of costly inspections, while increasing the true positive rate (TPR) to find as many NTL occurrences as possible. In Ref. 16, we propose to use a receiver operating characteristic (ROC) curve, which plots the TPR against the FPR. The area under the curve (AUC) is a performance measure between 0 and 1, where any binary classifier with an $AUC > 0.5$ performs better than random guessing. In order to assess a NTL prediction model using a single performance measure, the AUC was picked as the most suitable one in Ref. 16.

All works in the literature only use a fixed NTL proportion in the data set, for example in Refs. 17, 20, 22, 23, 31, 38 and 42. We think that it is necessary to investigate more into this topic in order to report reliable and imbalance-independent results that are valid for different levels of imbalance. This will allow to build models that work in different regions, such as in regions with a high NTL ratio as well as in regions with a low occurrence of NTLs. Therefore, we suggest to create samples of different NTL proportions and assess the models on the entire range of these samples. In the preliminary work of Ref. 16, we also noticed that the precision usually grows linearly with the NTL proportion in the data set. It is therefore not suitable for low NTL proportions. However, we did not notice this for the recall and made observations of non-linearity similar to related work in Ref. 27, as depicted in Table 1. With the limitations of precision and recall, the F_1 score did not prove to work as a reliable performance measure.

Furthermore, we suggest to derive multi-criteria evaluation metrics for NTL detection and rank customers that cause a NTL with a confidence level, for example models related to the ones introduced in Ref. 51. For example, the criteria we suggest to include are the costs of inspections and possible increases in revenue.

4.2. Feature description and modeling temporal behavior

How can we describe features that accurately reflect NTL occurrence and can we self-learn these features from data? NTL of customers is a set of inherently temporal events where for example a fraud of customers excites themselves or other related customers to commit fraud as well. How can we extend temporal processes to model the characteristics of NTL?

Most research on NTL uses primarily information from the consumption time series. The consumption is from traditional meters, such as in Refs. 1, 16, 17, 19 and 20, or smart meters, such as in Refs. 6, 25–27, 30, 32 and 33. Both meter types will co-exist in the next decade and the results of those works are not easily interchangeable. Therefore, we suggest to shift to self-learning of features from the

consumption time series. This topic has not been explored in the literature on NTL detection yet. Deep learning allows to self-learn hidden correlations and increasingly more complex feature hierarchies from the raw data input as discussed in Ref. 52. This approach has led to breakthroughs in image analysis and speech recognition as presented in Ref. 53. One possible method to overcome the challenge of feature description for NTL detection is by finding a way to apply deep learning to it.

In a different vein, we believe that the neighborhood of customers contains information about whether a customer may cause a NTL or not. Our hypothesis is confirmed by initial work described in Ref. 21, in which also the average consumption of the residential neighborhood is used for classification of NTL. We have shown in Ref. 44 that features derived from the inspection ratio and NTL ratio in a neighborhood help to detect NTL.

A temporal process, such as a Hawkes process described in Ref. 54, models the occurrence of an event that depends on previous events. Hawkes processes include self-excitement, meaning that once an event happens, that event is more likely to happen in the near future again and decays over time. In other words, the further back the event in the process, the less impact it has on future events. The dynamics of Hawkes processes look promising for modeling NTL: Our first hypothesis is that once customers were found to steal electricity, finding them or their neighbors to commit theft again is more likely in the near future again and decays over time. A Hawkes process allows to model this first hypothesis. Our second hypothesis is that once customers were found to steal electricity, they are aware of inspections and subsequently are less likely to commit further electricity theft. Therefore, finding them or their neighbors to commit theft again is more likely in the far future and increases over time as they become less risk-aware. As a consequence, we need to extend the Hawkes process by incorporating both, self-excitement in order to model the first hypothesis, as well as self-regulation in order to model the second hypothesis. Only few works have been reported on modeling anomaly detection using self-excitement and self-regulation, such as faulty elec-

trical equipment in subway systems reported in Ref. 55.

The neighborhood is essential from our point of view as neighbors are likely to share their knowledge of electricity theft as well as the outcome of inspections with their neighbors. We therefore want to extend this model by optimizing the number of temporal processes to be used. In the most trivial case, one temporal process could be used for all customers combined. However, this would lead to a model that underfits, meaning it would not be able to distinguish among the different fraudulent behaviors. In contrast, each customer could be modeled by a dedicated temporal process. However, this would not allow to catch the relevant dynamics, as most fraudulent customers were only found to steal once. Furthermore, the computational costs of this approach would not be feasible. Therefore, we suggest to cluster customers based on their location and then to train one temporal process on the customers of each cluster. Finally, for each cluster, the conditional intensity of its temporal process at a given time can then be used as a feature for the respective customers. In order to find reasonable clusters, we suggest to solve an optimization problem which includes the number of clusters, i.e. the number of temporal processes to train, as well as the sum of prediction errors of all customers.

4.3. Correction of spatial bias

Previous inspections may have focused on certain neighborhoods. How can we reduce the covariate shift in our training set?

The customers inspected are a sample of the overall population of customers. However, that sample may be biased, meaning it is not representative for the population of all customers. A reason for this is that previous inspections were largely focused on certain neighborhoods and were not sufficiently spread among the population. This issue has not been addressed in the literature on NTL yet. All works on NTL detection, such as Refs. 1, 16, 17, 20, 22, 23, 31, 33, 38 and 42, implicitly assume that the customers inspected are from the distribution of all customers. Overall, we think that the topic of bias correction is currently not receiving the

necessary attention in the field of machine learning as a whole. For about the last ten years, the paradigm followed has been labeled in Ref. 56: “It’s not who has the best algorithm that wins. It’s who has the most data.” However, we are confident to also show that having more *representative* data will help rather than just having a lot of more data for NTL detection.

Bias correction has initially been addressed in the field of computational learning theory, see Ref. 57, which also calls this problem covariate shift, sampling bias or sample selection bias in Ref. 58. For example, one promising approach is resampling inspection data in order to be representative for the overall population of customers. This can be done by learning the hidden selection criteria of the decision whether to inspect a customer or not. Covariate shift can be defined in mathematical terms as introduced in Ref. 58:

- Assume that all examples are drawn from a distribution D with domain $X \times Y \times S$,
- where X is the feature space,
- Y is the label space
- and S is $\{0, 1\}$.

Examples (x, y, s) are drawn independently from D . $s = 1$ denotes that an example is selected, whereas $s = 0$ does not. The training is performed on a sample that comprises all examples that have $s = 1$. If $P(s|x, y) = P(s|x)$ holds true, we can imply that s is independent of y given x . In this case, the selected sample is biased but the bias only depends on the feature vector x . This bias is called covariate shift. An unbiased distribution can be computed as follows:

$$\widehat{D}(x, y, s) := P(s = 1) \frac{D(x, y, s)}{P(s = 1|x)}. \quad (2)$$

Spatial point processes surveyed in Ref. 59 build on top of Poisson processes. They allow to examine a data set of spatial locations and to conclude whether the locations are randomly distributed in a space or if they are skewed. Eq. (2) requires $P(s = 1|x) > 0$ for all possible x . In order to compute this non-zero probability for spatial locations x , we

suggest to use and amend spatial point processes in order to reduce the spatial covariate shift of inspection results. This will in turn allow to train more reliable NTL predictors.

4.4. Scalability to smart meter profiles of millions of customers

How can we efficiently implement the models in order to scale to Big Data sets of smart meter readings?

Experiments reported in the literature range from data sets that have up to a few hundred customers in Refs. 1, 27 and 30 through data sets that have thousands of customers in Refs. 24 and 29 to tens of thousands of customers in Refs. 19 and 22. The world-wide electricity grid infrastructure is currently undergoing a transformation to smart grids, which include smart meter readings every 15 or 30 minutes. The models reported in the literature that work on smart meter data use only very short periods of up to a few days for NTL, such as in Refs. 24–26 and 27. Future models must scale to millions of customers and billions of smart meter readings. The focus of this objective is to perform the computations efficiently in a high performance environment. For this, we suggest to redefine the computations to be computed on GPUs, as described in Ref. 60, or using a map-reduce architecture introduced in Ref. 61.

4.5. Creation of a publicly available real-world data set

How can we compare different models?

The works reported in the literature describe a wide variety of different approaches for NTL detection. Most works only use one type of classifier, such as in Refs. 1, 22, 24 and 27, whereas some works compare different classifiers on the same features, such as in Refs. 29, 31 and 35. However, in many cases, the actual choice of classification algorithm is less important. This can also be justified by the “no free lunch theorem” introduced in Ref. 62, which states that no learning algorithm is generally better than others.

We are interested in not only comparing classification algorithms on the same features, but instead

in comparing totally different NTL detection models. We suggest to create a publicly available data set for NTL detection. Generally, the more data, the better for this data set. However, acquiring more data is costly. Therefore, a tradeoff between the amount of data and the data acquisition costs must be found. The data set must be based on real-world customer data, including meter readings and inspection results. This will allow to compare various models reported in the literature. For these reasons, it should reflect at least the following properties:

- Different types of customers: the most common types are residential and industrial customers. Both have very different consumption profiles. For example, the consumption of industrial customers often peaks during the weekdays, whereas residential customers consume most electricity on the weekends.
- Number of customers and inspections: the number of customers and inspections must be in the hundreds of thousands in order to make sure that the models assessed scale to Big Data sets.
- Spread of customers across geographical area: the customers of the data set must be spread in order to reflect different levels of prosperity as well as changes of the climate. Both factors affect electricity consumption and NTL occurrence.
- Sufficiently long period of meter readings: due to seasonality, the data set must contain at least one year of data. More years are better to reflect changes in the consumption profile as well as to become less prone to weather anomalies.

5. Conclusion

Non-technical losses (NTL) are the predominant type of losses in electricity power grids. We have reviewed their impact on economies and potential losses of revenue and profit for electricity providers. In the literature, a vast variety of NTL detection methods employing artificial intelligence methods are reported. Expert systems and fuzzy systems are traditional detection models. Over the past years, machine learning methods have become more popular. The most commonly used methods are support vector machines and neural networks, which

outperform expert systems in most settings. These models are typically applied to features computed from customer consumption profiles such as average consumption, maximum consumption and change of consumption in addition to customer master data features such as type of customer and connection type. Sizes of data sets used in the literature have a large range from less than 100 to more than one million. In this survey, we have also identified the six main open challenges in NTL detection: handling imbalanced classes in the training data and choosing appropriate evaluation metrics, describing features from the data, handling incorrect inspection results, correcting the covariate shift in the inspection results, building models scalable to Big Data sets and making results obtained through different methods comparable. We believe that these need to be accurately addressed in future research in order to advance in NTL detection methods. This will allow to share sound, assessable, understandable, replicable and scalable results with the research community. In our current research we have started to address these challenges with the methodology suggested and we are planning to continue this research. We are confident that this comprehensive survey of challenges will allow other research groups to not only advance in NTL detection, but in anomaly detection as a whole.

Acknowledgments

We would like to thank Angelo Migliosi from the University of Luxembourg and Lautaro Dolberg, Diogo Duarte and Yves Rangoni from CHOICE Technologies Holding Sàrl for participating in our fruitful discussions and for contributing many good ideas. This work has been partially funded by the Luxembourg National Research Fund.

References

1. Jawad Nagi, Keem Siah Yap, Sieh Kiong Tiong, Syed Khaleel Ahmed, and Malik Mohamad. Nontechnical loss detection for metered customers in power utility using support vector machines. *IEEE transactions on Power Delivery*, 25(2):1162–1171, 2010.

2. Saurabh Amin, Galina A Schwartz, and Hamidou Tembine. Incentives and security in electricity distribution networks. In *International Conference on Decision and Game Theory for Security*, pages 264–280. Springer, 2012.
3. Abhishek Chauhan and Saurabh Rajvanshi. Non-technical losses in power system: A review. In *Power, Energy and Control (ICPEC), 2013 International Conference on*, pages 558–561. IEEE, 2013.
4. Thomas B Smith. Electricity theft: a comparative analysis. *Energy Policy*, 32(18):2067–2076, 2004.
5. Bhavna Bhatia and Mohinder Gulati. Reforming the power sector, controlling electricity theft and improving revenue. *Public Policy for the Private Sector*, 2004.
6. Soma Shekara Sreenadh Reddy Depuru, Lingfeng Wang, Vijay Devabhaktuni, and Robert C Green. High performance computing for detection of electricity theft. *International Journal of Electrical Power & Energy Systems*, 47:21–30, 2013.
7. Jawad Nagi, Keem Siah Yap, Sieh Kiong Tiong, Syed Khaleel Ahmed, and Farrukh Nagi. Improving svm-based nontechnical loss detection in power utility using the fuzzy inference system. *IEEE Transactions on power delivery*, 26(2):1284–1285, 2011.
8. MS Alam, E Kabir, MM Rahman, and MAK Chowdhury. Power sector reform in bangladesh: Electricity distribution system. *Energy*, 29(11):1773–1783, 2004.
9. CCB de Oliveira, N Kagan, A Meffe, S Jonathan, S Caparroz, and JL Cavaretti. A new method for the computation of technical losses in electrical power distribution systems. In *Electricity Distribution, 2001. Part 1: Contributions. CIRED. 16th International Conference and Exhibition on (IEE Conf. Publ No. 482)*, volume 5, pages 5–pp. IET, 2001.
10. Stuart Jonathan Russell and Peter Norvig. Artificial intelligence: a modern approach (3rd edition), 2009.
11. Richard J Bolton and David J Hand. Statistical fraud detection: A review. *Statistical science*, pages 235–249, 2002.
12. Yufeng Kou, Chang-Tien Lu, Sirirat Sirwongwattana, and Yo-Ping Huang. Survey of fraud detection techniques. In *Networking, sensing and control, 2004 IEEE international conference on*, volume 2, pages 749–754. IEEE, 2004.
13. Maryam Kazerooni, Hao Zhu, and Thomas J Overbye. Literature review on the applications of data mining in power systems. In *Power and Energy Conference at Illinois (PECI), 2014*, pages 1–8. IEEE, 2014.
14. Rong Jiang, Rongxing Lu, Ye Wang, Jun Luo, Changxiang Shen, and Xuemin Sherman Shen. Energy-theft detection issues for advanced metering infrastructure in smart grid. *Tsinghua Science and Technology*, 19(2):105–120, 2014.
15. Stephen McLaughlin, Dmitry Podkuiko, and Patrick McDaniel. Energy theft in the advanced metering infrastructure. In *International Workshop on Critical Information Infrastructures Security*, pages 176–187. Springer, 2009.
16. Patrick Glauner, Andre Boechat, Lautaro Dolberg, Radu State, Franck Bettinger, Yves Rangoni, and Diogo Duarte. Large-scale detection of non-technical losses in imbalanced data sets. In *Innovative Smart Grid Technologies Conference (ISGT), 2016 IEEE Power & Energy Society*. IEEE, 2016.
17. J Nagi, AM Mohammad, KS Yap, SK Tiong, and SK Ahmed. Non-technical loss analysis for detection of electricity theft using support vector machines. In *Power and Energy Conference, 2008. PCon 2008. IEEE 2nd International*, pages 907–912. IEEE, 2008.
18. J Nagi, KS Yap, SK Tiong, SK Ahmed, and AM Mohammad. Detection of abnormalities and electricity theft using genetic support vector machines. In *TENCON 2008-2008 IEEE Region 10 Conference*, pages 1–6. IEEE, 2008.
19. Cyro Muniz, Marley Maria Bernardes Rebuszi Velasco, Ricardo Tanscheit, and Karla Figueiredo. A neuro-fuzzy system for fraud detection in electricity distribution. In *IFSA/EUSFLAT Conf.*, pages 1096–1101. Citeseer, 2009.
20. Cyro Muniz, Karla Figueiredo, Marley Vellasco, Gustavo Chavez, and Marco Pacheco. Irregularity detection on low tension electric installations by neural network ensembles. In *2009 International Joint Conference on Neural Networks*, pages 2176–2182. IEEE, 2009.
21. Eduardo Werley S Angelos, Osvaldo R Saavedra, Omar A Carmona Cortés, and André Nunes de Souza. Detection and identification of abnormalities in customer consumptions in power distribution systems. *IEEE Transactions on Power Delivery*, 26(4):2436–2442, 2011.
22. Breno C Costa, Bruno LA Alberto, André M Portela, W Maduro, and Esdras O Eler. Fraud detection in electric power distribution networks using an ann-based knowledge-discovery process. *International Journal of Artificial Intelligence & Applications*, 4(6):17, 2013.
23. Josif V Spirić, Slobodan S Stanković, Miroslav B Dočić, and Tatjana D Popović. Using the rough set theory to detect fraud committed by electricity customers. *International Journal of Electrical Power & Energy Systems*, 62:727–734, 2014.
24. Jose E Cabral, Joao OP Pinto, and Alexandra MAC Pinto. Fraud detection system for high and low voltage electricity consumers based on data mining. In *2009 IEEE Power & Energy Society General Meeting*, pages 1–5. IEEE, 2009.
25. Soma Shekara Sreenadh Reddy Depuru, Lingfeng Wang, and Vijay Devabhaktuni. Support vector ma-

- chine based data classification for detection of electricity theft. In *Power Systems Conference and Exposition (PSCE), 2011 IEEE/PES*, pages 1–8. IEEE, 2011.
26. J Nagi, KS Yap, F Nagi, SK Tiong, SP Koh, and SK Ahmed. Ntl detection of electricity theft and abnormalities for large power consumers in tnb malaysia. In *Research and Development (SCORED), 2010 IEEE Student Conference on*, pages 202–206. IEEE, 2010.
 27. Sanujit Sahoo, Daniel Nikovski, Toru Muso, and Kaoru Tsuru. Electricity theft detection using smart meter data. In *Innovative Smart Grid Technologies Conference (ISGT), 2015 IEEE Power & Energy Society*, pages 1–5. IEEE, 2015.
 28. Soma Shekara Sreenadh Reddy Depuru, Lingfeng Wang, and Vijay Devabhaktuni. Enhanced encoding technique for identifying abnormal energy usage pattern. In *North American Power Symposium (NAPS), 2012*, pages 1–6. IEEE, 2012.
 29. Caio César Oba Ramos, Andre Nunes De Souza, Danilo Sinkiti Gastaldello, and João Paulo Papa. Identification and feature selection of non-technical losses for industrial consumers using the software weka. In *Industry Applications (INDUSCON), 2012 10th IEEE/IAS International Conference on*, pages 1–6. IEEE, 2012.
 30. Caio CO Ramos, André N Souza, Joao P Papa, and Alexandre X Falcao. Fast non-technical losses identification through optimum-path forest. In *Intelligent System Applications to Power Systems, 2009. ISAP'09. 15th International Conference on*, pages 1–5. IEEE, 2009.
 31. Caio César Oba Ramos, André Nunes de Sousa, João Paulo Papa, and Alexandre Xavier Falcao. A new approach for nontechnical losses detection based on optimum-path forest. *IEEE Transactions on Power Systems*, 26(1):181–189, 2011.
 32. Caio CO Ramos, André N Souza, João P Papa, and Alexandre X Falcão. Learning to identify non-technical losses with optimum-path forest. In *Proceedings of the 17th International Conference on Systems, Signals and Image Processing (IWSSIP 2010)*, pages 154–157, 2010.
 33. Anisah Hanim Nizar, Jun Hua Zhao, and Zhao Yang Dong. Customer information system data pre-processing with feature selection techniques for non-technical losses prediction in an electricity market. In *2006 International Conference on Power System Technology*, pages 1–7. IEEE, 2006.
 34. Christopher M Bishop. *Neural networks: a pattern recognition perspective*. 1996.
 35. AH Nizar, ZY Dong, and Y Wang. Power utility nontechnical loss analysis with extreme learning machine method. *IEEE Transactions on Power Systems*, 23(3):946–955, 2008.
 36. Vladimir N Vapnik. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999.
 37. Léon Bottou. Stochastic learning. In *Advanced lectures on machine learning*, pages 146–168. Springer, 2004.
 38. Estevão Costa, Fábio Fabris, Alexandre Rodrigues Loureiros, Hannu Ahonen, and Flávio Miguel Varejão. Optimization metaheuristics for minimizing variance in a real-world statistical application. In *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, pages 206–207. ACM, 2013.
 39. JE Cabral, JOP Pinto, EM Gontijo, and J Reis. Rough sets based fraud detection in electrical energy consumers. In *WSEAS International Conference on MATHEMATICS AND COMPUTERS IN PHYSICS, Cancun, Mexico*, volume 2, pages 413–416, 2004.
 40. Nathalie Japkowicz and Shaju Stephen. The class imbalance problem: A systematic study. *Intelligent data analysis*, 6(5):429–449, 2002.
 41. Yuchun Tang, Yan-Qing Zhang, Nitesh V Chawla, and Sven Krasser. Svms modeling for highly imbalanced classification. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(1):281–288, 2009.
 42. Matías Di Martino, Federico Decia, Juan Molinelli, and Alicia Fernández. Improving electric fraud detection using class imbalance strategies. In *ICPRAM (2)*, pages 135–141, 2012.
 43. Pedro Domingos. A few useful things to know about machine learning. *Communications of the ACM*, 55(10):78–87, 2012.
 44. Patrick Glauner, Jorge Meira, Lautaro Dolberg, Radu State, Franck Bettinger, Yves Rangoni, and Diogo Duarte. Neighborhood features help detecting non-technical losses in big data sets. In *3rd IEEE/ACM International Conference on Big Data Computing Applications and Technologies (BDCAT 2016)*, 2016.
 45. Andrew Y Ng. Feature selection, l1 vs. l2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning*, page 78. ACM, 2004.
 46. Christopher M Bishop. *Pattern recognition. Machine Learning*, 128, 2006.
 47. B Boser Le Cun, John S Denker, D Henderson, Richard E Howard, W Hubbard, and Lawrence D Jackel. Handwritten digit recognition with a back-propagation network. In *Advances in neural information processing systems*. Citeseer, 1990.
 48. Tim Harford. Big data: are we making a big mistake? ft magazine, 2014.
 49. Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27,

- 2011.
50. Roger Penrose. A generalized inverse for matrices. In *Mathematical proceedings of the Cambridge philosophical society*, volume 51, pages 406–413. Cambridge Univ Press, 1955.
 51. Carlos López-Vázquez. A protocol for the ranking of interpolation algorithms based on confidence levels. *International Journal of Remote Sensing*, 37(19):4683–4697, 2016.
 52. Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
 53. Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97, 2012.
 54. Patrick J Laub, Thomas Taimre, and Philip K Pollett. Hawkes processes. *arXiv preprint arXiv:1507.02822*, 2015.
 55. Şeyda Ertekin, Cynthia Rudin, Tyler H McCormick, et al. Reactive point processes: A new approach to predicting power failures in underground electrical systems. *The Annals of Applied Statistics*, 9(1):122–144, 2015.
 56. Michele Banko and Eric Brill. Scaling to very very large corpora for natural language disambiguation. In *Proceedings of the 39th annual meeting on association for computational linguistics*, pages 26–33. Association for Computational Linguistics, 2001.
 57. Corinna Cortes and Mehryar Mohri. Domain adaptation and sample bias correction theory and algorithm for regression. *Theoretical Computer Science*, 519:103–126, 2014.
 58. Bianca Zadrozny. Learning and evaluating classifiers under sample selection bias. In *Proceedings of the twenty-first international conference on Machine learning*, page 114. ACM, 2004.
 59. Adrian Baddeley, Imre Bárány, and Rolf Schneider. Spatial point processes and their applications. *Stochastic Geometry: Lectures given at the CIME Summer School held in Martina Franca, Italy, September 13–18, 2004*, pages 1–75, 2007.
 60. Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
 61. Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. Spark: cluster computing with working sets. *HotCloud*, 10:10–10, 2010.
 62. David H Wolpert. The lack of a priori distinctions between learning algorithms. *Neural computation*, 8(7):1341–1390, 1996.