

Identifying Irregular Power Usage by Turning Predictions into Holographic Spatial Visualizations

Patrick Glauner*, Niklas Dahringer*, Oleksandr Puhachov*, Jorge Augusto Meira*,
Petko Valtchev†, Radu State* and Diogo Duarte‡

*Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg
2721 Luxembourg, Luxembourg

Email: {first.last, jorge.meira}@uni.lu

†Department of Computer Science, University of Quebec in Montreal,
H3C 3P8 Montreal, Canada

Email: valtchev.petko@uqam.ca

‡CHOICE Technologies Holding Sàrl

2453 Luxembourg, Luxembourg

Email: diogo.duarte@choiceholding.com

Abstract—Power grids are critical infrastructure assets that face non-technical losses (NTL) such as electricity theft or faulty meters. NTL may range up to 40% of the total electricity distributed in emerging countries. Industrial NTL detection systems are still largely based on expert knowledge when deciding whether to carry out costly on-site inspections of customers. Electricity providers are reluctant to move to large-scale deployments of automated systems that learn NTL profiles from data due to the latter’s propensity to suggest a large number of unnecessary inspections. In this paper, we propose a novel system that combines automated statistical decision making with expert knowledge. First, we propose a machine learning framework that classifies customers into NTL or non-NTL using a variety of features derived from the customers’ consumption data. The methodology used is specifically tailored to the level of noise in the data. Second, in order to allow human experts to feed their knowledge in the decision loop, we propose a method for visualizing prediction results at various granularity levels in a spatial hologram. Our approach allows domain experts to put the classification results into the context of the data and to incorporate their knowledge for making the final decisions of which customers to inspect. This work has resulted in appreciable results on a real-world data set of 3.6M customers. Our system is being deployed in a commercial NTL detection software.

Index Terms—Critical infrastructure, non-technical losses, time series classification, Microsoft HoloLens, spatial hologram.

I. INTRODUCTION

Critical infrastructure refers to assets that are essential for the functioning of a society and economy. They include power generation, transmission and distribution facilities. Losses in power grids can be grouped into technical losses, which appear naturally due to internal electrical resistance, and non-technical losses (NTL), which appear during power distribution. NTL include, but are not limited to, the following causes [1], [2]:

- Meter tampering in order to record lower consumptions
- Bypassing meters by rigging lines from the power source
- Arranged false meter readings by bribing readers
- Faulty or broken meters

NTL can range up to 40% of the total electricity distributed in countries such as Brazil, India, Malaysia or Lebanon [3]. As

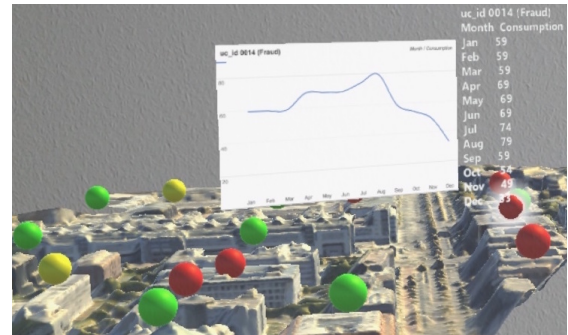


Fig. 1. Example usage of our NTL detection system: Customers are classified as either regular (green), irregular (red) or suspicious (yellow) by a machine learning system. Holographic spatial visualization of customers allows domain experts at the electricity providers to gather information about the customers as well as their neighborhood in order to decide which customers to inspect. The figure depicts the profile of an irregular customer whose consumption has significantly dropped in the last few months.

a consequence, electricity providers face financial losses as well as a decrease of stability and reliability in their power networks. It is for these reasons that electricity providers aim to reduce NTL in their networks by carrying out on-site inspections of customers that have potentially irregular behavior. To date, most NTL detection systems deployed in industry are based on expert knowledge rules [2]. In contrast, the predominant research direction reported in the recent research literature is the use of machine learning/data mining methods, which learn from customer data and known irregular behavior that was reported through inspection results. Due to the high costs per inspection and the limited number of possible inspections, electricity providers aim to maximize the return on investment (ROI) of inspections.

In this paper, we combine both worlds in a spatiotemporal approach that allows domain experts to visualize the prediction results of NTL classifiers in a holographic spatial visualization. An example of this outcome is depicted in Figure 1.

The main contributions of this paper are:

- We propose a novel and flexible framework to compute a large number of domain-specific features and generic features from the noisy consumption time series of customers for NTL detection.
- We retain the statistically meaningful features extracted from the noisy consumption data and optimize different classifiers to predict NTL.
- We present a novel approach to put the prediction results into context by visualizing them in a 3D hologram that contains information about customers and their spatial neighborhood. This hologram can be visualized in a Microsoft HoloLens.

The entire process in our proposed NTL detection system is depicted in Figure 2. As an outcome, domain experts can put the results generated by the classifiers into the context of the data in order to make the final decisions of whether to inspect specific customers. We are confident that this approach will lead to an increase of both stability and reliability of power grids by making better use of the limited number of inspections as well as lead to a greater ROI of the limited number of inspections.

II. RELATED WORK

State-of-the-art surveys on NTL detection are provided in [2], [4]. A data set of ~22K customers is used in [5] 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. Consumption profiles of 5K Brazilian industrial customer profiles are analyzed in [6]. 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. We have discussed the class imbalance and evaluation metric selection of NTL detection in [7] and shown that a large-scale machine learning approach outperforms rule-based Boolean and fuzzy logic expert systems. 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. We have shown in [8] that the sample of inspected customers may be biased, i.e. it does not represent the population of all customers. As a consequence, machine learning models trained on these inspection results may be biased as well and therefore may lead to unreliable predictions of whether customers cause NTL or not. Furthermore, we have shown that the neighborhoods of customers yield significant information in order to decide whether a customer causes a NTL or not [9], [10].

In the literature, different approaches for visualization of NTL are reported. In order to support the decision making, the visualization of the network topology on feeder level as

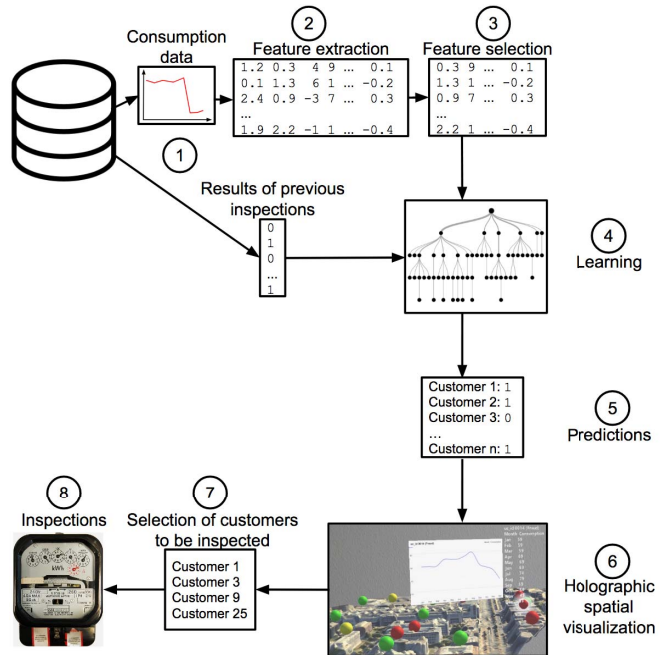


Fig. 2. Proposed NTL detection system: First, the data of previously inspected customers is loaded, which consists of their consumption data as well the inspection result. Second, a vast number of features are extracted from the customers’ noisy consumption data. Third, these features are reduced in order to retain the statistically meaningful ones. Fourth, using the set of reduced features and the results of previously carried out inspections, classifiers are trained in order to recognize NTL. Fifth, these classifiers are then used to predict for customers whether they should be inspected for NTL or not. Sixth, domain experts visualize the customers, their neighbors, inspection results and other data such as the consumption data in a spatial 3D hologram. Seventh, using their expert knowledge, they can review and amend the recommendations made by the classifiers in order to choose the customers for which an inspection appears to be justified from an economic point of view. Last, the inspections are carried out by technicians.

well as load curves on transformer level is proposed in [11]. In addition, the density of NTL in a 2D map is visualized in [12]. For analytics in power grids as a whole, the need for novel and more powerful visualization techniques is argued in [13]. The proposed approaches include heat maps and risk maps. All methods for visualization of NTL proposed in the literature focus only on 2D representations.

We are currently undergoing a paradigm shift in data visualization from not only 2D to 3D, but rather to augmented reality using holographic projections [14]. This shift allows to better understand and experience data [15]. Users are not constrained to looking at data on a screen, as they can interact with the data, e.g. walking around holograms to get a better understanding of Big Data sets. This comes with the benefit of increased productivity as users can use their hands to turn and manipulate objects rather than getting distracted caused by a change of focus from the screen to the input devices such as keyboards or mice [16]. A number of successful applications of holographic projections have been described in the literature including guided assembly instructions [17] as well as a combination of different geographical information

data sources in city management [16]. The literature also discusses the limitations of 3D visualizations, such as that users mistakenly may have greater confidence in the quality of the data [18].

III. DETECTION OF NTL

A. Data

The data used in this paper comes from an electricity provider in Brazil and consists of 3.6M customers. The data contains 820K inspection results, such as inspection date, presence of fraud or irregularity, type of NTL and inspection notes. 620K customers have been inspected at least once and the remaining ~ 3 M customers have never been inspected. Third, there are 195M meter readings from 2011 to 2016 such as consumption in kWh, date of meter reading and number of days between meter readings. From the 620K customers for which an inspection result is available, only the most recent inspection result is used in the experiments in Section IV. The available data per customer m is a complete time series of monthly meter readings of electricity consumption in kWh over the last N months before the most recent inspection, described as follows:

$$C^{(m)} = [C_0^{(m)}, \dots, C_{N-1}^{(m)}],$$

where $C_{N-1}^{(m)}$ is the most recent meter reading before the inspection. For greater N , less customers with a complete time series are available. In contrast, for smaller N , less information per customer is available.

B. Features

In this section, we describe the features that we compute from a customer's consumption time series $C^{(m)}$ for the detection of NTL.

1) *Difference Features*: The intra year difference

$$\text{intra_year}_d^{(m)} = C_d^{(m)} - C_{d-K}^{(m)},$$

for $K = 12$, is the change of consumption to the consumption in the same month of the previous year. In total, there are $N - 12$ intra year difference features.

The intra year seasonal difference

$$\text{intra_year_seasonal}_d^{(m)} = C_d^{(m)} - \frac{1}{3} \times \sum_{k=d-K-1}^{d-K+1} C_k^{(m)},$$

for $K = 12$, is the change of consumption to the mean of the same season in the previous year. In total, there are $N - 13$ intra year seasonal difference features.

The fixed interval

$$\text{fixed_interval}_d^{(m)} = C_d^{(m)} - \frac{1}{K} \times \sum_{k=d-K}^{d-1} C_k^{(m)},$$

for $K \in \{3, 6, 12\}$, is the change of consumption to the mean consumption in a period of time directly before a meter reading. In total, there are $3 \times (N - 12)$ fixed interval features. These features are inspired by [10], in which they are proposed only for the most recent meter reading. Instead, we compute these features for the entire consumption time series.

2) *Daily Averages*: A daily average consumption feature during month d for customer m in kWh is:

$$\text{daily_avg}_d^{(m)} = \frac{C_d^{(m)}}{R_d^{(m)} - R_{d-1}^{(m)}}. \quad (1)$$

$C_d^{(m)}$ is the consumption between the meter reading $R_d^{(m)}$ of month d and the previous one $R_{d-1}^{(m)}$ in month $d - 1$. $R_d^{(m)} - R_{d-1}^{(m)}$ is the number of days between both meter readings of customer m . In total, there are $N - 1$ daily average consumption features. This feature type is successfully used in a number of publications on NTL detection [7], [19]–[22]. It is therefore also relevant to our research.

3) *Generic Time Series Features*: In order to catch more characteristics of the consumption time series, we compute 222 generic time series features from it, comprising:

- Summary statistics, such as maximum, variance or kurtosis.
- Characteristics from sample distribution, such as absolute energy, whether a distribution is symmetric or the number of data points above the median.
- Observed dynamics, such as fast Fourier transformation coefficients, autocorrelation lags or mean value of the second derivative.

The full list of features is provided in [23].

C. Feature Selection

In total, 304 features are computed. In the subsequent learning phase, only the meaningful features should be used. One common dimensionality reduction method is the principal component analysis (PCA). However, time series, and in particular real-world data sets, are noisy, which can lead to poor performance of PCA [24]. It is for that reason that we do not use PCA for the feature selection. Instead, we employ hypothesis tests to the features in order to retain the ones that are statistically relevant. These tests are based on the assumption that a feature x_k is meaningful for the prediction of the binary label vector y if x_k and y are not statistically independent [25]. For binary features, we use Fisher's exact test [26]. In contrast, for continuous features, we use the Kolmogorov-Smirnov test [27].

D. Classifiers

1) *Decision Tree*: Decision tree learners such as ID3 or C4.5 [28] recursively split the input space by choosing the remaining most discriminative feature of a data set. To predict, the learned tree is traversed top-down.

2) *Random Forest*: A random forest [29] is an ensemble estimator that comprises a number of decision trees. Each tree is trained on a subsample of the data and feature set in order to control overfitting. In the prediction phase, a majority vote is made of the predictions of the individual trees.

3) *Gradient Boosted Tree*: A gradient boosted tree [30] is also an ensemble of decision trees. The ensemble is boosted by combining weak classifiers (i.e. classifiers that work little better than a random guess) into a strong one. The ensemble is built by optimizing a loss function.

TABLE I
MODEL PARAMETERS.

Parameter	Values	DT	RF	GBT	LSVM
Learning rate	[0.0001, 1] (log space)			✓	
Loss function	{AdaBoost, deviance}			✓	
Max. number of leaves	[2, 1000)	✓	✓	✓	
Max. number of levels	[1, 50)	✓	✓	✓	
Measure of the purity of a split	{entropy, gini}	✓	✓		
Min. number of samples required to be at a leaf	[1, 1000)	✓	✓	✓	
Min. number of samples required to split a node	[2, 50)	✓	✓	✓	
Number of estimators	20		✓	✓	
L_2 regularization	[0.001, 10] (log space)				✓

TABLE II
NUMBER OF FEATURES BEFORE AND AFTER SELECTION.

Name	#Features	#Retained features
Daily average (AVG)	23	18
Fixed interval	36	34
Generic time series (GTS)	222	162
Intra year difference	12	12
Intra year seasonal difference	11	11
Total	304	237

4) *Support Vector Machine*: A support vector machine (SVM) [31] is a maximum margin classifier, i.e. it creates a maximum separation between classes. Support vectors hold up the separating hyperplane. In practice, they are just a small fraction of the training examples. Therefore, a SVM is often less prone to overfitting than other classifiers, such as a neural network [32]. The training of a SVM can be defined as a Lagrangian dual problem having a convex cost function. By default, the separating hyperplane is linear. Training of SVMs using a kernel to map the input to higher dimension is only feasible for several dozens of thousands of training examples in a realistic amount of time [33]. Therefore, for Big Data sets only a linear implementation of SVMs is practically usable [34].

IV. EVALUATION

A. Metric

The performance measure used in the following experiments is the area under the receiver-operating curve (AUC) [35]. It plots the true positive rate or recall against the false positive rate. It is particularly useful for NTL detection, as it allows to handle imbalanced datasets and puts correct and incorrect inspection results in relation to each other. The superiority of the AUC over other metrics such as accuracy, precision or recall with respect to the problem of NTL detection has been witnessed in the literature [7].

B. Experimental Setup

We experimentally determined $N = 24$ months to work the best for the following experiments. Using $N = 24$ allows

the consumption data to reflect seasonality in the experiments. As a consequence, $M = 150,700$ customers are retained for the experiments. This data set is imbalanced: 100,471 have a negative label (non-NTL), whereas 50,229 have a positive one (NTL). Therefore, 33.33% of the customers used in the following experiments have been found to cause NTL.

We train the decision tree (DT), random forest (RF), gradient boosted tree (GBT) and linear support vector machine (LSVM) classifiers as follows:

- Handling class imbalance: We handle the class imbalance during training by assigning class weights to the examples of both classes in the training set:

$$w_0 = \frac{\# \text{ examples}}{\# \text{ examples}_{C=0}}, \quad (2)$$

$$w_1 = \frac{\# \text{ examples}}{\# \text{ examples}_{C=1}}. \quad (3)$$

- Performing model selection: We want to find the model which is able to distinguish between NTL and non-NTL customers the best. For this, we optimize various parameters for every classifier. The complete list of parameters and considered values per classifier is depicted in Table I. We use randomized grid search, which samples from the joint distribution of model parameters. In contrast to grid search, randomized grid search does not try out all parameter values. We use 100 sampled models in every model selection.
- Handling overfitting: We also employ model selection that splits the data set into $k = 10$ folds. This leads to a more reliable model for NTL detection. The AUC reported per model is the average of the AUCs of the k test sets.

C. Implementation

All computations were run on a server with 80 cores and 128 GB of RAM. The entire code was implemented in Python using `scikit-learn` [34] for machine learning. `scikit-learn` allows to distribute the training of the numerous classifiers among all cores. Using this infrastructure, the extraction of features took 6 hours. The feature selection took only 1 minute. The extensive model selection of classifiers took 4 days. In deployment, the training of classifiers will

TABLE III
TEST PERFORMANCE OF CLASSIFIERS ON FEATURES FROM MEASURED CONSUMPTION DATA.

Cif.	GTS		AVG		DIF		GTS+AVG		GTS+DIF		AVG+DIF		GTS+AVG+DIF	
	X_{all}	X_{ret}	X_{all}	X_{ret}	X_{all}	X_{ret}	X_{all}	X_{ret}	X_{all}	X_{ret}	X_{all}	X_{ret}	X_{all}	X_{ret}
DT	0.64544	0.64625	0.64037	0.63985	0.63730	0.63792	0.64712	0.64705	0.64638	0.64647	0.64348	0.64312	0.64646	0.64765 ^f
RF	0.65665 ^c	0.65726 ^c	0.65083 ^c	0.65248 ^c	0.65529 ^c	0.65459 ^c	0.65800 ^c	0.65835 ^c	0.65911 ^c	0.65896 ^c	0.65858 ^c	0.65755 ^c	0.65747 ^c	0.65977^{c,f}
GBT	0.63149	0.63125	0.63234	0.63186	0.62869	0.63019	0.63262	0.63322	0.63319	0.63358 ^f	0.63261	0.63245	0.63354	0.63355
LSVM	0.63696	0.63656	0.54982	0.54933	0.55749	0.55843	0.63725	0.63689	0.63731	0.63693	0.57173	0.57432	0.63728	0.63760 ^f

Test AUC for combinations of decision tree (DT), random forest (RF), gradient boosted tree (GBT) and linear support vector machine (LSVM) classifiers trained on sets composed of general time series (GTS), daily average (AVG) and difference (DIF) features.

The best overall combination of classifier and feature set is **highlighted**.

Per combination of classifier and feature set, the better result on either a full feature set (X_{all}) or retained feature set (X_{ret}) is **highlighted**.

^c denotes the best classifier per feature set.

^f denotes the best feature set per classifier.

perform significantly faster as the extensive model selection needs to be performed only when a new data set is used. We have also noticed that about 90% of the training time was spent on the gradient boosted tree. Therefore, a significant speedup can be achieved in deployment when skipping the training of this classifier.

D. Feature Selection

We first compute the features described in Section III and then perform the feature selection. In summary, there are three types of features: (1) generic time series (GTS) features, (2) daily average features (AVG) and (3) difference features (DIF) composed of fixed interval, intra year difference and intra year seasonal difference features. The numbers of features before and after selection are depicted in Table II. In total, 237 out of the 304 features are retained. The relevance of our hand-crafted difference features is confirmed: All intra year difference and intra year seasonal difference features are retained. In addition, 34 out of 36 fixed interval features are retained. The 2 features are not retained for $K = 3$, which is most likely due to the too short span of time they reflect. As a matter of fact, daily average features are widely used in the research literature on NTL detection. However, only 18 out of 23 daily average consumption features (i.e. 78%) are retained. The 5 daily average consumption features that are not retained are the ones for the first - i.e. the oldest - 6 months of the 24-month window. The statistical feature check leads to the conclusion that this type of feature is only useful for about 1.5 years of our data for NTL detection. In addition, 73% of the generic time series features are retained after the statistical relevance check. As these features are generic and not particularly made for NTL detection, it is to no surprise that the retention rate for these features is the lowest.

E. Classification Results

We train the four classifiers on each of the GTS, AVG and DIF feature sets as well as on all combinations thereof. The test performance of the best model per experiment returned by the model selection is depicted in Table III. The best test AUC of 0.65977 is achieved for training the random forest classifier on the combination of the retained GTS, AVG and

DIF features. In general, the random forest classifier works the best for every feature set. In total, we report the results of 28 experiments in Table III, both for the full feature sets as well as the retained feature sets. In 16 experiments, the feature selection leads to better results over using all features. Our observation can be explained by the “no free lunch theorem”, which states that no model is generally better than others [36]. However, our best result of 0.65977 is achieved for the retained feature set.

Generally, we observe that a combination of two or three feature sets leads to a better test result than for any of the respective single feature sets. An example to demonstrate this observation is as follows: The random forest classifier achieves test AUCs of 0.65726, 0.65248 and 0.65459 for the retained GTS, AVG and DIF features, respectively. It then achieves test AUCs of 0.65835, 0.65896, 0.65755 and 0.65977 for the retained GTS+AVG, GTS+DIF, AVG+DIF and GTS+AVG+DIF feature sets, respectively. Therefore, the test AUCs for each of the combined feature sets are greater than the test AUCs for any of the single feature sets.

F. Discussion

Previous works that employ the widely-used daily average features established a baseline that only achieved an AUC of slightly above 0.5 [7], [37], i.e. slightly above chance, on real-world NTL detection data sets using linear classifiers. First and foremost, we want to highlight that increasing the performance of machine learning models on noisy real-world data sets is far more challenging than doing so on academic data sets that were created and curated in controlled environments. Furthermore, a small increase of the performance of a real-world model can lead to a major increase of the market value of a company. Our framework presented in this paper significantly outperforms the baselines established in the literature. As a consequence, our models lead to a better detection of NTL and thus to an increase of revenue and profit for electricity providers as well as an increase of stability and reliability in their critical infrastructure. Our NTL detection framework allows other electricity providers to apply our extensive feature extraction, feature selection and



Fig. 3. Microsoft HoloLens [38].

model selection techniques to their data sets, which can lead to potentially greater improvements of NTL detection in their power networks.

It is to our surprise that the gradient boosted tree classifier performs consistently worse than the random forest classifier in our experiments. In the literature, the gradient boosted tree is reported to often lead in a wide range of classification problems [30]. However, our observation can also be explained by the “no free lunch theorem”.

V. HOLOGRAPHIC VISUALIZATION OF NTL

The NTL detection approach presented in Section III and evaluated in Section IV allows to predict whether customers cause NTL or not. It can then be used to trigger possible inspections of customers that have irregular electricity consumption patterns. Subsequently, technicians carry out inspections, which allow them to remove possible manipulations or malfunctions of the power distribution infrastructure. Furthermore, the fraudulent customers can be charged for the additional electricity consumed. Generally, carrying out inspections is costly, as it requires physical presence of technicians. In order to increase both the ROI of the limited number of inspections and the reliability and stability of the power grid, electricity providers in practice strongly rely on expert knowledge for making the decision of whether to inspect a customer or not [7]. As a consequence, electricity providers are reluctant to move to large-deployments of NTL detection systems based on machine learning. We therefore aim to combine automated statistical decision making for generating inspection proposals with incorporating knowledge of the domain experts at the electricity providers for making the final decisions of which customers to inspect.

A. HoloLens

Mixed reality smartglasses such as the Microsoft HoloLens [38] depicted in Figure 3 allow users to combine holographic projections with the real world. The HoloLens offers their user a new perception of 3D models and, perhaps, can provide a new meaning to it. Visualization of data through holograms has found its application in many areas. In medicine, future doctors can study human anatomy by looking at a representation of the human body and navigate through muscles, organs and skeletons [39]. The HoloLens has the ability to perform the holoportation, which allows to virtually place users to remote locations to see, hear and interact with others. Users can walk



Fig. 4. Gesture interactions with the spatial hologram allow to select customers as well as to zoom into or rotate holograms. We also provide a future yellow label that depicts a borderline case, which requires a manual check by domain experts.

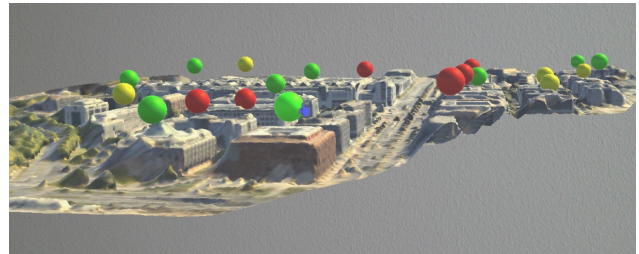


Fig. 5. Zoomed and rotated view on the spatial hologram.

around holograms and interact with them using gaze, gestures or voice in the most natural way. Spatial sound allows hearing holograms even if they are behind the user, considering its position and direction of the sound. Spatial mapping features provide a real-world representation of surfaces, creating convincing holograms in augmented reality.

B. Implementation

We created a 3D model using Google Earth Pro and Blender. Our model allows us to visualize customers and their neighborhood in a 3D spatial hologram that is depicted in Figure 4. A movie was recorded to capture the scene and all its objects from the different angles through Google Earth Pro. Afterwards, images were extracted in Windows Movie Maker from that movie at the best experimentally determined rate of 1 frame/sec. Then, those images were loaded in Blender, which in turn created a 3D FBX model. This model was exported to Unity. Holographic effects were implemented through HoloToolkit-Unity [40]. The GameObjects that handle input events implement the `IInputHandler` interface for tap and hold gestures. Classes that implement the `IManipulationHandler` interface handle manipulation gestures such as moving and rotating actions.

C. Results

This application is used by domain experts at the electricity providers who perceive that customers are classified as either regular (green) or irregular (red). Domain experts can walk around a spatial hologram and observe the data from different directions. Using their hand, they can also interact with the

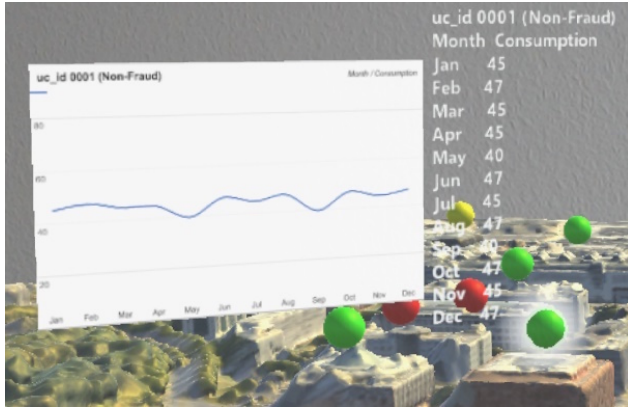


Fig. 6. Detailed view of a customer depicted by a green dot predicted to have a regular power consumption pattern.

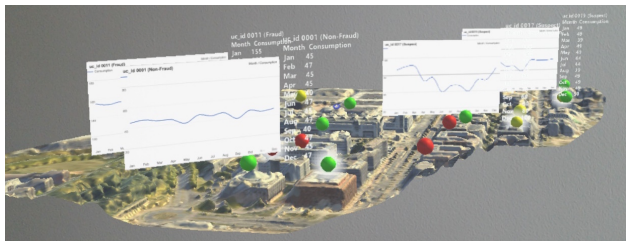


Fig. 7. Multi-view on multiple customers' power consumption history.

hologram in different ways, such as zooming into or rotating the hologram as depicted in Figure 5.

Domain experts can also learn more about a customer by tapping on it with their finger. The spatial hologram then also depicts the consumption profile of the respective customer over a selected period of time such as the previous 12 months. A customer with a predicted regular consumption profile is depicted in Figure 6. This customer's consumption has only changed very little in the last 12 months. As a consequence, the machine learning system classified this customer as non-NTL (green). A customer with an irregular consumption profile is depicted in Figure 1. This customer's consumption has undergone a significant drop over the last few months. Therefore, the machine learning system classified this customer as NTL (red). In both cases, domain experts can compare their observations with the prediction made by the machine learning system. If the prediction is not plausible, domain experts can choose not to follow the recommendation and therefore decide whether to inspect a customer. Our visualization allows domain experts to take the neighborhood of customers into account in order to decide which customers to inspect. Aside from the actual spatial visualization of satellite images of a neighborhood, domain experts can also visualize the consumption profile of neighbors as visualized in Figure 7 for comparing customers in order to decide whether to inspect a customer.

D. Discussion

Our holographic spatial visualization of customers and their neighborhood comes with the benefit of increased productivity. It has previously been shown that the neighborhoods of customers yield significant information in order to decide whether a customer causes a NTL or not [9], [10]. There are many interpretations of this fact. For example, fraudulent customers may share their knowledge with neighbors or there may be a correlation between electricity theft and the level of prosperity of a neighborhood. Our system allows to increase the ROI of inspections as well as to increase both the reliability and stability of the power grid by incorporating expert knowledge in the decision making process. Also, domain experts can use their hands to turn and manipulate objects rather than getting distracted by a change of focus from the screen to the input devices such as a keyboard or mouse.

VI. CONCLUSION AND FUTURE WORK

In this work, we have proposed a novel system for detecting non-technical losses (NTL) for a real-world data set of 3.6M customers. In the first stage, a machine learning system learns to predict whether a customer causes NTL or not. In order to do so, we have proposed to extract a number of domain-specific features from the noisy consumption data. We have shown the statistical relevance of these features over generic time series features. As a consequence, our machine learning system allows to detect NTL better than previous works described in the literature. In the second stage, we put the prediction results into context by visualizing further data of the customers and their neighborhoods in a spatial hologram using a Microsoft HoloLens. Using this hologram, domain experts can then review and amend the suggestions of which customers to inspect. As a result, they can make the final decisions of which customers to inspect in order to increase the ROI of the limited number of inspections.

We have previously referred to the main challenges to solve in order to advance NTL detection. We believe that covariate shift is one of the main impediments in advancing NTL detection. It has been argued that covariate shift is currently one of the main impediments in a wide range of real-world Big Data machine learning problems [41]. Therefore, reducing the covariate shift in the data should be a future priority in the detection of NTL. We therefore expect our models to perform better after the reduction of covariate shift. In order to make the visualization ready for large-scale production, we are planning to integrate other 3D data sources as well as to add compression algorithms such that large maps can be transferred to the HoloLens. We are also interested in visualizing other quantities, such as prosperity levels or credit worthiness, in the spatial holograms. We are also interested in exploring unsupervised methods in the future in order to build system that perform with only the availability of consumption data.

ACKNOWLEDGMENT

We would like to thank Tamer Aidek and Yves Rangoni from CHOICE Technologies Holding Sàrl for contributing many good ideas to our discussions. This work has been partially funded by the Luxembourg National Research Fund.

REFERENCES

- [1] A. Chauhan and S. Rajvanshi, "Non-technical losses in power system: A review," in *Power, Energy and Control (ICPEC), 2013 International Conference on*. IEEE, 2013, pp. 558–561.
- [2] P. Glauner, J. A. Meira, P. Valtchev, F. Bettinger *et al.*, "The challenge of non-technical loss detection using artificial intelligence: A survey," *International Journal of Computational Intelligence Systems*, vol. 10, no. 1, pp. 760–775, 2017.
- [3] S. S. R. Depuru, L. Wang, V. Devabhaktuni, and R. C. Green, "High performance computing for detection of electricity theft," *International Journal of Electrical Power & Energy Systems*, vol. 47, pp. 21–30, 2013.
- [4] J. L. Viegas, P. R. Esteves, R. Melício, V. Mendes, and S. M. Vieira, "Solutions for detection of non-technical losses in the electricity grid: A review," *Renewable and Sustainable Energy Reviews*, vol. 80, pp. 1256–1268, 2017.
- [5] B. C. Costa, B. L. Alberto, A. M. Portela *et al.*, "Fraud detection in electric power distribution networks using an ann-based knowledge-discovery process," *International Journal of Artificial Intelligence & Applications*, vol. 4, no. 6, p. 17, 2013.
- [6] C. C. O. Ramos, A. N. De Souza, D. S. Gastaldello *et al.*, "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*. IEEE, 2012, pp. 1–6.
- [7] P. Glauner, A. Boechar, L. Dolberg *et al.*, "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.
- [8] P. Glauner, A. Migliosi, J. Meira, P. Valtchev, R. State, and F. Bettinger, "Is big data sufficient for a reliable detection of non-technical losses?" in *19th International Conference on Intelligent System Applications to Power Systems (ISAP 2017)*, 2017.
- [9] P. Glauner, J. Meira, L. Dolberg *et al.*, "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.
- [10] J. Meira, P. Glauner, R. State *et al.*, "Distilling provider-independent data for general detection of non-technical losses," in *Power and Energy Conference at Illinois (PECI), 2017*. IEEE, 2017.
- [11] A. Abaide, L. Canha, A. Barin, and G. Cassel, "Assessment of the smart grids applied in reducing the cost of distribution system losses," in *Energy Market (EEM), 2010 7th International Conference on the European*. IEEE, 2010, pp. 1–6.
- [12] J. Porras, H. Rivera, F. Giraldo, and B. Correa, "Identification of non-technical electricity losses in power distribution systems by applying techniques of information analysis and visualization," *IEEE Latin America Transactions*, vol. 13, no. 3, pp. 659–664, 2015.
- [13] F. C. Trindade, L. F. Ochoa, and W. Freitas, "Data analytics in smart distribution networks: Applications and challenges," in *Innovative Smart Grid Technologies-Asia (ISGT-Asia), 2016 IEEE*. IEEE, 2016, pp. 574–579.
- [14] E. Olshannikova, A. Ometov, Y. Koucheryavy, and T. Olsson, "Visualizing big data with augmented and virtual reality: challenges and research agenda," *Journal of Big Data*, vol. 2, no. 1, p. 22, 2015.
- [15] M. A. Hoffman, "The future of three-dimensional thinking," *Science*, vol. 353, no. 6302, pp. 876–876, 2016.
- [16] Z. Lv, X. Li, B. Zhang, W. Wang, Y. Zhu, J. Hu, and S. Feng, "Managing big city information based on webrgis," *IEEE Access*, vol. 4, pp. 407–415, 2016.
- [17] G. Evans, J. Miller, M. I. Pena, A. MacAllister, and E. Winer, "Evaluating the microsoft hololens through an augmented reality assembly application," in *SPIE Defense+ Security*. International Society for Optics and Photonics, 2017, pp. 101970V–101970V.
- [18] S. Zanola, S. I. Fabrikant, and A. Çöltekin, "The effect of realism on the confidence in spatial data quality in stereoscopic 3d displays," in *Proceedings of the 24th international cartography conference (ICC 2009), Santiago, Chile, 2009*, pp. 15–21.
- [19] J. Nagi, K. S. Yap, S. K. Tiong, S. K. Ahmed, and M. Mohamad, "Nontechnical loss detection for metered customers in power utility using support vector machines," *IEEE transactions on Power Delivery*, vol. 25, no. 2, pp. 1162–1171, 2010.
- [20] J. Nagi, K. S. Yap, S. K. Tiong *et al.*, "Improving svm-based nontechnical loss detection in power utility using the fuzzy inference system," *IEEE Transactions on power delivery*, vol. 26, no. 2, pp. 1284–1285, 2011.
- [21] J. Nagi, A. Mohammad, K. Yap, S. Tiong, and S. Ahmed, "Non-technical loss analysis for detection of electricity theft using support vector machines," in *Power and Energy Conference, 2008. PECon 2008. IEEE 2nd International*. IEEE, 2008, pp. 907–912.
- [22] J. Nagi, K. Yap, S. Tiong, S. Ahmed, and A. Mohammad, "Detection of abnormalities and electricity theft using genetic support vector machines," in *TENCON 2008-2008 IEEE Region 10 Conference*. IEEE, 2008, pp. 1–6.
- [23] M. Christ, A. W. Kempa-Liehr, and M. Feindt, "Distributed and parallel time series feature extraction for industrial big data applications," *arXiv preprint arXiv:1610.07717*, 2016.
- [24] T.-c. Fu, "A review on time series data mining," *Engineering Applications of Artificial Intelligence*, vol. 24, no. 1, pp. 164–181, 2011.
- [25] P. Radivojac, Z. Obradovic, A. K. Dunker, and S. Vucetic, "Feature selection filters based on the permutation test," in *ECML*. Springer, 2004, pp. 334–346.
- [26] R. A. Fisher, "On the interpretation of χ^2 from contingency tables, and the calculation of p," *Journal of the Royal Statistical Society*, vol. 85, no. 1, pp. 87–94, 1922.
- [27] F. J. Massey Jr, "The kolmogorov-smirnov test for goodness of fit," *Journal of the American statistical Association*, vol. 46, no. 253, pp. 68–78, 1951.
- [28] J. R. Quinlan, "C4. 5: Programming for machine learning," *Morgan Kaufmann*, p. 38, 1993.
- [29] T. K. Ho, "Random decision forests," in *Document Analysis and Recognition, 1995., Proceedings of the Third International Conference on*, vol. 1. IEEE, 1995, pp. 278–282.
- [30] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. ACM, 2016, pp. 785–794.
- [31] V. N. Vapnik, "An overview of statistical learning theory," *IEEE transactions on neural networks*, vol. 10, no. 5, pp. 988–999, 1999.
- [32] L.-J. Cao and F. E. H. Tay, "Support vector machine with adaptive parameters in financial time series forecasting," *IEEE Transactions on neural networks*, vol. 14, no. 6, pp. 1506–1518, 2003.
- [33] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011.
- [34] F. Pedregosa, G. Varoquaux, A. Gramfort *et al.*, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [35] W. B. van den Hout, "The area under an roc curve with limited information," *Medical decision making*, vol. 23, no. 2, pp. 160–166, 2003.
- [36] D. H. Wolpert, "The lack of a priori distinctions between learning algorithms," *Neural computation*, vol. 8, no. 7, pp. 1341–1390, 1996.
- [37] J. Nagi, K. Yap, F. Nagi *et al.*, "Ntl detection of electricity theft and abnormalities for large power consumers in tn malaysia," in *Research and Development (SCoREd), 2010 IEEE Student Conference on*. IEEE, 2010, pp. 202–206.
- [38] M. Corporation, "Microsoft hololens," <https://www.microsoft.com/en-us/hololens>, 2016, [Online; accessed July 3, 2017].
- [39] A. G. Taylor, "Creating a holographic teaching tool," in *Develop Microsoft HoloLens Apps Now*. Springer, 2016, pp. 185–193.
- [40] M. Corporation, "Holotoolkit-unity," <http://github.com/Microsoft/HoloToolkit-Unity>, 2017.
- [41] A. L'Heureux, K. Grolinger, H. F. ElYamany, and M. Capretz, "Machine learning with big data: Challenges and approaches," *IEEE Access*, 2017.