

Predicting Short 802.11 Sessions from RADIUS Usage Data

Anisa Allahdadi*, Ricardo Morla*, Ana Aguiar[†]

and Jaime S. Cardoso*

*Faculty of Engineering University of Porto and INESC TEC

Porto, Portugal

Email: anisa.allahdadi@inescporto.pt, ricardo.morla@fe.up.pt, jaime.cardoso@fe.up.pt

[†]Faculty of Engineering University of Porto and Instituto de Telecomunicaes, Porto, Portugal

Email: ana.aguiar@fe.up.pt

Abstract—The duration of 802.11 user sessions has been widely studied in the context of analyzing user behavior and mobility. Short (smaller-than-5-minutes) sessions are never used or characterized in these analyses as they are unrelated to user behavior and considered as artifacts introduced by the wireless network. In this paper we characterize short 802.11 sessions as recorded through RADIUS authentication. We show that 50% of access points have 70% of smaller than 5 minutes sessions in a 5 months trace from the Eduroam academic wireless network in the University of Porto. Exactly because they are artifacts introduced by the network, short sessions are an important indicator for network management and the quality of the wireless access. Network managers typically do not collect and process session information but rely on SNMP to provide summaries of 802.11 usage data. We develop a modeling framework to provide predictions for the number of short sessions from SNMP data. We model the data stream of each access point using two methods of regression and one classification technique. We evaluate these models based on short session prediction accuracy. The models are trained on the 5 months data and the best results show prediction accuracy of 95.27% in polynomial regression at degree of 3.

I. INTRODUCTION

The IEEE 802.11 standard is being extensively implemented in most university campuses and organizations to provide wireless connectivity to their users. There are two fundamental aspects in understanding 802.11 technology deployments: 1) how and when the wireless network is used and 2) what performance it provides at a given usage. Large collections of wireless network data (a noteworthy example of which is the Dartmouth repository [1]) have been key in analyzing 802.11 usage and user behavior. This results in a better understanding of human mobility patterns on-off campus and between different parts of a campus, of the diversity of application usage, and of the in/out-bound nature of traffic (see [8], [14], [3] for examples and further references on 802.11 usage analysis).

The performance of deployed 802.11 networks has not been as widely studied. One reason for this is that network researchers and managers tend to have strong intuitions on how the network and network protocols behave, that help them manage the network. In large deployments with varying usage, channel conditions, and operational constraints, this intuition may sometimes fail. Another reason is that in order

to understand network performance, we often need to capture and analyze large amounts of data. This implies higher monitoring and storage costs and severely limits our data analysis capability. One example of such data – and the one we use in this paper – is 802.11 session data.

The 802.11 mobile stations connect to the wireless network by associating and disassociating to an access point (AP). An 802.11 session is the period of time during which a mobile station is associated to an access point. Sessions of interest to 802.11 user behavior and network usage studies are long and typically only sessions longer-than-5-minutes are used in these studies. This is reasonable, as user behavior is not expected to change drastically at time scales smaller than 5 minutes. It is also fortunate as this kind of data can be easily collected via SNMP queries, say every 5 minutes or more.

Shorter sessions are often not directly related to user behavior and are considered artifacts of the network that can hinder performance and indicate poor planning or misbehaving network devices. The ping-pong effect, for example, is the result of a greedy choice of which access point to connect to. Under heavy utilization and dense access point coverage, this greedy behavior may result in many small sessions with a mobile station continuously alternating between access points. Other examples are persistently resource-scarce access points that can force mobile stations to continuously disassociate, also resulting in a series of small sessions. The consequences of small sessions are the overhead of transmitting association frames and the potential loss of connectivity, between the time when a station is disassociated from an access point and the time when it is able to associate to another access point. Small session data, however, is not easily obtainable from SNMP sources. Moreover, the portion of smaller-than-5-minutes session records are potentially greater than longer-than-5-minutes sessions [4].

In this paper we propose to use indicators that can be obtained with typical SNMP queries to identify APs with excessive occurrence of short sessions. Specifically, we use machine learning techniques to predict the number of short sessions in a time-bin from the number of users, number of sessions and number of user sightings in that time-bin. The latter data can easily be obtained via SNMP polling.

We define three categories of short sessions that we intend to predict: zero length, below 30 seconds, and between 30 seconds and 2 minutes. We use two regression methods and a classification technique to learn a single model using data from all the access points and a AP-level model using data from each of the access points. We define mis-prediction and mis-classification errors and the best outcome of regression is compared with the best classification result evaluating the accuracy of the proposed models.

The rest of the paper is organized as follows. Section II presents the related work. Section III identifies the RADIUS authentication and the dataset we utilized. In section IV some analyses regarding the short sessions are provided. Feature specific metrics, the principal methodology and model construction are described in section V. Section VI explains the experimental setups and presents the obtained results. Finally in section VII we provide conclusions and disclose some future lines of work.

II. RELATED WORK

While the most recent work on modeling 802.11 wireless networks focuses on characterizing the network usage and user behavioral pattern [11], [12], [13], [5], this paper propose a peculiar performance estimation metric and provide some learning techniques not only to predict the target performance variable, but also to discern the discriminative structure of the data in various parts of the network.

Highlighting the infrastructure rather than the user mobility behavior, a number of studies [10], [6], [16] focused on characterizing the workload, access pattern, access duration of wireless APs, and total infrastructure usage. Characterizing and modeling the traffic load of each AP, Hernandez-Campos et al. [10] observed correlations between the number of association, total association time and overall traffic per AP. Papadopoulou et al. [16] in a relevant direction of work, conducted an extensive measurement study of wireless users and their association patterns on a university campus. Their study reveals some relation between the wireless access patterns and several parameters such as mobility, sessions and visit duration. In [7], Ergin et al. discussed some 802.11 deployment issues to preserve the required coverage and access quality for mobile users providing novel insights into the TCP behavior in multi-cell WLAN.

In a more related modeling approach, Campos et al. [9] proposed system-wide and AP-level models of traffic demands by focusing on wireless session and network flow to capture the network-independent characteristics of the traffic workload. A very recent line of work presented in [15], dealt with 802.11 connectivity problems by inspecting the usage pattern. Massa and Morla identified a 802.11 usage pattern termed as "abrupt ending" that happens when a large number of sessions in an AP terminate within a one second window. They provided a number of anomaly related patterns such as AP halt/crash, AP overload, interference across the vicinity of an AP, AP persistent interference as well as user authentication failure and intermittent connectivity. They made use of radius

authentication log data as we did, however the issue of short sessions is not considered in any of the recognized anomaly patterns of this work.

To the best of our knowledge, this phenomena is never addressed in such modeling efforts as it is not directly relevant to the user behavioral patterns and not considered in typical SNMP records. However, according to the aforementioned premises and subsequent statistical analyses (in IV), we propose this network usage attribute as a significant indicator of the performance and quality of access in 802.11 networks.

III. DATASET DESCRIPTION

The raw dataset consists of daily summary of connections between 338 access points and their corresponding stations in one of the busiest Eduroam hotspots at the Faculty of Engineering of the University of Porto (FEUP). The University hotspots are part of the Eduroam European wireless academic network. The data used for the current work belongs to a 5 month period of tracing since 2011-7-22 to 2011-12-22. The entire dataset incorporates records of 802.11 mobile stations association to APs stored at a Radius authentication server. When a client associate/disassociate to an 802.11 AP, a "Start"/"Stop" event is recorded. A probing log event "Alive" is generated every 15 minutes while the client is still connected to the network [2]. Moreover, each record includes association time, number of input and output bytes and packets, session ID, and session duration in case the log event is a "Stop".

For this work the user sessions are key, therefore given the session duration, the "Stop" records are sufficient to mark the beginning and ending of user sessions. The number of unique users observed during the 5 months of tracing is over 23 thousands and the corresponding number of "Stop" sessions is around 2.5 millions.

IV. SHORT SESSIONS

A. Mobile Stations

Figure 1 shows the Cumulative Distribution Function (CDF) of the duration of all sessions in our dataset. More than 75% of all sessions last less than the typical SNMP polling time of 5 minutes. More than 25% of all sessions last less than the reasonably small time of 15 seconds for opening the laptop lid to check emails. More than 10% of all sessions last below 3 seconds. As these session durations are too small to be caused by user behavior, we interpret them as indicators of network anomalies, induced by unexpected deficiency or poor planning of network resources. Hence, they can be utilized to alert network managers.

Figure 2 shows the CDF of the percentage of short sessions per user. Close to 90% of users have just 5% zero length sessions. And more than 45% of the users have 50% of their sessions smaller than 5 minutes. Another interesting statistic refers to the short sessions below 30 seconds. It is observed that half of the user sessions for more than 90% of the users fall under 30 seconds. As discussed earlier smaller than 30 seconds sessions are more likely to be generated as a result of

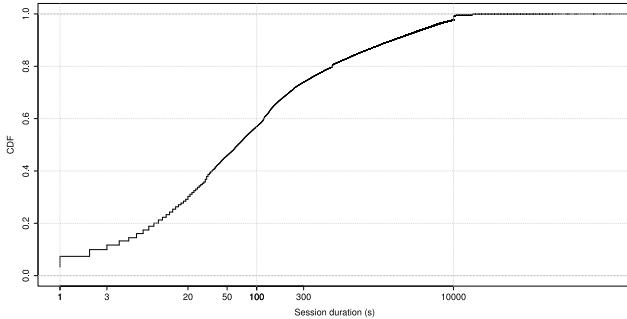


Fig. 1. CDF of session duration

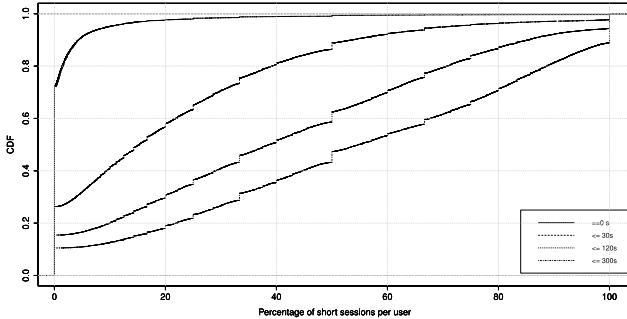


Fig. 2. CDF of the percentage of short sessions per user

poor access quality rather than mobility of users or termination of sessions on purpose.

B. Access Points

Figure 3 shows the CDF of the number of sessions per AP in our dataset. Almost 30% of the APs are considered lightly used with less than 1000 sessions recorded during the 5 months of this experiment. Figure 4 shows the CDF of the percentage of sessions per AP in our dataset. Close to 90% of APs have just 5% zero length sessions. And more than 30% of the APs have 60% of their sessions smaller than 5 minutes. We observed 31 APs with more than half of their sessions smaller than 30 seconds. This is clearly not related to user behavior, and points to poor network performance and large overhead for setting up and tearing down sessions.

Figure 5 shows the relation of AP session count with the percentage of short sessions. APs with higher portion of

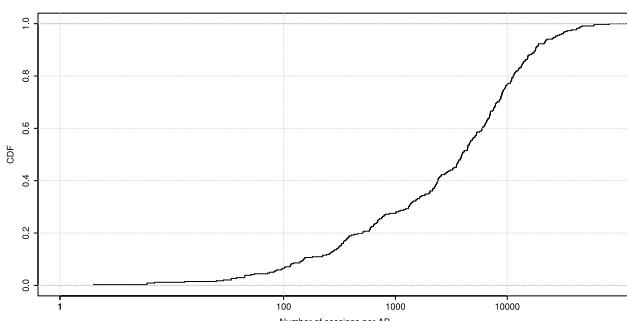


Fig. 3. CDF of the number of sessions per AP

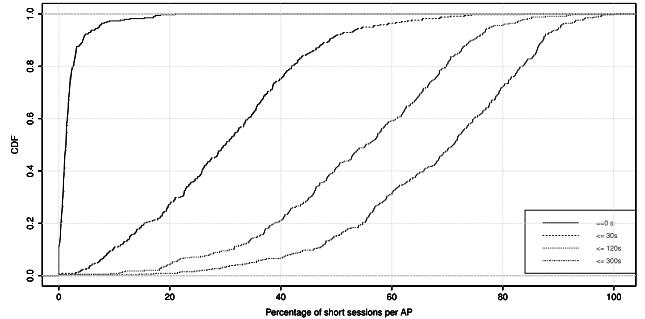


Fig. 4. CDF of the percentage of short sessions per AP

short sessions are not always those with larger session count. Taking total session count of an AP as an implication of its activity and liveliness, it is perceived that such active APs do not normally generate the highest amount of short sessions. However, these results are based on the statistics of the entire dataset and do not consider the daily usage pattern of APs. In the following sections we further explore the relation between AP usage, and the amount/group of short sessions.

V. PREDICTION METHODOLOGY

A. Metrics

Accomplishing an enumeration process on the raw dataset, a number of metrics emerge with application to model building and data analysis. The definition of these metrics can be found in the following paragraphs.

Suppose that L is a location in the campus (indicated by an AP), U is the set of all users, and S is the set of all session IDs.

User Count: the number of unique users observed in a specific location in a temporal period.

$$U_L := \{u \in U : u \text{ observed in } L\}$$

$$UCount(L) := |U_L|$$

Sessions: the raw count of active sessions during a time-bin regardless of the owner. This attribute shows the number of attempts have been made to associate with an AP by a number of users.

$$S_L := \{s \in S : s \text{ observed in } L\}$$

$$Sessions(L) := |S_L|$$

Sightings: the probing event occurring every 5 minutes as an indicator of users' duration of stay in a location, and a counter for sessions longer than 5 minutes. For example, if a user u maintain his connectivity for 12 continuous minutes, the corresponding *Sight* parameter will be 2, showing the presence of user u at least for 10 minutes.

$$Sight_L(u) := \sum_{s \in S_{L,u}} \left\lfloor \frac{\text{EndTime}(s) - \text{BeginTime}(s)}{5} \right\rfloor$$

$$Sightings(L) := \sum_{u \in U_L} Sight_L(u)$$

Short Sessions: This feature has been divided into three categories, zero length sessions, below 30s and between 30s and 2min sessions.

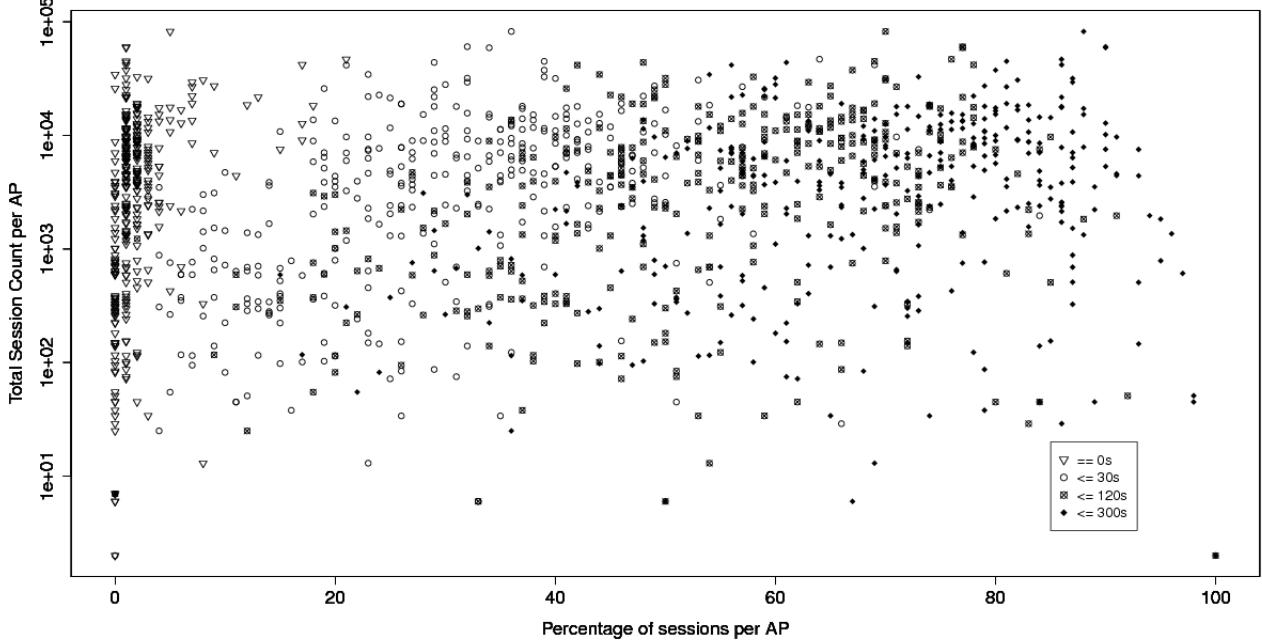


Fig. 5. AP short session percentage vs. AP session count

$$\begin{aligned}
 SS_{L1} &:= \{s \in S : s \text{ observed in } L \text{ and } \text{length}(s) == 0\} \\
 SS_{L2} &:= \{s \in S : s \text{ observed in } L \text{ and } \text{length}(s) \leq 30\} \\
 SS_{L3} &:= \{s \in S : s \text{ observed in } L \text{ and } 30 < \text{length}(s) < 120\} \\
 SSessions(L) &:= |(VSS_{L1}, VSS_{L2}, VSS_{L3})|
 \end{aligned}$$

A preliminary perspective of the variable relations and dependency is provided in Figure 6. This Figure illustrates pair-wise correlation of variables, both predictors and target variables, and shows how each predictor relate to different types of short sessions. Among the predictors, there is a high dependency between user count and sightings, as the correlation coefficient is very close to 1 (0.911), and then between user count and sessions (0.786).

The highest amount of correlation between the set of predictors and the target short sessions is observed between session count and two types of short sessions, below 30s (0.742) and between 30s and 2min (0.712). Hence, the number of short sessions are more probably to increase when the session count increment.

Another relation in this regard is the existence of rather high correlation between user count and short sessions between 30s and 2min (0.493). More users can provoke the congestion condition and difficulty of access to the shared medium, consequently the number of short sessions will excessively increase.

B. Supervised Model Construction

Based on the availability of data labels, learning techniques can operate in supervised, semi-supervised or unsupervised approaches. In the current work, supervised learning has been

utilized to predict the counts of short sessions and is achieved from two different directions (regressions and classification).

1) Regression Methods: The purpose of regression is to predict the value of one or more continuous *target* variables y given the value of a D -dimensional vector x of *input* variables. In this study, the target variable is the number of short sessions in three levels and we round the predicted continuous value of the target variable to estimate this quantity. We made use of linear and polynomial regression (up to degree of 9) for this purpose. In this work the predictor (x) and the target variable (y) are both a vector of integer-valued features: $x = \langle \text{user count, sightings, sessions} \rangle$ and $y = \langle \text{zero length, below 30s, between 30s and 2min} \rangle$.

a) Mis-Prediction Error: One of many ways of measuring the difference between values implied by an estimator and the true value of the quantity being estimated is Mean Squared Error (MSE). We check the accuracy of the provided regression model by comparing the MSE in the test set to simply measure the distance of the predicted value with the true ones and we call this difference mis-prediction error. In this context, MSE of a model shows the average squared distance of the predicted value for the short sessions with the real quantity. Lower MSEs denote better estimation of the target variable.

$$MSE = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2 \quad (1)$$

In equation 1, n is the total number of data instances, \hat{y} is the estimated value by the model and y is the true corresponding value.

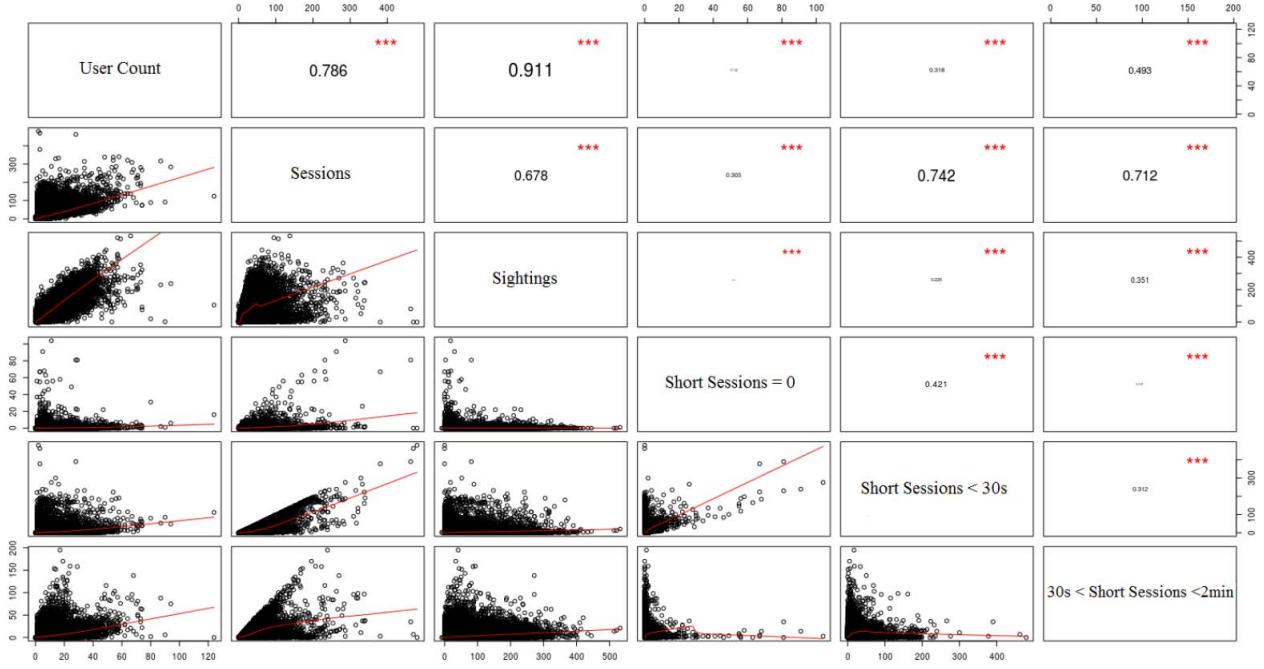


Fig. 6. Relation of feature specific metrics and the target variables

2) *Classification Methods:* In this section an analogous class of models for solving classification problems is dealt with. The goal in classification methods is to take an input vector x and assign it to one of K discrete classes C_k where $k = 1, \dots, K$. In the most common scenario, the classes are taken to be disjoint, so that each input is assigned to one and only one class. For this study, logistic regression has been selected as the classification technique. We provided 3 to 8 classes applying k-means clustering algorithm on the target feature transforming the count prediction to classification settings. This is achieved by vector quantization on the count output in a predefined number of type counts (using k-means). The outcome of the logistic regression represents the probabilities of the possible classes among which, one is more likely to describe the data instance in D -dimensional space.

a) *Mis-Classification Error:* In network context, recognition of the exact amount of target features is not of great importance in many cases. On the contrary, the network managers need to be aware of the range of safety or risk to conduct mitigation tasks or required adjustments. Having considered this, we performed clustering on the target feature to define various zones (e.g. low, medium, high). Distinguishing the anomalous zone, however, needs further assumptions and differs from context to context.

We measure the average of Euclidean distance between the estimated class and the true class. As the class labels are ordered for this problem, we penalize the wrong estimations by evaluating the distance between the predicted class and the initial one. In this case, as the labels do not inherently come with the data, like an expert knowledge, the process of labeling is assigning the cluster labels to the target variable.

TABLE III
AVERAGE MCE AND ACCURACY OF POLYNOMIAL REGRESSION VS.
LOGISTIC REGRESSION (ENTIRE DATASET FOR 3 TO 8 CLUSTERS)

Model	Zero Length Short Sessions	Short Sessions below 30s	Short Sessions between 30s and 2min	Total Accuracy
Polynomial Regression (degree:3)	0.0014	0.0131	0.0328	95.27%
Logistic Regression	0.0049	0.0318	0.0890	87.43%

$$MCE = \frac{1}{n} \sum_{i=1}^n (class(\hat{y}_i) - class(y_i))^2 \quad (2)$$

MCE or mis-classification error in equation 2 denotes the average distance of the predicted class of y (\hat{y}) from the true class of y . Note that k-means clustering algorithm do not provide the class labels in order, while our clusters inherently indicate an ordinal set. Hence, we assort the cluster centers and re-label the classes to evaluate the real distance of the predicted and primary classes.

VI. EXPERIMENTAL RESULTS

A. Single Model

In this case, the models are built based on the entire set of data concerning the statistics of all APs for 5 months of experiment. The motivation of building two models, single and AP-level, is to inspect whether there is a specific pattern in each AP or a single model is able to characterize all the APs with a same norm. In this experiment two sampling methods

TABLE I
MSE OF LINEAR REGRESSION VS. POLYNOMIAL REGRESSION (ENTIRE DATASET)

Model	Zero Length Short Sessions	Short Sessions below 30s	Short Sessions between 30s and 2min	Sum of Errors
Linear Regression (Hold-out)	0.1318	3.9003	2.6229	6.655
Polynomial Regression (Hold-out)	0.1317	2.324	2.0134	4.4691 (degree: 3)
Linear Regression (10-fold CV)	0.1605	4.0494	2.8288	7.0387
Polynomial Regression (10-fold CV)	0.1605	4.0494	2.8288	7.0387 (degree: 1)

TABLE II
MCE OF LOGISTIC REGRESSION FOR VARIOUS NUMBER OF CLUSTERS (ENTIRE DATASET)

Number of Clusters	Zero Length Short Sessions	Short Sessions below 30s	Short Sessions between 30s and 2min	Sum of Errors
3	0.0003	0.0058	0.0160	0.0221
4	0.0004	0.0061	0.0352	0.0417
5	0.0108	0.0186	0.0595	0.0889
6	0.0126	0.0308	0.0796	0.1230
7	0.0023	0.0577	0.1793	0.2393
8	0.0027	0.0717	0.1643	0.2387

have been used, hold-out and n-fold cross validation. In hold-out, the entire set of data is divided into train and test sets, so the statistics of two thirds of APs are randomly selected for training and the rest for testing. The model is then trained using the train data only and applied on the test data to predict the output values. To provide robust set of results, we repeated the hold-out method 4 times, and the final error is the average of 4 times modeling with this approach. In n-fold cross validation, the complete set of data is divided into k segments (or folds), each time one fold is selected as the test set and the others as training sets. The model is built based on the train sets and evaluated on the test set. The overall errors are calculated by taking the average over all iterations.

1) *Regression:* The results of the first set of experiments is shown in Table I for hold-out and 4-fold cross validation. Although the errors are quite similar in both cases, in hold-out experiment the overall error in polynomial regression method for three types of short sessions is less than the errors in linear Regression. Note that polynomial regression checks all the possible models from degree of 1 to 9 and the best result for this test is achieved at degree of 3. In the cross validation experiment, the best result of polynomial regression occurred at degree of 1 which is the same as linear regression. So in cross validation case no better result is achieved when the regression model is built with higher polynomial order.

2) *Classification:* The best classification results for the single model, is obtained in logistic regression for 3 clusters. The target value has been divided into a number of clusters (from 3 to 8), and the best outcome of each cluster (hold-out) has been presented in Table II. The initial number of clusters is considered 3 because 2 clusters do not really divide the data in a fair way. As the target feature is very sparse (mostly zero) and the short sessions only appear in special circumstances, dividing the data in two clusters will be highly biased (a large cluster of zeros and another cluster for the other values).

In Table II it can be seen that the sum of errors grows as the number of clusters increases, so to provide more reliable MCE

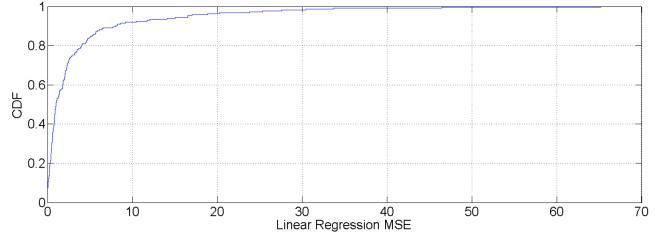


Fig. 7. Linear Regression CDF for the AP-level Model

of the classification method, we presented the final error as the average over all different number of clusters. Hence, the results provided in Table III are the average of mis-classification error for 3 to 8 clusters.

3) *Comparison between Best Results of Regression and Classification:* In order to compare the results of regression and classification, we need to simulate the prediction error of one method against the other one. Hence, we put the predicted values of the regression method in the same number of clusters obtained from the classification technique and check if the predicted values and the expected ones are located in the same cluster. In other words, we count the number of times the predicted value and the initial value are not in the same cluster as mis-classification error for the regression method. In such a way, the two mis-classification errors could be compared.

Table III shows that polynomial regression (at degree of 3) outperforms the logistic regression with accuracy of 95.27%. It means that 95.27% of the predicted values has remained in the same cluster with their corresponding true value changing the number of clusters from 3 to 8.

B. AP-level Models

Having built a single model of regression and classification for the entire dataset, the above experiments are reiterated in AP-level, as well. Note that, as the experiments in section IV affirm, not all the APs have the same, even similar usage patterns. For example there exist 135 APs with half of their

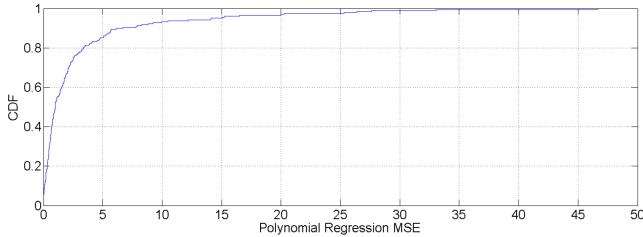


Fig. 8. Polynomial Regression CDF for the AP-level Model

sessions smaller than 2 minutes while there are just 31 APs with more than half of their sessions smaller than 30 seconds. Furthermore, as Figure 5 indicates, APs with higher percentage of short sessions are not always those with larger session counts. This inherent distinctiveness in the usage pattern of APs, conceives the idea of building separate models for each AP to inspect whether it is rational to make a more precise prediction in AP-level rather than the single model.

1) *Regression:* Figure 7 and 8 are the CDF of mis-prediction errors per AP. Figure 7 shows that more than 92% of APs obtain MSE lower than 10 in linear regression and the maximum MSE is 65. Figure 8 shows that the highest MSE in polynomial regression is 47 and almost 93% of APs get MSE less than 10. As the single model mis-prediction errors also indicate, polynomial regression at proper degree, can be a better option when predicting the target variable by regression. Note that in the AP-level approach, the most precise model is selected for each AP. For example in polynomial regression method, for some APs the best model achieved at degree of 3 and for some others at degree of 6. So in separate modeling, for each AP, the best set of parameters has been selected to estimate the target feature with highest precision.

In order to compare the single and AP-level models, the former is evaluated with the test set of each AP. CDF of Figure 9 shows the outcome of this assessment. In this Figure almost 92% of the APs get MSE lower than 10 while the maximum MSE is 155. Having compared the Figures 7 and 8 with Figure 9, the mis-prediction error of more than 92% of APs is below 10 in both single and AP-level models. But the maximum MSE in separate models is around 50 while in the single model it is almost three times greater (155). It shows that for the majority of APs it makes no difference whether the single or the AP-level model is employed. Inspecting the model data of these APs, it is conceived that they have very low probability of occurrence of short sessions. While for the remaining 8% of APs, with higher probability of short sessions, which are more critical for the network managers as well, the AP-level models are more accurate in estimating the target values causing much less mis-prediction errors.

2) *Classification:* Figure 10 shows the CDF of mis-classification errors per AP and Figure 11 shows the same when applying the single model on each AP test data. Comparing the single model and AP-level models in classification approach, we observed that almost 90% of APs in both models got 0.5 for the sum of mis-classification error. Note that the

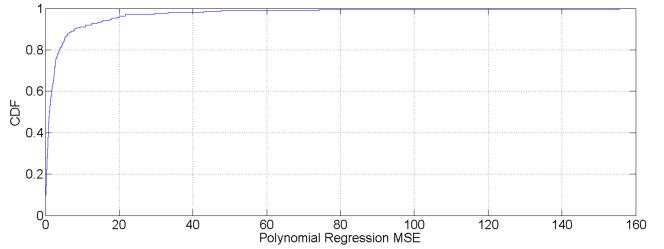


Fig. 9. Polynomial Regression CDF for the Single Model

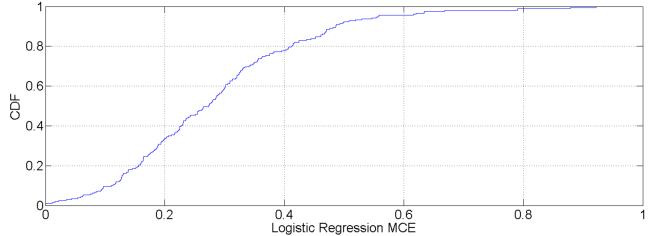


Fig. 10. Logistic Regression CDF for the AP-level Model

results in Figure 10 and 11 are depicted for 5 clusters. The maximum MCE for the single model is 1.4 while for the AP-level model is slightly higher than 0.9. So, although the differences of MCEs between these two models are trivial, AP-level model still achieves better results and therefore is recommended for predicting the classes of the short sessions.

VII. CONCLUSION AND FUTURE WORK

In this paper we focused on the short sessions in a deployed wireless network, as a network artifact and an indicator of the quality of wireless access which is normally overlooked in typical SNMP accounting. We applied machine learning techniques, namely regression and classification methods, to predict three types of short sessions (zero length, below 30s and between 30s and 2min) in RADIUS authentication data, based on the number of users, sightings and sessions per hour. We built a single model for the entire set of data and AP-level models specifically for each AP. In the single model, the best regression result is achieved in polynomial regression at degree of 3, and in logistic regression the average of classification accuracy for 3 to 8 classes is 87.44%. To make a comparison between results of regression and classification, we proposed a method simulating the regression errors and it is conceived that

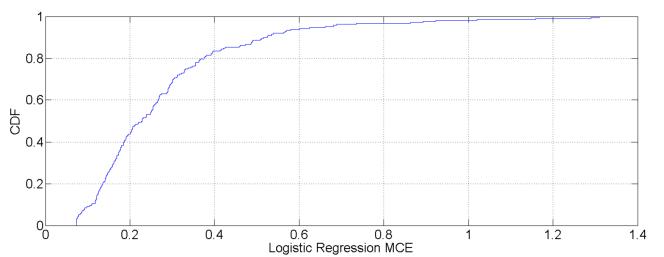


Fig. 11. Logistic Regression CDF for the Single Model

polynomial regression with accuracy of 95.27% outperforms the logistic regression.

Having built AP-level and single models, both in regression and classification techniques, we concluded that separate models per AP characterize the usage pattern and predict the short sessions more precisely. In particular, for APs with a higher probability of short sessions (and those that are more critical from the network management point of view), AP-level models are more accurate and provide better estimates for the number of short sessions.

An interesting direction of future work will be considering more complex models, Hidden Markov Models for example, to capture the evolving structure of APs regarding daily usage patterns in addition to make specific models for locations per hour. In a relevant line of work, similar APs could be clustered and one model could be built for each cluster to investigate whether the group model surpass the AP-level model in acquiring the underlying pattern of collections of APs.

Additionally, as the problem of this work was not exactly regression nor classification, another promising future line of research could be applying counting algorithms specifically for such sort of problems and making meaningful comparisons with regression and classification outcomes.

ACKNOWLEDGMENT

This work is financed by the ERDF European Regional Development Fund through the COMPETE Programme (operational programme for competitiveness) and by National Funds through the FCT Fundao para a Ciéncia e a Tecnologia (Portuguese Foundation for Science and Technology) within project Place Characterization from Sensing and Acting/Reference EXPL/EEI-SII/0902/2012 (FCOMP-01-0124-FEDER-029831)

REFERENCES

- [1] (2013) A community resource for archiving wireless data at dartmouth. [Online]. Available: <http://crawdad.cs.dartmouth.edu/>
- [2] (2013) Rfc 2139 radius authentication. [Online]. Available: <http://tools.ietf.org/html/rfc2139>
- [3] M. Afanasyev, T. Chen, G. Voelker, and A. Snoeren, "Usage patterns in an urban wifi network," *Networking, IEEE/ACM Transactions on*, vol. 18, no. 5, pp. 1359 –1372, oct. 2010.
- [4] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan, "Characterizing user behavior and network performance in a public wireless LAN," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 30, 2002, p. 195205. [Online]. Available: <http://dl.acm.org/citation.cfm?id=511359>
- [5] Y.-C. Chen, E. Rosensweig, J. Kurose, and D. Towsley, "Group detection in mobility traces," in *Proceedings of the 6th International Wireless Communications and Mobile Computing Conference*, ser. IWCMC '10. New York, NY, USA: ACM, 2010, pp. 875–879. [Online]. Available: <http://doi.acm.org/10.1145/1815396.1815597>
- [6] F. Chinchilla, M. Lindsey, and M. Papadopoulou, "Analysis of wireless information locality and association patterns in a campus," in *INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 2, 2004, pp. 906–917 vol.2.
- [7] M. A. Ergin, K. Ramachandran, and M. Gruteser, "An experimental study of inter-cell interference effects on system performance in unplanned wireless lan deployments," *Computer Networks*, vol. 52, no. 14, pp. 2728 – 2744, 2008.
- [8] T. Henderson, D. Kotz, and I. Abyzov, "The changing usage of a mature campus-wide wireless network," *Computer Networks*, vol. 52, no. 14, p. 26902712, 2008.
- [9] F. Hernández-Campos, M. Karaliopoulos, M. Papadopoulou, and H. Shen, "Spatio-temporal modeling of traffic workload in a campus wlan," in *Proceedings of the 2nd annual international workshop on Wireless internet*, ser. WICON '06. New York, NY, USA: ACM, 2006. [Online]. Available: <http://doi.acm.org/10.1145/1234161.1234162>
- [10] F. Hernandez-Campos and M. Papadopoulou, "A comparative measurement study the workload of wireless access points in campus networks," in *Personal, Indoor and Mobile Radio Communications, 2005. PIMRC 2005. IEEE 16th International Symposium on*, vol. 3, 2005, pp. 1776–1780 Vol. 3.
- [11] W.-j. Hsu, D. Dutta, and A. Helmy, "Mining behavioral groups in large wireless lans," in *Proceedings of the 13th annual ACM international conference on Mobile computing and networking*, ser. MobiCom '07. New York, NY, USA: ACM, 2007, pp. 338–341. [Online]. Available: <http://doi.acm.org/10.1145/1287853.1287899>
- [12] W. jen Hsu and A. Helmy, "On nodal encounter patterns in wireless lan traces," *Mobile Computing, IEEE Transactions on*, vol. 9, no. 11, pp. 1563–1577, 2010.
- [13] U. Kumar and A. Helmy, "Extract: mining social features from wlan traces—a gender-based case study," in *Proceedings of the 13th ACM international conference on Modeling, analysis, and simulation of wireless and mobile systems*, ser. MSWIM '10. New York, NY, USA: ACM, 2010, pp. 240–247. [Online]. Available: <http://doi.acm.org/10.1145/1868521.1868561>
- [14] A. Mahanti, C. Williamson, and M. Arlitt, "Remote analysis of a distributed wlan using passive wireless-side measurement," *Performance Evaluation*, vol. 64, pp. 909 – 932, 2007. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0166531607000545>
- [15] D. Massa and R. Morla, "Abrupt ending of 802.11 ap connections," in *Proceedings of the 18th IEEE Symposium on Computers and Communications*, ser. ISCC '13, 2013.
- [16] M. Papadopoulou, H. Shen, and M. Spanakis, "Characterizing the duration and association patterns of wireless access in a campus," in *Wireless Conference 2005 - Next Generation Wireless and Mobile Communications and Services (European Wireless), 11th European*, 2005, pp. 1–7.