

Behavioural biometrics using electricity load profiles

M. Bicego, F. Recchia, A. Farinelli
University of Verona
Italy

S. D. Ramchurn
University of Southampton
UK

E. Grosso
University of Sassari
Italy

Abstract—Modelling behavioural biometric patterns is a key issue for modern user centric applications, aimed at better monitoring users' activities, understanding their habits and detecting their identity. Following this trend, this paper investigates whether the electrical energy consumption of a user can be a distinctive behavioural biometric trait. In particular we analyse daily and weekly load profiles showing that they are closely related to the identity of the users. Hence, we believe that this level of analysis can open interesting application scenarios in the field of energy management and it provides a good working framework for the continuous development of smart environments with demonstrable benefits on real-world implementations.

I. INTRODUCTION

Biometric technology [1] is a powerful tool to identify individuals upon some distinctive physiological characteristics, such as face, fingerprints or iris patterns [2], [3], [4] or even EEG signals, footprints, ears, and many others [5], [6].

Recently, a number of biometric approaches have gone beyond the typical physiological characteristics of the users and consider also behavioral aspects, such as gait [7], signature [8], or even HCI-related behaviors [9], such as keystrokes, mouse dynamics [10], [11] or Internet browsing histories [12]. All these approaches fall under the context of *behavioral* biometrics [9] and provide key tools to build accurate models of the users. Such models are particularly important when considering the growing popularity of ubiquitous, context-aware computing and user-centric application developments, which aim at improving and personalizing existing services to specific users' needs [13].

Against this background, in this paper, we analyse and assess the suitability of a novel trait: the electrical energy consumption profile of a user (or load profile), i.e., the users' consumption of electrical energy over a given period. In fact, the recent large scale deployment of smart metering infrastructure coupled with the energy market liberalization in many countries around the world, result in an increasing availability of fine-grained electrical consumption data and in an increasing interest in analysing users' load profiles. Consequently, analysts have begun to use these data with the aim of discovering and categorizing groups of users, to optimize energy allocation and purchasing, or to help retailers designing new pricing models for implementing more accurate demand and supply profiles [14], [15], [16], [17]. For example, in [14] authors propose a framework to characterize users based on their energy consumption so as to provide decision support tools for defining adequate market strategies. In this context, users' load profile are typically represented by a set of features called load shape indexes (e.g., Load Factor, Lunch Impact and Night Impact) [14], [16]. In contrast, other approaches [15],

[17] characterize the load profile by considering each profile as time series of load measurements, and employ different unsupervised learning techniques to cluster data (such as Self Organizing Maps, K-means and Hidden Markov Models).

Here, we take a different perspective with respect to such previous work, and investigate whether load profiles can be considered a distinctive biometric trait, hence allowing the identification of a specific user. To the best of our knowledge, this has never been investigated in the literature, and this paper is a first step in this direction.

In more detail, our goal is to identify a specific user (i.e., an electricity meter) among several users, considering the users' load profiles (i.e., a series of consumption data, acquired at fixed interval during a day), over a given period of time (i.e., one month of data in our case). To do so, we investigate different representations for the load profiles, including raw measurements, frequency characterizations and typical load shape indexes [14], [16]. Based on these representations we experiment two classification approaches. The former, based on the classical Nearest Neighbor rule, assigns an unknown object to the class of the nearest neighbor point; the second is a classical Bayesian scheme, based on Hidden Markov Models (HMM – [18]), a probabilistic tool widely used to characterize sequential data, which has been recently applied also for load profiles grouping (in particular to characterise relationships between electricity consumption and consumers' preferences or behaviors – e.g., lifestyle – [17]).

We empirically evaluate our approach on a database of real energy consumption data acquired in the UK over a single continuous month, considering up to 400 meters. Our results suggest that the energy load profile is indeed a promising behavioural trait: in a problem with 100 users, our best classification scheme can correctly identify a given user within the first 20 answers of the system in 89.36% of the cases. Such rate significantly increases if we aggregate more days to characterize a given user, reaching 94.5% when using 7 days. This is even more interesting considering that a random classifier would have a recognition rate of 20% in such setting. Finally, we also assess the scalability of the method by showing that the performances of our classification schemes do not degrade when increasing the number of users (up to 400 in our experiments).

Notice that the data acquired by a single meter can incorporate the presence of more people (i.e., all the inhabitants of the households that the meter is monitoring). Hence, we are not strictly identifying the behaviour of one specific person (as in most biometric approaches), but rather the behaviour of a group of people. Nevertheless, this is a natural consequence of the peculiar type of data that we are considering (i.e.,

load profiles) which naturally leads to a mapping between a meter and a user. In fact, this mapping has been used in previous work on energy profile segmentation [17] and is the typical mapping performed by utility companies, which are not interested in the composition of the group that is served by a specific meter. Moreover, we believe that this level of analysis (i.e., considering a group of people living together as a single user) leads to interesting application scenarios as it allows to track and monitor the coherence of the behaviours of a user: if we can accurately characterize the behaviour of a group of inhabitants based on their energy consumption, we can detect drastic changes in such behaviours (e.g., due to a change in the group of people living in the households).

The rest of the paper is organized as follows: Section II describes our methodology for user recognition, detailing the different types of representation and the classification schemes we adopt. Section III discusses our empirical evaluation of the method and Section IV concludes the paper providing suggestions and directions for future investigations.

II. PROPOSED METHODOLOGY

In this section we detail our methodology for user identification. The starting point is the “electrical load profile”, which represents the energy consumption of each user throughout a day (every half hour in our experimental setting). Hence each profile is a vector of T elements (where $T = 48$ in our case), see Figure 1 for a visual representation of a series of load profiles.

Such profile can be characterized in different ways. In particular here we employed three representations based on time, frequency and load shape indexes. Given these representations, we then derive two different classification schemes, one based on the nearest neighbor rule and the other based on the HMM plus the Bayesian rule. In this last case, since HMM are models usable for sequential data, we restricted their employment to the sole time representation.

A. Load profile representations

Time representation. We use the raw signal, namely the T -dimensional vector as given by the acquisition process. This is a rather common choice in energy profile clustering – e.g. [15], [17]. We also applied a small level of Gaussian smoothing, in order to remove some fluctuations due to the sampling intervals. We decided not to normalize the signals (e.g., with z-score normalization), in order to maintain all possible information on the energy of the signal (e.g., the absolute scale of the consumption profile) which might be useful to differentiate between different users¹.

Frequency representation. We apply a Fast Fourier Transform to the original signal (smoothed or not), in order to investigate whether the frequential content of the signal may be useful to discriminate between users. After a careful evaluation of the frequency content of the profiles, only the first 10 coefficients have been retained.

¹Actually, we empirically evaluated different types of normalization, including the z-score normalization; however in all our experiments (not shown here), the normalization schemes we tried decreased the performance of the classification approach.

Normalized load shape indexes. These features represent a classical set of values directly extracted from the load profile: they have been shown to be adequate to characterize groups of users [14], [19]; since they are widely used in many studies, we decided to include them in our work, even if in our case the task is different – i.e. classification rather than clustering. In particular, in our work we employed the selection made in [14], where the following three indexes were used:

1) Load Factor:

$$LF = \frac{P_{av,day}}{P_{max,day}} \quad (1)$$

where $P_{av,day}$ is the average of the profile during the whole day, whereas $P_{max,day}$ represents the highest peak.

2) Night Impact:

$$NI = \frac{1}{3} \frac{P_{av,night}}{P_{av,day}} \quad (2)$$

where $P_{av,night}$ is the average of the profile during the night, assumed to last for 8 hours from 11pm to 7am.

3) Lunch Impact:

$$LI = \frac{1}{8} \frac{P_{av,lunch}}{P_{av,day}} \quad (3)$$

where $P_{av,lunch}$ is the average of the profile during the lunch time, assumed to last for 3 hours from 11.30am to 2.30pm.

Such features are computed for every profile, each one thus represented as a 3 dimensional feature vector.

B. Classification

The main goal of our study is to assess the potentials of the electrical load profile as a biometric trait for a given user. Specifically, in this first study, we aim at assessing identification performances (i.e. recognition), rather than authentication (i.e. verification). Therefore, we need to design a classification scheme, which, given an unknown profile, is able to assign it to one over a set of predetermined users. While the relevant literature offers a wide selection of classification schemes, since our primary focus is to investigate the use of load profiles as a biometric trait, we decided to use two basic, well known approaches for classification: the former is based on the classical Nearest Neighbor rule, which assigns an unknown object to the class of the nearest neighbor point. The latter is a classical Bayesian scheme, based on Hidden Markov Models (HMM – [18]): in the training phase, a single HMM is trained for every user, using all its training profiles; in the classification case, an unknown profile is assigned to the user whose HMM shows the highest likelihood on such profile.

In what follows we further describe such classification schemes and their application for load profile classification.

Nearest Neighbor Scheme. As described before, the classifier assigns an object to the class of its nearest neighbor; the definition of a proper proximity measure (either similarity or dissimilarity) is therefore crucial. In the following we

will introduce the ones used in our study, together with the motivations.

The first two distances that we employ are the L1 and the L2 norms (i.e. the Manhattan and the Euclidean distances, respectively). These are very common distances, widely applied in the Pattern Recognition community, and also specifically in the energy domains (as basis for clustering, as in [15]). Given two profile representations $\mathbf{x} = [x_1, \dots, x_n]$ and $\mathbf{y} = [y_1, \dots, y_n]$, they are defined as follows:

L1 norm (Manhattan distance):

$$L1(\mathbf{x}, \mathbf{y}) = \sum_i |x_i - y_i| \quad (4)$$

L2 norm (Euclidean distance):

$$L2(\mathbf{x}, \mathbf{y}) = \sqrt{\sum_i (x_i - y_i)^2} \quad (5)$$

While these two distances make sense for all the representations, we also investigate some other measures specifically designed for the time signal representation, based on the concept of signal correlation [20], a standard and well applied method to compare time series. The first is the standard Zero Lag Cross Correlation, which, given two profiles $\mathbf{p} = p_1, \dots, p_T$ and $\mathbf{q} = q_1, \dots, q_T$, is defined as:

Zero Lag Cross Correlation:

$$CC_0(\mathbf{p}, \mathbf{q}) = \frac{\sum_i (p_i \cdot q_i)}{\|\mathbf{p}\| \|\mathbf{q}\|} \quad (6)$$

The second class of correlation measures is based on the assumption that activities for a user in different days may not be completely overlapped, but there can be a small lag (consider for example when people have dinner). In order to capture this behaviour, we compute the cross correlation measure by allowing 1 and 2 time steps of lag, retaining at the end the maximum of the correlations. This reasoning leads to the following two measures:

Max (1-lag) Cross Correlation:

$$CC_1(\mathbf{p}, \mathbf{q}) = \max_{m \in \{-1, 0, 1\}} \frac{\sum_i (p_i \cdot (q_i + m))}{\|\mathbf{p}\| \|\mathbf{q}\|} \quad (7)$$

Max (2-lag) Cross Correlation:

$$CC_2(\mathbf{p}, \mathbf{q}) = \max_{m \in \{-2, -1, 0, 1, 2\}} \frac{\sum_i (p_i \cdot (q_i + m))}{\|\mathbf{p}\| \|\mathbf{q}\|} \quad (8)$$

Finally, we also tried to model the fact the lags displacement in the daily activities can be differently displaced over the day: in order to capture this behaviour we repeated the computation of the Max 2-lag cross correlation (as defined before) in small overlapping windows of 4 hours (with an overlap of two hours) – this allows to best align different parts of the day – taking at the end the mean or the max value. In formula:

Mean Windows Max (2-lag) Cross Correlation:

$$CC_{MeanW2}(\mathbf{p}, \mathbf{q}) = \frac{1}{z} \sum_{i=1}^z CC_2(\mathbf{w}_i(\mathbf{p}), \mathbf{w}_i(\mathbf{q})) \quad (9)$$

Max Windows Max (2-lag) Cross Correlation:

$$CC_{MaxW2}(\mathbf{p}, \mathbf{q}) = \max_{i \in \{1, \dots, z\}} CC_2(\mathbf{w}_i(\mathbf{p}), \mathbf{w}_i(\mathbf{q})) \quad (10)$$

where $\mathbf{w}_1(\cdot), \dots, \mathbf{w}_z(\cdot)$ are all possible overlapped sub-windows lasting 4 hours extracted for a given profile.

Hidden Markov Model-based Bayesian classification. This scheme is a classical Bayesian scheme, based on Hidden Markov Models, a probabilistic approach whose usefulness has been successfully applied in many different pattern recognition scenarios. More in detail, a discrete-time first order HMM [18] is a probabilistic model that describes a stochastic sequence² $\mathbf{O} = (O_1, O_2, \dots, O_T)$ as being an indirect observation of a hidden Markovian random sequence of states $\mathbf{Q} = (Q_1, Q_2, \dots, Q_T)$, where for $t = 1, \dots, T$, $Q_t \in \{1, 2, \dots, N\}$ (the set of states). Each state has an associated probability function that specifies the probability of observing each possible symbol, given the state. A HMM is thus fully specified by a set of parameters $\lambda = \{\mathbf{A}, \mathbf{B}, \boldsymbol{\pi}\}$ where $\mathbf{A} = (a_{ij})$ is the transition matrix, i.e., $a_{ij} = P(Q_t = j | Q_{t-1} = i)$; $\boldsymbol{\pi} = (\pi_i)$ is the initial state probability distribution, i.e., $\pi_i = P(Q_1 = i)$, and $\mathbf{B} = (b_i)$ is the set of emission probability functions. In our case, since the observations are continuous, we assume that each b_i is a Gaussian probability density function (leading to the so called Continuous Gaussian HMM).

The training of the model, given a set of sequences $\{\mathbf{o}^{(i)}\}$, is usually performed using the standard Baum-Welch re-estimation technique [18], which determines the parameters $(\mathbf{A}, \mathbf{B}, \boldsymbol{\pi})$ that maximize the probability $P(\{\mathbf{o}^{(i)}\} | \lambda)$. The evaluation step, i.e. the computation of the log probability $\log P(\mathbf{o} | \lambda)$, given a model λ and a sequence \mathbf{o} to be evaluated, is performed using the *forward-backward procedure* [18].

Given a C -class problem, the standard classification scheme (i.e. the Bayes rule) is realized in the following way: for every class c , a HMM λ_c is trained, using only the training sequences belonging to such class – obtaining at the end the set of C models $\lambda_1, \dots, \lambda_C$. Then, in the testing phase, an unknown sequence $\mathbf{o} = (o_1, \dots, o_T)$ is assigned to the class whose model shows the highest likelihood (assigning to each class the same prior probability), namely the label $\ell(\mathbf{o})$ is determined as

$$\ell(\mathbf{o}) = \arg \max_c \log P(\mathbf{o} | \lambda_c) \quad (11)$$

III. EXPERIMENTAL EVALUATION

In this section we describe our experimental evaluation. First, we introduce the experimental details. Then, a large set of experiments are presented, analysing the different representations and the different classification schemes. Subsequently, an analysis of the scaling capabilities of the methods will be described. A final experiment, aimed at investigating the possible aggregation of more days, is finally proposed.

²We adopt the common convention of writing stochastic variables with upper case and realizations thereof in lower case.

A. Empirical details

The dataset is derived from data typically employed in the collective energy purchasing domain, where energy consumers form groups to purchase energy at better prices [21]. In more detail, a profile records the energy consumption of a household at fixed intervals (every half hour in our case) collected, over a month in 2009, from different households in UK. Please note that all profiles are recorded during the same month, in order to avoid the possible presence of discriminative information derived from seasonality and not from person-specific characteristics. Examples of profiles (from different users and relative to different days) are displayed in Fig. 1.




Fig. 1. Some profiles: every column contains profiles from a different user; every row contains a different day.

In all experiments, in order to maintain separated training and testing sets, we used the first 14 days for training the system, testing it with the remaining 15 days. Instead of simply computing classification errors, we computed the more expressive Cumulative Match Curve (CMC), a common performance measure in the field of biometrics [22]: given a test profile, we compute its proximity to all the users. The curve then tells the rate at which the correct user is found within the first k matches (namely within the first k nearest users), with all possible k spanned on the x-axis. An useful measure that can be extracted from the CMC is the normalized Area Under the Curve (nAUC), which, as in the ROC curve, represents an aggregate measure able to express how good is the recognition rate (the higher the better).

In all the experiments HMM training has been performed using the standard Baum-Welch procedure, stopping after likelihood convergence. Initialization has been carried out, as in many applications involving Continuous Gaussian HMM, with a clustering based on Gaussian mixture models. The number of states, which represents the free parameter of the model, has been set in an automatic way starting only from the training data, following the procedure described in [23].

B. Experiment 1: different variants of the scheme

As a first set of experiments we analysed a dataset containing 100 different users, investigating the different representations and the different variants of the two classification schemes. Here our main goal is to investigate the best configuration for both schemes, to be used in the subsequent

TABLE I. RESULTS ON THE DATASET WITH 100 SUBJECTS

Nearest Neighbor Classifiers		
Representation	Classifier	norm AUC
Time Rep.	L1 + Nearest Neighbor	0.855
Time Rep.	L2 + Nearest Neighbor	0.798
Time Rep.	CC_0 + Nearest Neighbor	0.725
Time Rep.	CC_1 + Nearest Neighbor	0.752
Time Rep.	CC_2 + Nearest Neighbor	0.755
Time Rep.	CC_{MeanW2} + Nearest Neighbor	0.780
Time Rep.	CC_{MaxW2} + Nearest Neighbor	0.699
Smoothed Time Rep.	L1 + Nearest Neighbor	0.868
Smoothed Time Rep.	L2 + Nearest Neighbor	0.839
Smoothed Time Rep.	CC_0 + Nearest Neighbor	0.759
Smoothed Time Rep.	CC_1 + Nearest Neighbor	0.766
Smoothed Time Rep.	CC_2 + Nearest Neighbor	0.763
Smoothed Time Rep.	CC_{MeanW2} + Nearest Neighbor	0.781
Smoothed Time Rep.	CC_{MaxW2} + Nearest Neighbor	0.749
Frequency Rep.	L1 + Nearest Neighbor	0.848
Frequency Rep.	L2 + Nearest Neighbor	0.845
Load Shape Indexes	L1 + Nearest Neighbor	0.722
Load Shape Indexes	L2 + Nearest Neighbor	0.722
HMM based Bayesian classifiers		
Representation	Classifier	norm AUC
Time Rep.	HMM + Bayes Rule	0.927
Smoothed Time Rep.	HMM + Bayes Rule	0.905

experiments. As for the smoothing of the raw load profile (when used), we employ a simple Gaussian filtering, with sigma varying in the range [0.6 - 2.2]. We compute the CMC curves for all the parametrizations and all classification schemes, reporting the normalized Area Under the Curve (nAUC) in table I. As a general comment, we can observe that the system works reasonably well, reaching in the best case an nAUC of 0.927 – quite high nAUC if we consider that we are dealing with a behavioral biometrical trait. Concerning the two classification approaches, it seems evident that the Bayesian scheme outperforms the nearest neighbor rule: in this case, the learning phase present in the HMM scheme seems to be essential to capture and model the unique characteristics of every user. Nonetheless, also the NN rule works quite well, especially with L1 and L2 measures for the Time and the Frequency representations. On the contrary, the correlation based measures do not work as well as the others: probably the flexibility introduced by the different versions captures more the variations between classes rather than the intraclass one.

Concerning representations, we have to notice that Load Shape Indexes represent a poor choice: it seems evident that this representation, which was adequate for clustering and general data-mining [14], [19], is not very informative when trying to identify the differences between users. As a final remark, it is important to observe that smoothing the signal is almost always beneficial in all the versions of the Nearest Neighbor rule, whereas it is not for the HMM scheme.

C. Experiment 2: scalability

In this section we tried to investigate whether the proposed strategies scale well with the number of users, i.e. if the performances do not significantly degrade when increasing the dataset size. In particular we selected from the results of the previous experiment the best configurations for both

TABLE II. RESULTS ON THE DATASET WITH INCREASING NUMBER OF SUBJECTS.

Classes	Nearest Neighbor Scheme	HMM Bayesian Scheme
100	0.868	0.927
200	0.866	0.931
300	0.860	0.932
400	0.860	0.929

TABLE III. RESULTS ON THE DATASET WITH INCREASING NUMBER OF DAYS CONSIDERED FOR EVERY SUBJECT

Days	Scheme	100 classes		200 classes	
		Aver.	Concat.	Aver.	Concat.
1	L1 + NN	0.868	0.868	0.866	0.866
	HMM	0.927	0.927	0.931	0.931
2	L1 + NN	0.893	0.876	0.888	0.872
	HMM-Bayes	0.906	0.945	0.905	0.948
3	L1 + NN	0.904	0.872	0.901	0.869
	HMM-Bayes	0.901	0.953	0.905	0.960
4	L1 + NN	0.921	0.886	0.915	0.878
	HMM-Bayes	0.893	0.957	0.897	0.965
5	L1 + NN	0.927	0.875	0.920	0.866
	HMM-Bayes	0.905	0.960	0.902	0.967
6	L1 + NN	0.939	0.872	0.937	0.875
	HMM-Bayes	0.900	0.971	0.904	0.972
7	L1 + NN	0.941	0.892	0.933	0.877
	HMM-Bayes	0.910	0.963	0.901	0.970

the classification schemes: L1 norm on the Smoothed Time Representation (for the nearest neighbor scheme) and the HMM scheme on the original Time Representation. Then we tested the two configurations with 200, 300 and 400 subjects. Results are reported in Table II. For sake of clarity we also reported the results obtained with 100 subjects. As can be noted from the table, the performance of the classification schemes do not vary significantly when increasing the database size.

D. Experiment 3: enlarging the biometric trait

The main goal for this experiment is to investigate the possibility of aggregating more days to characterize every user: in particular we perform experiments by considering as a biometrical trait X consecutive days (with X ranging from 2 to 7). The rationale behind this is that energy consumption can significantly vary across different days of the week (e.g., week-days or week-ends). Given X days, we consider two simple ways to aggregate them: averaging them, thus obtaining again a T -dimensional profile, and concatenating them, obtaining a X times T long load profile. In the first case we are removing some noise / attenuating intraclass variation, while in the second case we can consider a richer set of (possibly noisy) data. To run the experiments we used the dataset with 100 subjects and the one with 200 subjects selecting, as in the previous section, the best configuration as suggested by the previous experiments: the nearest neighbor scheme with the L1 norm on the Smoothed Time Representation and the HMM scheme on the original Time Representation. Results are reported in Table III (for sake of clarity we also reported results with 1 day). As we can see, there is a beneficial impact in the performances obtained when aggregating more days. In particular, with 100 users, the nAUC of Nearest Neighbor scheme increases from

0.868 to 0.941, whereas it increases from 0.927 to 0.971 for the HMM-based scheme. In the former case, the best improvement is obtained while averaging different signals, whereas in the second case the best improvement is obtained while concatenating. This is somehow expected: the NN rule strongly depends on the goodness of single instances (since the comparison is pairwise), which is possibly enriched and improved via the averaging operation – some instance specific noise may be removed; the HMM-scheme, on the contrary, is able to build a model by simultaneously considering all the training sequences, thus being less sensitive to the noise of the single profile. On the contrary, such scheme can suffer from the lower cardinality of the training set which is obtained when averaging: this is confirmed by the experiments, where the concatenation approach largely outperforms the averaging scheme.

To provide a clearer idea of the performances of the proposed approach, we reported in Fig. 2 the CMC curves for the different configurations, using the averaging scheme and the L1 distance for the Nearest Neighbor rule and the concatenation scheme for the HMM-based approach. From the figure it is possible to note the full potentials of the proposed trait: using just one day, in the 89.4% of the cases we can correctly identify a given subject in the first 20 answers of the HMM-based rule. Such rate increases if we aggregate more days, in particular reaching 94.5% when using 7 days. This is even more interesting considering that a random classifier would have a recognition rate of 0.2. The same interesting behaviour can be observed for the experiment with 200 classes (here the random classifier would have a recognition rate of 0.1).

IV. CONCLUSIONS AND FUTURE WORK

This paper investigated the possibility of identifying users by monitoring their load profiles. The analysis was carried out over a significant data set of 400 users, by applying different types of representation and two different classification schemes: one based on the Nearest Neighbour rule with different proximity measures and one based on Hidden Markov Models. The results clearly show that load profiles can be considered as a distinctive biometric behavioural trait³. In our view, this is a very encouraging result that opens up several possibilities for further investigations of such novel distinctive trait, not only in biometrics but more in general for user centric applications.

REFERENCES

- [1] A. Jain, L. Hong, and S. Pankanti, "Biometric identification," *Commun. ACM*, vol. 43, no. 2, pp. 90–98, Feb. 2000.
- [2] S. Li and A. Jain, *Handbook of Face Recognition*. Springer, 2005.
- [3] D. Maltoni, D. Maio, A. Jain, and S. Prabhakar, *Handbook of Fingerprint recognition*, 2nd ed. Springer, 2009.
- [4] M. Burge and K. Bowyer, *Handbook of Iris Recognition*. Springer, 2013.
- [5] A. Abaza, A. Ross, C. Hebert, M. Harrison, and M. Nixon, "A survey on ear biometrics," *ACM Comput. Surv.*, vol. 45, no. 2, pp. 22:1–22:35, 2013.

³Actually, it would be really interesting to investigate how robust is the proposed approach with respect to the seasonability – is a model trained on Spring profiles usable also for Summer profiles? Unfortunately, at the moment we do not have data permitting such evaluation.




Fig. 2. CMC curves for 100 (left column) and 200 (right column) subjects. The vertical dotted line indicates the 20-rank recognition rate. (first row): the Nearest Neighbor rule; (second row): the HMM scheme.

- [6] A. Riera, M. Soria-Frisch, C. Caparrini, C. Grau, and G. Ruffini, "Unobtrusive biometric system based on electroencephalogram analysis," *EURASIP Journal on Advances in Signal Processing*, vol. 1, pp. 18–25, 2008.
- [7] M. Nixon, T. Tan, and R. Chellappa, *Human Identification Based on Gait*. Springer-Verlag, New York Inc., 2006.
- [8] G. Bailador, C. Sanchez-Avila, J. Guerra-Casanova, and A. de Santos Sierra, "Analysis of pattern recognition techniques for in-air signature biometrics," *Pattern Recognition*, vol. 44, no. 1011, pp. 2468 – 2478, 2011.
- [9] R. Yampolskiy and V. Govindaraju, "Behavioural biometrics: a survey and classification," *Int. J. Biometrics*, vol. 1, no. 1, pp. 81–113, 2008.
- [10] F. Bergadano, D. Gunetti, and C. Picardi, "User authentication through keystroke dynamics," *ACM Trans. on Inf. Syst. Secur.*, vol. 5, no. 4, pp. 367–397, Nov. 2002.
- [11] M. Pusara and C. Brodley, "User re-authentication via mouse movements," in *ACM workshop on Visualization and data mining for computer security*. ACM, 2004, pp. 1–8.
- [12] P. Eckersley, "How unique is your web browser?" in *Proc. Privacy Enhancing Technologies*, 2010, p. 118.
- [13] P. Rashidi and D. Cook, "Keeping the resident in the loop: adapting the smart home to the user," *IEEE Trans. on Systems, Man and Cybernetics, Part A: Systems and Humans*, vol. 39, no. 5, p. 949959, 2009.
- [14] V. Figueiredo, F. Rodrigues, Z. Vale, and J. Gouveia, "An electric energy consumer characterization framework based on data mining techniques," *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 596–602, 2005.
- [15] T. Zhang, G. Zhang, J. Lu, X. Feng, and W. Yang, "A new index and classification approach for load pattern analysis of large electricity customers," *IEEE Transactions on Power Systems*, vol. 27, no. 1, pp. 153–160, 2012.
- [16] G. Chicco, R. Napoli, P. Postolache, M. Scutariu, and C. Toader, "Customer characterization options for improving the tariff offer," *Power Systems, IEEE Transactions on*, vol. 18, no. 1, pp. 381–387, 2003.
- [17] A. Albert and R. Rajagopal, "Smart meter driven segmentation: What your consumption says about you," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4019–4030, 2013.
- [18] L. Rabiner, "A tutorial on Hidden Markov Models and selected applications in speech recognition," *Proc. of IEEE*, vol. 77, no. 2, pp. 257–286, 1989.
- [19] G. Chicco, O.-M. Ionel, and R. Porumb, "Electrical load pattern grouping based on centroid model with ant colony clustering," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 1706–1715, 2013.
- [20] S. Orfanidis, *Optimum Signal Processing: an Introduction*, 2nd ed. Prentice-Hall, 1996.
- [21] M. Vinyals, F. Bistaffa, A. Farinelli, and A. Rogers, "Coalitional energy purchasing in the smart grid," in *Proc. of IEEE Int. Energy Conference and Exhibition (ENERGYCON 12)*, 2012, pp. 848 –853.
- [22] H. Moon and P. Phillips, "Computational and performance aspects of pca-based face-recognition algorithms," *Perception*, vol. 30, pp. 303 – 321, 2001.
- [23] M. Bicego, V. Murino, and M. Figueiredo, "A sequential pruning strategy for the selection of the number of states in Hidden Markov Models," *Pattern Recognition Letters*, vol. 24, no. 9–10, pp. 1395–1407, 2003.