

AN ONLINE FEATURE SELECTION ARCHITECTURE FOR HUMAN ACTIVITY RECOGNITION

Katerina Karagiannaki^{*†} *Athanasia Panousopoulou*^{*} *Panagiotis Tsakalides*^{*†}

^{*} Institute of Computer Science, Foundation for Research and Technology-Hellas, Greece

[†] Department of Computer Science, University of Crete, Greece

{karayan, apanouso, tsakalid}@ics.forth.gr

ABSTRACT

Human Activity Recognition (HAR) must currently face up to the challenge of rethinking analytics from the perspective of real-time operation, wherein biophysical sensing streams are efficiently intertwined at close vicinity to the point of sensing. As such, feature selection techniques, traditionally employed for off-line data processing, should be evaluated with respect to their ability to filter out redundant information in real-time. In this work, we propose an online architecture for implementing feature selection on mobile devices, and we evaluate popular feature selection methods against constantly alternating activity labels. We perform a qualitative analysis to determine the dominant sensing modality that dictates the activities of a certain time duration. The results indicate that online feature selection performance changes among consecutive data partitions, leading to the conclusion that the type of available activity influences significantly the feature selection procedure.

Index Terms— Feature Selection, Human Activity Recognition, Ubiquitous Computing, Online Implementations

1. INTRODUCTION

The last two decades have witnessed pioneering efforts on establishing Human Activity Recognition (HAR) as a multi-disciplinary research field that can shape the future of personalized healthcare and well-being services [1, 2]. Signal processing [3, 4] and wearable computing [5, 6] have been empowering these efforts, providing the necessary means for monitoring the physiological performance of individuals and extracting contextual information from raw observations in an offline, centralized fashion. Nevertheless, latest trends in HAR applications [7] are focusing on addressing the challenge of moving towards online processing architectures, dictated by the need to analyse and interpret complex activities while in data capture.

Representative online learning frameworks in the HAR arena expand the set of modalities beyond typical inertial sensing [8, 9], or elaborate on implementing reliable classification techniques (e.g., Support Vector Machines) on mobile, computationally constrained, devices [10, 11]. These approaches address adequately both data sampling and online pre-filtering, as well as real-time aspects of landmark classification algorithms. Even so, limited consideration is given on how we can intelligently exploit the inherent correlations of diverse raw sensing streams for achieving dynamic data compression at no information loss for the classification procedure.

Recent empirical studies [12] emphatically conclude that considering feature selection [13] into the HAR chain, can yield similar or improved performance than the one achieved when the entire feature set is used. Feature selection is in principle a search problem, responsible for the automatic calculation of the data attributes

that are considered to have sufficient information for inferring the labels of different classes. Feature selection algorithms (FSA) either rank individual attributes based on their discriminating power between different classes (filter methods), or iteratively evaluate the complementarity of different attributes (wrapper methods) [14]. In addition, the theoretical framework of graph feature selection has been recently shaped [15, 16, 17], according to which features are modelled as the vertices of a graph, while their inter-similarity is reflected on the weighted edges of the graph.

FSA trends in the HAR domain propose frameworks that in their majority address off-line perspectives of the problem, wherein the sensing dataset is a priori available for processing. Zhang and Sawchuk [18] introduce the concept of physical features in architectures that consider a wired inertial sensor and present the performance between different supervised FSA. The benchmark studies recently introduced in [19] further extend such efforts, evaluating the performance of supervised, unsupervised, and graph methods on expanded datasets in terms of types of activity, as well as deployment on the human body. Moving towards on-line schemes, the authors in [20], combine the energy cost of individual features that are calculated in real-time and explore the efficacy of graph models to represent correlation and computing complexity of the features. Notwithstanding, recent literature reports a substantial gap on how FSA performs in dynamic environments, wherein dominant features are calculated while biophysical sensing streams are in capture.

In this work, we address this gap by studying online aspects of popular FSA [21, 22, 16]. We extend our recent work [19], where we evaluated the offline performance of FSA for HAR, by proposing a software framework for performing online feature selection in parallel with data acquisition, and implement it on modern mobile devices. Instead of performing calculations on the entire dataset [19], limited windows of sensing streams are available at any time for catering feature extraction and selection. Going well beyond the current state of the art, we herein yield analytical results on the dynamic FSA performance with respect to the variation of human activities, sampled at different parts of the human body. We introduce the concept of dominant streams, defined as the sensing modalities that convey dominant features, and we highlight how they impact the performance of different FSA. Finally, we address real-time aspects, elaborating on execution time and energy consumption.

2. THE ONLINE FEATURE SELECTION ARCHITECTURE

A common concept for most real-life HAR deployments [23, 24] is that wireless sensing devices are mounted on different parts of the human body and periodically transmit biophysical streams of data to a mobile device, which is responsible for storing and processing. Considering such architectures, one can think of sensor data arriv-

ing in small chunks at extremely frequent time intervals, belonging into the magnitude of milliseconds. As such, a substantial volume of raw data streams can become available within seconds for enabling real-time analytics on human movements. The herein proposed architecture, presented in Figure 1 exploits this rationale for extending the modular design of the FORTH-TRACE library [19], in order to perform online feature selection on mobile devices.

Specifically, the Data Gathering module, responsible for the data

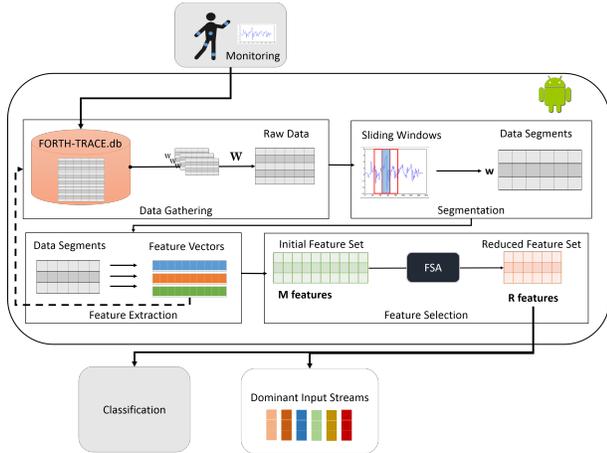


Fig. 1: The architecture of the online feature selection scheme.

preprocessing, is enriched with a database (FORTH-TRACE.db), which enables the dynamic storage and retrieval of data streams as they arrive from the wearable sensors to the mobile device. Depending on both the sampling rate and the memory and computational limitations of the device, these streams are retrieved in temporal windows of length equal to W and preprocessed for synchronization and removal of outliers. The resulting data chunks are fed into the Segmentation component for their further tessellation into sliding windows of observation length equal to w . The segmented data are employed for the calculation of M statistical features [18] and the concatenation of all respective feature vectors into a feature data matrix at the Feature Extraction module. The extracted features are stored to FORTH-TRACE.db by means of an independent, non-blocking background thread, and they are employed by the Feature Selection component for calculating the reduced feature set, comprised of R ($R < M$) dominant attributes, which represent the human activities that correspond to the initial data chunk.

The above procedure is periodically repeated, since the activities involved in such applications alternate at a frequent pace. The reduced set of features can be further utilized for the classification process, as well as for inferring qualitative information on the characteristics of the reduced feature set contents. Specifically, we categorize the dominant features based on the sensing modality they represent (e.g., acceleration, angular velocity), the type of attribute (e.g., statistical), and the pairwise correlations they express, both within the same modality (intra-modality), e.g., between data captured from different axes, as well as across different modalities (inter-modality), e.g., between accelerometer and gyroscope data. As highlighted in Section 3, the resulting *dominant streams* provide valuable insight into the current dominating sensor modality and its relationship to the performed human activity.

3. EVALUATION STUDIES

In this work, we investigate the influence of both the occurring activity labels and the resulting dominant streams on the FSA

performance. Towards this direction, the herein proposed architecture for online FSA has been implemented at a Samsung Galaxy Tab4, featuring Android KitKat (4.4.2). Without loss of generality, we emulate sampling from wearable sensors by employing a pseudo real-time procedure that periodically feeds our online feature selection library with consecutive sensing streams. For this purpose, we consider the FORTH-TRACE dataset [19], which employs Shimmer wearable devices [25], deployed on five different body locations. The dataset contains 3-axial inertial data (acceleration, angular velocity, variations of the magnetic field) collected from 15 healthy individuals while they are performing a series of 16 short activities within a 20 min time span. The set of activities includes conventional ones (e.g., climbing stairs), postural transitions (e.g., sit \rightarrow stand), and combined activities (e.g., walk and talk). In total, streams of 9 sensing channels, i.e. 3 per modality, are fed into the Feature Extraction component for calculating 135 statistical features for each temporal window. These attributes can be further grouped into the following categories of dominant streams: statistical \rightarrow {accelerometer, gyroscope, magnetometer}; pairwise intra-modality \rightarrow {accelerometer-accelerometer, gyroscope-gyroscope, magnetometer-magnetometer}; and pairwise inter-modality \rightarrow {accelerometer-gyroscope, accelerometer - magnetometer, gyroscope-magnetometer}.

We evaluate the performance of feature selection methods that yield the optimal results in the off-line studies [19], namely the unsupervised Feature Selection based on Feature Similarity (FSSA, wrapper method) [22], the supervised Relief-F (filter method) [21, 13], and the unsupervised Graph Clustering with Node Centrality (GCNC, filter method) [16]. The length w of the segmentation window is set to 2s, while the key experimental parameter is the length W of the initial temporal window and the number P of subsequent partitions it generates for performing FSA on a different subset of raw data and respective performed activities. Specifically, we consider $W \in \{2, 3, 4, 5\}$ (min), thereby generating $P \in \{9, 6, 5, 4\}$ partitions, respectively. Table 1, indicates the number of activity labels that occur in each $W - P$ combination, while Table 2 highlights the distribution of the different activities in successive data partitions when $W = 2$ min.

W	P								
	1	2	3	4	5	6	7	8	9
2	3	5	5	3	3	5	5	5	3
3	4	6	5	5	7	5	-	-	-
4	6	6	6	9	3	-	-	-	-
5	6	7	8	5	-	-	-	-	-

Table 1: No. of activity labels in each $W - P$ combination.

Activity	Partition Index								
	1	2	3	4	5	6	7	8	9
stand	41.88	18.64	18.64	3.39	16.1	16.95	3.39	16.1	64.47
sit	52.14	41.53	-	-	-	-	-	-	-
sit & talk	-	26.27	66.95	-	-	-	-	-	-
walk	-	-	5.08	94.92	-	-	-	49.15	-
walk & talk	-	-	-	-	81.36	22.88	-	28.81	31.58
climb	-	-	-	-	-	55.93	44.07	-	-
climb & talk	-	-	-	-	-	-	-	50.85	-
stand \rightarrow sit	5.98	-	-	-	-	-	-	-	-
sit \rightarrow stand	-	6.78	-	-	-	-	-	-	-
stand \rightarrow sit & talk	-	6.78	-	-	-	-	-	-	-
sit & talk \rightarrow stand	-	-	6.78	-	-	-	-	-	-
stand \rightarrow walk	-	-	2.54	-	2.54	-	-	-	-
walk \rightarrow stand	-	-	-	1.69	-	1.69	1.69	-	3.95
stand \rightarrow climb	-	-	-	-	-	2.54	-	2.54	-
climb \rightarrow walk	-	-	-	-	-	-	1.69	-	-
climb & talk \rightarrow walk & talk	-	-	-	-	-	-	-	1.69	-

Table 2: Percentage (%) of occurrence of activities detected in each data partition for $W = 2$ min.

The FSA performance is evaluated in a post-processing component, in which the normalized value $H_r \in [0,1]$ of the Representa-

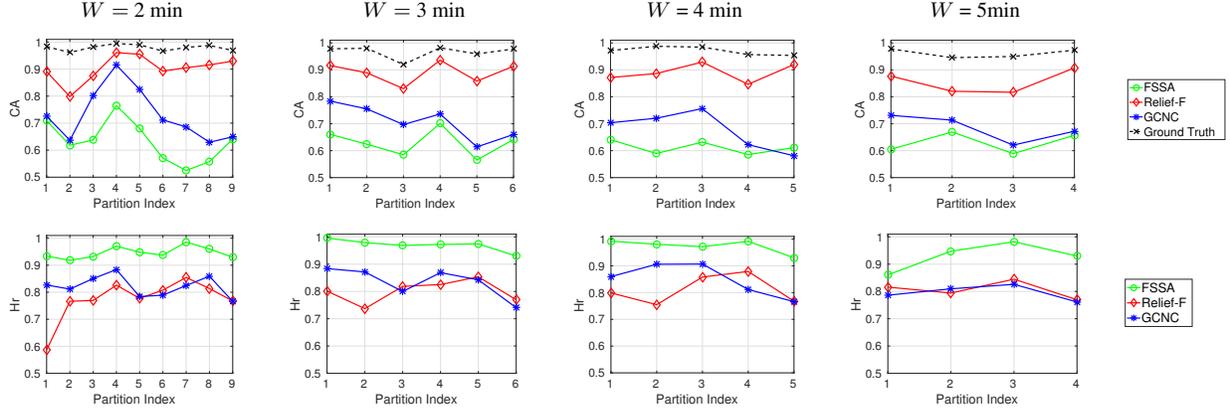


Fig. 2: Mean values of CA and \overline{H}_r for FSSA, Relief-F and GCNC from all sensor locations of the FORTH-TRACE dataset.

tion Entropy [22] is utilized to measure the quality of compression. Specifically, when $\overline{H} \rightarrow 0$, the data patterns are highly relevant to each other, and when $\overline{H} \rightarrow 1$ the information is equally distributed along the features. In addition, we employ in an off-line fashion the Gaussian-Kernel SVM classifier [26] for calculating the estimated activity labels. The respective classification accuracy (CA) metric quantifies the classification ability of the reduced feature set; when $CA \rightarrow 1$, optimal predictability is achieved. The run-time performance of the online feature selection library is quantified in terms of execution time per component and the Android energy requirements.

3.1. Dynamic Performance of the FSA

Figure 2 depicts the mean values of CA and \overline{H}_r of GCNC, FSSA and Relief-F with respect to the partition index for all cases of W and sensor locations. We also present the CA results for the case when no feature selection algorithm is applied (ground truth). We observe that short temporal windows W are accompanied by increased variations on the FSA performance across different data partitions. For example, when $W = 2$ min, the CA metric attains the values 0.61, 0.79, 0.63 in the 2nd data partition, and then increases to 0.76, 0.96 and 0.91 during the 4th data partition for FSSA, Relief-F, and GCNC, respectively. The same remark stands even for the ground truth case where CA transitions from 0.96 to 0.99 in the 2nd to 4th partition. This variation is aligned to the activities involved in the specific partitions, as presented in Table 2; during the second partition, the activities are mainly split into “sit” and “sit and talk”, which are considered less discriminative from each other compared to those primarily involved in the 4th data partition (“walk”). Similarly, the value of \overline{H}_r differentiates across different partitions especially regarding the Relief-F and GCNC algorithms, showing a peak in the 4th (0.82, 0.88 for Relief-F and GCNC, respectively), the 7th (0.85 for Relief-F), and the 8th (0.85 for GCNC) data partitions.

The performance variation across different partitions is less prominent, as the size of W increases. For instance, when $W = 4$ min, the CA metric attains the values 0.64, 0.87, 0.7 and 0.97 in the first partition and equals to 0.61, 0.91, 0.58 and 0.95 as it reaches the 5th partition for FSSA, Relief-F, GCNC, and the ground truth, respectively. Likewise, \overline{H}_r becomes $\{0.99, 0.85, 0.76\} \rightarrow \{0.92, 0.79, 0.76\}$ between the first and the last partition when FSSA, Relief-F, and GCNC are respectively employed. When data chunks corresponding to $W=5$ min are consecutively fed into the online FSA library, the maximum absolute difference of the value of CA (\overline{H}_r) between different partitions equals to 0.06 (0.07). These re-

sults highlight how the volume of the data streams can improve FSA performance both in terms of compression as well as in terms of the resulting predictability of the labels of the classes, regardless of the fact that the number of respective activities (Table 1) increases.

The differences in the performance of the three feature selection methods is aligned to their inherent characteristics. Specifically, Relief-F considers the class label in the calculation of the dominant features, and as such it exhibits better and more consistent performance in CA terms than the one observed when the unsupervised FSSA and GCNC are employed. Furthermore, the knowledge of the class labels leads Relief-F in convergence to the optimal CA results, which are indicated by the ground truth. Nevertheless, Relief-F does not consider the complementarity between different data attributes, and thereby has a limited consideration on yielding highly compressed reduced feature sets. As such, its performance in terms of \overline{H}_r is inferior when compared to the remaining approaches. Considering FSSA and GCNC, the different variations of activities and the lack of the respective labels in different data partitions have a greater impact on their performance both in terms of CA , as well as \overline{H}_r . The respective deviations are more consistent for smaller window sizes. Hence the remaining of this analysis elaborates on the case of having $W = 2$ min.

Performance among different data partitions and sensor locations. Table 3 presents the performance of the GCNC, Relief-F, and FSSA algorithms across different partitions for $W=2$ min for two locations of the dataset, namely the left wrist and the right thigh. The performance of all FSA varies with respect to the sensor location; activities inferred from the lower body (right thigh) offer consistently a more optimal performance in CA terms than the one resulting when data collected from the upper body are employed, for all feature selection techniques considered. This is related to the set of the activities incorporated (e.g., “walk” versus “stand”), which are more discriminative when data sampled from the lower body parts are involved. Among the herein examined FSA, the FSSA algorithm is highly influenced by the difference in the signals recorded from the different body parts; during the 4th partition, the CA decreases from 0.84 to 0.58 as we transit from the right thigh to the left wrist data streams. Nonetheless, FSSA, due to its clustering criterion, yields optimal \overline{H}_r values for all partitions of both locations. GCNC shows steep variations in terms of \overline{H}_r , among partitions 4-8 of the right thigh location; however this does not apply in the left wrist location. This observation is aligned to the limitation of the GCNC to systematically avoid the contribution of each feature to the overall entropy of the feature data set, hence causing fluctuations in the performance

Left Wrist									
Partition Index	1	2	3	4	5	6	7	8	9
GCNC	0.76, 0.87	0.66, 0.78	0.77, 0.82	0.81, 0.87	0.84, 0.75	0.78, 0.78	0.70, 0.96	0.68, 0.99	0.58, 0.49
Relief-F	0.81, 0.49	0.78, 0.66	0.90, 0.77	0.97, 0.79	0.97, 0.75	0.87, 0.79	0.92, 0.86	0.88, 0.81	0.92, 0.75
FSSA	0.62, 0.99	0.70, 0.99	0.46, 0.95	0.58, 0.98	0.71, 0.99	0.55, 0.96	0.39, 0.98	0.62, 0.99	0.52, 1

Right Thigh									
Partition Index	1	2	3	4	5	6	7	8	9
GCNC	0.67, 0.80	0.69, 0.92	0.74, 0.97	0.96, 0.80	0.80, 0.46	0.78, 0.65	0.58, 0.47	0.67, 0.93	0.70, 0.85
Relief-F	0.89, 0.54	0.72, 0.80	0.90, 0.83	0.96, 0.80	0.95, 0.77	0.89, 0.81	0.91, 0.84	0.95, 0.80	0.94, 0.83
FSSA	0.93, 0.99	0.62, 0.91	0.67, 0.87	0.84, 0.95	0.74, 0.82	0.63, 0.91	0.49, 0.99	0.54, 0.97	0.68, 0.82

Table 3: Values of $[CA, \overline{H}_r]$ of GCNC, Relief-F and FSSA algorithms across different partitions for $W = 2$ min.

GCNC, Relief-F	Left Wrist									Right Thigh								
	1	2	3	4	5	6	7	8	9	1	2	3	4	5	6	7	8	9
Partition Index	1	2	3	4	5	6	7	8	9	1	2	3	4	5	6	7	8	9
statistical acc	25, 20	12.5, 32	33.3, 20	44.4, 33.3	50, 37.5	12.5, 29.1	25, 29.1	11.1, 45.8	42.8, 20.8	14.2, 16.6	37.5, 12.5	0, 16.6	11.1, 29.1	12.5, 37.5	22.2, 20.8	22.2, 29.1	22.2, 20.8	22.2, 20.8
statistical gyro	62.5, 40	0, 28	0, 20	0, 45.8	12.5, 4.1	25, 25	12.5, 29.1	44.4, 0	0, 14.2	28.5, 50	12.5, 20.8	22.2, 16.6	55.5, 33.3	37.5, 16.6	22.2, 16.6	22.2, 25	22.2, 20.8	22.2, 16.6
statistical mag	12.5, 12	50, 16	33.3, 8	22.2, 8.3	37.5, 25	37.5, 20.8	25, 25	11.1, 16.6	42.8, 16.6	28.5, 8.3	37.5, 4.1	44.4, 25	22.2, 20.8	37.5, 25	44.4, 16.6	44.4, 16.6	33.3, 37.5	44.4, 12.5
p/w intra acc	0, 0	0, 0	0, 4	0, 4.1	0, 4.1	0, 0	0, 4.1	0, 0	0, 12.5	0, 8.3	0, 0	0, 4.1	11.1, 0	0, 4.1	0, 12.5	0, 4.1	0, 0	0, 8.3
p/w intra gyro	0, 0	0, 4	11.1, 0	0, 4.1	0, 4.1	0, 8.3	12.5, 4.1	0, 0	14.2, 0	0, 0	0, 0	0, 8.3	0, 0	0, 0	0, 0	0, 4.1	11.1, 0	0, 12.5
p/w intra mag	0, 0	0, 0	0, 0	0, 0	0, 0	0, 0	0, 0	0, 4.1	0, 0	14.2, 4.1	0, 4.1	0, 0	0, 0	0, 0	11.1, 0	11.1, 0	0, 0	0, 0
p/w inter acc-gyro	0, 12	12.5, 12	11.1, 20	0, 0	0, 8.3	12.5, 8.3	12.5, 4.1	11.1, 4.1	0, 20.8	0, 4.1	0, 33.3	0, 12.5	0, 12.5	0, 8.3	0, 20.8	0, 12.5	0, 4.1	0, 16.6
p/w inter acc-mag	0, 8	0, 8	0, 20	22.2, 4.1	0, 12.5	0, 8.3	0, 4.1	0, 0	0, 0	0, 8.3	0, 8.3	0, 8.3	0, 4.1	12.5, 4.1	0, 12.5	0, 0	11.1, 4.1	0, 4.1
p/w inter gyro-mag	0, 8	25, 0	11.1, 8	11.1, 0	0, 4.1	12.5, 0	12.5, 0	22.2, 0	0, 4.1	14.2, 0	12.5, 16.6	33.3, 8.3	0, 0	0, 4.1	0, 0	0, 8.3	0, 12.5	11.1, 8.3

Table 4: Percentage (%) of occurrence of Dominant Input Streams for (GCNC, Relief-F) across different data partitions for $W = 2$ min.

of the graph feature selection technique.

Impact of activity labels to the Classification Accuracy. The combination of results presented in Tables 2 and 3, verifies the claim that the distribution of activity labels within a single partition clearly affects the outcome of the CA . For instance, the value of CA at the 4th and 5th partitions, which correspond to a single primary activity (“walk”: 94.92%, “walk and talk”:81.36% respectively) attains optimal results (GCNC ≥ 0.8 , Relief-F ≥ 0.95). By contrast, the 2nd and 3rd partitions, wherein the data is shared between more activities, yield inferior performance regarding CA (GCNC ≤ 0.69 , Relief-F ≤ 0.72 , FSSA ≤ 0.7). Furthermore, it occurs that the proportion of transitional activities in a single partition has an effect on the FSA performance as well. Hence, the value of CA of Relief-F in the 8th partition, which contains $\sim 4.2\%$ of transitional activities, is approximately 20% larger than its CA value in the 2nd partition, which contains $\sim 13.5\%$ of transitional activities, for both locations. However, this observation does not apply for the unsupervised FSA.

Dominant Streams. Table 4 demonstrates the percentage of dominant streams selected by GCNC and Relief-F for the left wrist and the right thigh locations. The majority of dominant streams selected by GCNC belongs to the statistical domain, while the top rankings of Relief-F are distributed among the statistical and the pairwise correlation features. Moreover, the pairwise inter-modality streams dominate against the pairwise intra-modality ones for all cases herein presented. Notably, the dominant streams in partitions, which yield optimal CA results vary among the different locations; the majority of the input streams selected by GCNC in partitions 4-5 are from the statistical domain. By contrast, regarding partition 4 of both locations, Relief-F chooses a combination of statistical and pairwise correlation domain features. Finally, partitions yielding similar performance (e.g., partition 4, $CA = 0.96$) are described by different dominant input streams when either Relief-F or GCNC are employed. This remark stands for consecutive partitions, wherein each feature selection method shows similar performance (e.g. left wrist, partitions 4 and 5, Relief-F $CA = 0.97$), as well as for those that yield poor performance (e.g. partitions 2, 8).

3.2. Run-time aspects

Table 5 indicates the execution time of the online library modules in the Android environment. Large W possess large amounts of available to process data, thus, the separate modules are more time-consuming. Feature extraction and feature selection have a con-

siderable impact on the run-time performance of the whole framework, showing measurements reaching the minute scale as the size of W increases. With regard to feature selection, Relief-F reports the fastest execution ($< 12s$), however its supervised learning nature does not facilitate it for a real-time realization of the framework. Concerning the unsupervised FSA, GCNC is considered suitable for a real-time feature selection architecture, showing a maximum execution of $\sim 28s$, contrary to FSSA which requires more than 40s. Finally, we measure the energy requirements of the online framework using the PowerTutor utility [27]. Our library requires 20.3 J, outperforming both the Adaptive Accelerometer Activity Recognition (A3R) algorithm [24] and the Shimmer3 Gesture Recognition application [25], which require 100 J and 29.2 J respectively.

Execution time (s) per W	2	3	4	5
Acquisition	1.2	1.36	1.73	1.7
Segmentation	0.73	0.95	1.14	1.16
Feature Extraction	46.62	84.28	116.71	146.23
FS: FSSA	49.57	69.81	85.74	105.49
FS: Relief-F	4.3	7.36	9.75	11.71
FS: GCNC	21.52	21.74	23.06	28.25

Table 5: Execution time(s) of the library components per W .

4. CONCLUSIONS

In this work, we examined the online behaviour of feature selection in the HAR domain. We proposed a software architecture for enabling the transition from offline to online feature selection schemes. We implemented the herein proposed architecture on conventional mobile devices, and performed extensive studies on how the dynamic catering of sensing streams, corresponding to different activities and wearable sensor locations, affects the performance of popular FSA. We observed a diversity in the feature selection performance along different partitions of data, and different dominating sensor modalities. We conclude that the feature selection performance is affected by the intrinsic properties and learning nature of each distinct FSA and by the type of activity labels that characterize the given input data. In terms of real-time behaviour, the extraction and selection of features are the most time-consuming operations, especially when the a priori-knowledge of the activity labels is not considered for the calculation of the dominant streams. A future extension of this work will address the aforementioned remarks into the design of an unsupervised, yet dynamic and activity-adaptive FSA to fill the gap in online HAR realizations.

5. REFERENCES

- [1] Agusti Solanas, Constantinos Patsakis, Mauro Conti, Ioannis S Vlachos, Victoria Ramos, Francisco Falcone, Octavian Postolache, Pablo A Pérez-Martínez, Roberto Di Pietro, Despina N Perrea, et al., “Smart health: a context-aware health paradigm within smart cities,” *IEEE Communications Magazine*, vol. 52, no. 8, pp. 74–81, 2014.
- [2] Ya Tian, Hongxing Wei, and Jindong Tan, “An adaptive-gain complementary filter for real-time human motion tracking with marg sensors in free-living environments,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 21, no. 2, pp. 254–264, 2013.
- [3] S. Savvaki, G. Tsagkatakis, A. Panousopoulou, and P. Tsakalides, “Effects of matrix completion on the classification of undersampled human activity data streams,” in *2016 European Signal Processing Conference (EUSIPCO)*, 2016.
- [4] W. Xu, M. Zhang, A. A. Sawchuk, and M. Sarrafzadeh, “Robust human activity and sensor location corecognition via sparse signal representation,” *IEEE Transactions on Biomedical Engineering*, vol. 59, no. 11, pp. 3169–3176, Nov 2012.
- [5] O. D. Lara and M. A. Labrador, “A survey on human activity recognition using wearable sensors,” *IEEE Communications Surveys Tutorials*, vol. 15, no. 3, pp. 1192–1209, Third 2013.
- [6] Muhammad Shoaib, Stephan Bosch, Ozlem Durmaz Incel, Hans Scholten, and Paul JM Havinga, “Complex human activity recognition using smartphone and wrist-worn motion sensors,” *Sensors*, vol. 16, no. 4, pp. 426, 2016.
- [7] Seyed Ali Rokni and Hassan Ghasemzadeh, “Plug-n-learn: Automatic learning of computational algorithms in human-centered internet-of-things applications,” in *Proceedings of the 53rd Annual Design Automation Conference*, New York, NY, USA, 2016, DAC ’16, pp. 139:1–139:6, ACM.
- [8] G. Filios, S. Nikolettseas, C. Pavlopoulou, M. Rapti, and S. Ziegler, “Hierarchical algorithm for daily activity recognition via smartphone sensors,” in *Internet of Things (WF-IoT), 2015 IEEE 2nd World Forum on*, Dec 2015, pp. 381–386.
- [9] Asanga Wickramasinghe and Damith C. Ranasinghe, “Recognising activities in real time using body worn passive sensors with sparse data streams: To interpolate or not to interpolate?,” in *12th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, ICST, Brussels, Belgium, Belgium, 2015, pp. 21–30, ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).
- [10] Jorge Luis Reyes Ortiz, *Linear SVM Models for Online Activity Recognition*, pp. 95–112, Springer International Publishing, Cham, 2015.
- [11] Muhammad Shoaib, Stephan Bosch, Ozlem Durmaz Incel, Hans Scholten, and Paul JM Havinga, “A survey of online activity recognition using mobile phones,” *Sensors*, vol. 15, no. 1, pp. 2059–2085, 2015.
- [12] Nicole A. Capela, Edward D. Lemaire, and Natalie Baddour, “Feature selection for wearable smartphone-based human activity recognition with able bodied, elderly, and stroke patients,” *PLoS ONE*, vol. 10, no. 4, pp. 1–18, 04 2015.
- [13] Huan Liu and Hiroshi Motoda, *Computational methods of feature selection*, CRC Press, 2007.
- [14] Isabelle Guyon and André Elisseeff, “An introduction to variable and feature selection,” *The Journal of Machine Learning Research*, vol. 3, pp. 1157–1182, 2003.
- [15] Zhihong Zhang and Edwin R. Hancock, “A graph-based approach to feature selection,” in *Graph-Based Representations in Pattern Recognition*, Xiaoyi Jiang, Miquel Ferrer, and Andrea Torsello, Eds., vol. 6658 of *Lecture Notes in Computer Science*, pp. 205–214. Springer Berlin Heidelberg, 2011.
- [16] Parham Moradi and Mehrdad Rostami, “A graph theoretic approach for unsupervised feature selection,” *Engineering Applications of Artificial Intelligence*, vol. 44, pp. 33–45, 2015.
- [17] Guodong Zhao, Yan Wu, Fuqiang Chen, Junming Zhang, and Jing Bai, “Effective feature selection using feature vector graph for classification,” *Neurocomputing*, vol. 151, Part 1, pp. 376–389, 2015.
- [18] Mi Zhang and Alexander A Sawchuk, “A feature selection-based framework for human activity recognition using wearable multimodal sensors,” in *Proceedings of the 6th International Conference on Body Area Networks*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2011, pp. 92–98.
- [19] Katerina Karagiannaki, Athanasia Panousopoulou, and Panagiotis Tsakalides, “A benchmark study on feature selection for human activity recognition,” in *ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. 2016, ACM.
- [20] H. Ghasemzadeh, N. Amini, R. Saeedi, and M. Sarrafzadeh, “Power-aware computing in wearable sensor networks: An optimal feature selection,” *IEEE Transactions on Mobile Computing*, vol. 14, no. 4, pp. 800–812, April 2015.
- [21] Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H Witten, “The weka data mining software: an update,” *ACM SIGKDD explorations newsletter*, vol. 11, no. 1, pp. 10–18, 2009.
- [22] Pabitra Mitra, CA Murthy, and Sankar K Pal, “Unsupervised feature selection using feature similarity,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 3, pp. 301–312, 2002.
- [23] Oresti Banos, Rafael Garcia, Juan A Holgado-Terriza, Miguel Damas, Hector Pomares, Ignacio Rojas, Alejandro Saez, and Claudia Villalonga, “mhealthdroid: a novel framework for agile development of mobile health applications,” in *Ambient Assisted Living and Daily Activities*, pp. 91–98. Springer, 2014.
- [24] Zhixian Yan, Vigneshwaran Subbaraju, Dipanjan Chakraborty, Archan Misra, and Karl Aberer, “Energy-efficient continuous activity recognition on mobile phones: An activity-adaptive approach,” in *2012 16th international symposium on wearable computers*. Ieee, 2012, pp. 17–24.
- [25] “Shimmer sensing,” <http://www.shimmersensing.com/>, 2015.
- [26] Chih-Chung Chang and Chih-Jen Lin, “Libsvm: A library for support vector machines,” *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 3, pp. 27:1–27:27, May 2011.
- [27] Lide Zhang, Birjodh Tiwana, Zhiyun Qian, Zhaoguang Wang, Robert P Dick, Zhuoqing Morley Mao, and Lei Yang, “Accurate online power estimation and automatic battery behavior based power model generation for smartphones,” in *Proceedings of the eighth IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis*. ACM, 2010, pp. 105–114.