

Tracking agitation in people living with dementia in a care environment

Shehroz S. Khan^{†◇,*}, Thaejaesh Sooriyakumaran[‡], Katherine Rich[◇] Sofija Spasojevic[†] Bing Ye[◇]

Kristine Newman[§] Andrea Iaboni^{†◇} Alex Mihailidis[◇]

[†] KITE, Toronto Rehabilitation Institute, University Health Network, Canada

[‡] McMaster University, Canada

[◇] University of Toronto, Canada

[§] Ryerson University, Canada

Abstract

Agitation is a symptom that communicates distress in people living with dementia (PwD), and that can place them and others at risk. In a long term care (LTC) environment, care staff track and document these symptoms as a way to detect when there has been a change in resident status to assess risk, and to monitor for response to interventions. However, this documentation can be time-consuming, and due to staffing constraints, episodes of agitation may go unobserved. This brings into question the reliability of these assessments, and presents an opportunity for technology to help track and monitor behavioural symptoms in dementia. In this paper, we present the outcomes of a 2 year real-world study performed in a dementia unit, where a multi-modal wearable device was worn by 20 PwD. In line with a commonly used clinical documentation tool, this large multi-modal time-series data was analyzed to track the presence of episodes of agitation in 8-hour nursing shifts. The development of a baseline classification model (AUC=0.717) on this dataset and subsequent improvement (AUC= 0.779) lays the groundwork for automating the process of annotating agitation events in nursing charts.

Keywords: multi-modal sensors, agitation, dementia, long-term care, machine learning, long time-series

1. Introduction

Many older adults in the advanced stages of dementia live in long-term care (LTC) or nursing homes settings to receive the needed care. In Canada, about one-third of PwD younger than 80 years and 42% above 80 years live in long-term care homes [1]. However, this sector is poorly resourced and under-staffed, and this can impact on the ability of staff support the well-being of residents [2]. A common challenge in these settings are the behavioral and psychological symptoms of dementia representing a heterogeneous group of non-cognitive type of symptoms and behaviours such as apathy, depression, irritability, agitation and anxiety [3]. These behaviours are often in response to unmet needs, and their expression can place the residents and staff at risk. Lack of staffing resources have an important impact on the ability to monitor the health status of residents. With limited staff time and heavy workloads, the frequency, severity, and context around episodes of agitation are not always reliably documented [1].

Monitoring behavioural symptoms in dementia is clinically important. Changes in behaviour, such as increases in agitation, can signal a change in the health status of the resident, such as seen with delirium or worsening pain. Tracking changes in behaviour over time is also important when interventions are being trialed, to help determine their effectiveness. There is thus an important opportunity to use technology to develop objective measures of behavioural symptoms of dementia.

*corresponding_shehroz.khan@uhn.ca

In this paper, we present the outcomes of a study performed at the Specialized Dementia Unit, Toronto Rehabilitation Institute, Canada. In this 2-year study, 600 days worth of wearable multi-modal sensor data was collected from 20 patients. Using this data, it has been shown previously that multi-modal wearable sensor data can be used to detect incidents of agitation in PwD with high accuracy [4]. For this paper, we examine the predictive ability of this multi-modal sensor data to replicate a common behaviour clinical documentation tool, the Pittsburgh Agitation Scale [5], by rating the presence or absence of agitation events in over an 8 hour nursing shift. To approach this problem, we first reviewed nursing documentation and compared it to our research documentation of events of agitation. Next, we describe the data analysis and machine learning challenges associated with this problem, show baseline results and present improvements in achieving an Area Under the Curve (AUC) of Receiver Operating Characteristics (ROC) = 0.78 on this unique dataset.

2. Literature Review

From a machine learning perspective, the problem studied in this paper can be termed as classification of long time-series data. This classification problem is a challenging task because of unreasonably high training / testing time and degenerating performance in the presence of noise. A common state-of-the-art benchmark is the 1-nearest-neighbor dynamic time warping (DTW) that many methods compare to [6]. Schäfer [7] presented a method based on bag-of-words approach. This method is suited for long-term time-series because it extracts subsequences and compares two time-series based on their structural similarity. The method is shown to be accurate, fast and robust to noise in comparison to DTW for both short and long time-series. Sharbiani et al. [8] presented an efficient dynamic time warping (DTW) technique by transforming the time-series into a set of segments with aggregated values and duration forming a reduced 3-D vector. The speed of the method is up to two orders of magnitude faster while providing similar performance for time-series of greater than 500 points. Lahreche et al. [9] presented a method based on local-extrema and DTW for long time-series classification that is much faster and accurate than distance-based methods.

In the deep learning field, recurrent neural network (RNN) can be used for time-series classification. However, RNN suffers from the vanishing gradient problem due to training on long time-series. [10]. Time Warping Invariant Echo State Network has shown competitive accuracy on several long time-series datasets [11] as it can solve the vanishing gradient problem, especially when learning from long time-series.

Although the above stated research works are related, our problem is different due to the very large size of the time-series sequences (in the order of hundreds of thousands points) and the multi-modal nature of the sensor data (more details in Section 3.3). Therefore, working directly on raw sensor points and applying approaches based on DTW are not feasible due to their quadratic time complexity. Our focus is to build a baseline model on this unique data, develop insights and investigate possible improvements.

3. Agitation Detection Study

The data collection for this study took place between November 2017 - October 2019. Over this 2-year study, 600 days worth of wearable multi-modal sensor data was collected from 20 participants who were people with dementia admitted to a Specialized Dementia Unit, Toronto Rehabilitation Institute, Toronto, Canada. The full protocol can be found at [12]; it is described in brief below.

3.1. Data Collection Protocol

An Empatica E4 wearable watch was placed on the dominant hand of the PwD. This watch can record data on the device that includes accelerometer, blood volume pulse, electrodermal activity and skin temperature. Fifteen cameras were also installed in the common areas of the unit to fine tune the annotations of specific agitation events. In this paper, we do not consider the video data, so it is not discussed further. Once a participant was identified by the geriatric psychiatrist, informed consent was obtained from their substitute decision-maker. As per the study protocol, each participant was recruited in the study for a maximum of two months. If they did not show any symptoms of agitation for two consecutive weeks, they were removed from the study. The E4 device was removed from the participants wrist in the evening before bed and left to charge overnight. Therefore, the wearable sensor data for patients' sleeping time is not available. Every morning, the researcher would upload the data from E4 device to the cloud and then replace the E4 device on the participants wrist.

The time, duration, and details of agitation events was determined through review of documentation in the nursing charts. Nurses were trained to record details of any agitation events and to flag these events by putting a green dot in the chart. These events were validated by review of research video footage from the unit. In addition, nurses complete shift-by-shift clinical documentation of behavioural symptoms using the Pittsburgh Agitation Scale (PAS) [5]. For this study, we examined the PAS scores completed by the morning shift (0700 – 1500 hours) and evening shift (1501 – 2300 hours) as the E4 data for the overnight shift was not collected. In practice, the E4 device was mostly placed on the patient after 0700 hours in the morning and could be removed before 2300 hours, if the participant retired to bed early. Therefore, for a given nursing shift, a full 8 hour data may or may not be available for either of the morning or evening shifts.

3.2. Clinical Documentation of Behavioural Symptoms

The PAS [5] is a clinical documentation tool used to record PwDs severity of behavioural symptoms. The PAS is based on direct observation of PwD. The PAS rates the severity of agitation on the scale ranging from 0 to 4 in four behaviour group Aberrant Vocalization (AV), Motor Agitation (MA), Aggression (AG) and Resisting Care (RC). Depending on needs, the period of observation could range from 1 to 8 hours. In this study, the PAS scores were completed at the end of each of the 8-hour morning and evening shifts. The overnight PAS scores were not used in this analysis because the corresponding multimodal wearable sensor data is not available (as discussed above).

When examining the PAS scores completed for the study participants, we noticed some common issues. There were missing data for PAS scores (14.72% for AV and AG, and 15.58% for MA and RC). We also discovered that many shifts in which agitation had been recorded in the research data were scored as absent of agitation. To understand the distribution of PAS scores reported in the nursing charts, we plotted the scores for each type (AV, MA, AG, RC) for both the known agitation and non-agitation shifts based on our research data collection. An agitation shift is defined as the one that contained at least one mention of agitation event in the nursing chart. Similarly, a non-agitation shift is defined as the one that contained no mention of agitation event in the nursing chart. The general expectation is that for agitation shifts there will be more PAS scores with higher values and for non-agitation shifts there will be more PAS scores with lower values.

Figure 1 shows the histogram of PAS scores for AV, MA, AG and RC for non-agitation shifts. As expected, majority of the scores are below or equal to 3. This indicates that the patients were mostly calm in the non-agitation shifts. Figure 2 shows the histogram of PAS scores for AV, MA, AG and RC for agitation shifts. We observed that for AV, AG and

RC the PAS scores were on the lower side, whereas for MA ($\approx 56\%$) of the score was less than 3. This indicated a surprising contrast that on a shift when agitation occurred (and noted in the nursing chart), the PAS scores given may not be high (e.g. more than 3). This under-reporting of agitation events in the clinical documentation gave rise to our research question.

The focus of this analysis is thus to improve the accuracy of documentation of behavioural symptoms on a shift-by-shift basis, by using sensor data to classify each shift as agitation or no agitation. To answer this question, we used as our gold standard annotation, the presence (or absence) of agitation on any given shift which was collected from a detailed review of nursing charts and validated using research video review.

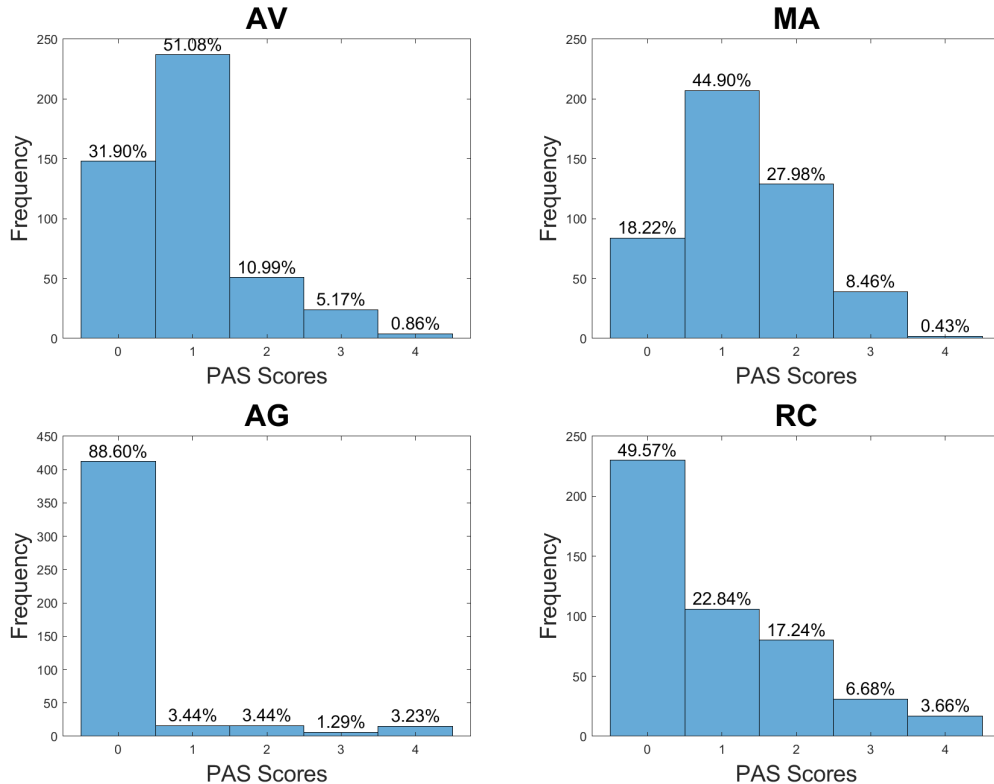


Figure 1. Distribution of PAS scores on non-agitation shifts.

3.3. Machine Learning Challenges

The data collected in this study poses several machine learning challenges:

- The agitation label is only available after a long shift or sequence of multi-modal sensor data. In an ideal case, the shift should be 8 hours. The sensors are sampled at 64Hz (more details in Section 4); therefore, an 8 hour shift will result in $64 \times 60 \times 60 \times 8 = 1,843,200$ points per sensor. There are four sensors in the Empatica E4 watch – accelerometer, blood volume pulse, electrodermal activity and skin temperature. These points are followed by one label – 0 for absence and 1 for presence of agitation in that shift. We observed sequences less than 8 hours as well (see Section 5.2); however, the order of data remains similar.

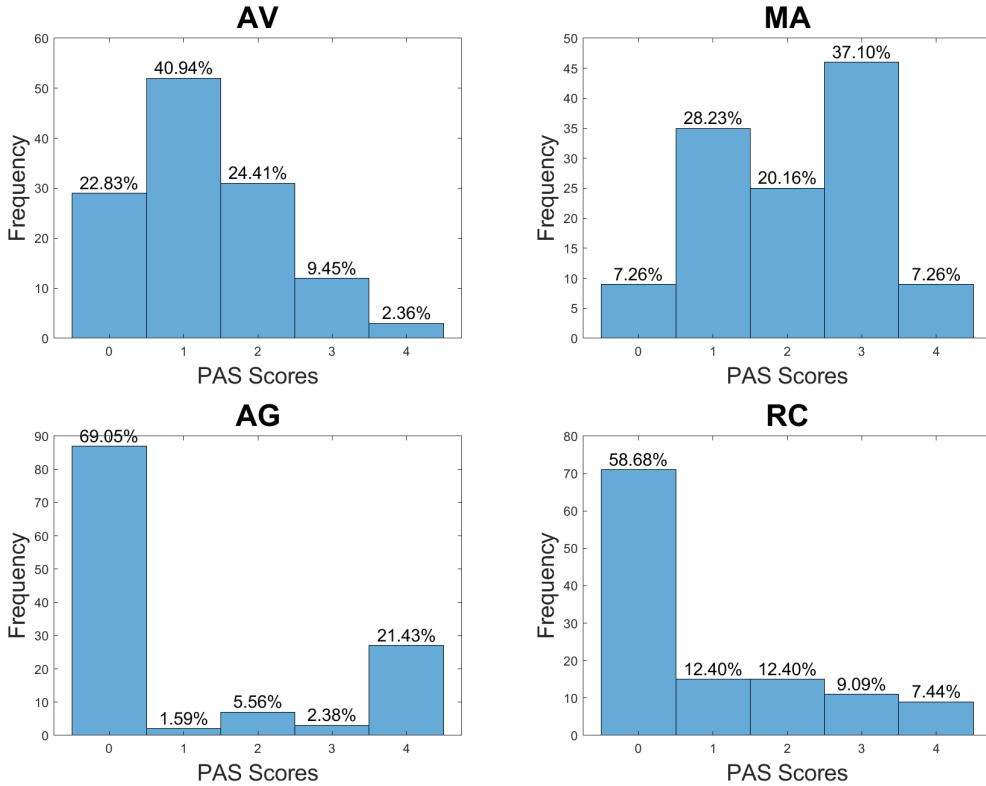


Figure 2. Distribution of PAS scores on agitation shifts.

- The duration of agitation events can have important effect in terms of extracting relevant features from the long multi-modal sensor data sequence. Previous work suggests that agitation events can vary from one minute to up to three hours [12]. However, the actual start/end time or duration of an agitation event is not available for this problem.
- The length of the multi-modal sensor sequences is not the same. The primary reason is that the sensors may not capture the entire data corresponding to the 8-hour nursing shift (see more details in Section 5.2). Therefore, building any temporal predictive model or strategy for feature extraction is not straight forward.

4. Data Processing and Experimental Setup

Different sensors in the Empatica E4 device sample data at a different sampling frequency (Accelerometer at 32Hz, Blood Volume Pulse at 64Hz, Electrodermal Activity and Skin Temperature at 4Hz). To avoid data loss, all the sensors were re-sampled to 64Hz to match with the maximum sampling rate of Blood Volume Pulse [12]. Then, the multi-modal sensor data is split into shifts based on the timestamps (available with the sensor data) – morning shift (0700 – 1500 hours) and evening shift (1501 – 2300 hours). In total, 693 shifts were extracted that contained sensor data. Then, for each shift the following 10 features [13] were extracted from each of the sensor modality – mean, minimum, maximum, standard deviation, interquartile range, difference of maximum and minimum, number of abrupt changes in the data, greatest abrupt change in the data, greatest gradient value in the data and coefficient of variation (mean divided by standard deviation). Additionally,

for electrodermal activity, phasic and tonic signal were extracted. Four features each were extracted for tonic and phasic signals – trapezoidal numerical integration, total number of peaks, their maximum and minimum. Two more features were extracted from tonic signal – maximum of difference and approximate derivative, and mode of the signal. In some cases, the coefficient of variation for temperature gave a *NaN* values. This could happen when the standard deviation is zero, meaning no change in skin temperature during the recording session. Therefore, to avoid numerical computation problems, this feature was removed and overall 49 features were considered for each shift for classification purposes.

We trained three classifiers to detect the presence or absence of agitation in a given shift – Logistic Regression (LR), Random Forest (RF) and Support Vector Machine (SVM). We performed 10-fold cross validation. An internal 5-fold cross validation is performed to tune the parameters of these classifiers. The tuned parameters for LR was lambda (regularization term) in the range of [0.01, 0.05, 0.1, 1, 10, 100], for RF were Number of Trees ([10, 30, 50, 70, 90]) and Number of Predictors ($[f/5, 2 * f/5, 3 * f/5, 4 * f/5]$, where f is the number of features), and for SVM were Box Constraint and Kernel Scale (both in the range [0.01, 0.1, 1, 10, 100, 1000]).

Out of the 693 sensors shifts, 140 contained agitation, i.e. approximately 20% contained agitation shifts and rest of them did not. Clearly, the data labels are skewed that may favour the majority class. Therefore, we introduced a misclassification ‘cost’ parameter in the RF, LR and SVM classifiers. The value of cost is calculated within each fold of the cross-validation step as follows. Let L be a variable that contains the N labels $\in (0, 1)$, where 0 means non-agitation shift and 1 means agitation shift, and w is the ratio of total number of agitation shifts and total number of shifts, defined as:

$$w = \frac{\sum_i^N (L_i \neq 0)}{\sum_i^N L_i} \quad (4.1)$$

and the cost matrix is defined as:

$$cost = \begin{bmatrix} 0 & \frac{1}{(1-w)} \\ \frac{1}{w} & 0 \end{bmatrix} \quad (4.2)$$

This cost matrix shows that when a minority class, i.e. agitation is predicted wrongly as non-agitation it will be penalized heavily as $\frac{1}{w}$ in comparison to when a non-agitation is wrongly classified as agitation $\frac{1}{1-w}$ (considering $w < 0.5$) It should be noted that in a classifier without additional misclassification cost, the off-diagonal elements are 1. The folds in each cross-validation step were stratified, so that data from both majority (non-agitation) and minority (agitation) class is spread equally in each fold. Since the cross-validation is stratified, the cost parameter will remain almost constant in each fold. During each fold of the cross-validation, the scores/probabilities on the test set are concatenated. The final concatenated vector of scores is used to calculate the area under the curve (AUC) of the receiver operating characteristic (ROC), which is the reported performance metric of different classifiers. The AUC of a random classifier is 0.5 and for the best classifier is 1. A higher value of AUC means more confidence is detecting the class of interest (agitation in our case).

5. Results

5.1. Baseline Model

Table 1 shows the classification results on each of the classifier without and with cost (with a subscript). We observe that including cost in the classification slightly improves the AUC in both RF and SVM classifiers (shown as RF_{cost} and SVM_{cost}); however, not so

Table 1. Performance of different classifiers on baseline data, both with and without cost.

Classifier	AUC
RF	0.688
RF _{cost}	0.690
LR	0.717
LR _{cost}	0.705
SVM	0.601
SVM _{cost}	0.699

much in the case of LR. Although, LR without cost gives highest AUC, the AUCs of RF_{cost} and SVM_{cost} classifiers are also very similar.

5.2. Improving the Baseline

As mentioned earlier, a typical nursing shift is 8 hours long. However, there can be situations where the sensor data collected for those shifts may be less than 8 hours. Some of these scenarios are:

- The watch was attached to the patient late in the morning shift (i.e. not at 0700 hours).
- The watch was removed from the patient before the shift ends (i.e. before 2300 hours) or anytime before that).
- The data collection was stopped either by the patient or nursing staff during the daily care (e.g. bathing).

Therefore, we define *sensors shift* as the time duration for which the sensor data is available to the corresponding nursing shift. If the sensor shift contains 8 hours of uninterrupted data, then it is equivalent to a nursing shift.




Figure 3. Distribution of multi-modal sensor data across sensors shifts of various duration.

We plot a histogram to verify the distribution of multi-modal sensor data across various shifts. Figure 3 shows the distribution of multi-modal sensor data corresponding to sensors shifts of different sizes. The histogram clearly shows that majority of the sensor data does

not correspond to the usual 8 hours nursing shift because of the reasons discussed above. Table 2 shows the distribution of sensors shift data across different bins, where Bin x means it contains x hours of multi-modal sensor data. The first row of Table 2 shows that out of total 693 shifts, only 110 shifts have data corresponding to 8 hours nursing shift (approximately 15.87%) and rest of the sensor data is distributed in sensors shifts of duration 1 to 7 hours. This imposes another challenge that the multi-modal sensor data is not equal in length; hence feature extraction becomes non-trivial. Another observation was that the distribution of agitation shifts in these sensors shifts is almost similar (third row of Table 2).

Table 2. Distribution of multi-modal sensor data with respect to different shifts.

	Bin1	Bin2	Bin3	Bin4	Bin5	Bin6	Bin7	Bin8
All Shifts	9	15	77	118	151	131	82	110
Agitation Shifts	2	3	17	25	37	23	13	20
% of Agitation Shifts	22.22	20	22.08	21.19	24.50	17.56	15.85	18.18

In this dataset, a very long multi-modal sensor sequence is given, the only information available is whether agitation event occurred or not. The actual timing and duration of the agitation events are not provided. Therefore, the features extracted must reflect the occurrence of agitation. In a very large sequence (e.g. 8 hours sensors shift), it is very challenging to extract meaningful features without knowing their actual timing and number of occurrences during a nursing shift. The larger the sequence is, the more difficult is to extract discriminatory features. Therefore, we hypothesize that considering shorter sequences for feature extraction will be more useful than larger sequences. To test this hypothesis, we performed the following two experiments:

- (1) Larger Sequences Experiment (LSE) – In this experiment, we only consider the sequences that are larger than a specific threshold, s.t. $SequenceLength \geq b$, where b is varied from 0 to 4 hours (or 240 minutes) with an equal increments of 15 minutes. Therefore, LSE will result in selecting sequences of larger length from the overall sensor data.
- (2) Smaller Sequences Experiment (SSE) – In this experiment, we only consider the sequences that are smaller than a specific threshold, s.t. $SequenceLength \leq b$, where b is varied from 8 to 4 hours with an equal decrements of 15 minutes. Therefore, SSE will result in selecting sequences of smaller length from the overall sensor data.

A threshold of 4 hours is selected in both LSE and SSE because too large or too small sensors sequences may not be very useful for extracting features. Our results confirm this choice, which is described next. For both LSE and SSE, we performed classification with and without misclassification costs (as discussed in Section 5.1).

Figures 4, 5 and 6 show the result of LSE and SSE with and without cost using RF, LR and SVM classifiers. It can be observed that SSE performed better than LSE without cost. With cost SSE performed better than LSE in all the classifiers except for one sensor shift point in SVM. This confirms our hypothesis that considering shorter sequences for feature extraction are more useful than larger sequences. The best AUC achieved with SSE without cost was 0.753 for RF, 0.772 for LR and .774 for SVM, and SSE with cost was 0.770 for RF, 0.779 for LR and 0.761 for SVM. These results are a significant improvement over the the baselines AUCs reported in Section 5.1.

5.3. Sub-Sampling

One of the challenges in this dataset was the very large size of time-series / sequential multi-modal sensor data. It is very hard to extract features that are meaningful and representative of the presence or absence of agitation events in it. One possibility to reduce raw




Figure 4. LSE and SSE results with RF as the base classifier.




Figure 5. LSE and SSE results with LR as the base classifier.

sensor data before performing feature extraction is to downsample the sensor data. However, this could also lead to loss of information due to downsampling. Some features described in Section 4 require minimum number of points; we set this number to 5. If downsampled data has less than 5 points, they are discarded to avoid numerical computation problems. Therefore, as the downsampling rate increases, number of sensors shift can reduce from the original number of 693 shifts. To verify the performance on reduced data, we downsampled the data from 1 minute (or 60 seconds) to an hour (or 3600 seconds) with an increment of 1 minute. We did not downsample the sensor data beyond one hour because that would lead to very few points and they may not capture any useful information in the sensors data. Figures 7, 8 and 9 show the results for RF, LR and SVM classifier with and without cost when the sensor data was downsampled from 1 minute to 60 minutes. We observe that the cost version slightly performs better than without cost. The performance improved in each the case around 20 minutes and then gradually decreased. It was expected that the performance would decrease at higher sampling rate due to loss of information from the raw sensor data. The interesting observation is that the best AUC on the downsampled




Figure 6. LSE and SSE results with SVM as the base classifier.




Figure 7. Performance of RF classifier at different downsampling rates (with and without cost).

data for each classifier (0.700 for RF+cost, 0.683 for LR+cost, and 0.673 for SVM is almost equivalent to the baselines results with full sensor data for RF (see Table 1). This result suggests that downsampling on very long multi-modal sensor sequences where the actual timing/duration of agitation event is not known, could lead to similar results as compared to the full sensor data. Agitation events can be gradual and can extend for several minutes to hours. Therefore, despite downsampling the sensor data, useful features can still be extracted. We could not perform a similar experiment as discussed in Section 5.2 of choosing longer and shorter length sequences, because that value can change for every sampling rate and would lead to an increased number of hyperparameters.

6. Conclusions and Future Work

In this paper, we presented the results of a real-world study that collected wearable multimodal sensor data from PwD in a care environment. The task was to detect the




Figure 8. Performance of LR classifier at different downsampling rates (with and without cost).

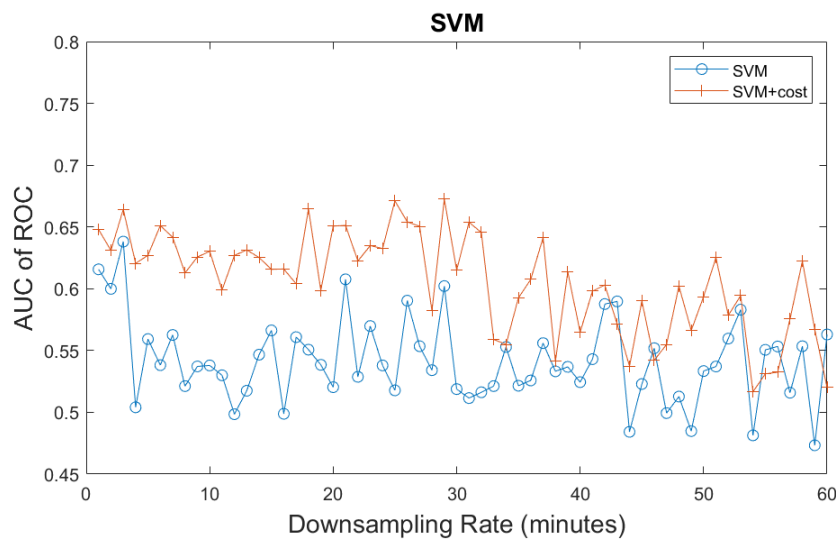


Figure 9. Performance of SVM classifier at different downsampling rates (with and without cost).

presence or absence of agitation during a nursing shift. This is a challenging problem due to the large number of points to analyze and standard techniques (e.g. DTW) fail due to quadratic time complexity. The time-series sequences were of unequal lengths and the actual time, duration and number of occurrences of agitation events are not available, and only one label is given at the end of very long sequence of data. We extracted features from these large multi-modal time-series data and presented a baseline results. Based on the observation that majority of sensors shifts were not equal to nursing shifts in terms of data collected, we chose shorter sequences to build models and improved the baseline results. An interesting experimental observation was that downsampling the sensor data by up to 21 minutes gave equivalent results in comparison to full sensor data. This work

also highlighted the importance of collecting high-quality ground truth labels for clinical studies and the limitations of using clinical documentation for this purpose. In future, we will extract localized features within a time window; however, it could also lead to variable feature length due to unequal length of sensor sequences. Encouraged by our results on downsampled data, we are currently exploring the use of Temporal Convolution Network [14] that can model long-range dependencies, and much faster than LSTM-based methods.

Acknowledgements

This project was funded through the Alzheimer’s Society Research Grant.

References

- [1] *Dementia in long-term care*. URL: <https://www.cihi.ca/en/dementia-in-canada/dementia-care-across-the-health-system/dementia-in-long-term-care>.
- [2] *Report Shows Understaffing in Long Term Care and How to Fix It*. 2021. URL: <https://www.qualicare.com/toronto-on/report-shows-understaffing-in-long-term-care-and-how-to-fix-it/>.
- [3] J. Cerejeira, L. Lagarto, and E. Mukaetova-Ladinska. “Behavioral and psychological symptoms of dementia”. In: *Frontiers in neurology* 3 (2012), p. 73.
- [4] S. S. Khan, S. Spasojevic, J. Nogas, B. Ye, A. Mihailidis, A. Iaboni, A. Wang, L. S. Martin, and K. Newman. “Agitation detection in people living with dementia using multimodal sensors”. In: *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2019, pp. 3588–3591.
- [5] J. Rosen, L. Burgio, M. Kollar, M. Cain, M. Allison, M. Fogleman, M. Michael, and G. S. Zubenko. “The Pittsburgh agitation scale: a User-Friendly instrument for rating agitation in dementia patients”. In: *The American Journal of Geriatric Psychiatry* 2.1 (1994), pp. 52–59.
- [6] G. A. Susto, A. Cenedese, and M. Terzi. “Time-series classification methods: Review and applications to power systems data”. In: *Big data application in power systems*. Elsevier, 2018, pp. 179–220.
- [7] P. Schäfer. “Scalable time series classification”. In: *Data Mining and Knowledge Discovery* 30.5 (2016), pp. 1273–1298.
- [8] A. Sharabiani, H. Darabi, A. Rezaei, S. Harford, H. Johnson, and F. Karim. “Efficient classification of long time series by 3-d dynamic time warping”. In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 47.10 (2017), pp. 2688–2703.
- [9] A. Lahreche and B. Boucheham. “A fast and accurate similarity measure for long time series classification based on local extrema and dynamic time warping”. In: *Expert Systems with Applications* (2020), p. 114374.
- [10] R. Pascanu, T. Mikolov, and Y. Bengio. “On the difficulty of training recurrent neural networks”. In: *International conference on machine learning*. 2013, pp. 1310–1318.
- [11] H. I. Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P.-A. Muller. “Deep learning for time series classification: a review”. In: *Data Mining and Knowledge Discovery* 33.4 (2019), pp. 917–963.
- [12] S. Spasojevic, J. Nogas, A. Iaboni, B. Ye, A. Mihailidis, A. Wang, S. J. Li, L. S. Martin, K. Newman, and S. S. Khan. “A pilot study to detect agitation in people living with dementia using multi-modal sensors”. In: *Under Review* (2020).
- [13] S. S. Khan, M. E. Karg, D. Kuli, and J. Hoey. “X-factor HMMs for detecting falls in the absence of fall-specific training data”. In: *International Workshop on Ambient Assisted Living*. Springer. 2014, pp. 1–9.
- [14] C. Lea, M. D. Flynn, R. Vidal, A. Reiter, and G. D. Hager. “Temporal convolutional networks for action segmentation and detection”. In: *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, pp. 156–165.