Detecting Behaviour Changes in Accelerometer Data

Claudio Diaz, Kalina Yacef
School of Information Technologies
The University of Sydney
cdia0348@uni.sydney.edu.au, kalina.yacef@sydney.edu.au

Abstract

How can the impact of Health Education programs promoting physical activity be analysed? One common way with learning programs is to conduct pre- and post-tests and measure whether/how target knowledge has evolved. In the case of physical activity, unobtrusive accelerometers can capture detailed data about people's movements, but the challenge is to extract information from these raw data to investigate whether/how physical activity behaviours have evolved. This paper presents a methodology to do so, by extracting bouts of physical activity of specific intensity levels and of various lengths, and by using these as features to cluster students' daily behaviours before and after intervention. This approach enables a more insightful analysis of the physical activity behaviours of the participants, and point to the nature of behaviour changes, if present. We illustrate this methodology with pre- and post-test data collected in the context of an e-learning program aimed at educating school children about healthy behaviours, with a focus on reaching recommended levels of daily physical activity: the pre- and post-tests were carried out by measuring unobtrusively and continuously their physical activity for five consecutive school days using research-grade accelerometers (GENE-Activ).

1 Introduction and Related Work

Obesity and sedentarity in children has increased in the last three decades [Ng *et al.*, 2014]. In order to reverse this trend, countries and organisations worldwide implement health education programs for seniors, adults and children, in order to promote behaviour changes and raise awareness with regards to diet and physical activity, two major factors linked to obesity and non-communicable diseases. In particular, studies suggest that physical activity is positively associated with many health benefits, and that in children should accumulate at least 60 mins per day of moderate to vigorous physical activity [Janssen and LeBlanc, 2010].

The use of technology in health promotion interventions has shown great potential to improve health behaviours and provide insights on how to improve their effectiveness [Krebs *et al.*, 2010]. With increasingly available wearable technologies, researchers more routinely use sensors for measuring physical activity unobtrusively and continuously [Plasqui *et al.*, 2013]. Accelerometers provide objective, continuous data of real daily life physical activity, replacing or complementing self-reported data (often inaccurate and coarse). This is especially important when studying children because their self-reported data and/or parent reports can be very inaccurate [Kelly *et al.*, 2007].

Whilst the most frequent use of accelerometers in Health Education is to quantify physical activity, much deeper information can be captured from their data, such as activity recognition [Ravi *et al.*, 2005] and changes in everyday physical activity [Sprint *et al.*, 2016]. Detecting changes in learning behaviour is not new: Specialised data science fields such as Educational Data Mining (EDM) [Baker and Yacef, 2009] and Learning Analytics and Knowledge (LAK)[Siemens, 2013] have developed techniques to extract learning behaviour changes which can certainly be explored for Health Education contexts using accelerometer data. There is indeed an emerging interest in using sensors to better understand complex behaviours in education: for example, in learning kinaesthetic skills like martial arts, dancing or use of clinical equipment [Martinez-Maldonado *et al.*, 2017], or sometimes using several sensors such as, for example, in the analysis of hand movements for engineering building activities [Worsley, 2014], leading to the added complexity of dealing with multimodal data sources [Ochoa, 2017] requiring the creation of different analytics and data mining techniques to extract meaningful information from multi-sensor data [Blikstein and Worsley, 2016]. However the techniques for extracting learning-useful information from sensor data are still in infancy.

In this paper we are concerned with modelling and comparing physical activity behaviours between two sets of accelerometer data, captured before and after a learning intervention, in order to understand its impact. The contribution of this paper is a clustering-based approach for a more insightful analysis of the physical activity behaviour of the participants, and of the nature of physical activity behaviour changes, if present. The paper is structured as follows. Section 2 presents our data and its context. Section 3 describes the methodology, and Section 4 presents the results of this approach on our dataset. Section 5 concludes the paper and suggests avenues of future work.

2 Data and Overall Analysis

The data was collected from the iEngage project [Yacef *et al.*, 2018]: iEngage aims to educate 10-13 year old school children about healthy behaviours, with a focus on reaching recommended levels of daily Physical Activity (PA). PA can fall into one of four different categories: sedentary time (therefore absence of physical activity), light, moderate and vigorous PA. The recommendations are that children should do at least 60 minutes of moderate to vigourous PA (shortened to MVPA). The elearning program also raises awareness about sedentary time, encourages children to limit it, and break them up on a regular basis by some light activity at least. As shown in Figure 1, we conducted a controlled study with two groups of children. The experimental group followed the iEngage learning sessions over 5 weeks, whilst the control group did not. Pre and post-tests were carried out on both groups measuring unobtrusively and continuously their physical activity with GENEActiv [Activinsights Ltd., 2017] activity trackers for five consecutive school days.

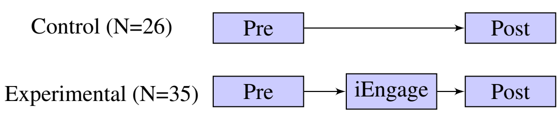

Figure 1: High-level protocol of the intervention

The GENEActiv accelerometers were worn on the wrist of their non-skilled hand and captured acceleration in three axes (x,y,z) with a sample frequency of 60Hz. At the end of each 5 day period (pre and post, for each group), the GENEActiv trackers were collected and their data downloaded to a computer, hence generating two five-day datasets per child, for a total of 61 children.

Overall analysis of the sum of minutes spent in PA showed that pre-intervention, the control and experimental groups spent similar time doing PA at each intensity (p-values of 0.63, 0.62,0.76, 0.29 for Sedentary, Light, Moderate and Vigorous intensities respectively). However, the experimental group post intervention did significantly more PA, especially in MVPA levels (p-values of 0.12, 0.003, 0.017 respectively for L, M and V). While this is consistent with the intervention reaching the desired effect (at least short term) on this population, we are seeking to get more insights on how this activity is distributed throughout the day, and how it evolved: for instance, an important question is whether the additional MVPA occurred in longer bouts of activity (which would suggest more sustained intentional activity), or was it scattered in minuscule amounts throughout the day (which is more likely to be more incidental)? This led us to explore bouts of PA in terms of intensity level, length and frequency.

3 Methodology for Extracting Daily Physical Activity Behaviours

We devised a methodology for characterising daily behaviours of PA at a coarse level, yet capturing essential elements of how the PA is distributed throughout the day. Indeed, two days (for 2 different children, or 2 days for the same child) can show the same total quantity of MVPA (e.g. 40 minutes), but one will contain a lot of sedentary time and long sessions of MVPA, whilst another can show more broken down MVPA but less sedentary time (hence more light activity). The idea is to be able to identify the types of distributions of activity that are present in the cohort data, and to distinguish these distributions.

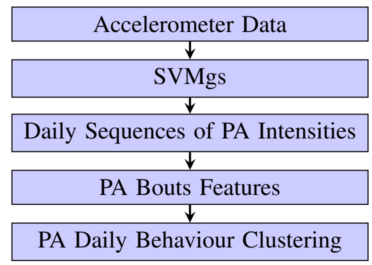

Figure 2: Methodology

Our methodology, shown in Figure 2, can be summarised as follows. First, we processed and categorised students' raw GENEActiv accelerometer data into sequences of PA intensity levels for both datasets (pre and post intervention). We then extracted the bouts of PA, and used their characteristics as features for clustering all the data, to identify types of daily PA behaviours. As we will show in section 4, these, we were then able to follow students' movements across these clusters before and after the learning intervention. The next sub-sections will detail these steps.

3.1 Data Pre-processing

The data pre-processing was done using R [Ihaka and Gentleman, 1996], which has a specific library to manipulate GENEActiv trackers data [Fang and Langford, 2013]. From this point onward, as we are interested in analysing the changes in the experimental population, we worked with the data from the experimental group (N=35). First, we converted the accelerometer binary files to data frames. Next, as we focus here on daily PA behaviours, we filtered out the sleeping times, thus extracting 12-hour daytime records (from 8:00 to 20:00 hrs). To ensure that the daily records were all comparable, weexcluded days where the tracker was not used the whole day, thus excluding the Monday and Friday which were incomplete. DUe to absence or sickness,not all children wore the trackers before and after the intervention. Therefore, from the initial 35 children in the experimental group, we ended up with 30 pre intervention children with three daytime records and 24 post intervention children with three daytime records, thus 54 (30+24) three-day PA records all up.

3.2 From Accelerometer to SVMgs

The next step translated the three dimensional 60 Hz acceleration data into quantities of physical activity within a 1 second epoch. We took the data frames from the binaries and extracted the triaxial acceleration records with timestamps of every child to calculate gravity-subtracted Signal Vector Magnitudes (SVMgs) [Esliger *et al.*, 2011], with gravity approximated to 1, for each 1 second epoch (see Formula

1). This process produced a long vector of physical activity SVMgs per second for each child over the 3 days, thus 54 vectors in total, each being 129,600 second long (3 days x 12 hours x 60 minutes x 60 seconds).

$$SVMgs = \sum_{i=1}^{60} |\sqrt{x_i + y_i + z_i} - 1| \quad (1)$$

3.3 From SVMgs to PA Intensity Levels

We then categorised the SVMgs at each second of data into a PA intensity level, using cutoffs scientifically validated for assessment of physical activity intensity in children [Phillips *et al.*, 2013]. These cutoffs are shown in Table 1.

Table 1: SVMgs Cut Off Levels

Physical Activity Intensity Levels	**SVMgs Cut Off**
Sedentary	[0, 4.5[
Light	[4.5, 16.5[
Moderate	[16.5, 42[
Vigorous	≥ 42

Figure 3 displays an example of SVMgs time series over one day for one student. The red horizontal line represents the cutoff from sedentary to light, the blue line the cutoff from light to moderate, and the green line the cutoff from moderate to vigorous.

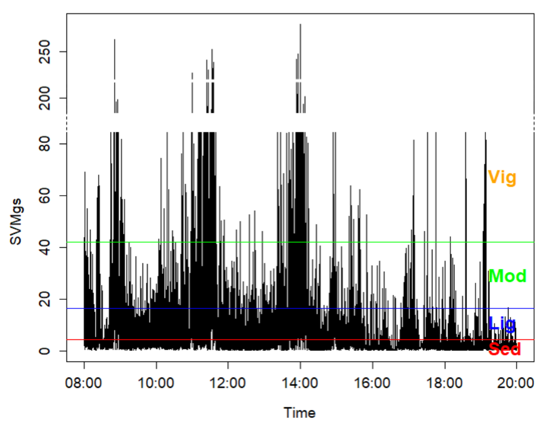

Figure 3: SVMgs time series of one student over one day (the figure is truncated between 80-200 SVMgs for better presentation clarity)

Using the cutoffs above, each second was coded as follows: S for sedentary time, L for light PA, M for moderate PA and V for vigorous PA. As an example, a piece of 5 seconds length of this string can be LLLVV, which can be read as 3 seconds of light activity followed by 2 seconds of vigorous activity. This step therefore produced 54 strings of 129,600 characters, where each character represents the PA intensity level for one second of PA.

3.4 Bouts of PA

As mentioned earlier, we are interested in assessing daily PA behaviours by looking at how their MVPA and sedentary times are distributed throughout the day. Therefore we chose to explore the intensity level, length and frequency of each bout of MVPA and sedentary times. Let us introduce some definitions.

- A **bout** is a continuous episode of physical activity at a specific range of intensity level.
- The **length of a bout** is the number of seconds spent during that bout.
- The **bout frequency** is the number of occurrences of all bouts of a certain length during a day.

We focused on bouts in the range of Moderate to Vigorous Physical Activity (MVPA) and Sedentary Activity (SED), as the aim of the health program is to increase MVPA and decrease SED. Therefore we merged M and V into one category "MVPA". For instance, a sequence of 11 seconds spent in M, 8 seconds in V, and 12 seconds in M preceded and followed by L's would generate one bout of MVPA that would be 31 seconds long.

One of the first questions we explored was: was the increased MVPA that was observed overall after the intervention done in longer bouts? As a first step, we analysed the total time spent in MVPA done in bouts of at least x seconds. Formula 2 shows the reverse cumulative sequence, where t is the bout threshold and b is the number of seconds spent in bouts of length of at least t. For t=1, this is equivalent to the total number of seconds spent in MVPA. For t=2, the total number of seconds spent in bouts of at least 2 seconds (therefore excluding the 1 second-long bouts), and so on.

$$Bouts\, Cum\, Sum_t = \sum_{i=t}^{n} b_i \quad (2)$$

Figure 4 shows a sample of the result of these calculations, where every line shows the average daily MVPA cumulative bout length for a particular student. Over 10 seconds the lines start to flatten as bout length increases.

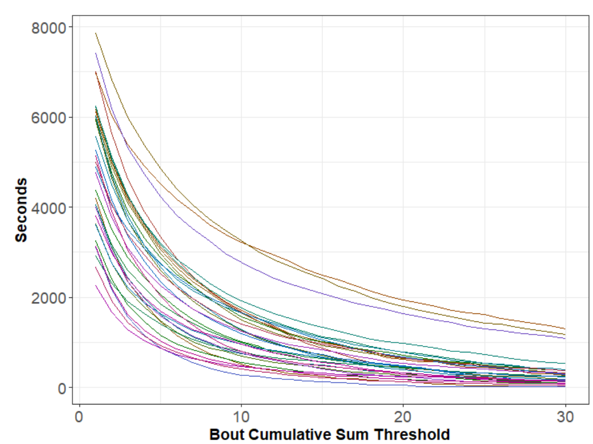

Figure 4: Reverse Cumulative Bout Lengths

A paired T-Test on the before and after cumulative series reveals that overall, students increased MVPA bouts length (p-value=6.883e-10), increased MVPA bout frequency (p-value=0.007814), decreased SED bout length (p-value=2.2e-16) and decreased SED bout frequency (p-value=2.2e-16). This therefore suggests an overall positive effect of the learning program.

3.5 Clustering of PA Behaviours

To explore how students changed their PA patterns before and after, we averaged the daily behaviours of the children pre- and post-intervention and clustered these average daily behaviours using bout characteristics as features: the average time per day spent in bouts of at least a specific length and the average frequency of bouts per day. We selected the daily thresholds of MVPA and SED bouts not only based on our exploration above but also following the established literature [Schaefer *et al.*, 2014]. In particular, meaningful MVPA detected by GENEactivs starts at 3 seconds, as any shorter activity is likely to be noise. The thresholds are shown in Table 2.

Table 2: Clustering Features

Physical Activity Intensity	Bouts Threshold
MVPA	3,10,30
SED	60,120,300

Using these features, we generated daily PA behaviour clusters with all the 54 three-day long records (30 pre-intervention + 24 post-intervention). This means children can be present in up to 2 clusters: one from their daily PA behaviour before the intervention, and the other from their PA behaviour after the intervention. Of course, both their PA behaviours could fall into the same cluster. The features were standardised and a k-means unsupervised algorithm [Macqueen, 1967] with k=6 was applied. This number of clusters was determined by analysing when including another cluster does not improve enough the total within-cluster sum of square (see Figure 5).

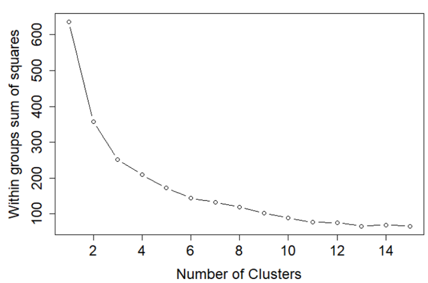

Figure 5: Total within-cluster sum of square by cluster

The cluster centroids are shown in Table 3. We can see that, from a MVPA point of view, the centroids of clusters C4, C5 and C6 fulfil the minimum recommendation of 60 minutes daily of MVPA [Janssen and LeBlanc, 2010], but those of C1, C2 and C3 do not. Also, from a SED point of view we can see that C2, C3 and C1 has the longest and more frequent SED. In detail, C1 shows the lowest medium/long MVPA and the third highest short SED, C2 shows the lowest short bouts of MVPA, longest short SED, C3 shows the third lowest short MVPA and the second highest short SED, C4 shows third highest short MVPA and the lowest short SED, C5 shows the second highest MVPA and the third lowest SED and finally C6 shows the highest MVPA and the second lowest short SED.

Table 3: Clusters Centroids

MVPA Inten.	Measure	1 (N=8)	2 (N=12)	3 (N=5)	4 (N=14)	5 (N=9)	6 (N=6)
>= 3 Secs							
	Tot. Time (min)	32.6	31.4	49.2	61.8	65.8	92.3
	Num. of Bouts	352.6	308.7	472.4	605.3	594.7	698.4
>= 10 Secs							
	Tot. Time (min)	9.5	11.5	18.7	22.1	27.9	49.3
	Num. of Bouts	35.7	40.7	66.1	81.6	97.2	143.6
>= 30 Secs							
	Tot. Time (min)	1.3	2.1	4.0	3.7	5.8	18.2
	Num. of Bouts	1.8	3.2	5.5	5.4	8.5	22.4

SED Inten.	Measure	1 (N=8)	2 (N=12)	3 (N=5)	4 (N=14)	5 (N=9)	6 (N=6)
>= 60 Secs							
	Tot. Time (min)	136	226.3	217.9	70.6	118.4	96.6
	Num. of Bouts	63.9	76.8	39.7	39.6	61.8	41.9
>= 120 Secs							
	Tot. Time (min)	73.2	156.1	179.8	29.8	59.8	57.5
	Num. of Bouts	17.8	24.9	11.2	9.1	18.3	12.6
>= 300 Secs							
	Tot. Time (min)	29.8	101.1	158.2	7.0	14.3	26.8
	Num. of Bouts	2.2	6	3.7	0.8	1.9	1.8

Given these observations, we ordered the clusters in increasing level of PA behaviour, from the lowest activity student cluster (C1) to the highest activity one (C6), and characterise them as seen in Table 4.

Table 4: Cluster Descriptions (Those meeting the daily recommendation of MVPA are flagged with *)

Cluster	Summary Description
1	Not very active cluster (Half of MVPA recommended amounts) but average amount of sedentary times
2	Not very active cluster (A little over half of MVPA recommended amounts) combined with high amount of sedentary time but broken down in many bouts
3	Fairly low MVPA (11 mins short of recommended levels) and very high amount of long sedentary bouts
4*	Active cluster (meeting the recommended amounts of MVPA) combined with little sedentary time, and even fewer long sedentary bouts
5*	Active cluster, slightly more MVPA than cluster 4 but contrasted with higher amounts of short sedentary bouts, and reasonable long bouts of sedentary time
6*	Active cluster, with highest amount of MVPA and low sedentary bouts, but more longer sedentary bouts than the 2 other active clusters.

4 Behaviour Change

The clusters above capture the daily behaviours for all children, before and after, with regards to MVPA and sedentary times. We can now look at whether and how the children from the experimental population moved from one cluster to another, or stayed in the same cluster, as this can be a sign of behaviour change. We can do so only for those children who wore the GENEactivs in both periods (N=22).

Table 5 shows the movement matrix between daily PA behaviour clusters before and after the intervention. The green area shows the top desirable moves (from a low PA cluster to a higher PA cluster), light green shows acceptable moves (from any PA cluster that already meets the daily recommendations to any cluster that also meets them). Yellow shows unimproved moves (from a low PA cluster to a similar PA cluster), and red area shows undesirable moves (from a high PA cluster to a low PA one, or from a low PA one to an even lower PA one).

Table 5: Cluster movement matrix

		To Cluster					
		1	**2**	**3**	**4**	**5**	**6**
From Cluster	**1**	0	2	0	1	1	0
	2	1	2	1	0	1	0
	3	0	0	0	1	1	1
	4	0	1	0	4	0	0
	5	0	0	0	1	2	0
	6	0	0	0	0	0	2

We observe four different behaviour changes,

- Children that moved to clusters with higher MVPA (9 children).
- Children that moved to clusters with lower MVPA clusters (3 children).
- Children who were already in a cluster with MVPA above the daily recommended guidelines, and remained in the same high MVPA cluster (8 children)
- Children who were in a cluster that did not meet the recommended guidelines of MVPA and remained in the same one (2 children)

In particular we can see that over half of the students who were in the cluster with the least MVPA (C1) have moved up to more active clusters, and that all the students who were in the average/fairly low MVPA cluster (C3) have moved to more active clusters. Students who were already active (in C4, C5 and C6) remained active, except for one student who became more sedentary (moved to C2).

5 Conclusion

We presented a methodology to extract aspects of children PA behaviour and how these changed before and after an intervention. First we calculated from accelerometers the SVMgs, then later use them to calculate the PA intensities bouts length and frequency, who were later used as features to cluster their behaviour and monitor changes before and after the intervention.

This methodology helps understand the impact of the intervention from a general and individual level. Whilst we focused here on MVPA and SED intensity levels, a similar approach can be used to also include sleep for instance. The advantage of this methodology is that it provides an aggregated analysis (via the clusters), but capturing important and essential aspects of the activity (the length and frequency of bouts).

With our small sample data, clusters revealed six groups. The first three (C1, C2 and C3) where under the daily recommendations and the other three (C4, C5 and C6) were above these, but each had different characteristics with regards to the occurrence of the MVPA and sedentary times. Cluster movement analysis enables to see students behaviour change in different ways.

Future work will include applying this methodology to larger datasets, exploring varying some of the thresholds used and combine it with more refined sequential pattern analysis.

Acknowledgements

This project was funded by Diabetes Australia Research Trust. We acknowledge all the iEngage team. C. Diaz thanks Universidad Adolfo Ibáñez for their support.

References

[Activinsights Ltd., 2017] Activinsights Ltd. GENEActiv Original - Wrist-Worn Actigraphy Device — GENEActiv Accelerometers, 2017.

[Baker and Yacef, 2009] Ryan S.J.D. Baker and Kalina Yacef. The State of Educational Data Mining in 2009 : A Review and Future Visions. *Journal of Educational Data Mining*, 1(1):3–16, 2009.

[Blikstein and Worsley, 2016] Paulo Blikstein and Marcelo Worsley. Multimodal learning analytics and education data mining: using computational technologies to measure complex learning tasks. *Journal of Learning Analytics*, 3(2):220–238, 2016.

[Esliger *et al.*, 2011] Dale W. Esliger, Ann V. Rowlands, Tina L. Hurst, Michael Catt, Peter Murray, and Roger G. Eston. Validation of the GENEA accelerometer. *Medicine and Science in Sports and Exercise*, 43(6):1085–1093, 2011.

[Fang and Langford, 2013] Zhou Fang and Maintainer Joss Langford. Package ' GENEAread ', 2013.

[Ihaka and Gentleman, 1996] Ross Ihaka and Robert Gentleman. Interface Foundation of America R: A Language for Data Analysis and Graphics R: A Language for Data Analysis and Graphics. *Source Journal of Computational and Graphical Statistics*, 5(3):299–314, 1996.

[Janssen and LeBlanc, 2010] Ian Janssen and Allana G LeBlanc. Systematic review of the health benefits of physical activity and fitness in school-aged children and youth. *International Journal of Behavioral Nutrition and Physical Activity*, 7(1):40, 2010.

[Kelly *et al.*, 2007] Louise A. Kelly, John J. Reilly, Diane M. Jackson, Colette Montgomery, Stanley Grant, and James Y. Paton. Tracking physical activity and sedentary behavior in young children. *Pediatric exercise science*, 19(1):51–60, 2 2007.

[Krebs *et al.*, 2010] Paul Krebs, James O. Prochaska, and Joseph S. Rossi. Defining what Works in Tailoring: A Meta-Analysis OF Computer Tailored Interventions for Health Behavior Change. *Prev Med*, 51(3-4):214–221, 2010.

[Macqueen, 1967] J.B. Macqueen. Some methods for classification and analysis of multivariate observations. *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 1(233):281–297, 11 1967.

[Martinez-Maldonado *et al.*, 2017] Roberto Martinez-Maldonado, Kalina Yacef, Augusto Dias Pereira Dos Santos, Simon Buckingham Shum, Vanessa Echeverria, Olga C. Santos, and Mykola Pechenizkiy. Towards Proximity Tracking and Sensemaking for Supporting Teamwork and Learning. In *Proceedings - IEEE*

17th International Conference on Advanced Learning Technologies, ICALT 2017, pages 89–91. IEEE, 2017.

[Ng *et al.*, 2014] Marie Ng, Tom Fleming, Margaret Robinson, Blake Thomson, Nicholas Graetz, Christopher Margono, ..., and Emmanuela Gakidou. Global, regional, and national prevalence of overweight and obesity in children and adults during 1980-2013: A systematic analysis for the Global Burden of Disease Study 2013. *The Lancet*, 384(9945):766–781, 2014.

[Ochoa, 2017] Xavier Ochoa. Multimodal Learning Analytics. *Handbook of Learning Analytics*, pages 129–141, 2017.

[Phillips *et al.*, 2013] Lisa R S Phillips, Gaynor Parfitt, and Alex V. Rowlands. Calibration of the GENEA accelerometer for assessment of physical activity intensity in children. *Journal of Science and Medicine in Sport*, 16(2):124–128, 2013.

[Plasqui *et al.*, 2013] G. Plasqui, A. G. Bonomi, and K. R. Westerterp. Daily physical activity assessment with accelerometers: New insights and validation studies. *Obesity Reviews*, 14(6):451–462, 2013.

[Ravi *et al.*, 2005] Nishkam Ravi, Nikhil Dandekar, Preetham Mysore, and Ml Michael L Littman. Activity Recognition from Accelerometer Data. In *Proceedings of the Seventeenth Conference on Innovative Applications of Artificial Intelligence(IAAI)*, volume 5518 LNCS, pages 1541–1546. 2005.

[Schaefer *et al.*, 2014] Christine A. Schaefer, Claudio R. Nigg, James O. Hill, Lois A. Brink, and Raymond C. Browning. Establishing and evaluating wrist cutpoints for the GENEActiv accelerometer in youth. *Medicine and Science in Sports and Exercise*, 46(4):826–833, 4 2014.

[Siemens, 2013] George Siemens. Learning Analytics: The Emergence of a Discipline. *American Behavioral Scientist*, 57(10):1380–1400, 2013.

[Sprint *et al.*, 2016] Gina Sprint, Diane J. Cook, and Maureen Schmitter-Edgecombe. Unsupervised detection and analysis of changes in everyday physical activity data. *Journal of Biomedical Informatics*, 63:54–65, 2016.

[Worsley, 2014] Marcelo Worsley. Multimodal learning analytics as a tool for bridging learning theory and complex learning behaviors. *3rd Multimodal Learning Analytics Workshop and Grand Challenges, MLA 2014*, pages 1–4, 2014.

[Yacef *et al.*, 2018] Kalina Yacef, Corinne Caillaud, and Olivier Galy. Supporting Learning Activities with Wearable Devices to Develop Life-Long Skills in a Health Education App. In *Artificial Intelligence in Education Conference*, 2018.