Journal of Biomedical Informatics 63 (2016) 54–65

Contents lists available at ScienceDirect

Journal of Biomedical Informatics

journal homepage: www.elsevier.com/locate/yjbin

Unsupervised detection and analysis of changes in everyday physical activity data

Gina Sprint [a,*], Diane J. Cook [a], Maureen Schmitter-Edgecombe [b]

[a] School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA, United States
[b] Department of Psychology, Washington State University, Pullman, WA, United States

ARTICLE INFO

Article history:
Received 23 February 2016
Revised 8 June 2016
Accepted 22 July 2016
Available online 25 July 2016

Keywords:
Physical activity monitoring
Wearable sensors
Unsupervised learning
Change point detection
Data mining

ABSTRACT

Sensor-based time series data can be utilized to monitor changes in human behavior as a person makes a significant lifestyle change, such as progress toward a fitness goal. Recently, wearable sensors have increased in popularity as people aspire to be more conscientious of their physical health. Automatically detecting and tracking behavior changes from wearable sensor-collected physical activity data can provide a valuable monitoring and motivating tool. In this paper, we formalize the problem of unsupervised physical activity change detection and address the problem with our Physical Activity Change Detection (PACD) approach. PACD is a framework that detects changes between time periods, determines significance of the detected changes, and analyzes the nature of the changes. We compare the abilities of three change detection algorithms from the literature and one proposed algorithm to capture different types of changes as part of PACD. We illustrate and evaluate PACD on synthetic data and using Fitbit data collected from older adults who participated in a health intervention study. Results indicate PACD detects several changes in both datasets. The proposed change algorithms and analysis methods are useful data mining techniques for unsupervised, window-based change detection with potential to track users' physical activity and motivate progress toward their health goals.

1. Introduction

In recent years, sensors have become ubiquitous in our everyday lives. Sensors are ambient in the environment, embedded in smartphones, and worn on the body. Data collected from sensors form a time series, where each sample of data is paired with an associated timestamp. This sensor-based time series data is valuable when monitoring human behavior to detect and analyze changes. Such analysis can be used to detect seasonal variations, new family or job situations, or health events. Analyzing sensor-based time series data can also be used to monitor changes in human behavior as a person makes progress toward a fitness goal. Making a significant lifestyle change often takes weeks or months of establishing new behavior patterns [1], which can be challenging to sustain. Automatically detecting and tracking behavior changes from sensor data can provide a valuable motivating and monitoring tool.

Recently, wearable sensors have increased in popularity as people aspire to be more conscientious of their physical health. Many consumers purchase a pedometer or wearable fitness device in order to track their physical activity (PA), often in pursuit of a goal such as increasing cardiovascular strength, losing weight, or improving overall health. Physical activity is estimated by pedometers and fitness trackers in terms of the steps taken by the wearer [2]. To track different types of changes in physical activity data, two or more time periods, or windows, of PA data can be quantitatively and objectively compared. If the two time windows contain significantly different sensor data then this may indicate a significant behavior change. Existing off-the-shelf change point detection methods are available to detect change in time series data, but the methods do not provide context or *explanation* regarding the detected change. For PA data, algorithmic approaches to change detection require additional information about what type of change is detected and its magnitude to potentially report progress to users for motivation and encouragement purposes. Furthermore, existing approaches often do not provide a method for determining if a detected change is *significant*, meaning the magnitude of change is high enough to suspect it likely resulted from a lifestyle alteration. A personalized, data-driven approach to significance testing for fitness tracker users is a necessary feature of physical activity change detection.

* Corresponding author.
E-mail addresses: gsprint@eecs.wsu.edu (G. Sprint), cook@eecs.wsu.edu (D.J. Cook), schmitter-e@wsu.edu (M. Schmitter-Edgecombe).

http://dx.doi.org/10.1016/j.jbi.2016.07.020
1532-0464/

Currently, there is no clear consensus regarding which change detection approaches are best for detecting and analyzing changes in PA data. Consequently, we formalize the problem of unsupervised physical activity change detection and address the problem with our Physical Activity Change Detection (PACD) approach. PACD is a framework that (1) segments time series data into time periods, (2) detects changes between time periods, (3) determines significance of the detected changes, and (4) analyzes the nature of the significant changes. We review recently proposed change detection methods and we evaluate the ability of four different change detection approaches to capture pattern changes in synthetic PA data. Next, we illustrate how the change approaches are used to monitor, quantify, and explain behavior differences in Fitbit data collected from older adults who participated in a health behavior intervention. Finally, we conclude with discussions about the limitations of current approaches and suggestions for continued research on unsupervised sensor-based change detection.

2. Related work

In the literature, a few studies have aimed to detect change specifically in human behavior patterns. These approaches have quantified change statistically [3,4], graphically [4–6], and algorithmically [5,7–9]. Recently, Merilahti et al. [3] extracted features derived from actigraphy data collected for at least one year. Each feature was individually correlated with a component of the Resident Assessment Instrument for insights into how longitudinal changes in actigraphy and functioning are associated. While this approach provides insight into the relationship between wearable sensor data and clinical assessment scores, this study does not directly quantify sensor-based change.

Wang et al. [5] introduced another activity-based change detection approach in which passive infrared motion sensors were installed in apartments and utilized to estimate physical activity in the home and time away from home. The data were converted into co-occurrence matrices for computation of image-based texture features. Their case studies suggest the proposed texture method can detect lifestyle changes, such as knee replacement surgery and recovery. Though the approach does not provide explanation of the detected changes over time, visual inspection of the data is suggested with activity density maps. More recently, Tan et al. [6] applied the texture method to data from Fitbit Flex sensors for tracking changes in daily activity patterns for elderly participants. Another approach for activity monitoring is the Permutation-based Change Detection in Activity Routine (PCAR) algorithm [7]. PCAR researchers modeled activity distributions for time windows of size three months. Changes between windows were quantified with probabilities of change acquired via hypothesis testing.

The change detection algorithms described previously are intended for monitoring human activity behavior. There are several additional approaches that are not specific to activity data, but instead represent generic statistical approaches to detecting changes in time series data. Change point detection, the problem of identifying abrupt changes in time series data [10], constitutes an extensive body of research as there are many applications requiring efficient, effective algorithms for reliably detecting variation. There are many families of change detection algorithms that are suitable for different applications [11]. Algorithms appropriately handling two sample, unlabeled data are most relevant to the current study due to their data-driven change score computation and no need for ground truth information. Unsupervised change detection approaches include subspace models and likelihood ratio methods [8]. One particular subgroup of likelihood ratio methods, direct density ratio estimator methods, is used in various applications [12,13]. Relative Unconstrained Least-Squares Importance Fitting (RuLSIF) [8] is one such approach used to measure the difference between two samples of data surrounding a candidate change point. Other recent change point detection research includes work on multidimensional [14,15] and streaming time series data [11].

The above approaches are effective methods for detecting change between two samples of data; however, they are not explanatory methods as they only identify if two samples are different and do not provide information on how the samples are different. Once a change is detected and determined significant, additional analyses are required to explain the change that occurred. Hido et al. [9] formalized this problem as *change analysis*, a method of examination beyond change detection to explain the nature of discrepancy. Hido's solution to change analysis utilizes supervised machine learning algorithms, specifically virtual binary classifiers (VCs), to identify and describe changes in unsupervised data. Research by Ng and Dash [16] and Yamada et al. [10] have also explored methods for detecting and explaining change in time series data.

The aforementioned methods provide several options for change detection and analysis, each with their own suitability for various applications. In this paper, we evaluate the following methods for use in our PACD method: (1) RuLSIF [8], (2) texture-based dissimilarity [5,6], (3) our proposed adaptation of PCAR [7] to handle small window sizes (sw-PCAR), and (4) VC-based change analysis [9].

3. Methods

Physical activity is often defined as any bodily movement by skeletal muscles that results in caloric energy expenditure [17]. Physical activity consists of bouts of movement that are separated by periods of rest. Physical activity bouts are composed of four dimensions [17]:

1. Frequency: the number of bouts of physical activity within a time period, such as a day.
2. Duration: the length of time an individual participates in a single bout.
3. Intensity: the physiological effort associated with a particular type of physical activity bout.
4. Activity type: the kind of exercise performed during the bout.

To add exercise throughout the day, individuals can increase their number of bouts (frequency), increase the length of bouts (duration), increase the intensity of bouts, and vary the type of physical activity performed during the bouts. These four components of PA represent four distinct types of changes that can reflect progress toward many different health goals, such as increasing physical activity or consistency in one's daily routine.

We study the problem of detecting and analyzing change in physical activity patterns. More specifically, we introduce methods to determine if a significant change exists between two windows of time series step data sampled from a physical activity sensor. Algorithm 1, PACD, outlines this process. Let X denote a sample of time series step data segmented into days, $D = \{x_1, x_2, \ldots, x_t, \ldots, x_m\}$, where x_t is a scalar number of steps taken at time interval $t = 1, 2, \ldots, m$ and m is the number of equal-sized time intervals in a day. Let t_{mins} denote the number of minutes per time interval, t. For example, if the sampling rate of the wearable sensor device is one reading per minute, $t_{mins} = 1$ min and $m = 1440$ min/t_{mins} = 1400 intervals. Now, let W be a window of n days such that $W \subseteq X$. Furthermore, an aggregate window, $\widehat{W}$, represents the average of all days within the window W: