3-2016

Did you take a break today? Detecting playing foosball using your smartwatch

Sougata SEN
Singapore Management University, sougata.sen.2012@smu.edu.sg

Kiran K. RACHURI
Samsung Research America

Abhishek MUKHERJI
Samsung Research America

Archan MISRA
Singapore Management University, archanm@smu.edu.sg

Citation

SEN, Sougata; RACHURI, Kiran K.; MUKHERJI, Abhishek; and MISRA, Archan. Did you take a break today? Detecting playing foosball using your smartwatch. (2016). *2016 IEEE International Conference on Pervasive Computing and Communications Workshops: Sydney, March 14-18 March: Proceedings*. 7457165:1-6.
Available at: https://ink.library.smu.edu.sg/sis_research/3608

The Second IEEE International Workshop on Sensing Systems and Applications Using Wrist Worn Smart Devices, 2016
https://doi.org/10.1109/PERCOMW.2016.7457165

Did You Take a Break Today?: Detecting Playing Foosball Using Your Smartwatch

Sougata Sen‡, Kiran K. Rachuri†, Abhishek Mukherji†, Archan Misra‡
‡Singapore Management University and †Samsung Research America
‡{sougata.sen.2012,archanm}@smu.edu.sg, †{k.rachuri,a.mukherji}@samsung.com

***Abstract*—Prolonged working hours are a primary cause of stress, work related injuries (e.g, RSIs), and work-life imbalance in employees at a workplace. As reported by some studies, taking timely breaks from continuous work not only reduces stress and exhaustion but also improves productivity, employee bonding, and camaraderie. Our goal is to build a system that automatically detects *breaks* thereby assisting in maintaining healthy work-break balance. In this paper, we focus on detecting *foosball breaks* of employees at a workplace using a smartwatch. We selected foosball as it is one of the most commonly played games at many workplaces in the United States. Since playing foosball involves wrist and hand movement, a wrist-worn device (e.g., a smartwatch), due to its position, has a clear advantage over a smartphone for detecting foosball activity. Our evaluation using data collected from real workplace shows that we can identify with more than 95% accuracy whether a person is playing foosball or not. We also show that we can determine how long a foosball session lasted with an error of less than 3% in the best case.**

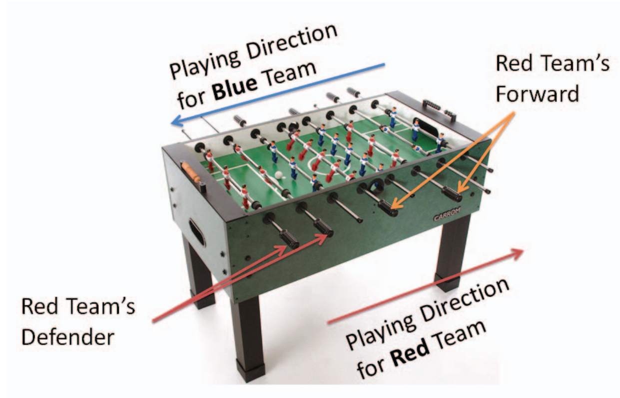

Fig. 1: Standard Foosball Table

I. Introduction

Prolonged working hours at workplaces has many adverse effects on employees. It is a primary cause of stress for many people impacting their health and well-being. Project deadlines and late night working hours further exacerbate work-life balance of employees. According to a work-stress survey conducted in 2013, 83% of American workers say that they feel stressed out by their jobs. Taking frequent breaks can reduce stress and exhaustion, further, it has a positive impact on the employees productivity as reported by multiple studies [4], [18]. Long continuous working hours also increase the risk of Repetitive Stress Injuries (RSI). As reported by The National Institute for Occupational Safety and Health (NIOSH) [1], musculoskeletal disorders are the most costly category of workplace injuries. The U.S. spends $100 billion on lost productivity, employee turnover, and other expenses, and $20 billion annually on workers' compensation costs due to RSIs. It is, therefore, important to encourage employees to take frequent breaks from continuous work to reduce their risk of RSIs and improve their overall well-being [2].

The rate of adoption of wearable devices such as smartwatches and wrist-bands has been rising continuously. Moreover, most of the sensors required for physical activity detection such as accelerometer, gyroscope, GPS are already included in smartwatches (e.g., Apple Watch, Samsung Gear S2). Wrist-worn wearables, due to their form factor and position, have many advantages over smartphones in determining certain activities. Physical activities that are primarily dependent on hand and wrist movement can be more accurately detected using a wrist-worn device. For example, smoking, playing foosball, table tennis, driving etc. involve unique hand movements that can be more accurately detected using the user's smartwatch than her smartphone. Therefore, in this work, we use a smartwatch to detect the physical activities performed during a break. During breaks, employees at a workplace might get a coffee, go for a walk, play a game of table tennis or foosball, or just chat with a colleague. Although an employee can perform any of these diverse set of activities, in this paper, we specifically focus on detecting *playing foosball* as it is one of the most commonly played games at workplaces in the United States. If a user performs any other activity, for example, walking or running, then our system can be extended using existing work [11], [13] to detect these activities. If an employee goes outside of her work location, this can be easily inferred using their location detected from Wi-Fi and GPS signals.

In this paper, we present a system that detects whether the user is playing foosball or not. The system comprises of a data acquisition component that captures data from smartwatch sensors such as accelerometer and gyroscope. The data is then split into windows (e.g., 2 seconds), and then passed into a feature extraction module. A binary classifier uses these extracted features to detect whether the user is playing foosball or not. A smoothing technique is applied on top of the classified output to smoothen the output to accurately

This work was supported partially by Singapore Ministry Education Academic Research Fund Tier 2 under research grant MOE2011-T2-1001 and partially by the National Research Foundation, Singapore under its Interactive Digital Media (IDM) Strategic Research Programme. All findings and recommendations are those of the authors and do not necessarily reflect the views of the granting agency, or Singapore Management University.

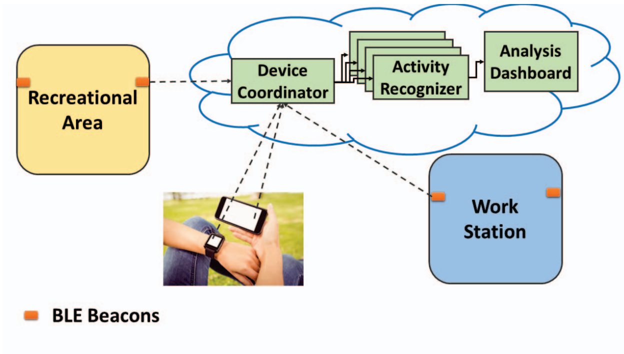

Fig. 2: Architecture of the System

infer the duration of a foosball break. The inferred foosball break events are logged into a database on the smartwatch for later retrieval. Although there is considerable work in the field of smartphone based physical activity detection [19], we believe that smartwatch based activity recognition provides unique advantages in determining activities such as *playing foosball* due to its position, which we exploit. It might not be possible to determine the *playing foosball* activity using only a smartphone. Our main contributions are the following:

- We present design, implementation, and evaluation of a smartwatch based system for detecting *playing foosball* activity.
- We show using dataset collected from a real workplace that with an accuracy of over 95%(precision=82.2%), we can determine whether a person is playing foosball.
- We also show that we can estimate the average duration of a foosball session with a 3% error.

II. System Architecture

To realize a system which can identify when a person is taking a break, coordination between multiple personal devices is necessary as different activities performed when a person is taking a break might be identifiable through separate devices. For example, gaming activities such as playing foosball or snooker/pool might be easier to identify using smartwatches, whereas activities such as cycling might be easier to identify using smartphones, if the smartphone is kept in the pant pocket. Other than a person's personal devices, infrastructure sensors can also be helpful in providing usable hints regarding individual's break times - e.g. using Wi-Fi based indoor localization to identify that the person is in the pantry.

Figure 2 shows the overall architecture of the system which can be helpful in identifying breaks. A central server is needed which can orchestrate between multiple devices to promptly identify if the user is taking a break. The processing pipeline of the central server is explained in Figure 3. The sensing pipeline consists of the following components:

- Framing stage: The raw accelerometer data is broken down into multiple frames. The frame length is tunable. For our dataset, we empirically determined the frame length to be 0.5 seconds. We used a sliding window with an overlap of 50% between two successive frames.
- Feature extraction: For each frame, features are extracted. These extracted features line in both time and frequency domain (similar to [19]). The summary of features extracted is listed in Table I.
- Gesture Classification: Once the features have been extracted from the data, the feature instances are passed through multiple classification models, where each model is a binary classifier to determine a particular *break activity*. Based on our empirical evaluation of multiple classifiers, we found that Random Forest usually has higher accuracy with tolerable latency. We have thus used Random Forest in all our evaluation.
- Smoothing: Finally output of the classifier has to be smoothed. A lot of false positives can appear when we consider a small window size. We thus have to smooth the output to remove these false positives. Currently we have used a dominant set labeling approach for smoothing. However in future, we can use more sophisticated smoothing techniques to have a better classification accuracy.

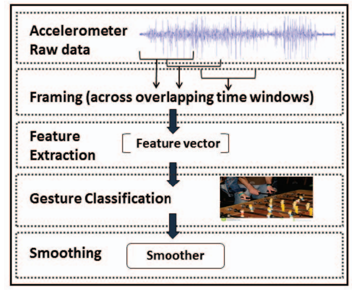

Fig. 3: Activity Recognition module pipeline

III. Activity Detection

We next provide some background of foosball and then discuss about the data that we collected and the gesture identification that we performed.

A. Background

Figure 1 represents a typical foosball table. A foosball table consists of 8 bars (4 on each side) to which 11 to 13 miniature figures representing a soccer player is attached. The bars have to be moved forward and backward to align the players with the ball and rotated to strike the ball. The game can be played as singles, where one person controls all four bars or doubles where two players control the four bars. In a singles game, the player has to switch bars, based on the position of the ball, while in the doubles game, a player acts either as a defender (control the first two rods from own goal) or as a forward (control rods three and four). The objective of the game is to score a certain number of goals (typically 8 to 10) before the opponent scores the same number of goals and a game is considered over when either of the teams has scored the pre-decided number of goals.

B. DataSet

We collected data from 5 individuals who were part of the same research unit in an organization. The data was collected in phases over a period of one month. A custom smartwatch application was deployed on the Samsung Gear 2 watch. This application would collect accelerometer and gyroscope data

Feature Name	Count	Description
Mean	4	The mean of the sensor readings for X,Y,Z axis and the magnitude for a framing window
Variance	4	The variance of the three axis of accelerometer and the magnitude in a framing window
Covariance	3	The covariance of each of the accelerometer axis w.r.t. the other axes
Energy	4	The energy of the X,Y,Z axis of accelerometer and the magnitude
Entropy	4	The entropy of the three axis of accelerometer and the magnitude

TABLE I: List of features extracted from the raw accelerometer readings

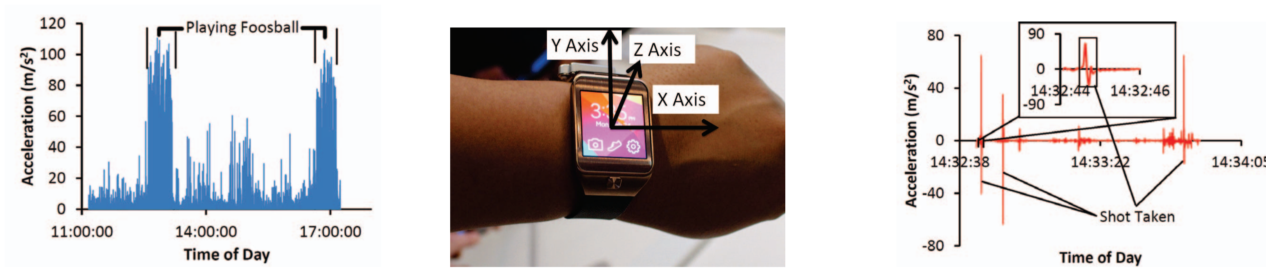

(a) Entire day's plot of magnitude of the 3 axis accelerometer

(b) Direction of different accelerometer axes

(c) Variation of Y axis of accelerometer when person is playing

Fig. 4: Visualizing variations in raw sensor readings

from the smartwatch. A total of 23 person-day data was collected using this application, where the *Start of day* was between 10 AM to 12 noon and the *End of Day* was any time between 4 PM and 6:30 PM. During these 23 person-days, a total of 81 hours data was obtained. On all these days the individuals played foosball atleast once during the day. Other than foosball, in 4-person-days data, the individuals also played table tennis. The table tennis data has not been removed from our dataset when we performed the evaluation. We did not try identifying table tennis gesture due to the limited amount of table tennis data. For the foosball data, the total number of times the individuals played foosball (foosball session) was 74 times which resulted in 709 minutes of data. All foosball games played were doubles, where four people were playing at a time. The playing four in any game was not restricted to the 5 individuals whose data was collected. Anyone in the organization was free to join in. Overall, in 30 cases, the data collected was from the individual who was playing as a forward, while in 44 cases, the individual was playing the role of the defender.

To collect the ground truth, the custom application on the watch had a user interface, which had options where the user would mark when they started playing a game, when they finished the game, whether they were playing offense or defense and when a goal was scored. Every gesture during the game was not marked individually. We next discuss about activities we are tying to detect.

C. Did You Play Foosball Today?

Since the data was collected on days when the users played foosball atleast once, we wanted to see how accurately we could predict the playing time. To understand if this identification was even possible, we plotted the magnitude of the 3 axis accelerometer to see if there was any variation during the playing period. Figure 4a shows the variation of the magnitude over the period of a day. When a person is playing foosball, we expect to see a lot of very fast hand movements. From the figure, we can clearly see the huge spikes when the person played foosball. We zoom into the data and plotted the Y axis data because the maximum variation was observed in the Y-axis as that is the direction in which hand is rotated (See Figure 4b). The variation is shown in Figure 4c. From this we can see that when a person is taking a shot there is initially positive acceleration and then negative acceleration in the Y-axis. This is because the hand has to be rolled in clockwise initially to raise the players for power in the shot and then it is rolled back anti clockwise to take the actual shot. Other than the huge variation in the Y axis during shot taking, there is continuous variation in the X-axis (not plotted) because during the game, the player has to continuously move the rod forward and backward to align the players to gain possession of the ball.

From the datasets, we wanted to answer two questions : (i) how accurately can we identify the foosball gesture and (ii) can we group all the foosball gestures together to declare a period of time as a foosball playing period. We next describe our approach in answering the two questions.

1) Identifying the playing gesture: We can see that when a person is playing foosball, there is a period of some short, bursty, high acceleration gestures, which is not visible during other times of the day. We wanted to identify this period and mark it as the period when the person took breaks. To identify this, we used a Random Forest classifier with *number of trees* set to 10. We extracted the features described in Table I. Based on empirical evaluation, we found that a Foosball shot gesture takes ≈ 1.3 seconds, where the actual shot taking part (highlighted in the zoomed Figure 4c takes less than 0.5 seconds). We had to ensure that we could identify this short high acceleration period. We varied the window size which was used to calculate features between 0.5 seconds to 5 seconds (with 50% overlap) to understand how accurately we could identify 'Playing Foosball'. Based on the data we found that the wrist rotation speed (which in turn resulted in acceleration) varied across individuals. We also found that for the same individual, there was variation across shots (as seen in Figure 4c).

2) How long did you play?: To identify playing sessions, we examining the temporal density of the fast hand gesture. Since a player will take multiple shots in a game, there should be multiple high acceleration readings in a small window. If we could identify the time period between the first high acceleration shot and the last one during a session, we could approximately identify the break taken by the individual. To identify the break session, we use the entire day's trace. We considered a time window w and counted the number of gestures classified as playing foosball in this window. If the count of gestures in window was above a threshold t, we considered this period as a 'playing foosball' period. From the ground truth data we found that on average, games in the dataset lasted for 9 minutes 35 seconds with a minimum of 6 minutes 10 seconds. Since a foosball session in our study was always more than 6 minutes, we varied w so that it was could cover even a 6 minute window. We considered the starting of a foosball session when t gestures in a window was determined as playing foosball and considered the session to end when a window was determined as not playing foosball.

D. What Was Your Role In The Game?

Till now we were trying to determine the playing moment. We next wanted to see if we could determine the role of the player in the game. Identifying this can be interesting because if we can help providing statistics to players at the end of a certain period of time regarding how they played in certain positions, against certain opponents, with certain team mates etc. To identify the role of a player in a game, we took the data from the actual playing period of the individuals, determined by the ground truth data. We used the features that were extracted from the best framing window size. We separated out data based on sessions. We then performed a leave one session out cross validation to determine the role of the player. For the evaluation, we used a random forest classifier with *number of trees* set to 10.

IV. Evaluation

We next describe the performance of our system.

A. Did You Play Foosball Today?

To identify the time of playing foosball, we had to identify (i) Foosball Playing Gesture: to determine if an individual is playing foosball and (ii) Playing Period Estimation: to determine how long the foosball game lasted.

1) Foosball Playing Gesture: To identify the gesture, we took the data 23 person-days data of the 5 users and performed a leave-one-person-day-out-cross validation. We took data from 22 days and created a training set, where the class labels in the training set was either *Playing* or *Not Playing*. The data from the 23^{rd} day was used for testing. Based on our evaluation we found that we could obtain a precision of 82.2% and a recall of 90.49% when we used a window length of 0.5 seconds, which decreases to Precision=54.7% and Recall=82.9% when we varied the window to 5 seconds. The plot for varied window lengths is plotted in Figure 5. From this we see that the precision and recall increases as

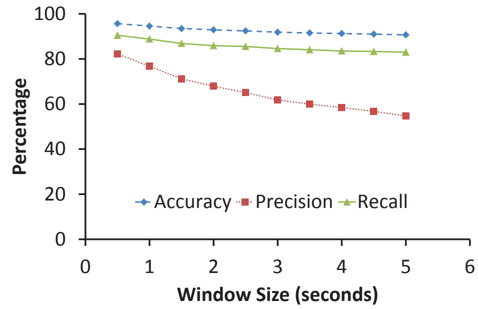

Fig. 5: Effect of varying the feature calculation window size on accuracy, precision and recall

we decrease the feature calculation window size. One reason for this can be since gestures in foosball are short and quick, it can be identified more accurately in smaller frames. Thus based on this, we decide to use a window size of 0.5 seconds for further calculations. Other than the precision and recall, we also computed the accuracy of foosball gesture identification and found the accuracy to be at 95.65%. The accuracy of the system will be high because of the imbalance in the two classes (plating foosball and not playing foosball). In our dataset, less than 10% of the data is from instances when the individual is playing foosball.

2) Playing Period Estimation: To identify a playing session, we considered windows of size $w = \{1,2,3,4,5\}$ minutes. To determine if w was a *Playing Foosball* window, we counted the number of foosball gestures that were identified in the window. If the number of foosball gestures were above a threshold $t = 1,2,\cdots,30$, then we declared the window to be a *Playing Foosball* window. Since we were using a 0.5 seconds framing window, we had 120 classifications in a minute. We discarded any session prediction if the predicted session was less than five minutes. From the data we found that setting the threshold of gestures too low results in over estimation in the number of sessions predicted in a day. While setting it too high led to under estimation. The reason for this is because there can be multiple false positives during the day while setting this threshold too high resulted in under estimation. We can compute the error in predicting the number of breaks as

$$Error = \frac{ActualNo.of breaks - PredictedNo.of breaks}{ActualNo.of breaks} \quad (1)$$

Here a value close to zero is what is desirable. A positive value indicates that we are underestimating, while a negative number indicates we are over estimating the number of breaks taken. Figure 6 shows the error in predicting the number of breaks that are taken when window size w is varied from 1 minute to 5 minutes. The X axis of the graph represents the number of gestures considered(t). When the number of gestures in a time window is low, there is always over estimation in the number of sessions. However the error gradually reduces to 0 as we increase the number of gestures.

We next wanted to see how accurately we could identify the total break time. Figure 7 shows the variation in prediction

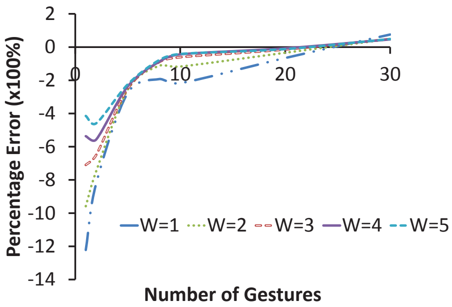

Fig. 6: Error in predicting number of breaks taken when window size (w) is varied

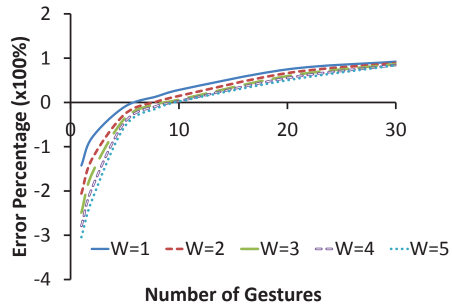

Fig. 7: Error in predicting the duration of a break when when window size (w) is varied

of break duration when the number of gestures threshold, considered is varied. From the figure we again find that with a low number of gestures, we always over estimate the duration of breaks taken. A lot of false positive in a day gets added to the break time. However if we considered too many gestures in *w*, we end up under estimating the duration of breaks. The reason for this is, there might not be a lot of foosball gesture in the entire break duration (e.g. when switching between games or after a goal has been scored) and thus the threshold might not be reached. From analysing the data, we find that we can get an average estimation error of less than 3% when we consider a time window of 2 minutes and we consider 8 gestures in those two minutes. However this value varies across individuals. A point to note here is that our gesture estimation is from only one hand (as the player wore only one watch). If we assume that a player plays equally with both the hands, then the total number of gestures in the 2 minutes window might be 16. In future if we have data from both the hands, we will try to predict this value more accurately.

B. Role in Game

We finally wanted to determine the role of a person (whether forward or defender) in the game. We performed a leave one session out cross validation to determine an individual's role in the game. During one session multiple games could have been played and it was not necessary that the role of a player had to be the same in all the games played in a session. Based on our evaluation we found that with 59.4% accuracy, we could identify the role of a player. The reason for the low accuracy in role determination is because a person's playing style remains almost invariant whether playing as a forward or as a defender. However, based on observation, we saw that the moving pattern for the goal keeper bar (primary role is to save the shot on target) was different from that of the first row of forward (primary role was to take a quick gestured shot to score a goal). However in our study, since each person was wearing the watch in their dominant hand, we do not have data from the hand controlling the goal keeper bar. In future, if we have data from both hands of an individual, we can correlate the data between the left and right hand, which might help in identifying a person's role more accurately.

V. Discussion and Future Work

We have shown the feasibility of implementing a system which can help in detecting one type of break. There are multiple possible extensions to the system which we are listing down here.

A. Break Determination

Our current implementation aims at identifying if a user is playing foosball. However, foosball might be one of many activities that can be performed when a user is taking a break. There are other recreational activities which an individual can engage in - e.g. table tennis or snooker/pool. Other non sport activities such as having coffee might also occur when a person takes a break. As a logical extension of this work, in future, we intend to identify other *break-taking* activities that are performed by individuals during working hours. For a robust estimation, data from multiple devices (smartphone or infrastructure sensor) might be needed. We foresee the use of a central coordinator in coordinating between various devices to decide if a person is taking a break.

B. Sensing Modality

We have currently used only the accelerometer. We did not use the gyroscope because the best possible sampling rate achievable for the gyroscope in the watch we used was quite low for any meaningful inference. In an activity like foosball, the gyroscope might provide a higher accuracy as there is rotation along one of the axis of the watch. However the gyroscope can drain out the battery very fast. In future we will investigate how we can use the accelerometer to trigger the gyroscope to capture such gestures more accurately with minimum increase in battery drain rate. We also intend to take advantage of other sensors (e.g. GPS) to improve the activity recognition. Infrastructure sensors can also be augmented with the sensors on the wearables for robust activity recognition.

C. End User Application

We currently have implemented a basic foosball gesture identifier which can help in determining when a person is taking a break. However, currently we have not explored possibilities of generating statistical data from the activity such as identifying how many goals an individual scored or during a month who did the individual partner with in the maximum number of occasions. In future we plan to explore predicting such statistics. Currently, we tried identifying the role of a player using simple gesture recognition. In future, we intend to study how the ball moves across the table and based on the sequence in which players interact with the ball, can the role of a player be determined.

Currently we have shown the feasibility of building one *break taking* activity detector. If we have multiple such activity recognition systems, we can use the identification of breaks for self reflection of the individual. Statistics such as amount of time spent at workstation in an office as compared to recreational areas can help individuals in maintaining a good work life balance. If the amount of time spent in each activity can be presented to the individual, it can help her in planning her days better.

VI. Related Work

Unobtrusive physical wellness monitoring has always been of interest to researchers. As smartwatches and wristbands are becoming common place, more people are embracing them for fitness tracking such as step counts during the day [9]. Efforts such as [15], [19] have established that continuous activity recognition is possible on consumer mobile devices, while work such as [5] shows how continuous activity recognition is possible in wearable devices accurately and in an energy efficient manner.

Activity Recognition using wrist worn devices to determine physical wellness has been a well-studied topic [6], [7]. Parate et al. [14] used a wrist worn device to determine smoking, while Amft et al. used it to understand food consumption [3].

Techniques such as [12], [16] have used the way a person is interacting with mobile devices to determine the individual's emotion and social behavior. Additionally there is scope for learning through wrist-worn devices how healthy (both physically and mentally) a lifestyle a person is leading [8], [17]. Yet another aspect of interest is a person's social interactions. Knighten et al. [10] recognize social interactions of the users through numerous gestures such as hand shake, fist bump, high five and pointing. Such social gestures may be indicative of person's mental and social well-being as well.

Our paper is orthogonal to these works and focuses on identifying breaks taken by the person during work hours in the form of activities such as foosball or table tennis. The aim is to monitor if the user takes sufficient breaks during work. The overreaching goal is to encourage timely breaks and interactions with colleagues through activities such as foosball that in turn improves productivity at work place [4], [18].

VII. Conclusion

In this work we have shown that by using off the shelf smartwatch, we can detect one of the most popular office recreational activity in the United States - playing foosball. We have shown how a system can be built where a break identification activity such as foosball can help in managing a healthy work life balance. With 23 day long data collected from 5 users, we show that our system can predict when an individual is playing foosball with over 95% accuracy and precision over 82%. We also show how in future we plan to extend this work to get more insights about the game as well as identify other *break* related activities.

References

[1] Prevention of job-related musculoskeletal disorders draws national audience at conference hosted by niosh, osha. http://www.cdc.gov/niosh/updates/ergprs.html.

[2] Repetitive strain injury. http://web.eecs.umich.edu/~cscott/rsi.html.

[3] Amft, O., Stäger, M., Lukowicz, P., and Tröster, G. Analysis of chewing sounds for dietary monitoring. *Proceedings of the 7th International Conference on Ubiquitous Computing*, UbiComp'05, pages 56–72, Berlin, Heidelberg, 2005. Springer-Verlag.

[4] Ariga, A. and Lleras, A. Brief and rare mental "breaks" keep you focused: deactivation and reactivation of task goals preempt vigilance decrements. *Cognition Journal*, 2011.

[5] Berlin, E. and Van Laerhoven, K. Detecting leisure activities with dense motif discovery. *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, UbiComp '12, pages 250–259, New York, NY, USA, 2012. ACM.

[6] Cheng, H.-T., Sun, F.-T., Griss, M., Davis, P., Li, J., and You, D. Nuactiv: Recognizing unseen new activities using semantic attribute-based learning. *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services*.

[7] De, D., Bharti, P., Das, S. K., and Chellappan, S. Multimodal wearable sensing for fine-grained activity recognition in healthcare.

[8] Dong, Y., Hoover, A., Scisco, J., and Muth, E. A new method for measuring meal intake in humans via automated wrist motion tracking. *Applied psychophysiology and biofeedback*, 37(3):205–215, 2012.

[9] Fitbit. http://www.fitbit.com/, December 2015.

[10] Knighten, J., McMillan, S., Chambers, T., and Payton, J. Recognizing social gestures with a wrist-worn smartband. *IEEE International Conference on Pervasive Computing and Communication Workshops, PerCom Workshops 2015*.

[11] Lu, H., Yang, J., Liu, Z., Lane, N. D., Choudhury, T., and Campbell, A. T. The jigsaw continuous sensing engine for mobile phone applications. *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*.

[12] Luo, L., Yang, J., Bao, X., Yan, Z., and Jiang, Y. Swan: A novel mobile system to track and analyze social well-being. *Adjunct Proceedings of the 2015 ACM International Symposium on Wearable Computers*.

[13] Nguyen, L. T., Zeng, M., Tague, P., and Zhang, J. Recognizing new activities with limited training data. *Proceedings of the 2015 ACM International Symposium on Wearable Computers*.

[14] Parate, A., Chiu, M.-C., Chadowitz, C., Ganesan, D., and Kalogerakis, E. Risq: Recognizing smoking gestures with inertial sensors on a wristband. *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services*.

[15] Park, T., Lee, J., Hwang, I., Yoo, C., Nachman, L., and Song, J. E-gesture: A collaborative architecture for energy-efficient gesture recognition with hand-worn sensor and mobile devices. *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*.

[16] Rachuri, K. K., Musolesi, M., Mascolo, C., Rentfrow, P. J., Longworth, C., and Aucinas, A. Emotionsense: a mobile phones based adaptive platform for experimental social psychology research. *UbiComp 2010: Ubiquitous Computing, 12th International Conference,*.

[17] Shoaib, M., Bosch, S., Scholten, H., Havinga, P. J. M., and Incel, Ö. D. Towards detection of bad habits by fusing smartphone and smartwatch sensors. *IEEEInternational Conference on Pervasive Computing and Communication Workshops, PerCom Workshops 2015*.

[18] Waber, B., Olguin, D. O., Kim, T., and Pentland, A. Productivity through coffee breaks: Changing social networks by changing break structure. *Proceedings of International Sunbelt Social Network Conference'10*, 2010.

[19] Yan, Z., Subbaraju, V., Chakraborty, D., Misra, A., and Aberer, K. Energy-efficient continuous activity recognition on mobile phones: An activity-adaptive approach. *Wearable Computers (ISWC), 2012 16th International Symposium on*.