Movement and Orientation Visualization using Wearable Inertial Sensors

Md Abu Sayeed Mondol, John A. Stankovic
University of Virginia
Charlottesville, Virginia 22904-4740
{mondol,stankovic}@virginia.edu

ABSTRACT

Activity recognition using wearable sensors has widespread and important applications in different domains including healthcare, safety and behavior monitoring. Almost all the solutions for gesture and activity recognition are data-driven, and so, understanding the data and their characteristics is fundamental toward developing effective solutions. Visualization is perhaps the most effective approach for getting insight into data as well as communicating the insights with others. A novel visualization method that would provide additional utility to the existing methods is of utmost desire. However, such novel methods for visualization are rarely invented, particularly in the area of wearable and mobile sensing. This paper presents novel methods for visualizing movement and orientation using inertial sensors. we demonstrate the use of the methods for visualizing several activities and gestures. We also developed an efficient method for smoking puff detection leveraging the visualization methods.

CCS CONCEPTS

•Computer systems organization →Sensors and actuators;

KEYWORDS

Activity Recognition, Wearable Sensor, Visualization

ACM Reference format:
Md Abu Sayeed Mondol, John A. Stankovic. 2020. Movement and Orientation Visualization using Wearable Inertial Sensors. In *Proceedings of 17th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services, Cyberspace, December 2020 (Mobiquitous 2020),* 10 pages.
DOI: 10.1145/nnnnnnn.nnnnnnn

Mobiquitous 2020, Cyberspace

DOI: 10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Smart wearable devices enriched with inertial sensors are widely used nowadays. The inertial sensors available in the devices can be used for activity and gesture recognition that have widespread applications in different areas including healthcare, safety, behavior monitoring, energy management, manufacturing, and elderly care. Almost all the solutions for activity and gesture recognition are data-driven, and so, understanding the data and their characteristics is fundamental toward developing effective solutions. Visualization is perhaps the most effective approach for getting insight into data as well as communicating the insights with others. A novel visualization method that provides additional utility to the existing methods is of utmost desire. However, such novel methods for visualization are rarely invented, particularly in the area of wearable and mobile sensing. Here, we present novel methods for visualizing movement and orientation using inertial sensors. Our visualization method is not a replacement to the existing methods; rather it provides additional utility toward better understanding the movements associated with different activities.

Inertial sensors provide data for three orthogonal directions, namely along X, Y, and Z-axes. Accelerometers, gyroscopes and magnetometers sense acceleration, rotation rate, and direction of the device, respectively. The data from these sensors are time-series in nature and usually visualized as a series of data points, often connected by straight line segments. These plots are useful in understanding the data. Our method complements the existing approaches and provides ways to visualize the data from different perspectives. Our methods are based on the orientation of the device where the time series orientation data are organized in 3-dimensional space, resulting in spatio-temporal representations. We extract the orientations of the device from quaternions that are usually constructed by fusing data from an accelerometer, a gyroscope, and a magnetometer. Our methods consist of creating four visualization-based primitives: Orientation Trace, Spatio-temporal Representation, Orientation Reachability, and Sphere Segmentation. A quaternion is decomposed into three unit vectors representing the orientation of the device with respect to earth gravity, north, and east. These vectors

are placed on the surface of a unit sphere to visualize orientation trace in the spatio-temporal dimension. To better visualize where the orientation reaches for some activities, the sphere is segmented into some nearly uniform cells.

The method presented here are not activity or gesture specific. It can be used to visualize movement and orientation for different activities. This paper presents examples of several activities of daily living to demonstrate the use of our method in understanding the underlying movement associated with the activities. Our visualization methods help to analyze and understand data that can further help in developing effective and efficient solutions. In this paper, we analyze a dataset related to smoking activity recognition and provide several insights through representing the data using our methods. We have also developed an efficient solution for smoking puff detection using the concept of orientation reachability. Finally, we illustrate how our method can be used to monitor movement patterns related to some rehabilitation exercises.

The major contributions of this paper are listed below.

- This paper presents novel methods for movement and orientation visualization using inertial sensors.
- It shows the usage of the methods in understanding the movement associated with several activities in daily living.
- It presents an efficient solution for smoking puff detection leveraging the visualization methods.
- It demonstrates the use of the methods for monitoring rehabilitation exercises.

2 VISUALIZATION METHODS

Our methods are based on device orientations that are expressed by quaternions. A quaternion is a representation of the orientation of a coordinate system with respect to another coordinate system. In the context of smart devices like smart phones and smart watches, a quaternion represents the orientation of the device with respect to the world coordinate system. The device coordinate system may differ based on how sensors are placed inside the device. For Android devices, the X axis usually points to the right of the device, the Y axis points up, and the Z axis points perpendicular to the display, as shown in Figure 1(b). For the world coordinate system (Figure 1(a)), the Z axis is perpendicular to the ground pointing opposite to the earth's gravity, the Y axis is tangential to the ground at the device's current location and points towards the magnetic north, and the X axis, defined as the vector product of Y and Z, is tangential to the ground at the device's current location and roughly points to the East. A quaternion q is formally defined by a scalar component (q_s) and a 3D vector (q_x, q_y, q_z) as:

$$q = q_s + q_x\hat{i} + +q_y\hat{j} + q_z\hat{k} \quad (1)$$

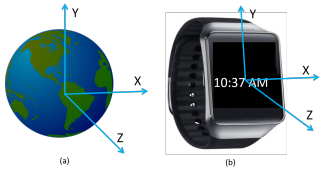

Figure 1: Axes of the coordinate systems for (a) the world and (b) an Android watch.

where i, j, and k are the imaginary basis elements. The quaternion is called a unit quaternion if its magnitude equals to one, i.e. if $|q| = \sqrt{q_s^2 + q_x^2 + q_y^2 + q_z^2} = 1$

The quaternion values depend on the orientation of the device with respect to both the direction of earth gravity and the direction of earth magnetic field. So, the quaternions for the same gestures are different when the gestures are performed facing in different magnetic directions (e.g, north or south). Consequently, it is difficult to interpret the movement and orientation of a device using quaternions. We use rotation matrix, an equivalent representation of the quaternion, to decompose the orientation with respect to the earth gravity and magnetic directions. The rotation matrix R, equivalent to the quaternion q, is defined as:

$$\begin{bmatrix} q_s^2 + q_x^2 - q_y^2 - q_z^2 & 2(q_xq_y - q_sq_z) & 2(q_xq_z + q_sq_y) \\ 2(q_sq_z + q_xq_2) & q_s^2 - q_x^2 + q_y^2 - q_z^2 & 2(q_xq_z + q_sq_y) \\ 2(q_sq_z - q_sq_2) & 2(q_yq_z + q_sq_x) & q_s^2 - q_x^2 - q_y^2 + q_z^2 \end{bmatrix}$$

The rotation matrix is orthogonal, and so the rows and columns of R are orthonormal i.e. orthogonal unit vectors. The first, second and third row of the matrix represent orientation of the device with respect to east (x-axis), north (y-axis), and gravity (z-axis) of the world, respectively. The first, second and third element of a unit vector (row) represent the projection of that vector along the the X, Y and Z axes of the device, respectively. It should be noted that multiplying the unit gravity vector (the third row) by the magnitude of gravitational acceleration (g) results the actual gravitational acceleration along the different axes of the device. Though the value of g is constant at a location, it is somewhat different at different locations on the earth. In contrast, the magnitude of the unit vector is fixed (i.e. 1) irrespective of the location on the earth.

The values of accelerations (from an accelerometer) and rotation rates (from a gyroscope) depend on the speed or intensity of the movement of the device. In contrast, quaternions or the unit vectors change only when the orientation of the device changes. For example, the accelerations and the unit gravity vectors along X axis from a wrist device for brushing teeth at different speeds are shown in Figure 2(a) and 2(b), respectively. Due to space limitation, we do not

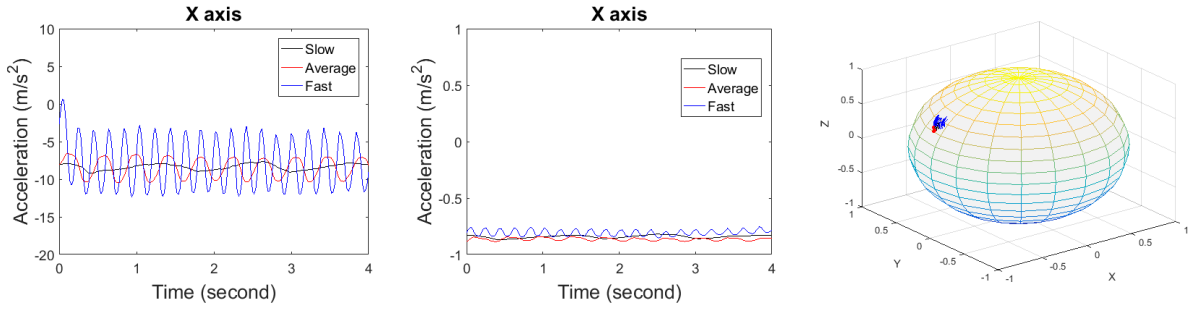

(a) (b) (c)

Figure 2: (a) Accelerations and (b) unit gravity vectors from a wrist device along X-axis for brushing teeth in a specific way at different speeds (slow, average, fast), (c) The unit gravity vectors (orientation trace) in 3D space.

show the accelerations and unit gravity vectors along the Y and Z axes, but their characteristics are similar to those for the X-axis. As shown, speed of movement results significant difference in acceleration signals. In contrast, the gravity vectors do not change significantly, indicating no significant change of the orientation of the wrist. Here, the unit vectors provide complementary insight about the wrist movements. In this section, we present a set of novel primitives for movement visualization using the unit vectors.

2.1 Primitive 1: Orientation Trace

Since the vectors are unit in length, they can be placed on the surface of a unit sphere. The trace of the vectors on the sphere helps to understand the movement and orientation of the device, and thus of the body part where the device is placed. Figure 2(c) shows the gravity vectors on the surface of a unit sphere for brushing teeth at different speeds, where as Figure 2(a) and 2(b) show data from X axis for accelerations and unit gravity vectors in 2-dimensional space, as mentioned earlier. From the trace, we see that the orientation of the wrist for brushing teeth in a specific way is nearly the same even though the intensity of the wrist movement is different. Here, the orientation traces on the sphere is limited to some specific area irrespective of the time duration of the data.

The magnetic direction is irrelevant for most of the human activities. For example, a person can eat facing in any direction, and so, the direction might not be useful in understanding and recognizing eating gestures. However, direction provides additional information in several contexts. For example, Figure 3(a) shows the unit gravity vectors from a wrist device, respectively, for opening two similar doors facing to different directions. Figure 3(b) shows corresponding unit north vectors. Here, the unit gravity vectors are similar and not suitable in distinguishing the doors. However, the doors can be easily distinguished by the trace of the unit north vectors. Though a magnetometer also provides direction information, its values are significantly distorted by magnetic materials (e.g. metal) and magnetic fields [11]. In contrast, quaternions are more stable as they are generated through fusion of multiple sensors.

2.2 Primitive 2: Spatio-temporal Representation

The spatio-temporal representation, as defined by the location of the unit vectors on the sphere and by the order of the unit vectors over time, is robust against the speed of a gesture. It is because the orientation trace of the device is nearly the same for similar gestures even though the speeds of performing the gestures are different. Due to difference in speed, the duration of an activity or gesture differs. Figure 4 (a) shows unit gravity data for two apple bites, one bite is slower and the other is faster, and it is not obvious that they are from similar activity. Here, bite means moving the apple to the mouth, biting the apple, and then moving the hand down. Figure 4(b) shows the sequences of the unit vectors on the sphere for those two bites. Here, the orientation traces are much more similar both in shape and length compared to the time series representation of 4(a). This representation gives the insight that the movement patterns are from similar activity which is hard to understand from Figure 4(a).

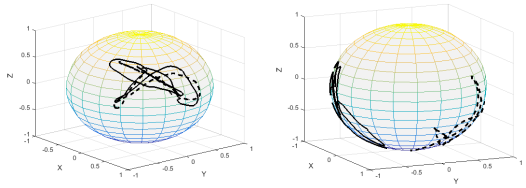

(a) (b)

Figure 3: Traces of (a) unit gravity vectors and (b) north vectors from a wrist device for opening two similar doors facing to different directions.

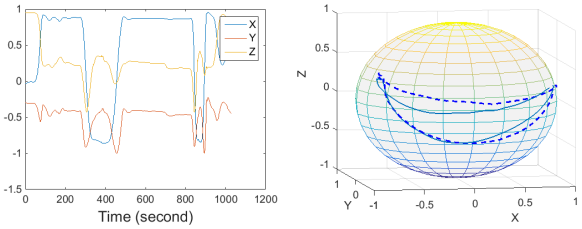

(a) (b)

Figure 4: For two apple bites, one slower and another faster, unit gravity vectors in the (a) time dimension, (b) spatio-temporal dimension

2.3 Primitive 3: Orientation Reachability

We define orientation reachability for an activity as the set of gravitational unit vectors that can be generated during the activity. It is represented by the area on the sphere that is covered by the orientation trace of the activity. The size and shape of the area depends on the activity. For example, Figure 2(c) shows the orientation reachability corresponding to the brushing teeth gestures mentioned before. Here, the trace is limited to a small area due to the fact that wrist orientation doesn't change significantly during brushing teeth in a specific way. On the other hand, Figure 4(b) shows the potential area on the sphere for apple bites which is quite wide. The traces for these bites are longer because they were taken while standing. The hand was moved from a downward position up to the mouth, and then again to the downward position. Nonetheless, the area of an activity is much less than the total area for all possible activities, and orientation reachability can be used to distinguish between different activities.

2.4 Primitive 4: Sphere Cell

The orientation traces can be at any places of the sphere. For visualization, we might need to rotate the sphere to bring the trace to the front. Though we get insight into the shape of the orientation trace, it is not easy to understand location of the orientation trace on the sphere because the sphere looks uniform. For example, we have rotated the sphere of Figure 2(c) to bring the traces to the front, and so the X and Y axes positions of this figure are different from those of Figure 3. It is difficult to understand exactly where the traces are on the sphere and how closer or further they are. To better understand the location of the traces on the sphere, we divide the sphere into some nearly uniform cells and assign an unique number for each of the cells.

The cells on the unit sphere are generated using a regular icosahedron, a polyhedron with 20 equilateral triangular faces and 12 vertices, as shown in Figure 5(a). The vertices of an icosahedron lie on the surface of a sphere. The cells

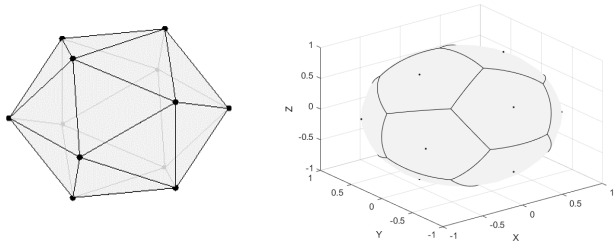

(a) (b)

Figure 5: (a) A regular icosahedron (b) The Voronoi cells generated from the regular icosahedron.

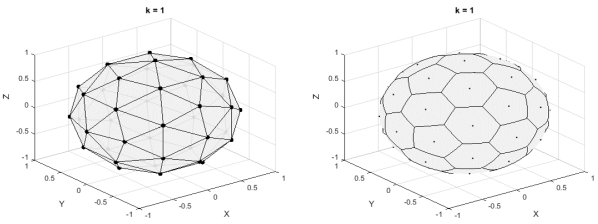

(a) (b)

Figure 6: (a) Polyhedron generated from dividing the regular icosahedron once, (b) Voronoi cells from the vertices of the polyhedron.

on the sphere are defined by the Voronoi diagram generated from the vertices of the icosahedron. Figure 5(b) shows the Voronoi diagram generated from the vertices of the regular icosahedron. The cells generated from the regular icosahedron are big. More fine grained cells are generated by dividing the icosahedron repeatedly. An icosahedron is divided by dividing each triangular face into four smaller triangular faces and connecting the mid-points of the edges of a face. The mid-points are then normalized so that they lie on the surface of the sphere of the icosahedron. For example, Figure 6(a) shows the polyhedron generated by dividing the regular icosahedron once, and Figure 6(b) shows the corresponding Voronoi cells. We can generate more fine grained cells by repeating the process. Icosahedron based division of sphere surface is widely used in areas like meteorology [7].

Rotating the sphere moves the cells from one place to another, and due to similarity it is very difficult to identify the individual cells. We assign unique numbers to each of the cells to identify them, and thus some specific areas on the sphere. The number of cells on the sphere depends on the size of the cells i.e., on how many times the regular icosahedron is divided to generate the cells. The number of cells generated from the regular icosahedron is 12, and there are $10 \times 4^k + 2$ cells when the cells are generated by dividing the regular

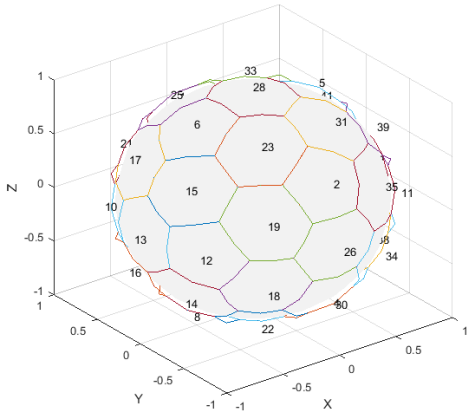

Figure 7: Unique cell numbers for k=1

icosahedron k times. For example, the number of cells for $k = 1$ (dividing the icosahedron once) and $k = 2$ (dividing twice) are 42 and 162, respectively. For identification, each of the cells is provided with a unique integer index in the range of [1, total number of cells]. When new cells are generated, the center of the old cells remain same, only size of the cells are reduced and new cells are generated. So, the cells whose centers are same to old cells retain the same indices of the old cells. This ensures consistency of the cell locations with different granularity. Figure 7 shows assignment of the numbers to the cells for $k = 1$. Such assignment of unique numbers to the cells helps to identify a cell or location easily even if the sphere is rotated. It also helps to distinguish the traces of different activities easily.

3 EXPERIMENTS

The methods presented in this paper are not activity or gesture specific. They can be used to visualize movement and orientation for different activities. The primitives are useful to understand different aspects of human activities including diversity in performing an activity, similarities between different activities, anomaly in data, quality of activity and so on. They are also useful in developing activity recognition solutions. In this section, we demonstrate the application of our method in understanding data and developing an efficient solution for smoking activity recognition. We also demonstrate how our method can be used to monitor rehabilitation exercises.

3.1 Smoking Activity Recognition

We have used our visualization methods to get insight into a dataset related to smoking activity recognition [13]. We also leverage orientation reachability to develop a very efficient solution for detecting smoking puffs. The dataset contains about 45 hours of data from 11 participants for smoking, drinking, eating, sitting, standing, and walking. The participants smoked while sitting, standing, walking, and group conversation. The dataset is weakly annotated where annotations are available for each type of activity instead of individual actions like a puff or a sip. It comes with data for total acceleration, linear acceleration and rotation rate from a smartwatch worn on the right wrist and a smartphone. In this experiment, we have used data from the wrist device only. We derive the gravity values by subtracting linear acceleration from total acceleration along each of the axes, and then scale the gravity values to unit vectors using L2 normalization.

3.1.1 Data Visualization. Our method provides better insight into data compared to traditional visualization methods. Figure 8 shows 10 minutes long data for the smoking (sitting), drinking (sitting) and eating using traditional 2-D plots (a, b, c), traditional 3-D plots (d, e, f) and our method (g, h, i). Using our method, we see that the area on the sphere for smoking is larger indicating that there is significant diversity in the smoking gestures, even from the same person. It is difficult to understand this phenomenon from the 2-D plots. Though the 3-D plots give some insights about this phenomenon, our method provides better insights by showing the orientation traces on a sphere with cell numbers.

One phenomenon related to activity or gesture is that different persons often perform the same activity differently, even same person does the same activity differently. We can easily visualize such phenomenon using our method. Figure 9 shows orientation traces for smoking (while standing) from three different participants. Though the orientation traces reach into or near cell 15 during the puffs, they are significantly different from person to person. It justifies why personalized models generally work better than person-independent models. Also, visualization helps us to understand why activity recognition is a very challenging task. Most of the works justify such facts using results (e.g., accuracy, f1-score) from classification algorithms where the results depend on the method used. Our method helps to better understand and communicate such insights through visualization of movement and orientation. We present few examples here for space limitations, but our methods can be used to get many different insights about the data and activities.

3.1.2 Smoking Puff Detection. We can leverage our methods for different activity recognition tasks. Here, we demonstrate how our methods can be used to detect smoking puffs efficiently. Like eating and drinking [8], the right hand is moved up to mouth and then down during a smoking puff resulting gravitational accelerations along the X axis go down

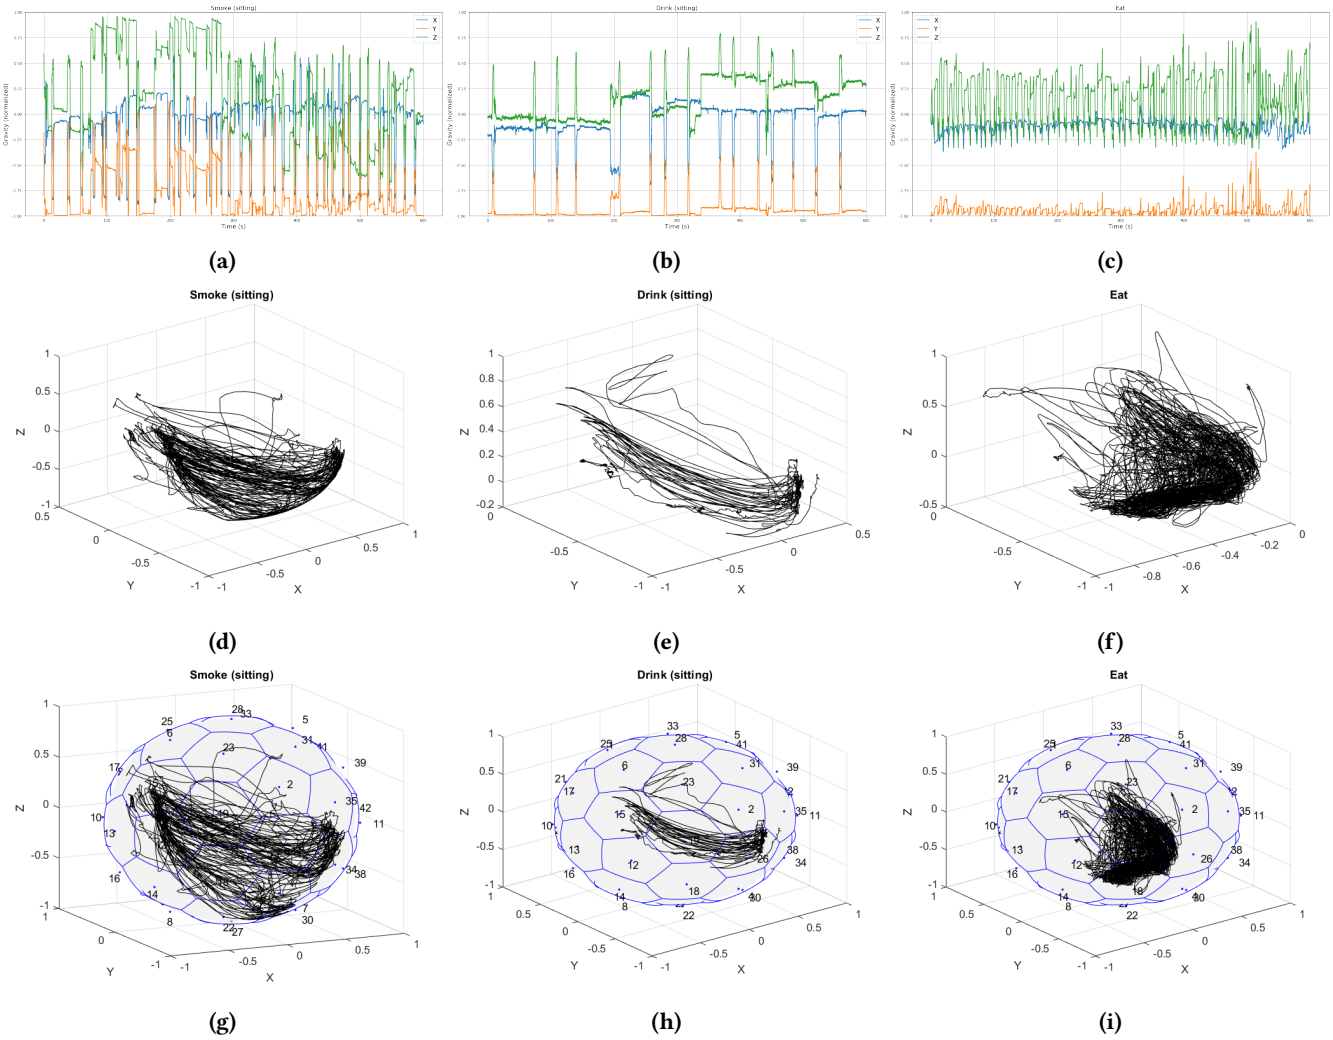

Figure 8: Plots of the gravity vectors for 10 minute long data using traditional 2-D plot (a, b, c), traditional 3-D plot (d, e, f) and our method (g, h, i) for Smoking (sit), Drinking (sit) and (c) Eating.

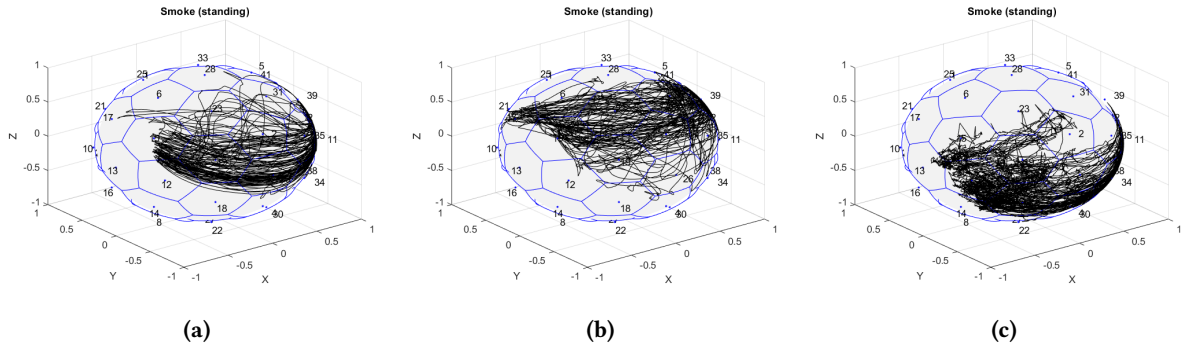

Figure 9: Smoking gestures from three different participants. These figures show that orientation of the wrist for smoking can differ significantly from person to person.

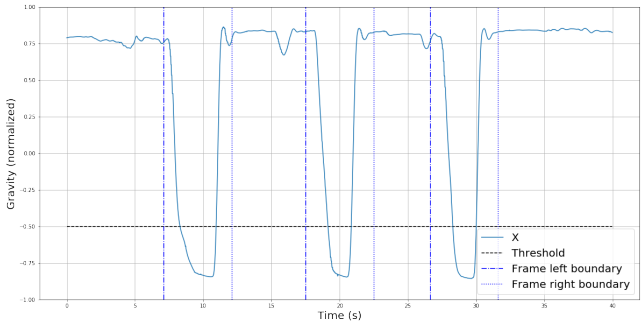

Figure 10: Detecting potential puffs by a threshold.

and then up. Figure 10 shows the gravity along the X axis for some puffs. We cut the X values using a threshold, and the segments that lie below the threshold contains the puffs. Such segments are also generated from other gestures like eating and drinking. The duration of a puff usually ranges between 1 to 2.4 seconds [14]. So, we only use the segments that are at least one second long and discard the others.

The midpoints of the segments approximately represent the orientation of the wrist when the hand is moved to the mouth. Figure 11 shows the gravity vectors (X, Y and Z) at the midpoints of the segments for smoking, drinking, and eating. We see that when the hand is moved to the mouth, the orientation of the wrist for smoking is significantly more similar to drinking than to eating. Also, we see that the orientation of the wrist is confined to a limited area for these activities. We use the orientation-reachability property of the activities to further filter-out non-smoking segments resulting reduction in overall computation. We leverage the distribution of the midpoints of the smoking segments to filter out the frames that are likely to be from non-smoking activities. The frames that remain after this filtering step are further classified using a more sophisticated classifier.

The pre-classification filter-out process is described in Figure 12. We calculate the mean ($\hat{\mu}$) and standard deviation ($\hat{\Sigma}$) of the midpoints of smoking segments. Then we define a distance threshold d_{th} in terms of the standard deviation so that the mid points for most of the smoking segments are within that threshold. If the distance is greater than the distance threshold, we consider it as non-smoking and discard it. Otherwise, we take a 5-second long frame around the midpoint and classify the frames using a classifier to determine whether it is a puff or not. As the usual duration of a puff is 1-2.4 seconds [14], the 5-second long frame captures a puff as well as corresponding hand movements. Our solution is not dependent on the classifier, and so we can use any effective classifier to classify the frames. We have used a convolutional neural network similar to that used for eating gesture detection in MFED [8]. It should be noted that only the X values of the the gravity vectors are used to detect the potential segments whereas data from all the axes (X, Y, and Z) of the gravity vectors are used for classification.

3.1.3 Result Analysis. The total duration of the data and the frames are about 45 hours and 9 hours, respectively. So, using the threshold on the gravitational acceleration along X axis, we can discard about 80% of the data in contrast to state-of-the-art methods that use computationally expensive classifiers to process all the data [2, 13]. Most of the segments come from smoking, drinking, and eating, with very few from other activities. The duration of smoking, drinking, and eating sessions in the dataset is about 72% of the total data. Usually, the ratio should be less in the real world, and so more portion of the data could be discarded.

For classification, we use Leave One Person Out (LOPO) approach. In this approach, data from one participant is tested by the model developed using the data from other participants. This process is repeated for all the participants, and results are calculated by combining the frames from all. We use the percentile of the distances of the smoking frames to determine the distance threshold. Figure 13 shows the rate of frames discarded at different percentiles. We see that more frames are discarded when the percentile is smaller and vice versa. For example, at the 90 percentile we discard about 22.8% of the non-smoking frames with the expense of only 5.2% of the smoking frames. It should be noted that 10 percent of the training frames are discarded at the 90 percentile threshold, but the discard rate of test frames is lower. It is because the distance threshold calculated from 10 persons (training) covers more puffs from the left-out person.

Discarding data using the distance threshold before using a computationally expensive classifier(e.g., a neural network) reduces overall computational overhead, but it also has implications for other performance metrics. Figure 14 shows the precision, recall, and F1-score at different percentiles. We see that as the percentile is decreased the recall decreases significantly, particularly for the lower percentiles. However, there is no significant gain in precision. Consequently, the F1-score reduces. However, the reduction in F1-score for higher percentiles is very low. For example, at 90 percentile threshold, the recall and the f1-score are reduced by 0.04 and 0.015 only. We can use our method to balance the trade-off between efficiency and accuracy metrics. Our solution can run on resource constraint devices without significant loss in accuracy.

3.2 Monitoring Rehabilitation Exercises

Rehabilitation programs, particularly after a surgery or an injury, generally involve a set of physical exercises that help the patients recover. The performance of the patients in performing the exercises is a good indicator of their health conditions. Monitoring the exercise patterns of the patients

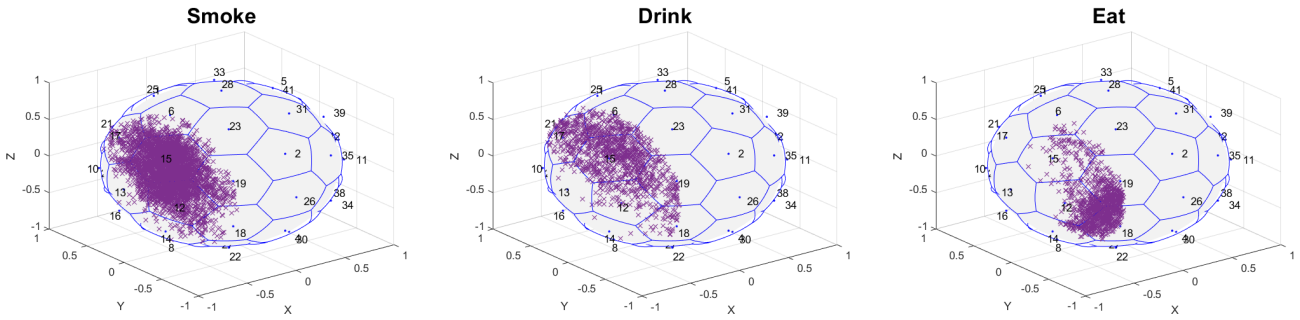

Figure 11: Mid points of the potential puff segments for (a) Smoking, (b) Drinking and (c) Eating.

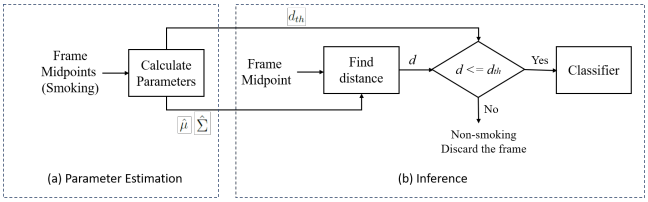

Figure 12: The process of filtering out frames before classification.

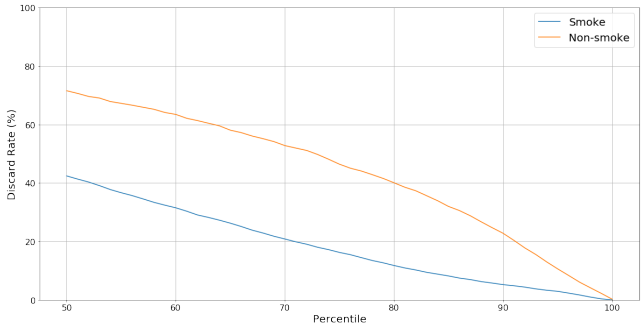

Figure 13: The rate of smoking and non-smoking frames discarded at different percentiles

would help the caregivers and physicians to better understand the patients' condition. In this section, we demonstrate how our methods can be used to monitor rehabilitation exercises. Due to space limitation, we demonstrate visualization for only 3 rehabilitation exercises (Figure 15).

For the purpose of this demonstration, the performance metric has been categorized into three categories: poor, average, and excellent. Poor performance for an exercise means very little hand movements during the exercise compared to the desired movements that is excellent. For this experiment, a participant performed all the exercises with poor, average and excellent performance following the guidelines in [1]. For each of the exercises, the participant first performs the exercise with little movement that corresponds to poor performance. Then the exercise is performed with more movement, and finally with desired movement that map to average and excellent performance, respectively. The participant is a healthy person, not a real patient. It should be noted that the purpose of this experiment is not to monitor real patients, but to show how our methods can be used for monitoring exercises.

Figures 15 shows the orientation traces of the exercises. The red, green and blue traces represent data for poor, average and excellent performances, respectively. The figures show that the orientation traces provides useful insights about the the quality of the exercises. The red traces have much less orientation reachability compared to that of the green traces. Smaller cell sizes might be used for fine-grained assessment of the quality of the exercises. Also the cell numbering is useful for better understanding and visualization of the orientation reachabilty. For example, it might seem that the gestures of Figure 15(a) and 15(c) have overlapping

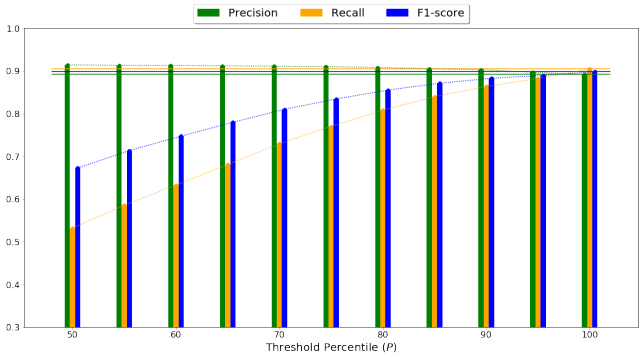

Figure 14: Precision, Recall and F1-score at different percentiles. The horizontal lines (dahsed) represent the metrics when no segment is discarded.

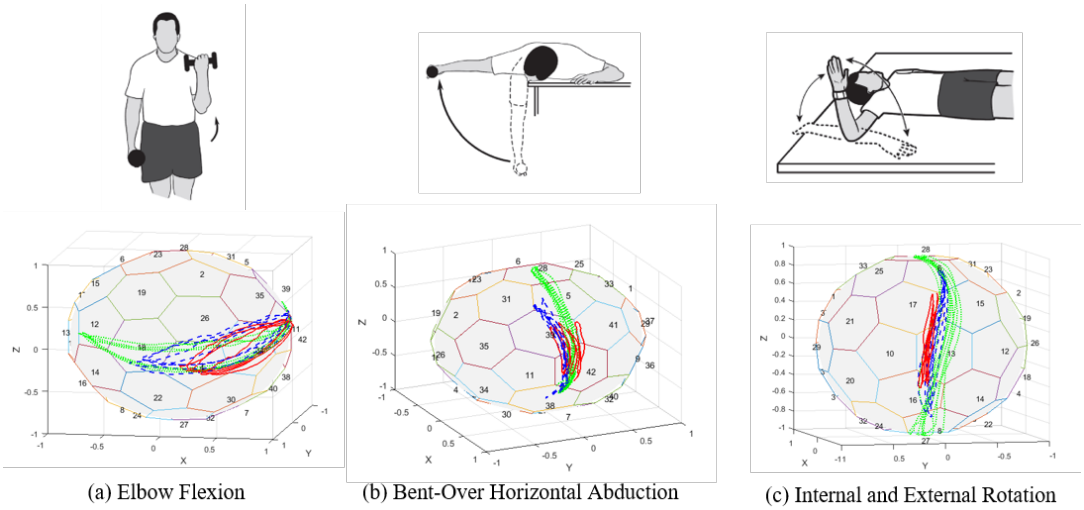

(a) Elbow Flexion (b) Bent-Over Horizontal Abduction (c) Internal and External Rotation

Figure 15: Plots for rehabilitation exercises (a) Elbow Flexion, (b) Bent-Over Horizontal Abduction, (c) Internal and External Rotation. The exercise images are adapted from [1].

reachability. But if we look into the cell numbers for these two exercises, they are different. This is due to the fact that the spheres have been rotated to bring the traces to the front. Thus, cell numbering helps in identifying the cells or location of the traces on the sphere easily.

4 RELATED WORKS

Our method places unit vectors from orientation on a unit sphere which is divided into some nearly uniform and numbered cells. The earth's shape resembles a sphere, and so spheres are widely used for visualization of geographical data. The sphere is often divided into uniform cells using a regular icosahedron and 2 of the 12 vertices of the icosahedron are aligned with the North and the South Poles [7] of the earth. A global weather prediction model, named GME, has been proposed based on the icosahedral–hexagonal grid [6]. These works focus on representing geographical data like weather for different regions of the earth.

State-of-the-art works on wearable and mobile sensing, including those for activity and gesture recognition, are generally focused on data transformation, feature engineering or classification. Existing techniques like line plot, scatter plot, pie chart, histogram, heat map, and bar chart are used for visualization of the sensor data. Some works transform data from one representation to another and visualize the transformed data using traditional plots mentioned above. Maaten et al. developed t-SNE [5] to visualize high-dimensional data in a two or three-dimensional space. For that purpose, the high-dimensional data are embedded into two or three dimensions using stochastic neighbor embedding. The embedded representation depends on different parameters used in the method including perplexity and number of iterations. The final embedding is usually visualized using scatter plots.

Though accelerations and rotation rates are the most widely used sensing modalities, orientation has also been used for activity recognition in several studies [3, 4, 12]. These works are focused on recognizing some specific activities, not visualization. RisQ [11] use inertial sensors on smart wristbands for recognizing smoking gestures. It uses quaternion to extract the orientation of the wrist. It computes the trajectory of the wrist relative to the elbow assuming that the elbow is stationary and in a fixed distance away from the wrist. This assumption is not often practical in free-living contexts. We focus more on visualization and the solution for smoking we present here does not have the above mentioned limitation.

HAWAD [10], a solution for hand washing detection, uses mean, standard deviation and threshold to filter out most of the data from activities other than handwashing. It uses out-of-distribution approach to address the problem of unknown activities. In contrast to HAWAD which focuses on accuracy and robustness, our solution for puff detection focuses on efficiency. Unlike HAWAD, we use the gravity vectors of the midpoints of the segments instead of the output of the penultimate layer of the neural network, and we filter out the instances before sending it to a classifier.

5 DISCUSSION

Visualization helps to understand data more effectively and easily. The methods presented here can help in developing better solutions for activity recognition. We present an efficient solution for smoking puff detection using orientation reachability. Our solution filters out most of the

non-smoking data using a threshold, and it can be used in combination with any classifier. Filtering out most of the activities other than the activities of interests not only reduces computation overhead but also it can improve accuracy by reducing class imbalance. This approach can be used to develop efficient and effective solutions for other activity recognition tasks. Our methods can be used to understand the orientation and movement of the wrist to improve user experience for smartwatch based interactive systems like MedRem [9].

The quality and quantity of data are crucial for most of the state-of-the-art solutions that are data-driven. However, developing datasets using wearable sensors, particularly with ground truth annotations, is very difficult and requires significant effort and time. The insights gained from movement visualization can be used to design data collection protocols. For example, activities with similar orientation trace and reachability are more likely to be misclassified into each other by a classifier. In order to build a more robust classification model for an activity, we need to provide the classifier with more data from those activities that have similar trace and reachability to the activity of interest than those activities that are different. We can detect such activities using our visualization methods. Though this paper mainly focuses on human activity and wearable sensors, the generic nature of the methods would make it possible to use them for different use cases beyond human activities. For example, the methods can be used to visualize movement and orientation of animals or objects if the inertial sensors are attached to them.

In this paper, we have demonstrated movement patterns of different rehabilitation exercises. We have not evaluated the effectiveness of monitoring the exercises involving real patients and physicians/therapists. It is possible to develop solutions that can measure the quality of exercise. However, our goal was to provide insights through visualization. It complements other solutions that provide some quantitative measures. A potential future work is to evaluate the effectiveness of our visualization method involving real patients.

6 CONCLUSION

This paper presents novel methods for movement and orientation visualization using inertial sensors. We have demonstrated the effectiveness and usefulness of the method to understand data, to develop a solution for smoking puff detection, and to monitor rehabilitation exercises. Methods for visualization using wearable sensors are rarely invented. Our methods for visualizing movement and orientation will be very useful for inertial sensor based research and development endeavours.

ACKNOWLEDGMENTS

This work was supported, in part, by NSF Grants CNS-1646470 and IIS-1521722.

REFERENCES

[1] 2020. Rehabilitation Exercises. https://orthoinfo.aaos.org/globalassets/pdfs/2017-rehab_shoulder.pdf. (2020). Last accessed: 2020-08-29.

[2] Fayez Alharbi and Katayoun Farrahi. 2018. A Convolutional Neural Network for Smoking Activity Recognition. In *2018 IEEE 20th International Conference on e-Health Networking, Applications and Services (Healthcom)*. IEEE, 1–6.

[3] Jonathan Alon, Vassilis Athitsos, Quan Yuan, and Stan Sclaroff. 2005. Simultaneous localization and recognition of dynamic hand gestures. In *Application of Computer Vision, 2005. WACV/MOTIONS'05 Volume 1. Seventh IEEE Workshops on*, Vol. 2. IEEE, 254–260.

[4] Luyang Liu, Cagdas Karatas, Hongyu Li, Sheng Tan, Marco Gruteser, Jie Yang, Yingying Chen, and Richard P Martin. 2015. Toward detection of unsafe driving with wearables. In *Proceedings of the 2015 workshop on Wearable Systems and Applications*. ACM, 27–32.

[5] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.

[6] Detlev Majewski, Dörte Liermann, Peter Prohl, Bodo Ritter, Michael Buchhold, Thomas Hanisch, Gerhard Paul, Werner Wergen, and John Baumgardner. 1998. The new global icosahedral–hexagonal grid point model GME of the Deutscher Wetterdienst. In *Proc. ECMWF Seminar on Recent Developments in Numerical Methods for Atmospheric Modelling, ECMWF, Reading, United Kingdom*. Citeseer.

[7] Detlev Majewski, Dörte Liermann, Peter Prohl, Bodo Ritter, Michael Buchhold, Thomas Hanisch, Gerhard Paul, Werner Wergen, and John Baumgardner. 2002. The operational global icosahedral–hexagonal gridpoint model GME: description and high-resolution tests. *Monthly Weather Review* 130, 2 (2002), 319–338.

[8] Md Abu Sayeed Mondol, Brooke Bell, Meiyi Ma, Ridwan Alam, Ifat Emi, Sarah Masud Preum, Kayla de la Haye, Donna Spruijt-Metz, John C Lach, and John A Stankovic. 2020. MFED: A System for Monitoring Family Eating Dynamics. *arXiv preprint arXiv:2007.05831* (2020).

[9] Md Abu Sayeed Mondol, Ifat Afrin Emi, and John A Stankovic. 2016. MedRem: an interactive medication reminder and tracking system on wrist devices. In *Wireless Health*. IEEE, 1–8.

[10] Md Abu Sayeed Mondol and John A Stankovic. 2020. HAWAD: Hand Washing Detection using Wrist Wearable Inertial Sensors. In *The 16th International Conference on Distributed Computing in Sensor Systems*.

[11] Abhinav Parate, Meng-Chieh Chiu, Chaniel Chadowitz, Deepak Ganesan, and Evangelos Kalogerakis. 2014. Risq: Recognizing smoking gestures with inertial sensors on a wristband. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*. ACM, 149–161.

[12] Didier Pittet, Benedetta Allegranzi, and John Boyce. 2009. The World Health Organization guidelines on hand hygiene in health care and their consensus recommendations. *Infection Control & Hospital Epidemiology* 30, 07 (2009), 611–622.

[13] Muhammad Shoaib, Hans Scholten, Paul JM Havinga, and Ozlem Durmaz Incel. 2016. A hierarchical lazy smoking detection algorithm using smartwatch sensors. In *2016 IEEE 18th International Conference on e-Health Networking, Applications and Services (Healthcom)*. IEEE.

[14] James P Zacny and Maxine L Stitzer. 1996. Human smoking patterns. *Smoking and tobacco control monograph* 7 (1996), 151–60.