

Monitoring and Coaching the Use of Home Medical Devices

Yang Cai, Yi Yang, Alexander Hauptmann, and Howard Wactlar

1 Introduction

Home medical devices (e.g. infusion pumps, inhaler, nebulizers, etc.) which are used by patients at home on their own, are becoming more and more prevalent, due to their cost-saving advantages. However, non-professional users, especially for elderly people with cognitive decline, may sometimes wrongly operate a medical device (e.g. not following required operating procedures). More seriously, the device-use error can lead to fatal results. For example, according to [13], during 2005–2010, 710 deaths linked to the use of one kind of home medical devices, the infusion pumps, which intravenously deliver life-critical drugs, food and other solutions to patients. Therefore, it's critical to have some external mechanisms to supervise patient's use of these devices and keep the use-error from happening. One straightforward solution to this problem would be let a professional person play the supervisor's role. However, because it contradicts to the main objective of home medical devices, i.e. reducing the cost, this solution is obviously infeasible. Instead of using "expensive" human's supervision, in this paper, we propose a cognitive assistive system whose objective is to automatically monitor the use of home medical devices.

The cognitive assistive system has two-fold functionalities: perception and recognition. On one hand, the system should be able to perceive user's operations. Since various advanced sensors been developed in the past decades, we are able to perceive user's operations from many different aspects. For example, the Kinect¹

¹<http://en.wikipedia.org/wiki/Kinect>.

Y. Cai • Y. Yang • A. Hauptmann (✉) • H. Wactlar
Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA
e-mail: yangcai1988@gmail.com; yee.i.yang@gmail.com; alex@cs.cmu.edu; hdw@cs.cmu.edu

that was invented recently can provide us not only the RGB information but also the depth that cannot be captured by traditional cameras. On the other hand, another important requirement for building a successful cognitive assistive system is to be able to recognize the operations so that use-errors can be identified [9]. However, unlike the well-developed perception module, it is still unknown whether the current techniques are mature enough to fulfill this requirement. Even though many techniques have been proposed to recognize human actions [1, 6, 15, 16, 20], none of them has been applied to recognize operations involved in using home medical devices and therefore, we are not sure if they are adequate in such scenario.

Since the lack of corresponding database is the main reason causing the situation, we construct a database (called PUMP) which was specially designed for studying the use of home medical devices and present it in this paper. Particularly, we take the example of patients using an infusion pump as a typical type of home medical device operation for collecting this database. An infusion pump is a device that infuses fluids, medication or nutrients into a patient's blood stream, generally intravenously. Because they connect directly into a person's circulatory system, infusion pumps are a source of major patient safety concerns. Because of this significant impact, we used an infusion pump as the sample device. To collect the data we first define an operation protocol for correct use of an infusion pump [7]. Then, each user was asked to simulate the use of infusion pump for several times. Their operations were recorded by three Kinect cameras from different views as shown in Fig. 1. Seventeen volunteers participated in data collection and 68 multi-view operation sequences were generated respectively. The operation sequences were then manually annotated by locating the temporal intervals of all operations in each sequence. The database will be released to public for research purpose.




Fig. 1 Examples of data recorded in PUMP database. The *first row* is the RGB data while the *second row* is the corresponding depth data. All user operations were captured by three Kinect cameras from different views




Fig. 2 Comparison between four actions and corresponding extracted MoSIFT features [4]. Only features in *green box* are relevant to actions by definition. (a) is “running” from [11], (b) is “jumping” from [3], (c) is “hula hoop” from [18] and (d) is “turning a device on” recorded by myself (Color figure online)

After building the database, we then evaluate the recognition performance of the existing approaches on the database. Even though using state-of-the-art approach which demonstrates near perfect performance in recognizing general human actions, we observe significant performance drop when applying it to recognize device operations. A subtle and overlooked unique characteristic of actions involved in using devices restrains the performance of the existing action recognition algorithms.

The uniqueness is illustrated in Fig. 2. It shows four actions and the corresponding extracted MoSIFT features [4]. The regions that are relevant to the target actions are indicated by green boxes. Figure 2a–c show the actions of “running”, “jumping” and “hula hoop” selected from three popular action recognition datasets [3, 11, 18]. For all three cases, most feature points in the whole frame are inside the green box. Because the “noise” points outside the box are relatively few, it is safe to use all features in the whole frame to model an action. However, as shown in Fig. 2d, for the action “turning a device on”, a typical action in using a home medical device, only a very small part of the features lies in the green box compared to all the features extracted from the whole frame. In this case, it is no longer reasonable to use all the features to represent actions, since the representation will be contaminated by substantial amount of essentially random noise. Such differences in feature distributions can be attributed to the fact that the relevant motion of the action in Fig. 2d is non-dominant and with a relatively small area compared to co-occurring non-relevant motion in the frame. We call this type of actions as *tiny actions*. Most of device operations are tiny actions, because we usually operate a device only with a body part, such as hand or foot, instead of the whole body.

To recognize tiny actions, it’s critical to focus on the local area where the target action happens, namely the region of interest (ROI). Therefore, in the second part of this paper, we introduce a simple but effective approach to estimating ROI for recognizing tiny actions. Specifically, the method learns the ROI for an

action by analyzing the correlation between the action and the sub-regions of the frame. The estimated ROI is then used as a filter for building more accurate action representations. The experiments show performance improvements over the traditional methods in terms of recognition precision. Note that, the proposed ROI estimation method can be used as a preprocessing step before applying any number of existing methods in the literature of action recognition.

The paper is organized as follows. We first give a review of the related works in Sect. 2 and then introduce the PUMP database in Sect. 3. In Sect. 4, we describe the ROI estimation method for recognizing tiny actions, followed by the experiments on the PUMP database in Sect. 5. We conclude our work and discuss the future work at Sect. 7.

2 Related Works

2.1 Existing Action Databases

We give a review of current existing action databases and compare them to our proposed one using three taxonomies: (1) RGB videos or RGB+Depth videos, (2) single camera or multiple cameras and (3) significant action or tiny action.

2.1.1 RGB Videos Versus RGB+Depth Videos

The videos of most current databases are RGB videos captured by traditional cameras, such as UCF50 [18], Hollywood2 [12], HMDB [10], KTH [21], Weizmann [3], UT-Interaction [19] and IXMAS [24]. Thanks to the greater availability of RGB+Depth cameras (e.g. Kinect), recently a few 3D databases which contain RGB+Depth videos have been proposed. For example, the MSR3D [23]. Compared to traditional RGB videos, the RGB+Depth videos preserve the additional depth information which could be useful for action analysis. PUMP is a RGB+Depth database.

2.1.2 Single Camera Versus Multiple Cameras

The single camera database refers to those recording actions only use one camera each time, while each action in multi-camera databases was simultaneously recorded by multiple cameras with overlapped views. Again, most of current databases are single camera ones (e.g. UCF50, Hollywood2, HMDB, KTH, Weizmann, UT-Interaction and MSR3D). There are only a few multi-camera action databases have been published, such as IXMAS where each action was captured by five cameras from different views. Since we use three cameras in PUMP, it is therefore a multi-camera database.

2.1.3 Significant Actions Versus Tiny Actions

We classify actions into two types in terms of their motion strength and relative area compared to whole motion region. The motion of significant actions is strong and dominant one compared to other co-occurred motion in a frame. In contrast, tiny actions' motion is weak, non-dominant and with relative small area. By this definition, the actions in KTH, Weizmann, UT-Interaction, MSR3D and IXMAS are mainly significant actions. For UCF50, Hollywood2 and HMDB, they contain videos with both significant actions and tiny actions. As shown in Fig. 4, in PUMP database, all actions are about hands operations, whose movement are weak and only taking a small area compared to co-occurred motion on body, head and etc. Therefore, they are all tiny actions.

2.2 Recognizing Human-Object Interaction

Accurately speaking, the action recognized in cognitive assistive system can be further categorized as human-object interaction which is an important group of actions recognition problems [1]. As [1] indicated, existing methods for recognizing interactions between human and object can generally be classified into two classes by judging if recognition of objects and actions are independently or collaboratively. For methods falling into first categories, objects recognition serves the following action recognition. For example, objects are usually recognized first and then actions are recognized by analyzing the object's motion. As for methods in second class, object and action are recognized in a collaborate fashion and the recognition of objects and actions serve each other. This work falls into the first category. However, different from previous works, we novelly use the object (device) as a cue for estimating the ROI.

2.3 Detecting Salient Region

Similar to ROI detection, salient region detection also finds a sub-region in an image or video that is considered to be salient. However, despite the similarity, their difference also worths noting. The saliency of a region is defined by its visual uniqueness, unpredictability, rarity and is caused by variations in image attributes, such as color, gradient, edges, and boundaries [5]. In other words, the saliency detection is not task-dependent by relies on the above rules. However, the ROI detection in visual-based coaching system is task-dependent. For example, the region including device operations may not be salient, due to the weak motion, but is of interest. Due to this difference, existing approaches for salient region detection cannot be directly applied to solve our problem.

3 PUMP Database

3.1 Data Collection Methodologies

We selected the Abbot Laboratories Infusion Pump (AIM Plus Ambulatory Infusion Manager) as an example of home medical devices for research. With the help of a medical devices expert, we first defined an operation protocol for correct use of infusion pump, as shown in Table 1. Then, as illustrated in Fig. 3, we set up a “workplace” for data recording. Specifically, we used three Kinect cameras to record user’s operation from three views: front, side, above. Before each time of recording, an infusion pump with off-state, two refilled syringes and a box of alcohol pads were prepared on the table.

Each user was asked to perform the operation following certain procedures for several times. In each time, they followed either the exact same procedures as described in Table 1 (correct operation protocol) or the predefined wrong procedures (to simulate the use-errors). Table 2 listed four types of predefined wrong procedures. They were different from the correct operation protocol by including steps disordering and steps missing. During the data recording, there were videos where users unintentionally deviated from the operation protocol they were asked to follow. Since these videos in fact reflected the use error that user made in real-life, we kept them in our database and provided additional error descriptions if the errors they made not belonged to any of the four predefined errors (Due to the limited space, we didn’t include the descriptions in our paper but kept them as an independent file in the database).

Table 1 The proposed operation protocol for correct use of infusion pump

1	Turn the pump on
2–5	Press buttons to set up infusion program
6	Uncap pump tube end
7	Clean pump tube end using alcohol pad
8	Open arm port
9	Clean arm port using alcohol pad
10	Flush arm port using syringe
11	Connect arm port and pump tube end
12	Press “START” button to start infusion
13	Press “STOP” button after infusion
14	Disconnect arm port and pump tube end
15	Clean pump tube end using alcohol pad
16	Cap pump tube end
17	Clean arm port using alcohol pad
18	Flush arm port using syringe
19	Clean arm port using alcohol pad
20	Cap arm port
21	Turn the pump off




Fig. 3 An illustration of data collection setting for PUMP database [8]

Table 2 The four types of predefined wrong operation protocols

W1	Switch step 6 and 7: Disorder flushing using syringe with cleaning arm port
W2	Remove step 4 and 6: Forget cleaning arm port and pump tube end
W3	Remove step 13 and 17: Forget cap arm port and pump tube end
W4	Remove step 9 and 10: Forget press buttons to start and stop infusion

Note that it only listed the differences between the correct protocol and wrong one and non-mentioned parts were same as correct protocol

Since some different steps in operation protocol were in fact the same actions, we therefore categorized them into one action class. We further combined the classes with same actions but operating different devices into one class, which finally led to seven action classes. We listed the aggregated classes in Table 3 and gave the snapshots of corresponding actions in Fig. 4.

Table 3 The action categories generalized from operation protocol for PUMP database

1	Turn the pump on/off
2	Press buttons
3	Uncap tube end/arm port
4	Cap tube end/arm port
5	Clean tube end/arm port
6	Flush using syringe
7	Connect/disconnect



Fig. 4 An illustration of seven action classes generalized from operation protocol of PUMP database. (a) Turn the pump on/off. (b) Press buttons. (c) Uncap tube end/arm port. (d) Cap tube end/arm port. (e) Clean tube end/arm port. (f) Flush using syringe. (g) Connect/disconnect tube end and arm port

Table 4 The inventory of PUMP database

Video type	OpenNI	RGB	Depth
View count	3		
Participant count	17		
Operation sequences	68		
File count	204		
File format	ONI	AVI	AVI
Frame rate	≈ 20	20	20
Resolution	640×480		
Average duration (m)	4.31		
Total duration (h)	14.64		
Disk storage (GB)	157	21	27

Note that due to the Kinect hardware issue, the frame rate of OpenNI videos was not constant but with little variation

Based on the generalized action categories, we manually annotated all sequences by locating the temporal intervals of all actions in each sequence.

3.2 Database Statistics and Recording Details

There were 17 volunteers participating in the data recording. Each user was asked to operate the infusion pump for four times with different appearances. Specifically, two times were correct operations while the others two came from the wrong operations. We finally constructed a database containing 68 operation sequences where each sequence had three synchronized RGB+Depth videos recorded from different views.

We adopted OpenNI² for Kinect recording and stored the raw data in ONI file format. To facilitate the use of this database, we also provided calibrated RGB videos and depth videos extracted from raw OpenNI data in our database. In Table 4, we show detail statistics of the PUMP database. The average video duration was 4.31 in and the total video duration was 14.64 h.

4 ROI Estimation for Recognizing Tiny Actions

4.1 Notations

Let (x, y, z) be the coordinates of the corner index of a cuboid region, A_j be an action of class j , M be the number of action classes. Let $T^{A_j}(x, y, z)$ be a density map that describes the probability for a region at (x, y, z) belonging to the ROI of A_j . Specifically, we call $T^{A_j}(x, y, z)$ as the *ROI template*.

²<http://openni.org/>.

4.2 Action-Region Correlation Estimation

Noticing the ROI for an action can be interpreted as regions that have strong correlation with the action, we estimate the correlation between an action A_j and each region at (x, y, z) in the “3D” frame recorded by Kinect cameras.

To represent each region, we generate sliding windows starting from the origin of “3D” frame and calculate the bag-of-words (BoW) [22] representation for each cuboid window. To reduce the computation cost in following steps, we then apply the principal component analysis (PCA) on the BoW of each window and only keep the dimensions corresponding to the top K largest eigenvalues. The fisher score [2] is used to estimate the correlations between each region and an action.

Let $D_b^{(x,y,z).A_j}$ and $D_w^{(x,y,z).A_j}$ be action class A_j 's *between class distance* and *within class distance* [2] respectively for all regions at (x, y, z) among all training data. Then, $D_b^{(x,y,z).A_j}$ and $D_w^{(x,y,z).A_j}$ are defined as:

$$D_b^{(x,y,z).A_j} = \sum_{k=1}^M \left(\mu_{\delta(k=j)}^{(x,y,z)} - \mu^{(x,y,z)} \right)^T \left(\mu_{\delta(k=j)}^{(x,y,z)} - \mu^{(x,y,z)} \right)$$

$$D_w^{(x,y,z).A_j} = \sum_{k=1}^M \sum_{b \in B_k^{(x,y,z)}} \left(b - \mu_{\delta(k=j)}^{(x,y,z)} \right)^T \left(b - \mu_{\delta(k=j)}^{(x,y,z)} \right)$$

where $\delta(z)$ is an indicator function that outputs 1 if z is true and 0 otherwise, $B_k^{(x,y,z)}$ is a set of BoWs of A_k at region (x, y, z) , $\mu^{(x,y,z)}$ is the mean of BoWs at region (x, y, z) for all action classes and $\mu_{\delta(k=j)}^{(x,y,z)}$ is the mean of BoWs at region (x, y, z) for action class A_j or action classes other than A_j (depending on if $k = j$). Then the fisher score $F^{(x,y,z).A_j}$ for an action class A_j and a region at (x, y, z) is simply:

$$F^{(x,y,z).A_j} = \frac{D_b^{(x,y,z).A_j}}{D_w^{(x,y,z).A_j}}.$$

If one region (x, y, z) is highly correlated to action A_j , it will then have relatively small *within class distance* $D_w^{(x,y,z).A_j}$ and large *between class distance* $D_b^{(x,y,z).A_j}$, which gives large fisher score $F^{(x,y,z).A_j}$. Therefore, fisher score can be an indicator of correlation between an action and regions.

By normalizing the fisher score at different regions, we get the representation of the ROI template $T^{A_j}(x, y, z)$:

$$T^{A_j}(x, y, z) = \frac{F^{(x,y,z).A_j}}{\sum_{x,y,z} F^{(x,y,z).A_j}}.$$

4.3 ROI Adaption and Noise Filtering

For a given input video sequence, we simply attach the ROI template $T^{A_j}(x, y, z)$ to each frame of the video sequence. Then, all feature points with ROI score lower than threshold λ are removed. In this way, we model the action only based on features in ROI.

5 Experiments

5.1 Experimental Setting

The total 68 videos are divided into two-fold and we in turn use one fold as training and the other one as testing data. For each video, we extract the MoSIFT [4] as low-level features and encode them into visual words [22] using a codebook with vocabulary size of 1,000. Then three different methods are used for generating the action representations (see Sect. 5.2 for details). SVM classifier with RBF kernel is adopted for action classification and two-fold cross validation is used for classification model training. Specifically, the training and testing is done independently for each view. To evaluate the effectiveness of different methods, the mean average precision (MAP) [17] which is the average precision (AP) over all actions is computed.

5.2 Action Representation Methods

The bag-of-words model (BoW) is adopted for action representation. Each high dimensional local feature point (e.g. MoSIFT) is first mapped to the closest cluster center using the pre-trained codebook and then the cluster's id is assigned to the feature as "visual word". After that, a pooling step is applied to calculate the statistics of all the visual words in the video segment and represent it as vector with same dimension of the codebook. This vector representation is called the BoW of the video segment. In this paper, we experiment with three different pooling methods for action representation as introduced below.

5.2.1 Whole Frame Based Pooling (WF-BoW)

Most of existing BoW-based action recognition approaches [4, 21] use this pooling method. Namely, all visual words in the whole frame are aggregated together first and the frequency for each visual word is calculated then. Finally, the normalized frequency histogram is used as the final representation.

5.2.2 Depth-Layered Multi-Channel Based Pooling (DLMC-BoW)

Since the videos are recorded by Kinect camera, then each feature point extracted from the key frame has not only the x and y coordinates but also the depth z coordinate. Based on this observation, in [14], z axis is first divided into several depth-layered channels, and then features within different channels are pooled independently, resulting in a multiple depth channel histogram representation. In our implementation, we uniformly divide the depth space into five channels.

5.2.3 ROI Based Pooling (ROI-BoW)

In order to estimate the ROI, $200 \times 200 \times 100$ pixels (in the order of x , y and z axis) sliding cuboid windows with moving step of 40 pixels are generated and represented as BoW by aggregating all visual words inside the cuboid window. We then apply the PCA on the BoW of each window and only keep the dimensions corresponding to the top 100 largest eigenvalues. After that, the ROI template is calculated using the method described in Sect. 4.2. Note that, all these process can be done off-line. At online testing stage, all visual words with ROI score lower than 0.5 are filtered out and the final BoW representation is built only based on the left visual words belonging to the ROI.

5.3 Experimental Results and Analysis

5.3.1 ROI Visualizations

To qualitatively evaluate the proposed ROI estimation method, in Fig. 5, we visualize estimated ROIs for each action of all three views using density map where higher intensity means high probability of belonging to be the ROI. Because the cameras are put at different positions, the ROI for the same action but different

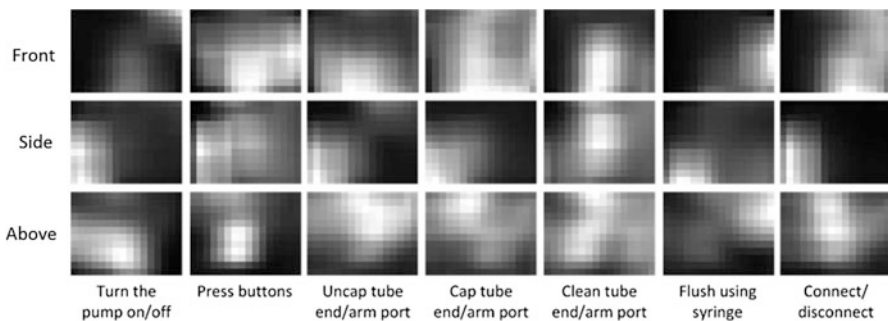


Fig. 5 The visualizations of estimated ROI for each action. Each row corresponds to one of the three views. The action examples of different views can be found in Fig. 3

views can be different. Comparing the action examples shown in Fig. 3 and the corresponding ROIs in Fig. 5, we can find the estimated ROIs make sense intuitively. For example, for the “front” view, the ROI is in the middle for action “press buttons” while is on the right for action “Flush using syringe”. This is because the “pressing buttons” always happen in the middle of the frame while for “flush using syringe” users always need to take out the syringe from syringe bag which is placed on the right side of the frame. Also, we can observe some of actions’ ROI are dense and sharp (e.g. “Connect/disconnect” in side view) while others are relative sparse and soft (e.g. “Cap tube end/arm port” in front view). This difference can be attributed to the different motion patterns of different actions. For example, “Connect/disconnect” concentrates in a narrow area but the motion of other actions like “Cap tube end/arm port” distributes in relatively larger areas.

5.3.2 Action Recognition Performance

In Table 5, we summarize the experimental results of three action representation methods. For each method, the first three columns correspond to the performance on three different views while the last column is the fusion results given by manually selecting the best performance among the three views. The fusion results can be interpreted as the best performance that the cognitive assistive system achieves using that method. Comparing the average fusion results of three action representation methods, we can see that the *ROI-BoW* achieves the best performance and improves the MAP for 3.33 % compared with *WF-BoW*. *DLMC-BoW* has almost the same performance as *WF-BoW*.

If we further compare three methods’ performance on each single view, we observe a different performance changing pattern. For the “front” and “side” views, both *ROI-BoW* and *DLMC-BoW* show significant improvements over the *WF-BoW*. Specifically, on average, *DLMC-BoW* improves for 4.42 and 10.20 % while *ROI-BoW* improves for 8.84 and 19.28 % on those two views. However, for the “above” view, both *ROI-BoW* and *DLMC-BoW* don’t show improvement but in fact slightly decrease the performance compared to *WF-BoW*. The reason causing the inconsistent performance changing pattern on different views is illustrated in Fig. 6. It visualizes of extracted MoSIFT features in example frames with the same time stamp but different views for action “Flush using syringe”. Again, we use green boxes to indicate the regions that are relevant to the target action. We can see that for “front” and “side” views, only a very small part of the features lies in the green box compared to all the features extracted from the whole frame, while most features are inside the green box for the “above” view. Therefore, due to the difference in camera positions, actions recorded by “front” and “side” cameras are always the most typical tiny actions. Because both *DLMC-BoW* and *ROI-BoW* can be interpreted as feature location based visual word weighting methods (*DLMC-BoW* does it implicitly by using SVM to weight features at different depth differently while *ROI-BoW* does it explicitly by hard weighting features inside ROI 1 and outside 0), they are most effective when actions are typical tiny actions.

Table 5 MAP comparison of three action representation methods on PUMP database

Actions	WF-BoW			DLMC-BoW			ROI-BoW					
	Front	Side	Above	Fusion	Front	Side	Above	Fusion	Front	Side	Above	Fusion
Turn the pump on/off	0.8356	0.5931	0.7222	0.8356	0.863	0.7761	0.8232	0.863	0.894	0.8348	0.8696	0.894
Press buttons	0.8044	0.5841	0.7703	0.8044	0.8312	0.7812	0.782	0.8312	0.8833	0.8058	0.8378	0.8833
Uncap tube end/arm port	0.4933	0.2679	0.6335	0.6335	0.6252	0.4034	0.5201	0.6252	0.6541	0.5215	0.5326	0.6541
Cap tube end/arm port	0.3466	0.2898	0.4356	0.4356	0.3724	0.3303	0.4352	0.4352	0.4455	0.398	0.3923	0.4455
Clean tube end/arm port	0.8684	0.6967	0.8676	0.8684	0.8834	0.7331	0.8445	0.8834	0.9202	0.8438	0.8401	0.9202
Flush using syringe	0.831	0.6718	0.8329	0.8329	0.8845	0.7589	0.8038	0.8845	0.948	0.9048	0.7931	0.948
Connect/disconnect	0.451	0.3361	0.6349	0.6349	0.48	0.3704	0.5302	0.5302	0.5037	0.4801	0.5335	0.5335
Average	0.6615	0.4914	0.6996	0.7208	0.7057	0.5933	0.677	0.7218	0.7498	0.6841	0.6856	0.7541

For each method, the performance is evaluated on each camera independently. The best performance among the three cameras is manually selected and shown in the “fusion” column




Fig. 6 The visualization of extracted MoSIFT features in frames with the same time stamp but of different views

However, even though *ROI-BoW* improves the performance significantly, if looking into the absolute performance of each action, we realize only actions “Turn the pump on/off”, “Press buttons”, “Clean tube end/arm port” and “Flush using syringe” achieve reasonable high average precision while the performance of left three actions is still low. To build an applicable cognitive assistive system, we have to accurately recognize all the actions. Therefore, it still requires special effort to further boost the performance of the difficult actions.

6 An Interaction Model for Coaching Use of Home Medical Devices

Separate from the detection problem is the interaction model for dealing with the user. The detection module, as described in the above sections, handles the analysis of the camera feeds and reports whether or not a step has been detected as correctly performed, and with what level of confidence. This is reported to an interaction module. The interaction module will have prior knowledge what the typical average confidence of the detection for the particular step is, what the acceptable sequences of steps are, and what the importance is for each step based on the scale which we will describe later. The user may also interact with the system to go through the instructions by using a mouse.

6.1 Importance Levels

We assign four different levels of importance to each step in the procedure. The interaction module functions differently for each step, depending on the level of importance of that step. An harmless or obvious step would be one where is unnecessary to warn the user if not performed, such as turning the system on. This would be categorized as level 0 importance. A very important step would be placed at Level 3 importance, which would require the system to make it hard for a user

Table 6 Example of warnings and reminders at different levels of error importance and confidence

Confidence	Importance			
	Level 0	Level 1	Level 2	Level 3
High	No action, green light	Audio-visual correction information	Flashing audio-visual correction information	Red, flashing visual correction with loud audio and pop-up warning
Moderate	Yellow light	Visual reminder with soft audio	Flashing audio-visual reminder with soft audio	Red, flashing visual reminder with soft audio and pop-up warning
Low	Yellow light	Visual reminder only	Flashing visual reminder with beep	Red, flashing visual reminder with soft audio and pop-up warning

to ignore the error warnings issued by the system, e.g. by combining loud, flashing audio-visual notification. Table 6 shows the different actions by the system given their importance levels.

6.2 Event Detection Certainty

One of the ways to prevent user frustration with system miss-recognitions is to have the system communicate what the certainty of its detection is, with appropriately adapted feedback. The quickest way for a user to begin doubting and ultimately ignoring the system, is to report an event as detected or not detected when (to the user) it is clearly the opposite. To combat this, the system incorporates an algorithm to decide whether an action is detected as incorrect or not detected. The warning for a user error with a high enough confidence rating will be shown to the user. Furthermore, the system adapts the presentation of the message to convey the confidence level.

- Generally, moderate to high confidence detection of a correct step is only indicated with a green light, without otherwise disrupting the user.
- Less confident detections of the correct step might be indicated in yellow.
- Moderate to high confidence detection of an incorrect action triggers a strong warning
- Low confidence detection of an incorrect action would trigger a suggestion or reminder, without commitment by the system that an actual error was made.

6.3 *Warning Levels*

Instead of confusing the user with percentages of confidence in the event detection, warnings will be presented appropriately to denote both error severity and confidence levels. This will prevent the system from making blatant statements that might contradict reality such as noting an event as not detected when to the user it clearly happened. These confidence and severity appropriate warnings allow the system to function usefully if it occasional events are misclassified.

7 Conclusion and Future Work

In this paper, we proposed a cognitive assistive system to monitor the use of home medical device. To build such a system, accurately recognition of user actions is one of the most essential problems. However, since few research has been done in this specific direction, it is still unknown if current techniques are adequate to solve the problem. In order to facilitate the research in this area, we made three contributions in this paper. First of all, we constructed a database where users were asked to simulate the use of infusion pump following predefined procedures. The operations were recorded by three Kinect cameras from different views. All the data was manually labeled for experimental purpose. Secondly, we performed a formal evaluation of some existing approaches on the proposed database. Because we realized current methods can hardly deal with tiny actions involved in using home medical devices, we made our third contribution by introducing an ROI estimation method and applying the ROI for building more accurate action representations. The experiments show significant performance improvements over the traditional methods by using the proposed methods.

Finally, we outline an interaction model how to handle different levels of errors recognized with varying certainty and appropriately provide feedback to the user reflecting the severity/importance of the mistake as well as the confidence of the system in the correctness of the recognition.

Currently, we treat all the actions in operating home medical devices as independent ones and recognize them separatively. However, it obvious they are mutually related because they are operations in a procedure. Therefore, in the future, we will focus on leveraging the inner relations between actions to further improve the recognition performance.

Acknowledgements This work was supported in part by the National Science Foundation under Grant No. IIS-0917072. Any opinions, findings, and conclusions expressed in this material are those of the author(s) and do not reflect the views of the National Science Foundation.

References

1. Aggarwal, J. K., & Ryoo, M. S. (2011). Human activity analysis: A review. *ACM Computing Surveys (CSUR)*, 43(3), 1–43.
2. Belhumeur, P. N., Hespanha, J. P., & Kriegman, D. J. (1997). Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7), 711–720.
3. Blank, M., Gorelick, L., Shechtman, E., Irani, M., & Basri, R. (2005). Actions as space-time shapes. In *International Conference on Computer Vision*.
4. Chen, M. Y., & Hauptmann, A. (2009). MoSIFT: Recognizing human actions in surveillance videos. In *CMU-CS-09-161*.
5. Cheng, M. M., Zhang, G. X., Mitra, N. J., Huang, X., & Hu, S. M. (2011). Global contrast based salient region detection. In *IEEE Conference on Computer Vision and Pattern Recognition*.
6. Dollár, P., Rabaud, V., Cottrell, G., & Belongie, S. (2005). Behavior recognition via sparse spatio-temporal features. In *Joint IEEE International Workshop on Visual Surveillance and Performance Evaluation of Tracking and Surveillance*.
7. Gao, Z., Chen, M. Y., Detyniecki, M., Wu, W., Hauptmann, A., Wactlar, H., et al. (2010). Multi-camera monitoring of infusion pump use. In *IEEE International Conference on Semantic Computing*.
8. Gao, Z., Detyniecki, M., Chen, M. Y., Hauptmann, A. G., Wactlar, H. D., & Cai, A. (2010). The application of spatio-temporal feature and multi-sensor in home medical devices. *International Journal of Digital Content Technology and Its Applications*, 4(6), 69–78.
9. Gao, Z., Detyniecki, M., Chen, M. Y., Wu, W., Hauptmann, A. G., & Wactlar, H. D. (2010). Towards automated assistance for operating home medical devices. In *International Conference of Engineering in Medicine and Biology Society*.
10. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., & Serre, T. (2011). HMDB: A large video database for human motion recognition. In *International Conference on Computer Vision*.
11. Laptev, I. (2005). On space-time interest points. *International Journal of Computer Vision*, 64(2/3), 107–123.
12. Marszałek, M., Laptev, I., & Schmid, C. (2009). Actions in context. In *IEEE Conference on Computer Vision and Pattern Recognition*.
13. Meier, B. (2010). F.D.A. Steps up oversight of infusion pumps. In *New York Times*.
14. Ni, B., Wang, G., & Moulin, P. (2011). RGBD-HuDaAct: A color-depth video database for human daily activity recognition. In *International Conference on Computer Vision Workshops*.
15. Niebles, J. C., Wang, H., & Fei-Fei, L. (2008). Unsupervised learning of human action categories using spatial-temporal words. *International Journal of Computer Vision*, 79(3), 299–318.
16. Parameswaran, V., & Chellappa, R. (2006). View invariance for human action recognition. *International Journal of Computer Vision*, 66(1), 83–101.
17. Philbin, J., Chum, O., Isard, M., Sivic, J., & Zisserman, A. (2007). Object retrieval with large vocabularies and fast spatial matching. In *IEEE Conference on Computer Vision and Pattern Recognition*.
18. Reddy, K., & Shah, M. (2012). Recognizing 50 human action categories of web videos. *Machine Vision and Applications Journal*, 25(5), 97–81.
19. Ryoo, M. S., Aggarwal, J. K. (2010). UT-Interaction dataset. In *ICPR Contest on Semantic Description of Human Activities (SDHA)*.
20. Ryoo, M. S., & Aggarwal, J. K. (2011). Spatio-temporal relationship match: Video structure comparison for recognition of complex human activities. In *International Conference on Computer Vision*.
21. Schuldt, C., Laptev, I., & Caputo B (2004) Recognizing human actions: A local sVM approach. In *International Conference on Pattern Recognition*.
22. Sivic, J., & Zisserman, A. (2003). Video google: A text retrieval approach to object matching in videos. In *International Conference on Computer Vision*.

23. Wang, J., Liu, Z., Wu, Y., & Yuan, J. (2012). Mining actionlet ensemble for action recognition with depth cameras. In *IEEE Conference on Computer Vision and Pattern Recognition*.
24. Weinland, D., Özuysal, M., & Fua, P. (2010). Making action recognition robust to occlusions and viewpoint changes. In *European Conference on Computer Vision*.