Textual Summarization of Events Leading to Health Alerts

Akshay Jain*, *Student Member, IEEE* and James M Keller, *Life Fellow, IEEE*

Abstract— **Extracting information from the sensors installed in the homes of elderly pose a unique set of challenges. Add to it the short amount of time the clinicians and nurses have to analyze this data, and the problem becomes more complicated. A system already in place at an "Aging in Place" facility monitors the activities of residents through multiple non-intrusive sensors and sends alerts on detecting an unusual event. We present an approach to generate textual summaries of events leading to the alerts. We analyze our system using four case studies and also list the comments provided by collaborators in healthcare domain. The system was then iterated to take some of those suggestions into account to give a glimpse of what an ideal system should look like.**

I. Introduction

In the recent times, there has been substantial growth in the development of health monitoring sensors. For instance, monitoring vital signs of patients in hospitals, living routines of elderly whose apartments are equipped with multiple sensors or even that of people just trying to monitor their health with the wearable sensors that have surfaced of late in the market. This leads to the generation of huge amounts of data which is believed to be rich in information content. To make sense of this information rich data, a popular method is to visualize it graphically and present it to people in the concerned domains, which can be medical personnel in hospitals, nurses etc. But in this method the healthcare professionals are subjected to an added responsibility of understanding the data along with inferring clinical information from it. Along with this, the visualization method pose challenges like getting familiar with it, axis scaling issues and the time to extract information from plots.

A different take to deal with this is to design systems that can identify important patterns in the data and present these in natural language to the concerned audience. This paradigm has been explored in the past by the joint efforts of the Natural Language Generation community and people in health domain. We list a few of them here for reference. The system presented in [1] combines the visualization and the text summarization techniques to efficiently present the medical histories of patients to clinicians. The visual interface presents a snapshot of the patient's history while the textual summaries of specific events are shown when queried by the user. Their data-to-text system has been evaluated in [2] by comparing the automatically generated textual summaries with raw data (consisting of test results, documents, etc.). In order to assess the usefulness of the system, it was presented to health domain experts, along with the medical histories of the patients. It was concluded that the clinicians were more inclined towards the automatic generated text summaries as compared to the patient history documentation, due to the consistency in the automatic summaries. The system in [3] generates natural language summaries of data in neo natal intensive care units and presents them along with the raw data to the nurses between the shifts. Then, based on a questionnaire posed to the nurses, they conclude that the summaries are understandable, accurate and helpful. Moreover, they show that their system can find "interesting" patterns and summarize them automatically which might be missed by not so experienced professionals. Another example is [4] in which the authors summarize medical histories for patient's personal use.

The major goal of this work is natural language summarization of sensor data obtained from the homes at TigerPlace, which is an "Aging In Place" facility for elderly at Columbia, MO. The apartments have various sensors installed to monitor the elderly in order to help nurses keep track of their health conditions and assist them if there is a possibility of a bad event. Linguistic Protoform Summaries (LPS) are used to compute the summaries of data which are then modified with natural language rules to appear more intuitive to the end user. A numeric system already in place produces real time alerts (notifications via email) in case a possible unusual activity is detected in any individual sensor data stream (motion sensor, bed sensor etc.). Feedback is then provided by the clinicians on the quality of the alerts. In this work we present four such cases and make a case for how linguistic summaries can help better understand the data. In order get an idea of the preferences of the clinicians, we presented the alert summaries to our collaborators having backgrounds in healthcare. We also list the comments provided by them and modify the summaries accordingly.

II. BACKGROUND

In this section we provide a brief description of the system and the techniques used in this work that have already been developed previously.

A. System at TigerPlace

The apartments at TigerPlace are equipped with motion sensors, bed sensors and Microsoft Kinect cameras to monitor the activities of elderly. In this work we focus on four such apartments which all have multiple motion sensors spanning across the apartment and a bed sensor to monitor restlessness in bed. Each sensor is attached with a numerical alert framework which signals the nurses and researchers about a possible unusual activity in the incoming data stream. An activity is deemed unusual if its distance from the mean of the

*Supported by Center For Eldercare and Rehabilitation Technology at the University of Missouri.

Akshay Jain (e-mail: aj4g2@mail.missouri.edu) and James M Keller (e-mail: kellerj@missouri.edu) are with Department of Electrical and Computer Engineering at the University of Missouri, Columbia, MO 65211 USA.

normal distribution comprising of past two weeks of data is greater than a threshold. The thresholds are determined experimentally. A more detailed study on the selection of these thresholds can be found in [5]. Also, on receiving the alert email, the clinicians dig into the sensor data and residents medical history, and provide a rating of the alert on a scale of 1 to 5 along with a textual feedback.

B. *Linguistic Protoform Summaries (LPS)*

Linguistic Protoform Summaries are short template based sentences providing a snapshot of data in textual form. For example, a summary focusing on the sizes of balls in a bag might look like, *Most of the balls in the bag are big*. In the sense of Yager [6], they are comprised of a Quantifier (*Most, Few* etc.) conveying information about the quantity of the attribute being summarized, a Summarizer (*big, small* etc.) measuring the feature in concern and a truth value (0 to 1) which signifies the validity of the summary with respect to the data. To generate summaries of the data, we define a set of Quantifiers and Summarizers which are modelled by fuzzy sets used to computed the truth values. Then for each summarizer, the summary with highest truth value is chosen as its best representation. The truth value is computed using the method presented recently in [7]. We now provide a brief description of this method. The truth value of the summaries of the form *A y's are P* is computed as shown in (1)

$$T(A\ y\text{'s are } P) = \max_{\alpha \in [0,1]} \left(\alpha \wedge A(P_\alpha)\right) \qquad (1)$$

where, ^ is the minimum operator, $P(x)$ is the membership function of the summarizer P, $P_\alpha = \frac{\left|\{y_i \in Y \mid P(y_i) \geq \alpha\}\right|}{N}$ is the proportion of objects whose membership in $P(x)$ is greater than or equal to α (varies from 0 to 1 in small intervals), $|\,.\,|$ denotes the cardinality of a set and $A(x)$ is a normal, convex and monotonically non-decreasing membership function of the quantifier A. For quantifiers whose membership function is not monotonically non-decreasing, it is split into two monotonically non-decreasing functions, $A_1(x)$ and $A_2(x)$ (which is used to compute $\overline{A_2}(x)$) and the truth value is computed as shown in (2). Please refer [7] for more details.

$$T(A\,y\text{'s are } P) = T(A_1\,y\text{'s are } P) \wedge (1 - T(\overline{A_2}\,y\text{'s are } P)) \qquad (2)$$

III. METHODOLOGY

Figure 1 presents the block diagram of our system.

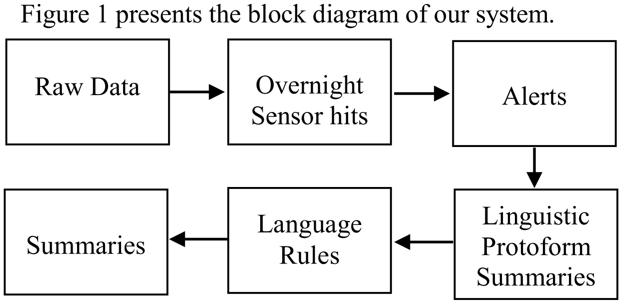

Figure 1: System Block diagram

The system already in place generates alerts for individual sensor parameters, such as motion density, bathroom motion, bed restlessness. However, on receiving an alert of one kind, it is useful to look at the other parameters to test the validity of the alerts as well as to have a snapshot of the person's routine. In fact, we have observed that when the clinicians look at a specific sensor alert, they go back and check the other parameters as well. Taking inspiration from this, we start by considering three parameters, namely, the motion density, the bed restlessness and the bathroom motion.

The motion sensors and bed sensor monitor the activity of the residents inside the apartments round the clock. However, in this study we only focus with activities during the night time, that is 12:00 AM to 06:00 AM. For each night, we measure the motion density, bathroom motion and bed restlessness. Motion density for each hour is computed by counting the total number of motion sensor hits in that hour divided by the fraction of time spent inside the apartment during that hour. To calculate motion density during the night time, motion density for each hour is accumulated together. The bathroom motion is computed by counting the total number of motion sensor hits in the bathroom during the night time. Bed restlessness is computed by accumulating all the readings of the four sensors installed under the mattress measuring the motion on bed.

The selection of bed restlessness and bathroom motion mentioned above enables to monitor the sleeping condition as well as how many times the person makes a bathroom visit. On the other hand, the motion density tells us whether a person moves around the apartment during the night time. In the following, we explain the different parts of our system to generate textual summaries of the parameters mentioned above which are accompanied with the alerts.

A. *Linguistic Protoform Summaries of the sensor parameters*

For each sensor parameter we generate summaries of the form: *The bed restlessness tonight is* ***a lot higher*** *than* ***most*** *of the nights in the* past *two weeks,* where, *most* and *a lot higher* is the quantifier and summarizer respectively. We start by defining their membership functions as shown in Figure 2 and 3, respectively.

Corresponding to the alert for which the summary is to be generated, we go back past two weeks and compute the difference between the sensor parameter tonight and each night in the last two weeks. These values are then used to compute the memberships in the summarizers (shown in Figure 3) of each sensor parameter. For each summarizer, we

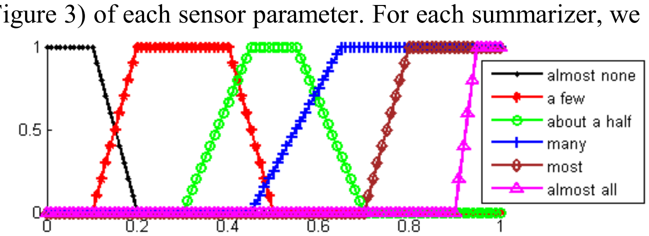

Figure 2: Membership functions of Quantifiers

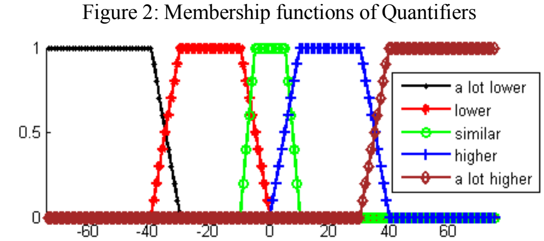

Figure 3: Membership functions of Summarizers

find the quantifier that best represents it, which is the one with highest truth value. In case of two summaries having the same truth value, the summary comprising of the quantifier on the right side in Figure 2 (that is, the one which represents more information) is chosen. The same procedure is carried out for all the three sensor modalities, irrespective of which one produced the alert. Table 1 shows the LPS for the motion density data shown in Figure 4. It is easy to see that the truth value for each summary is in line with the membership functions and data being summarized.

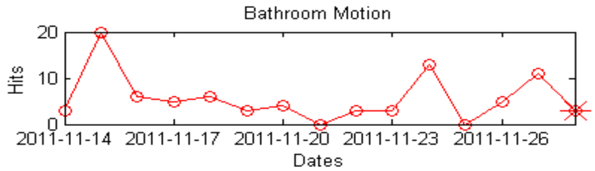

Figure 4: Overnight bathroom motion for a 2 week period

Table 1: LPS of the form: *The bathroom motion tonight is P than/to A of the nights in last 2 weeks,* where A and P are the quantifier and summarizer pair in each row.

Quantifier (A)	Summarizer (P)	Truth Value
'almost none'	'a lot lower'	1.0
'a few'	'lower'	0.8
'many'	'similar'	1.0
'almost none'	'higher'	0.7
'almost none'	'a lot higher'	1.0

B. *Language Rules*

Intuitiveness is a very important aspect of any Natural Language Generation (NLG) system. The linguistic protoform summaries presented in Table 1 convey the information about the data, however, reading it as it is, might be tiresome for the user. In the NLG domain, this stage where the information is modified to sound more intuitive is called the document planning and micro-planning and realization [8]. However, since our summaries are not too complex, we use a simple set of rules to make the presentation of textual summaries intuitive. In the following, we list some of those rules.

In the LPS of Table 1, the summaries with the quantifier *almost none* are useful when comparing two sets of summaries, for instance in [9]. However, when presented to a user, they might be redundant. Hence, for this application, we discard all the summaries with quantifier *almost none*. Another important feature that makes the language intuitive, is the use of appropriate conjunctions. For instance, the use of the conjunction '*however*' when joining two sentences representing information on the opposite end. Similarly, the rules for conjunctions, '*Also'* and *'And'* needs to be defined. Using such rules, the LPS presented in Table 1 produces the following summary: '*The bathroom motion tonight is* ***similar*** *to* ***many*** *of the past 14 nights. However, it is* ***lower*** *than a* ***few*** *nights.*'

IV. Case Studies

The experimental procedures involving human subjects described in this paper were approved by the Institutional Review Board. We presented four retrospective case studies with textual summaries of the form mentioned above to a group of health care professionals comprising of nurses, physicians, etc. to see whether they are helpful in analyzing the alerts. For each case there was a numeric alert, for one of the sensor parameters out of the motion density, bathroom motion and bed restlessness. All of the four alerts were rated very well by the group previously. For each case, the summaries of all three sensor modalities were presented.

The general consensus was that the textual summary of the past events of sensor modalities other than the one which generated the alert is helpful in analyzing the events that lead to it. However, they suggested that for individual summaries, the words representing the quantifier and summarizer membership functions should be more intuitive to a non-technical audience. For example, the use of phrase '*a lot*' is preferable to '*significantly*' or instead of saying '*past 14 days*', calling it '*past 2 weeks*' would be more helpful. Also, the inclusion of too many summarizers in one summary is not desirable. For instance, in cases where multiple summarizers have a quantifier other than *almost none*, it might be better to filter or combine some of the summaries based on the information they are conveying. Moving to the visualization, the group usually analyzes the sensor data using bar plots. However, we presented the textual summaries along with the line plots like the one shown in Fig 4. This led to a discussion on scale of the plot axis which is an important parameter in identifying patterns in the raw data. A single reading at one point can affect the scale of the complete plot. For example, a large scale may prohibit a user to interpret the consistency of sensor parameter. Since presenting the data in form of textual summaries involves a pre-processing step, we believe that they have an added advantage of solving this problem of scale. There were also a lot of suggestions on the inclusion of sensor patterns in the summaries.

Taking cue from the discussion, we modified the "language rules" presented earlier in this section to incorporate the suggestions. Moreover, we found out that many comments entered by the clinicians described the data trend that lead to the alerts. Therefore, we also developed a simple method to identify the pattern of the data that led to the alert and attached that along with the summary. To detect this pattern, we first median filter the data with a window size of 3 and then move backwards from the last day before the alert until when we find a change in direction. If the number of days with either increasing or decreasing trend are greater than or equal to 3, then we indicate the trend along with the alert summary.

Figure 5 through 8 presents the past 2 weeks data for the bathroom motion, motion density and bed restlessness that lead to the alerts, along with the text summaries for each sensor parameter. Note that the graphs may not accompany the summaries initially in a fielded system. The axis for each line plot is selected in order to have maximum resolution.

A. *Case I: Bathroom Motion Alert*

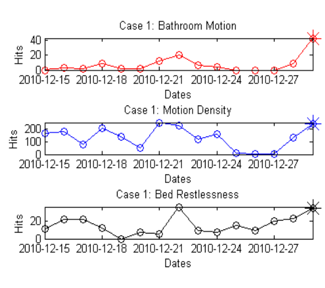

'The bathroom motion tonight is a ***lot higher*** than ***many*** of the nights in past 2 weeks with an ***increasing*** trend in the past 3 nights.'

'The motion density tonight is a ***lot higher*** than ***many*** of the nights in past 2 weeks.'

'The bed restlessness tonight is ***higher*** than ***most*** of the nights in past 2 weeks with an ***increasing*** trend in the past 6 nights.'

Figure 5: Case 1: The alert was generated by bathroom motion

The figure below shows the variation of the data that lead to the bathroom motion alert. It is clear to observe that the bathroom motion on the day of alert does not look normal as compared to the previous two weeks. Also the increasing trend in the last few days in bathroom motion and bed restlessness is evident which is also described by the summaries.

B. *Case II: Bed Restlessness Alert*

In this case also the bathroom motion and bed restlessness on the night of alert are on the higher side than usual and has been increasing for the past some nights. The same has been reported by the automatic text summaries. For the motion density, there are no significant number of days falling under the category of any of the summarizers. This might not be clear from the plot above since the scale is very large for most of the entries due to one very high value on December 13.

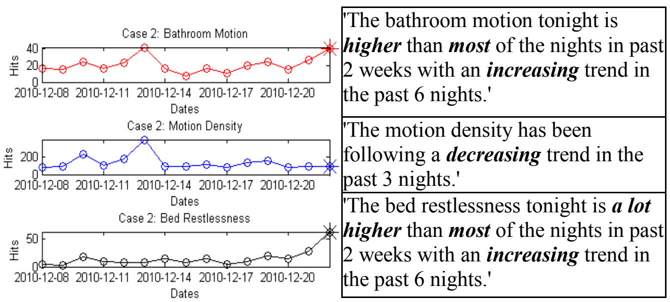

Figure 6: Case 2: The alert was generated by bed restlessness

C. *Case III: Bathroom Motion Alert*

In this case, the bathroom motion on the night of alert is *higher* than *about a half* and *a lot higher* than the other half of the nights in past two weeks. Therefore, we combine the two sentences in order to improve readability. However, it doesn't have any evident trend leading to the night of alert. The motion density in this case is been changing quite a bit, which ends up being summarized as *a lot higher* and *a lot lower* than *many* and *few* nights, respectively. Also, on smoothing the data we observe an increasing trend in motion density in the some of the last days leading up to the alert. The bed restlessness has been decreasing in the past few days, which is reported by the summary.

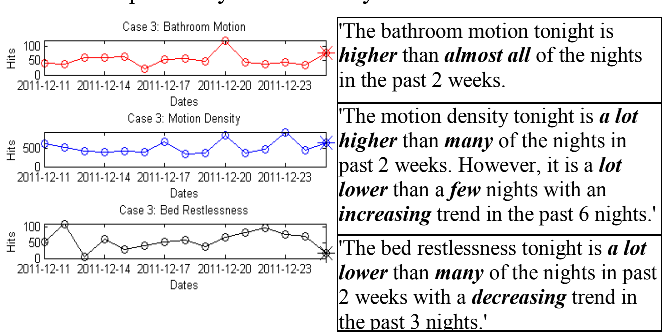

Figure 7: Case 3: The alert was generated by bathroom motion

D. *Case IV: Bed Restlessness Alert*

In this case the alert was generated due to high bed restlessness which is visible in its plot. However, there is no obvious trend in the past few days, therefore no trend summary has been reported. For motion density, there seems to be an increasing trend for the past three nights, but after filtering the data, this trend disappears, hence the system does not generate any trend summary for motion density. For the case of bathroom motion, the median filtering removes the spikes, after which it is concluded that there is an increasing trend. Also, the bathroom motion has been varying a lot in the past 2 weeks, which can be interpreted by the summary conveying that the bathroom motion is similar to many but lower than a few nights.

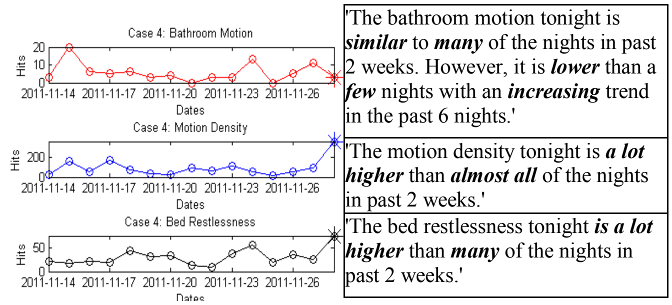

Figure 8: Case 4: The alert was generated by bed restlessness

V. CONCLUSION

With the help of four case studies we showed that our method has the potential to produce valid and useful textual summaries of the sensor data. Moreover, accompanying the summaries with patterns provides a deeper insight into data and saves the effort of identifying them manually. Another important aspect of such systems is the use of Natural Language that sounds intuitive to the user and can be tailored to clinicians, family members, or the resident. One of the areas of improvements might be to identify more patterns in the data and checking that if a similar pattern led to a health event before. Also, a system which adapts itself to each individual resident and/or sensor is something worth exploring and has the potential of getting over the scale issue faced while visualizing the data graphically.

REFERENCES

[1] C. Hallett, "Multi-modal presentation of medical histories," in *Proceedings of the 13th international conference on Intelligent user interfaces*, 2008, pp. 80-89.

[2] D. Scott, C. Hallett, and R. Fettiplace, "Data-to-text summarisation of patient records: Using computer-generated summaries to access patient histories," *Patient Education and Counseling,* vol. 92, pp. 153-159, 2013.

[3] J. Hunter, Y. Freer, A. Gatt, E. Reiter, S. Sripada, and C. Sykes, "Automatic generation of natural language nursing shift summaries in neonatal intensive care: BT-Nurse," *Artificial intelligence in medicine,* vol. 56, pp. 157-172, 2012.

[4] B. Di Eugenio, A. D. Boyd, C. Lugaresi, A. Balasubramanian, G. M. Keenan, M. Burton, *et al.*, "PatientNarr: Towards generating patient-centric summaries of hospital stays," *INLG 2014,* p. 6, 2014.

[5] M. J. Rantz, S. D. Scott, S. J. Miller, M. Skubic, L. Phillips, G. Alexander, *et al.*, "Evaluation of health alerts from an early illness warning system in independent living," *Computers, informatics, nursing: CIN,* vol. 31, p. 274, 2013.

[6] R. R. Yager, "A new approach to the summarization of data," *Information Sciences,* vol. 28, pp. 69-86, 1982.

[7] A. Jain and J. M. Keller, "On the Computation of Semantically Ordered Truth Values of Linguistic Protoform Summaries," *Accepted, FUZZ-IEEE*, 2015.

[8] E. Reiter and R. Dale, "Building Natural Language Generation Systems," 2006.

[9] A. Wilbik, J. M. Keller, and J. C. Bezdek, "Linguistic Prototypes for Data From Eldercare Residents," *Fuzzy Systems, IEEE Transactions on,* vol. 22, pp. 110-123, 2014.