Giovanni Acampora, Diane J. Cook, Parisa Rashidi, Athanasios V. Vasilakos

Healthcare Data Analytics

Contents

Chapter 1

Data Analytics for Pervasive Health

1.1 Introduction

1.1.1 Intelligent Health Systems

Nowadays, the majority of industrialized nations are facing significant complications regarding the quality and cost of various healthcare and wellbeing services. These difficulties will exacerbate even more due to an increasing aging population, which translates into a multitude of chronic diseases and tremendous demand for various health care services. As a result, the cost of the healthcare sector might not be sustainable and therefore industrialized countries need to find and plan policies and strategies to use the limited economical resources more efficiently and effectively. This need for sustainable healthcare systems translates into a range of challenges in science and technology which if solved, ultimately could benefit our global society and economy. In particular, the exploitation of information and communication technology for implementing autonomous and pro-active healthcare services will be extremely beneficial.

In face of such challenges, there has been an increasing interest in applying analytics techniques to health care problems. Analytics techniques can be applied in scenarios such as assisted living for individuals with disabilities, aging in place, and remote health monitoring. The need for developing healthcare applications based on analytics techniques is not just underscored by researchers, but also governments are trying to use such techniques to lower the cost of health care ~~In~~ US and elsewhere. ~~Especially,~~ with recent advances in analytics as well as sensor technology, we are embarking on the path of revolutionary low cost intelligent health systems embedded within the home and living environments (Black et al., 2004; Kameas and Calemis, 2010). Some examples include cognitive health monitoring systems based on activity recognition, persuasive systems for motivating users to change their health and wellness habits, and abnormal health condition detection systems. Figure 1.1 depicts how intelligent health systems might be used as cohesive services integrated into different environments and devices.

A wide variety of sensor modalities can be used when developing intelligent health systems, including wearable and ambient sensors. In the case of wearable sensors, sensors are attached to the body (Yang et al., 2009) or woven into garments (Harms et al., 2008; Metcalf et al., 2009). For example, 3-axis accelerometers distributed over an individual's body can provide information about the orientation and movement of the corresponding body part. Researchers commonly use these inertial measurement units to recognize ambulatory movements (e.g., walking, running, sitting, climbing, and falling) (Maurer et al., 2006; Srinivasan et al., 2010), posture (Lee and Mase, 2002) and gestures (Agrawal and Srikant, 1995; Junker et al., 2008; Krishnan et al., 2010; Mantyjarvi et al., 2001).

Ambient sensors such as infrared motion detectors, magnetic door sensors, break-beam sensors, and pressure mats (Barger et al., 2005; Wilson and Atkeson, 2005) also have been used to gather information about health status and user activities in indoor environments

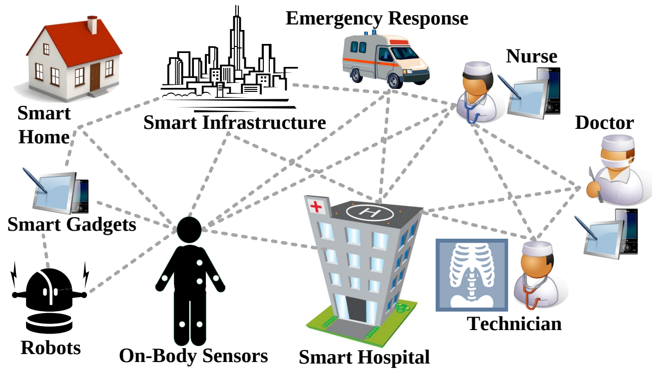

FIGURE 1.1: Interconnected world of intelligent health services.

(Cook, 2010; Nugent et al., 2009). Because this approach embeds sensors within environments, it is well suited to creating intelligent health systems such as smart environments and has been widely adopted for health monitoring and ambient assisted living (van Kasteren et al., 2010a). Other sensors include RFID tags (Buettner et al., 2009; Patterson et al., 2005), shake sensors (Philipose et al., 2004), video cameras (Brdiczka et al., 2009; Duan et al., 2010; Moeslund et al., 2006; Weinland et al., 2011), microphones (Hollosi et al., 2010; Lukowicz et al., 2004), and GPS locators (Liao et al., 2007; Patterson et al., 2003).

In this chapter, we will explore how different analytics techniques can be used for supporting the development of intelligent health systems and sensor data analysis.

1.2 Introduction to Analytic Techniques

Among data analytics methods, machine learning techniques, have been heavily used in pervasive health applications in recent years. Machine learning is a sub-discipline of artificial intelligence which allows real world systems to learn from data. Machine learning algorithms are able to generalize in case of new, unseen data examples by learning from a set of observed data examples called a *training set*. For example, after being trained on a training set of sample accelerometer data marked as `walking` or `jogging`, a machine learning algorithm will be able to classify the future data points into `walking` and `jogging` classes.

Machine learning methods are prevalent in modern applications, and they have been successfully implemented and deployed into many real world applications. In the health care domain, these methods also have gained vast popularity by being used in numerous applications such as electronic health record analysis for predicting life expectancy (Meystre et al., 2008), computer aided diagnosis (Doi, 2009), DNA sequence classification (Wang et al., 1999), medical imaging (Wernick et al., 2010), pharmacovigilance and post-marketing adverse drug reaction detection (Harpaz et al., 2013; Hauben et al., 2005), and early prediction of diseases (Pearce et al., 2006), among many other health care applications.

Machine learning also has become quite ubiquitous in many pervasive health applications. Many current pervasive health applications rely on machine learning to analyze the sensor data. These algorithms make it possible to obtain insights from the wealth of sensor data, and to support decision making in pervasive health applications. Some example applications include activity recognition algorithms for monitoring dementia patients (Rashidi and Cook, 2011a), or physical activity monitoring algorithms for fitness and wellbeing applications (Kwapisz et al., 2011). More examples can be found in section 1.2.3.

1.2.1 Supervised Techniques

Two common forms of machine learning techniques are *supervised* learning and *unsupervised* learning algorithms. The former focuses on predicting the known properties of data, while the second focuses on discovering unknown knowledge from data.

Supervised learning techniques construct an internal model of the observed data, which is used to predict the label of future examples. Supervised learning methods predict the label of a previously unseen data point based on the model's properties learned from the training data. Typically, if the predicted label is discrete, the supervised technique is called a *classifier*. If the predicted attribute is a continuous value, it is called a *regression* task.

Supervised algorithms require a labeled training set, where each training data point is annotated with its predictive label. Each data point itself is described in terms of a number of features. For example, if the goal is to distinguish between `walking` and `jogging` using accelerometer data, then each data point in the training set can be described by features such as the mean and the standard deviance of points in the surrounding window, peak value in this window, and other relevant features. Each one of the data points in the training set also needs to be annotated with one of the two labels: `jogging` or `walking`. The annotated training data set is then used as input in the training stage of the algorithm, thus allowing the machine learning algorithm to make generalization and predictions in case of future unseen data points. To test the performance of a machine learning algorithm, part of the annotated data is usually set aside, and is used to test the performance of the algorithm. If the performance is satisfactory, the machine learning algorithm can be deployed and integrated with the rest of the system.

Figure 1.2 shows how a supervised machine learning method works. In general, processing data is a multi-stage process. First, data are captured and annotated with their corresponding labels. Then, preprocessing tasks are performed on data. For example, accelerometer data might be filtered to remove high frequency noises and is segmented into shorter segments. Next, statistical and morphological features are extracted from data. This step might also reduce the number of features by applying feature selection and dimensionality reduction techniques. Finally, the classification step predicts the class of activity according to the features.

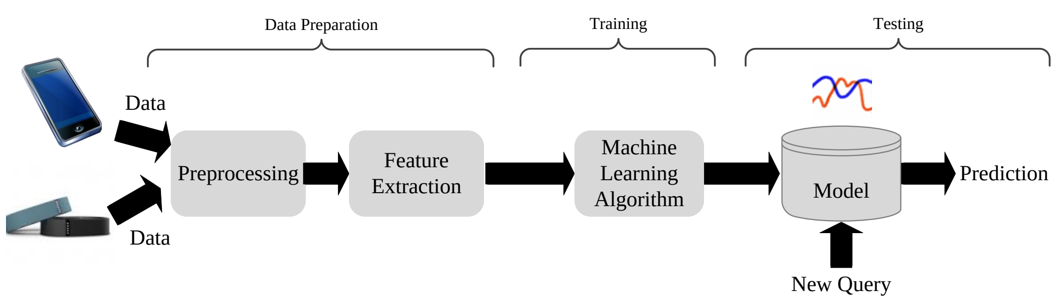

FIGURE 1.2: Machine Learning often involves a number of steps, including data preparation, training, and testing.

Classification techniques are one of the most widely used type of supervised machine learning techniques in the context of pervasive health applications. As mentioned before, if the predicted label is a discrete value, such as predicting whether the user is `walking` or `jogging`, it is called a *classification* task, and the algorithm that performs the classification is called a *classifier*. Some famous classification techniques include naive bayes (Lewis, 1998), decision trees (Quinlan, 1986; Safavian and Landgrebe, 1991), support vector machines (Cortes and Vapnik, 1995), logistic regression (Hosmer Jr et al., 2013), and neural networks (Hagan et al., 1996).

Figure 1.3 shows how a classification algorithm works for a toy problem. Here, each data point is represented in terms of two features (i.e. x and y axis). In real world, most datasets have tens, hundreds, or even thousands of features. As can be seen in figure 1.3, the decision boundary separates the first class from the second class. In reality, we do not know what the true boundary looks like, therefore machine learning algorithms try to reconstruct an approximate boundary based on extracting information from the training set. Some algorithms might only reconstruct a liner boundary such as the perceptron algorithm (Rosenblatt, 1958), or a recti-linear boundary such as decision trees (Quinlan, 1986), or a non-linear boundary such as the nearest neighborhood (Chang, 1974) and kernel machines (Williams and Seeger, 2001). Some algorithms might also assign a confidence value to each prediction, such as the logistic regression algorithm.

If the predicted value is a continuous value, the supervised machine learning technique is called a *regression* task. For example, if our goal is to predict blood pressure value in the next few days based on current food intake and activity level, then a regression algorithm is needed. Similar to classification techniques, regression algorithms might be linear or non-linear in nature.

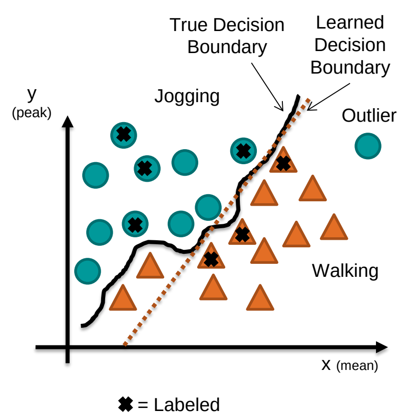

FIGURE 1.3: The true decision Boundary is typically unknown, machine learning algorithms try to reconstruct an approximate boundary.

1.2.2 Unsupervised Techniques

Unlike the supervised learning algorithms, unsupervised methods do not require any labeled data. Instead, they try to automatically find interesting patterns in unlabeled data, such as by grouping similar examples together into a cluster. For example, sequence mining can be used to discover user activities from ambient sensor data obtained in a smart home. Some of the unsupervised machine learning techniques includes cluster analysis (Kaufman and Rousseeuw, 2009), as well as a large class of data mining methods including association rule mining (Zhang and Zhang, 2002), frequent item-set mining (Borgelt, 2012), and sequence mining (Mooney and Roddick, 2013).

Frequent pattern mining is an area of data mining that involves finding frequently observed patterns in data. Pattern mining algorithms might look for different types of patterns: frequent itemsets, frequent sequences, frequent trees, or frequent graphs, among others. Frequent pattern mining algorithms can be useful in many pervasive health applications. For example, activity recognition is often used in monitoring the activities of elderly with dementia. While recognizing predened activities often relies on supervised learning techniques, frequent pattern mining is valuable for its ability to discover recurring sequences of unlabeled sensor activities that may comprise activities of interest.

The pioneering work of Agrawal's Aprori algorithm (Agrawal et al., 1994) was the starting point in this area. Apriori finds frequent itemsets and uses a bottom up approach, where frequent subsets are extended one item at a time in a step known as candidate generation, and then groups of candidates are tested against the data; the algorithm terminates when no further successful extensions are found. There have been many extensions and variations of frequent pattern mining algorithms (Garofalakis et al., 1999; Han et al., 2000, 2007; Aggarwal et al., 2009). For more information, refer to surveys on frequent pattern mining (Mooney and Roddick, 2013; Borgelt, 2012).

1.2.3 Example Applications

There are many examples of using both supervised and unsupervised machine learning techniques in pervasive health technology in the literature (Acampora et al., 2013), here we mention two prominent examples: continuous health monitoring and emergency detection.

Continuous Health Monitoring: Machine learning techniques can be used in continuous health monitoring applications, where a variety of noninvasive sensors monitor various physiological parameters such as ECG, EEG, respiration, and even biochemical processes such as wound healing. For example, Jin et al. (Jin et al., 2009) describe a cell phone-based real-time monitoring technology for cardiovascular disease (CVD) which automatically detects and classifies abnormal CVD conditions using neural networks. The training data is a combination of both an individual's cardiac characteristics and information from clinical ECG databases. Similar research has been done using sensors to monitor EEG to predict onset of epilepsy (Shoeb and Guttag, 2010), or in spirometry sensing using the built-in microphone of cell phones (Larson et al., 2012). In addition to monitoring physiological signs, one can also monitor and track activity as an indicator of physical and cognitive function. Activity might refer to activities of daily living (ADL) when monitoring dementia patients and elderly (Rashidi et al., 2011a), or to physical activity in the context of fitness and well-being applications (Kwapisz et al., 2011), or to online activity when monitoring patients with mental disorders (Burns et al., 2011). Both supervised and unsupervised approaches have been used quite frequently in activity recognition, especially supervised techniques (Acampora et al., 2013; Lane et al., 2010).

Emergency Detection: While it is valuable to monitor common normal events, we are also very interested in abnormal events. These abnormal events may indicate a crisis or an abrupt change in a regimen that is associated with health difculties. Classification techniques can be used to detect abnormal events from normal events based on incoming sensor data. There have been solutions for detecting emergency situations using PIR sensor network (Buckland et al., 2006), for detecting possible fall in elderly using ambient sound and sensors (Zhuang et al., 2009; Alwan et al., 2006), or for classifying cane usage and walking patterns, in case of high risk of falling (Wu et al., 2008).

TABLE 1.1: Advanced analytics techniques used in intelligent health systems.

Component	Example Techniques	Example Applications
Activity Recognition	Graphical models	Health monitoring
Behavior Discovery	Sequence mining	Behavior monitoring
Anomaly Detection	Statistical methods	Emergency detection
Planning	D-HTN	Prompting
Decision Support	Knowledge based	Communication of care personnel
Anonymization	K-Anonymity	Privacy preservation

1.3 Advanced Analytic Techniques

In this section we introduce the set of more advanced analytics techniques that enable us to develop sophisticated intelligent health care systems, as summarized in Table 1.1.

1.3.1 Activity Recognition

Intelligent health systems focus on the needs of a human and therefore require information about the activities being performed by the human (Singla et al., 2010a). At the core of such technologies is activity recognition, which is a challenging and well researched problem. The goal of activity recognition is to identify activities as they occur based on data collected by sensors. There exist a number of approaches to activity recognition (Chen et al., 2010; Rashidi et al., 2011b) that vary depending on the underlying sensor technologies that are used to monitor activities, the machine learning algorithms that are used to model the activities, and the complexity of the activities that are being modeled.

In terms of sensor technology, besides using wearable and ambient sensors, researchers have used a varity of other sensors for activity recognition. Some activities such as washing dishes, taking medicine, and using the phone are characterized by interacting with unique objects. In response, researchers have explored the usage of RFID tags (Buettner et al., 2009; Patterson et al., 2005) and accelerometers or shake sensors (Philipose et al., 2004) for tagging these objects and using the data for activity recognition. The challenge with this modality is deciding which objects to tag with sensors. One approach that has been investigated (Palmes et al., 2010; Perkowitz et al., 2004) is to mine web page description of activities to determine which objects are instrumental to the activity and help differentiate the activity from others. Other sensor modalities that have been researched for activity recognition include video cameras (Brdiczka et al., 2009; Duan et al., 2010; Moeslund et al., 2006; Weinland et al., 2011), microphones (Hollosi et al., 2010; Lukowicz et al., 2004), and GPS locators (Liao et al., 2007; Patterson et al., 2003). Each of these does face a unique challenge for use in health care applications. Cameras and microphones need to be carefully positioned and robust in the presence of occlusion. Furthermore, these technologies are not always well accepted because of privacy concerns. Smart phones are increasing in popularity for activity recognition (Győrbíró et al., 2009; Lane et al., 2011) because sensors in the phone collect all of the gyro, accelerometer, GPS, acoustic, and video data found the other methods, as long as they are on the individual while they perform activities.

1.3.1.1 Activity Models

The methods that are used to model and recognize activities are as varied as the sensor modalities used to observe activities. Existing methods can be broadly categorized into template matching / transductive techniques, generative, and discriminative approaches. Template matching techniques employ a nearest-neighbor classifier based on Euclidean distance or dynamic time warping (Bao and Intille, 2004; Stikic and Schiele, 2009). Generative approaches such as naïve Bayes classifiers where activity samples are modeled using Gaussian mixtures have yielded promising results for batch learning (Brdiczka et al., 2009; Tapia et al., 2004; van Kasteren and Krose, 2007). Generative probabilistic graphical models such as hidden Markov models (Brand et al., 1997; Tapia et al., 2004; van Kasteren et al., 2011; Cao and Liu, 2010) and dynamic Bayesian networks (Liao et al., 2007; Wilson and Atkeson, 2005) have been used to model activity sequences and to smooth recognition results of an ensemble classifier (Cook et al., 2012). Decision trees as well as bagging and boosting methods have been tested (Maurer et al., 2006). Discriminative approaches, including support vector machines (Brdiczka et al., 2009) and conditional random fields (Hsu et al., 2010; Mahdaviani and Choudhury, 2007; van Kasteren et al., 2010a; Vail et al., 2007) which attempt to maximally separate activity clusters, have also been effective.

1.3.1.2 Activity Complexity

Many of these methods analyze pre-segmented activity sequences that have been collected in controlled settings. More recently, attempts have been made to perform automatic segmentation of the data into sensor events that belong to the same activity class (Gu et al., 2010; Palmes et al., 2010; Rashidi and Cook, 2010). Still others have focused on recognizing activities in real time from continuous sensor streams (Rashidi and Cook, 2010). In addition, researchers have also investigated methods of leveraging information or models in one setting to boost activity recognition for a new sensor network (van Kasteren et al., 2010b), a new environmental setting (Rashidi and Cook, 2011a,b; Cook, 2010), or new activity labels (Iglesias et al., 2011). Another level of complexity for activity recognition is analyzing data for interwoven activities (Modayil et al., 2008; Singla et al., 2010b) or concurrent activities (Wang et al., 2009b). Humans often make efficient use of time by performing a step for one activity while still in the middle of another activity, causing the sensor streams to interweave. Concurrent activities may occur if a single sensor event contributes to more than one activity. This situation may also indicate that multiple residents are in the space, which can be a challenge for activity recognition algorithms (Crandall and Cook, 2009; Phua et al., 2011).

1.3.2 Behavioral Pattern Discovery

While recognizing predefined activities often relies on supervised learning techniques, unsupervised learning is valuable for its ability to discover recurring sequences of unlabeled sensor activities that may comprise activities of interest. Methods for activity discovery build on a rich history of discovery research, including methods for mining frequent sequences, mining frequent patterns using regular expressions (Barger et al., 2005), constraint-based mining (Pei et al., 2007), and frequent-periodic pattern mining (Rashidi and Cook, 2009; Heierman III and Cook, 2003).

More recent work extends these early approaches to look for more complex patterns.

Ruotsalainen et al. (Ruotsalainen et al., 2007) design the Gais genetic algorithm to detect interleaved patterns in an unsupervised learning fashion. Other approaches have been proposed to mine more complex discontinuous patterns (Rashidi et al., 2011b; Pei et al., 2007), from streaming data over time (Rashidi and Cook, 2010), in different types of sequence datasets and to allow variations in occurrences of the patterns (Pei et al., 2007). Discovered behavioral patterns are valuable to interpret sensor data, and models can be constructed from the discovered patterns to recognize instances of the patterns when they occur in the future.

1.3.3 Anomaly Detection

While it is value to characterize and recognize common normal events that account for the majority of the sensor events that are generated, for health applications we are also very interested in abnormal events. These abnormal events may indicate a crisis or an abrupt change in a regimen that is associated with health difficulties.

Abnormal event detection, or anomaly detection, is also important in security monitoring where suspicious activities need to be flagged and handled. Anomaly detection is most accurate when it is based on behaviors that are frequent and predictable. There are common statistical methods to automatically detect and analyze anomalies including the box plot, the chart, and the CUSUM chart (Tukey, 1977). Anomalies can be captured at different population scales. For example, while most of the population may exhibit condition A, one person might exhibit condition B, which pinpoints a condition needing further investigation (Dawadi et al., 2011). Anomalies may also be discovered at different temporal scales, including single events, days, or weeks (Song et al., 2010).

Little attention has been devoted to anomaly detection in intelligent health systems. This is partly because the notion of an anomaly is somewhat ill-defined. Many possible interpretations of anomalies have been offered and use cases have even been generated for intelligent health systems (Tran et al., 2010). Some algorithmic approaches have been suggested that build on the notion of expected temporal relationships between events and activities (Jakkula and Cook, 2008). Others tag events as anomalies if the occur rarely and they are not anticipated for the current context (Yin et al., 2008).

1.3.4 Planning and Scheduling

Automatic planning and scheduling can be useful in many intelligent health systems applications. Automatic planning techniques achieve a goal state by starting from an initial known state and choosing among possible actions at each state. Planning can be useful in a number of different intelligent health scenarios. For example, planning can be used to schedule daily activities in a flexible manner for reminding dementia patients about their daily activities. It also can be used in order to detect any possible deficiencies in task execution, and to help dementia patients to complete those steps. Another use of planning is in automating daily routines, in order to allow users with physical limitations to live a more independent lifestyle.

In the past, many planning techniques have been proposed. Some techniques include decision-theoretic techniques (e.g. Markov Decision Processes (Bellman, 1957)), search methods (e.g. forward and backward search (Bonet and Geffner, 2001)), graph-based techniques (e.g. GraphPlan (Blum and Furst, 1997)), hierarchal techniques (e.g. O-Plan (Tate et al., 2000)), and reactive planning techniques (e.g. (Firby, 1987)). For example, graph-based planning techniques represent search space of possible actions in form of a graph, hierarchal planning techniques use hierarchies to predefine groups of actions, and reactive planning techniques adjust the plan based on sensed information.

Intelligent health systems pose many new challenges to the classical planning techniques. For example, the planner has to be functional in a dynamic environment where the outcome of the actions and their duration is not deterministic. Also, the availability of resources might change due to user mobility or other factors. Therefore, more advanced planning techniques have been proposed by extending classical planning techniques (Simpson et al., 2006). One example is the distributed hierarchal task network (D-HTN) technique (Amigoni et al., 2005) which extends the hierarchal task network (HTN). It uses a centralized approach to manage the distributed capabilities provided by the distributed devices. The distributed devices might be available in a permanent or transient manner. D-HTN has been studied in the context of care for diabetic patients at home, where different home devices communicate and coordinate plans with each other in a distributed manner. For example data from monitoring devices might require actions such as adjusting the room temperature, suggesting insulin injection, or contacting medical help.

Several intelligent health systems have been reported in the literature which use automated planning and scheduling, especially to help dementia patients. COACH is one such system which provides task guidance to Alzheimer's disease patients (Mihailidis et al., 2004). It uses a hand-coded representation of detailed steps of hand-washing, and relies on vision techniques to recognize user steps. If user is unable to complete a particular step, detailed instructions are provided. Another example is PEAT which also provides task guidance to the user (Levinson, 1997). It maintains a detailed model of the daily plan in terms of hierarchal events, and tracks their execution. PEAT has the capability of rescheduling activities in case of unexpected events, however it lacks any real sensory information from the world, except for user feedback. Autominder by Pollack et al. (Pollack et al., 2003) is another system which provides users with reminders about their daily activities by reasoning about any disparities between what the client is supposed to do and what she is doing, and makes decisions about whether and when to issue reminders.

1.3.5 Decision Support

Decision Support Systems (DSS) (Efraim, 1996; Eom et al., 1998; Eom and Kim, 2006) have been widely used in the field of heath care for assisting physicians and other health care professionals with decision making tasks, for example for analyzing patient data (Omichi et al., 1984; Bakonyi et al., 1985; Bankowitz et al., 1992; Linnarsson, 1993; Snyder-Halpern, 1999; Romano and Stafford, 2011; Perwez et al., 2012). DSS systems are mainly based on two mainstream approaches: *knowledge-based* and *non knowledge-based*.

The knowledge-based DSS consists of two principal components: the knowledge database and the inference engine. The knowledge database contains the rules and associations of compiled data which often take the form of IF-THEN rules, whereas the inference engine combines the rules from the knowledge database with the real patients' data in order to generate new knowledge and to propose a set of suitable actions. Different methodologies have been proposed for designing health care knowledge databases and inference engines, such as the ontological representation of information (Kaptein et al., 2010).

The non knowledge-based DSS have no direct clinical knowledge about a particular health care process, however they learn clinical rules from past experiences and by finding patterns in clinical data. For example, various machine learning algorithms such as decision trees represent methodologies for learning health care and clinical knowledge.

Both of these approaches could be used in conjunction with intelligent health systems. Indeed, the sensitive, adaptive, and unobtrusive nature of intelligent health systems is particularly suitable for designing decision support systems capable of supporting medical staffs in critical decisions. In particular, intelligent health systems enable the design of the third generation of *telecare systems*. The first generation was the *panic-alarms gadgets*, often worn as pendants or around the wrist to allow a person to summon help in the case of a fall or other kinds of health emergency. The second generation of telecare systems uses sensors to automatically detect situations where assistance or medical decisions are needed. Finally, the third generation represents intelligent health systems which move away from the simple reactive approach and adopt a proactive strategy capable of anticipating emergency situations. As a result, DSSs could be used with *multimodal sensing* and *wearable computing* technologies for constantly monitoring all vital signs of a patient and for analyzing such data in order to take real-time decisions and opportunely support that people.

Finally, DSSs are jointly used with the intelligent health systems paradigm for enhancing communications among health personnel such as doctors and nurses. For example, Anya et al. have introduced a DSS system based on context aware knowledge modeling aimed at facilitating the communication and improving the capability to take decisions among healthcare personal located in different geographical sites (Anya et al., 2010).

1.3.6 Anonymization and Privacy Preserving Techniques

As intelligent health systems become more ubiquitous, more information will be collected about individuals and their lives. While the information is intended to promote the well being of individuals, it may be considered an invasion of privacy and, if intercepted by other parties, could be used for malicious purposes.

While some privacy concerns focus on the perception of intrusive monitoring (Demiris et al., 2008), many heavily-deployed Internet gadgets and current intelligent systems are nearly devoid of security against adversaries, and many others employ only crude methods for securing the system from internal or external attacks. The definition of privacy will continue to evolve as ambient intelligent systems mature (Hayes et al., 2007). This is highlighted by the fact that even if personal information is not directly obtained by an unwanted party, much of the information can be inferred even from aggregated data. For this reason, a number of approaches are being developed to ensure that important information cannot be gleaned from mined patterns (Laszlo and Mukherjee, 2005; Wang et al., 2009a).

1.4 Conclusions and Future Outlook

With the help of analytics techniques, intelligent health systems promise to enhance our health and well-being in many aspects in a radical manner by successful acquisition and interpretation of contextual information. We are aware that the goals set up for analytics techniques in intelligent health systems are not easily reachable and there are still many challenges to face and, consequently, this research field is getting more and more impetus. Researchers with different backgrounds are enhancing the current state of the art of intelligent health systems by addressing fundamental problems not only in analytics, but also related to human factors, design and implementation, security, as well as social and ethical issues. As a result, we are confident that this synergic approach will materialize the complete vision of intelligent health systems and its full application to health care and human well-being.

Bibliography

Acampora, G., Cook, D., Rashidi, P., Vasilakos, A., 2013. A survey on ambient intelligence in healthcare. Proceedings of the IEEE 101 (12), 2470–2494.

Aggarwal, C. C., Li, Y., Wang, J., Wang, J., 2009. Frequent pattern mining with uncertain data. In: Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, pp. 29–38.

Agrawal, R., Srikant, R., 1995. Mining sequential patterns. In: Data Engineering, 1995. Proceedings of the Eleventh International Conference on. IEEE, pp. 3–14.

Agrawal, R., Srikant, R., et al., 1994. Fast algorithms for mining association rules. In: Proc. 20th Int. Conf. Very Large Data Bases, VLDB. Vol. 1215. pp. 487–499.

Alwan, M., Rajendran, P., Kell, S., Mack, D., Dalal, S., Wolfe, M., Felder, R., 2006. A smart and passive floor-vibration based fall detector for elderly. In: Information and Communication Technologies, 2006. ICTTA'06. 2nd. Vol. 1. IEEE, pp. 1003–1007.

Amigoni, F., Gatti, N., Pinciroli, C., Roveri, M., 2005. What planner for ambient intelligence applications? Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on 35 (1), 7–21.

Anya, O., Tawfik, H., Amin, S., Nagar, A., Shaalan, K., july 2010. Context-aware knowledge modelling for decision support in e-health. In: Neural Networks (IJCNN), The 2010 International Joint Conference on. pp. 1 –7.

Bakonyi, P., Bekessy, A., Demetrovics, J., Kerekfy, P., Ruda, M., 1985. Microcomputer-network based decision support system for health-care organizations. In: IFAC Proceedings Series. pp. 1651–1658, cited By (since 1996) 0.

Bankowitz, R., Lave, J., McNeil, M., 1992. A method for assessing the impact of a computer-based decision support system on health care outcomes. Methods of Information in Medicine 31 (1), 3–10, cited By (since 1996) 12.

Bao, L., Intille, S., 2004. Activity recognition from user-annotated acceleration data. Pervasive Computing, 1–17.

Barger, T., Brown, D., Alwan, M., 2005. Health-status monitoring through analysis of behavioral patterns. Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on 35 (1), 22–27.

Bellman, R., 1957. A markovian decision process. Tech. rep., DTIC Document.

Black, J. P., Segmuller, W., Cohen, N., Leiba, B., Misra, A., Ebling, M. R., Stern, E., June 2004. Pervasive computing in health care: Smart spaces and enterprise information systems. In: Proceedings of the MobiSys 2004 Workshop on Context Awareness. Boston.

Blum, A., Furst, M., 1997. Fast planning through planning graph analysis. Artificial intelligence 90 (1-2), 281–300.

Bonet, B., Geffner, H., 2001. Planning as heuristic search. Artificial Intelligence 129 (1), 5–33.

Borgelt, C., 2012. Frequent item set mining. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery 2 (6), 437–456.

Brand, M., Oliver, N., Pentland, A., 1997. Coupled hidden markov models for complex action recognition. In: Computer Vision and Pattern Recognition, 1997. Proceedings., 1997 IEEE Computer Society Conference on. IEEE, pp. 994–999.

Brdiczka, O., Crowley, J., Reignier, P., 2009. Learning situation models in a smart home. Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on 39 (1), 56–63.

Buckland, M., Frost, B., Reeves, A., 2006. Liverpool telecare pilot: telecare as an information tool. Informatics in Primary Care 14 (3), 191–196.

Buettner, M., Prasad, R., Philipose, M., Wetherall, D., 2009. Recognizing daily activities with rfid-based sensors. In: Proceedings of the 11th international conference on Ubiquitous computing. ACM, pp. 51–60.

Burns, M. N., Begale, M., Duffecy, J., Gergle, D., Karr, C. J., Giangrande, E., Mohr, D. C., 2011. Harnessing context sensing to develop a mobile intervention for depression. Journal of medical Internet research 13 (3).

Cao, X.-Q., Liu, Z.-Q., 2010. Human motion detection using markov random fields. Journal of Ambient Intelligence and Humanized Computing 1 (3), 211–220, cited By (since 1996) 0.

Chang, C.-L., 1974. Finding prototypes for nearest neighbor classifiers. Computers, IEEE Transactions on 100 (11), 1179–1184.

Chen, L., Hoey, J., Nugent, C., Cook, D., Hu, Z., 2010. Sensor-based activity recognition. IEEE Transactions on Systems, Man, and Cybernetics, Part CIn Press.

Cook, D., 2010. Learning setting-generalized activity models for smart spaces. Intelligent Systems, IEEE (99), 1–1.

Cook, D., Feuz, K. D., Krishnan, N. C., 2012. Transfer learning for activity recognition: A survey. International Journal on Knowledge and Information Systems To appear.

Cortes, C., Vapnik, V., 1995. Support vector machine. Machine learning 20 (3), 273–297.

Crandall, A., Cook, D., 2009. Coping with multiple residents in a smart environment. Journal of Ambient Intelligence and Smart Environments 1 (4), 323–334.

Dawadi, P., Cook, D., Parsey, C., Schmitter-Edgecombe, M., Schneider, M., 2011. An approach to cognitive assessment in smart home. In: Proceedings of the 2011 workshop on Data mining for medicine and healthcare. ACM, pp. 56–59.

Demiris, G., Oliver, D., Dickey, G., Skubic, M., Rantz, M., 2008. Findings from a participatory evaluation of a smart home application for older adults. Technology and Health Care 16 (2), 111–118.

Doi, K., 2009. Computer-aided diagnosis in radiological imaging: current status and future challenges. In: Sixth International Symposium on Multispectral Image Processing and Pattern Recognition. International Society for Optics and Photonics, pp. 74971A–74971A.

Duan, L., Xu, D., Tsang, I., Luo, J., 2010. Visual event recognition in videos by learning from web data. In: Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on. IEEE, pp. 1959–1966.

Efraim, T., 1996. Implementing decision support systems: a survey. In: Proceedings of the IEEE International Conference on Systems, Man and Cybernetics. Vol. 4. pp. 2540–2545, cited By (since 1996) 6.

Eom, S., Kim, E., 2006. A survey of decision support system applications (1995-2001). Journal of the Operational Research Society 57 (11), 1264–1278, cited By (since 1996) 23.

Eom, S., Lee, S., Kim, E., Somarajan, C., 1998. A survey of decision support system applications (1988-1994). Journal of the Operational Research Society 49 (2), 109–120, cited By (since 1996) 30.

Firby, R., 1987. An investigation into reactive planning in complex domains. In: Proceedings of the sixth National conference on Artificial intelligence. Vol. 202. p. 206.

Garofalakis, M. N., Rastogi, R., Shim, K., 1999. Spirit: Sequential pattern mining with regular expression constraints. In: VLDB. Vol. 99. pp. 7–10.

Gregoski, M. J., Mueller, M., Vertegel, A., Shaporev, A., Jackson, B. B., Frenzel, R. M., Sprehn, S. M., Treiber, F. A., 2012. Development and validation of a smartphone heart rate acquisition application for health promotion and wellness telehealth applications. International journal of telemedicine and applications 2012, 1.

Gu, T., Chen, S., Tao, X., Lu, J., 2010. An unsupervised approach to activity recognition and segmentation based on object-use fingerprints. Data & Knowledge Engineering 69 (6), 533–544.

Győrbíró, N., Fábián, Á., Hományi, G., 2009. An activity recognition system for mobile phones. Mobile Networks and Applications 14 (1), 82–91.

Hagan, M. T., Demuth, H. B., Beale, M. H., et al., 1996. Neural network design. Pws Pub. Boston.

Han, J., Cheng, H., Xin, D., Yan, X., 2007. Frequent pattern mining: current status and future directions. Data Mining and Knowledge Discovery 15 (1), 55–86.

Han, J., Pei, J., Mortazavi-Asl, B., Chen, Q., Dayal, U., Hsu, M.-C., 2000. Freespan: frequent pattern-projected sequential pattern mining. In: Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, pp. 355–359.

Harms, H., Amft, O., Tröster, G., Roggen, D., 2008. Smash: A distributed sensing and processing garment for the classification of upper body postures. In: Proceedings of the ICST 3rd international conference on Body area networks. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), p. 22.

Harpaz, R., Vilar, S., DuMouchel, W., Salmasian, H., Haerian, K., Shah, N. H., Chase, H. S., Friedman, C., 2013. Combing signals from spontaneous reports and electronic health records for detection of adverse drug reactions. Journal of the American Medical Informatics Association 20 (3), 413–419.

Hauben, M., Madigan, D., Gerrits, C. M., Walsh, L., Van Puijenbroek, E. P., 2005. The role of data mining in pharmacovigilance.

Hayes, G., Poole, E., Iachello, G., Patel, S., Grimes, A., Abowd, G., Truong, K., 2007. Physical, social, and experiential knowledge in pervasive computing environments. IEEE Pervasive Computing, 56–63.

Heierman III, E., Cook, D., 2003. Improving home automation by discovering regularly occurring device usage patterns. In: Data Mining, 2003. ICDM 2003. Third IEEE International Conference on. IEEE, pp. 537–540.

Hollosi, D., Schroder, J., Goetze, S., Appell, J., 2010. Voice activity detection driven acoustic event classification for monitoring in smart homes. In: Applied Sciences in Biomedical and Communication Technologies (ISABEL), 2010 3rd International Symposium on. IEEE, pp. 1–5.

Hosmer Jr, D. W., Lemeshow, S., Sturdivant, R. X., 2013. Applied logistic regression. Wiley. com.

Hsu, K., Chiang, Y., Lin, G., Lu, C., Hsu, J., Fu, L., 2010. Strategies for inference mechanism of conditional random fields for multiple-resident activity recognition in a smart home. Trends in Applied Intelligent Systems, 417–426.

Iglesias, J., Angelov, P., Ledezma, A., Sanchis, A., 2011. Creating evolving user behavior profiles automatically. Knowledge and Data Engineering, IEEE Transactions on (99), 1–1.

Jakkula, V., Cook, D., 2008. Anomaly detection using temporal data mining in a smart home environment. Methods of information in medicine 47 (1), 70–75.

Jin, Z., Sun, Y., Cheng, A. C., 2009. Predicting cardiovascular disease from real-time electrocardiographic monitoring: An adaptive machine learning approach on a cell phone. In: Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE. IEEE, pp. 6889–6892.

Junker, H., Amft, O., Lukowicz, P., Tröster, G., 2008. Gesture spotting with body-worn inertial sensors to detect user activities. Pattern Recognition 41 (6), 2010–2024.

Kameas, A., Calemis, I., 2010. Pervasive Systems in Health Care. In: Nakashima, H., Aghajan, H., Augusto, J. C. (Eds.), Handbook of Ambient Intelligence and Smart Environments. Springer, p. 315.

Kaptein, M., Markopoulos, P., de Ruyter, B., Aarts, E., 2010. Persuasion in ambient intelligence. Journal of Ambient Intelligence and Humanized Computing 1 (1), 43–56, cited By (since 1996) 6.

Kaufman, L., Rousseeuw, P. J., 2009. Finding groups in data: an introduction to cluster analysis. Vol. 344. Wiley. com.

Krishnan, N., Lade, P., Panchanathan, S., 2010. Activity gesture spotting using a threshold model based on adaptive boosting. In: Multimedia and Expo (ICME), 2010 IEEE International Conference on. IEEE, pp. 155–160.

Kwapisz, J. R., Weiss, G. M., Moore, S. A., 2011. Activity recognition using cell phone accelerometers. ACM SIGKDD Explorations Newsletter 12 (2), 74–82.

Lane, N., Xu, Y., Lu, H., Hu, S., Choudhury, T., Campbell, A., Zhao, F., 2011. Enabling large-scale human activity inference on smartphones using community similarity networks (csn). Ubicomp11, 355–364.

Lane, N. D., Miluzzo, E., Lu, H., Peebles, D., Choudhury, T., Campbell, A. T., 2010. A survey of mobile phone sensing. Communications Magazine, IEEE 48 (9), 140–150.

Larson, E. C., Goel, M., Boriello, G., Heltshe, S., Rosenfeld, M., Patel, S. N., 2012. Spirosmart: using a microphone to measure lung function on a mobile phone. In: Proceedings of the 2012 ACM Conference on Ubiquitous Computing. ACM, pp. 280–289.

Laszlo, M., Mukherjee, S., 2005. Minimum spanning tree partitioning algorithm for microaggregation. Knowledge and Data Engineering, IEEE Transactions on 17 (7), 902–911.

Lee, S., Mase, K., 2002. Activity and location recognition using wearable sensors. Pervasive Computing, IEEE 1 (3), 24–32.

Levinson, R., 1997. The planning and execution assistant and trainer (peat). The Journal of Head Trauma Rehabilitation 12 (2), 85.

Lewis, D. D., 1998. Naive (bayes) at forty: The independence assumption in information retrieval. In: Machine learning: ECML-98. Springer, pp. 4–15.

Liao, L., Patterson, D., Fox, D., Kautz, H., 2007. Learning and inferring transportation routines. Artificial Intelligence 171 (5), 311–331.

Linnarsson, R., 1993. Drug interactions in primary health care. a retrospective database study and its implications for the design of a computerized decision support system. Scandinavian Journal of Primary Health Care 11 (3), 181–186, cited By (since 1996) 33.

Logan, B., Healey, J., Philipose, M., Tapia, E., Intille, S., 2007. A long-term evaluation of sensing modalities for activity recognition. In: Proceedings of the 9th international conference on Ubiquitous computing. Springer-Verlag, pp. 483–500.

Lukowicz, P., Ward, J., Junker, H., Stäger, M., Tröster, G., Atrash, A., Starner, T., 2004. Recognizing workshop activity using body worn microphones and accelerometers. Pervasive Computing, 18–32.

Mahdaviani, M., Choudhury, T., 2007. Fast and scalable training of semi-supervised crfs with application to activity recognition. Advances in Neural Information Processing Systems 20, 977–984.

Mantyjarvi, J., Himberg, J., Seppanen, T., 2001. Recognizing human motion with multiple acceleration sensors. In: Systems, Man, and Cybernetics, 2001 IEEE International Conference on. Vol. 2. Ieee, pp. 747–752.

Maurer, U., Smailagic, A., Siewiorek, D., Deisher, M., 2006. Activity recognition and monitoring using multiple sensors on different body positions. In: Wearable and Implantable Body Sensor Networks, 2006. BSN 2006. International Workshop on. IEEE, pp. 4–pp.

Metcalf, C., Collie, S., Cranny, A., Hallett, G., James, C., Adams, J., Chappell, P., White, N., Burridge, J., 2009. Fabric-based strain sensors for measuring movement in wearable telemonitoring applications. In: IET Conference on Assisted Living. The Institution of Engineering and Technology, pp. 1–4.

Meystre, S. M., Savova, G. K., Kipper-Schuler, K. C., Hurdle, J. F., 2008. Extracting information from textual documents in the electronic health record: a review of recent research. Yearb Med Inform 35, 128–44.

Mihailidis, A., Carmichael, B., Boger, J., 2004. The use of computer vision in an intelligent environment to support aging-in-place, safety, and independence in the home. Information Technology in Biomedicine, IEEE Transactions on 8 (3), 238–247.

Milosevic, M., Shrove, M. T., Jovanov, E., 2011. Applications of smartphones for ubiquitous health monitoring and wellbeing management. Journal of Information Technology and Application (JITA) 1 (1), 7–14.

Modayil, J., Bai, T., Kautz, H., 2008. Improving the recognition of interleaved activities. In: Proceedings of the 10th international conference on Ubiquitous computing. ACM, pp. 40–43.

Moeslund, T., Hilton, A., Krüger, V., 2006. A survey of advances in vision-based human motion capture and analysis. Computer vision and image understanding 104 (2), 90–126.

Mooney, C. H., Roddick, J. F., 2013. Sequential pattern mining–approaches and algorithms. ACM Computing Surveys (CSUR) 45 (2), 19.

Nugent, C., Mulvenna, M., Hong, X., Devlin, S., 2009. Experiences in the development of a smart lab. International Journal of Biomedical Engineering and Technology 2 (4), 319–331.

Omichi, M., Maki, Y., Ohta, T., Sekita, Y., S., F., 1984. A decision support system for regional health care planning in a metropolitan area. Japan-hospitals : the journal of the Japan Hospital Association 3, 19–23, cited By (since 1996) 0.

Palmes, P., Pung, H., Gu, T., Xue, W., Chen, S., 2010. Object relevance weight pattern mining for activity recognition and segmentation. Pervasive and Mobile Computing 6 (1), 43–57.

Patterson, D., Fox, D., Kautz, H., Philipose, M., 2005. Fine-grained activity recognition by aggregating abstract object usage. In: Wearable Computers, 2005. Proceedings. Ninth IEEE International Symposium on. IEEE, pp. 44–51.

Patterson, D., Liao, L., Fox, D., Kautz, H., 2003. Inferring high-level behavior from low-level sensors. In: UbiComp 2003: Ubiquitous Computing. Springer, pp. 73–89.

Pearce, C. B., Gunn, S. R., Ahmed, A., Johnson, C. D., 2006. Machine learning can improve prediction of severity in acute pancreatitis using admission values of apache ii score and c-reactive protein. Pancreatology 6 (1), 123–131.

Pei, J., Han, J., Wang, W., 2007. Constraint-based sequential pattern mining: the pattern-growth methods. Journal of Intelligent Information Systems 28 (2), 133–160.

Perkowitz, M., Philipose, M., Fishkin, K., Patterson, D., 2004. Mining models of human activities from the web. In: Proceedings of the 13th international conference on World Wide Web. ACM, pp. 573–582.

Perwez, M., Ahmad, N., Javaid, M., Ehsan Ul Haq, M., 2012. A critical analysis on efficacy of clinical decision support systems in health care domain. Advanced Materials Research 383-390, 4043–4050, cited By (since 1996) 0.

Philipose, M., Fishkin, K., Perkowitz, M., Patterson, D., Fox, D., Kautz, H., Hahnel, D., 2004. Inferring activities from interactions with objects. Pervasive Computing, IEEE 3 (4), 50–57.

Phua, C., Sim, K., Biswas, J., 2011. Multiple people activity recognition using simple sensors. In: International Conference on Pervasive and Embedded Computing and Communication Systems. pp. 313–318.

Pollack, M. E., Brown, L., Colbry, D., McCarthy, C. E., Orosz, C., Peintner, B., Ramakrishnan, S., Tsamardinos, I., 2003. Autominder: an intelligent cognitive orthotic system for people with memory impairment. Robotics and Autonomous Systems 44 (3-4), 273–282.

Quinlan, J. R., 1986. Induction of decision trees. Machine learning 1 (1), 81–106.

Rashidi, P., Cook, D., 2010. Mining sensor streams for discovering human activity patterns over time. In: Data Mining (ICDM), 2010 IEEE 10th International Conference on. IEEE, pp. 431–440.

Rashidi, P., Cook, D., 2011a. Activity knowledge transfer in smart environments. Pervasive and Mobile Computing.

Rashidi, P., Cook, D., 2011b. Domain selection and adaptation in smart homes. Toward Useful Services for Elderly and People with Disabilities, 17–24.

Rashidi, P., Cook, D., Holder, L., Schmitter-Edgecombe, M., 2011a. Discovering activities to recognize and track in a smart environment. Knowledge and Data Engineering, IEEE Transactions on 23 (4), 527–539.

Rashidi, P., Cook, D., Holder, L., Schmitter-Edgecombe, M., 2011b. Discovering activities to recognize and track in a smart environment. Knowledge and Data Engineering, IEEE Transactions on (99), 1–1.

Rashidi, P., Cook, D. J., 2009. the resident in the loop: Adapting the smart home to the user. IEEE Transactions on Systems, Man, and Cybernetics journal, Part A 39 (5), 949–959.

Romano, M., Stafford, R., 2011. Electronic health records and clinical decision support systems: Impact on national ambulatory care quality. Archives of Internal Medicine 171 (10), 897–903, cited By (since 1996) 31.

Rosenblatt, F., 1958. The perceptron: a probabilistic model for information storage and organization in the brain. Psychological review 65 (6), 386.

Ruotsalainen, M., Ala-Kleemola, T., Visa, A., 2007. Gais: A method for detecting interleaved sequential patterns from imperfect data. In: Computational Intelligence and Data Mining, 2007. CIDM 2007. IEEE Symposium on. IEEE, pp. 530–534.

Safavian, S. R., Landgrebe, D., 1991. A survey of decision tree classifier methodology. Systems, Man and Cybernetics, IEEE Transactions on 21 (3), 660–674.

Scully, C. G., Lee, J., Meyer, J., Gorbach, A. M., Granquist-Fraser, D., Mendelson, Y., Chon, K. H., 2012. Physiological parameter monitoring from optical recordings with a mobile phone. Biomedical Engineering, IEEE Transactions on 59 (2), 303–306.

Shoeb, A. H., Guttag, J. V., 2010. Application of machine learning to epileptic seizure detection. In: Proceedings of the 27th International Conference on Machine Learning (ICML-10). pp. 975–982.

Simpson, R., Schreckenghost, D., LoPresti, E., Kirsch, N., 2006. Plans and planning in smart homes. In: Designing Smart Homes. Vol. 4008 of Lecture Notes in Computer Science. pp. 71–84.

Singla, G., Cook, D., Schmitter-Edgecombe, M., 2010a. Recognizing independent and joint activities among multiple residents in smart environments. Journal of Ambient Intelligence and Humanized Computing 1 (1), 57–63, cited By (since 1996) 14.

Singla, G., Cook, D., Schmitter-Edgecombe, M., 2010b. Recognizing independent and joint activities among multiple residents in smart environments. Journal of ambient intelligence and humanized computing 1 (1), 57–63.

Snyder-Halpern, R., 1999. Assessing health care setting readiness for point of care computerized clinical decision support system innovations. Outcomes management for nursing practice 3 (3), 118–127, cited By (since 1996) 11.

Song, C., Koren, T., Wang, P., Barabási, A., 2010. Modelling the scaling properties of human mobility. Nature Physics 6 (10), 818–823.

Srinivasan, R., Chen, C., Cook, D., 2010. Activity recognition using actigraph sensor. In: Proc. of the Fourth Int. Workshop on Knowledge Discovery form Sensor Data (ACM SensorKDD'10), Washington, DC, July. pp. 25–28.

Stikic, M., Schiele, B., 2009. Activity recognition from sparsely labeled data using multi-instance learning. Location and Context Awareness, 156–173.

Tapia, E., Intille, S., Larson, K., 2004. Activity recognition in the home using simple and ubiquitous sensors. Pervasive Computing, 158–175.

Tate, A., Dalton, J., Levine, J., 2000. O-plan: A web-based ai planning agent. In: Proceedings of the National Conference on Artificial Intelligence. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, pp. 1131–1132.

Tran, A., Marsland, S., Dietrich, J., Guesgen, H., Lyons, P., 2010. Use cases for abnormal behaviour detection in smart homes. Aging Friendly Technology for Health and Independence, 144–151.

Tukey, J., 1977. Exploratory data analysis. Reading, MA.

Vail, D., Veloso, M., Lafferty, J., 2007. Conditional random fields for activity recognition. In: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems. ACM, p. 235.

van Kasteren, T., Englebienne, G., Kröse, B., 2010a. An activity monitoring system for elderly care using generative and discriminative models. Personal and ubiquitous computing 14 (6), 489–498.

van Kasteren, T., Englebienne, G., Kröse, B., 2010b. Transferring knowledge of activity recognition across sensor networks. Pervasive Computing, 283–300.

van Kasteren, T., Englebienne, G., Kröse, B., 2011. Hierarchical activity recognition using automatically clustered actions. Ambient Intelligence, 82–91.

van Kasteren, T., Krose, B., 2007. Bayesian activity recognition in residence for elders. In: Intelligent Environments, 2007. IE 07. 3rd IET International Conference on. IET, pp. 209–212.

Wang, J., Luo, Y., Zhao, Y., Le, J., 2009a. A survey on privacy preserving data mining. In: Database Technology and Applications, 2009 First International Workshop on. Ieee, pp. 111–114.

Wang, J. T., Rozen, S., Shapiro, B. A., Shasha, D., Wang, Z., Yin, M., 1999. New techniques for dna sequence classification. Journal of Computational Biology 6 (2), 209–218.

Wang, L., Gu, T., Tao, X., Lu, J., 2009b. Sensor-based human activity recognition in a multi-user scenario. Ambient Intelligence, 78–87.

Weinland, D., Ronfard, R., Boyer, E., 2011. A survey of vision-based methods for action representation, segmentation and recognition. Computer Vision and Image Understanding 115 (2), 224–241.

Wernick, M., Yang, Y., Brankov, J., Yourganov, G., Strother, S., 2010. Machine learning in medical imaging. Signal Processing Magazine, IEEE 27 (4), 25–38.

Williams, C., Seeger, M., 2001. Using the nyström method to speed up kernel machines. In: Advances in Neural Information Processing Systems 13. Citeseer.

Wilson, D., Atkeson, C., 2005. Simultaneous tracking and activity recognition (star) using many anonymous, binary sensors. Pervasive computing, 329–334.

Wu, W., Au, L., Jordan, B., Stathopoulos, T., Batalin, M., Kaiser, W., Vahdatpour, A., Sarrafzadeh, M., Fang, M., Chodosh, J., 2008. The smartcane system: an assistive device for geriatrics. In: International conference on Body area networks. pp. 1–4.

Yang, A., Jafari, R., Sastry, S., Bajcsy, R., 2009. Distributed recognition of human actions using wearable motion sensor networks. Journal of Ambient Intelligence and Smart Environments 1 (2), 103–115.

Yin, J., Yang, Q., Pan, J., 2008. Sensor-based abnormal human-activity detection. Knowledge and Data Engineering, IEEE Transactions on 20 (8), 1082–1090.

Zhang, C., Zhang, S., 2002. Association rule mining: models and algorithms. Springer-Verlag.

Zhuang, X., Huang, J., Potamianos, G., Hasegawa-Johnson, M., 2009. Acoustic fall detection using gaussian mixture models and gmm supervectors. In: Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on. IEEE, pp. 69–72.