Engineering Applications of Artificial Intelligence 41 (2015) 232–248

Contents lists available at ScienceDirect

Engineering Applications of Artificial Intelligence

journal homepage: www.elsevier.com/locate/engappai

Design and validation of a light-weight reasoning system to support remote health monitoring applications

Aniello Minutolo *, Massimo Esposito, Giuseppe De Pietro

National Research Council of Italy, Institute for High Performance Computing and Networking (ICAR), Via P. Castellino 111, 80131 Naples, Italy

A R T I C L E I N F O

Article history:
Received 4 April 2014
Received in revised form
20 November 2014
Accepted 30 January 2015
Available online 21 March 2015

Keywords:
Decision support systems
Knowledge engineering
Inference algorithms
Patient monitoring

A B S T R A C T

Recently, mobile devices have dramatically improved their communications and processing capabilities, so enabling the possibility of embedding knowledge-based decision support components within Remote Health Monitoring (RHM) applications for the ubiquitous and seamless management of chronic patients. According to these considerations, this paper presents a light-weight, rule-based, reasoning system, purposely designed and optimized to build knowledge-based Decision Support Systems efficiently embeddable in mobile devices. The key issues of such a system are both a domain-independent reasoning algorithm and knowledge representation capabilities, specifically thought for both computation intensive and real-time RHM scenarios. The performance evaluation of the proposed system has been arranged according to the Taguchi's experimental design and performed directly on a mobile device in order to quantitatively assess its effectiveness in terms of memory usage and response time. Moreover, a case study has been arranged in order to evaluate the effectiveness of the proposed system within a real RHM application for monitoring cardiovascular diseases. The evaluation results show that the system offers an innovative and efficient tool to build mobile DSSs for healthcare applications where real-time performance or computation intensive demands have to be met.

1. Introduction

Recent developments in telecommunication and computer science along with the availability of new and cheaper technologies are radically impacting the management of chronic diseases, broadly exploded with longer lives, by enabling the successful provision of services for Remote Health Monitoring (in the following, RHM). RHM of patients affected by chronic illness has been expected to increase the individual's comfort, enhance quality of life, and encourage patient empowerment, while reducing the number of needless transfers at hospitals and the cost of provided healthcare services (Kumar et al., 2013; Lee et al., 2009; Osmani et al., 2013; Palaniswamy et al., 2013). RHM paradigm strongly relies on a new generation of medical devices designed to enable long-term, continuous, and unobstructed monitoring of physiologic information and, contextually, provide more realistic indication of the patient's health status, otherwise inaccessible in clinical settings (de Toledo et al., 2006; Rajan and Rajamony, 2013; Logan, 2013; Tamrat and Kachnowski, 2012; Martínez-Pérez et al., 2013).

However, gathering, processing and correlating data coming from these new medical sensors in order to infer information about the patient's condition represent undoubtedly a very costly and complex process.

In this respect, the real added-value could be supplied by intelligent components, in the form of Decision Support Systems (in the following, DSSs), able to automatically process and correlate huge volumes of monitored physiologic parameters, detect suspicious changes and supply alarms as a response to a worsening of the patient's status, plus suggestions about the actions to take. The potential for RHM applications integrating DSSs is well-documented (Martínez-Pérez et al., 2014) and can ease personalized diagnoses and treatments, predict patient status and follow-up based on multi-level observations, and empower patients to actively participate in their health.

So far, a good number of intelligent RHM applications has been proposed (Babu et al., 2013; Lee et al., 2007; Prasad et al., 2013; Won-Jae Yi et al., 2014; Xiaohui et al., 2012). Some of them are characterized by a centralized architecture that requires permanent communication between local sensing-data sources and remote decision support components, so implying sophisticated and reliable mechanisms for handling data losses or delays in the information flow. The remote decision support components are typically based on desktop-oriented reasoning systems equipped

* Corresponding author. Tel.: +39 081 6139515; fax: +39 081 6139531.
E-mail address: aniello.minutolo@na.icar.cnr.it (A. Minutolo).

http://dx.doi.org/10.1016/j.engappai.2015.01.019
0952-1976/

with knowledge management tools, aimed at automating the inference process of complex sensed data depending on a knowledge base (in the following, KB). Other existing RHM applications rely on distributed mobile agents able to process and exchange information in a reconfigurable platform so increasing scalability and reliability of provided services (Su, 2008; Kulkarni et al., 2014; Su and Chiang, 2014). Such intelligent RHM applications have been reviewed for being applied to care delivery and management in cases of chronic diseases (Paré et al., 2007). A comparison of these applications has been also proposed in (Pawar et al., 2012) with respect to their communication facilities, e.g. sensors and back-end servers. Moreover, a study about how the reduction of information transmitted over networks and presented to healthcare professionals actually improves the quality of mobile health applications is diffusely described in (Varshney, 2014). Finally, a review of the literature about existing mobile systems and applications dedicated to the medical decision support, as well as mobile apps currently available on the most used app stores has been proposed in (Martínez-Pérez et al. 2014).

In this respect, even if, a centralized and/or cloud infrastructure can limit the resource consumption to locally perform the intelligent data processing, by demanding reasoning functionalities to some remote nodes, it cannot be proficiently applied to many real situations. First, the necessity to communicate with a remote node for processing users' contextual and situational changes introduces an undesirable overhead for properly facing the rapid response time challenges of highly dynamic scenarios (Steller and Krishnaswamy, 2009). In particular, on the one hand, the overhead involved in relaying raw unprocessed data can be very high, since such data flow through the network and can eventually cause congestion. On the other hand, also the intelligent data processing overhead at the remote node can act as a bottleneck in the RHM application (Tai et al., 2013). Furthermore, users' data are extremely sensitive in medical settings, and the needs of transmitting them to a remote node could also raise privacy concerns for the users (Kleemann, 2006). Indeed, users could (rightly) be uncomfortable with sharing certain information outside of the mobile device (Woensel et al., 2014).

Second, this paradigm relies on the availability of a reliable network connectivity with a remote reasoning infrastructure. However, in some critical health scenarios, network connections established to reach a remote node could be not adequate, irregular or not feasible. For instance, in emergency medical care situations, the RHM application can make use only of a GSM network connection. This connection should prioritize the transmission of critical data, such as the real-time change in patient status, in order to promptly support the ambulance crew in the process of extrapolating or inferring what they need as a guide to the decision that reflects the reality of the given situation. In these scenarios, indeed, the role of ambulances and paramedics in the provision of emergency medical care is recognized as fundamental (Anantharaman and Swee Han, 2001), especially in patients experiencing an acute episode of their chronic illness, since urgent medical treatments can create the potential for better outcomes and lesser mortality and morbidity. Nevertheless, existing GSM networks provide only “best effort” service and do not explicitly provide for prioritized traffic, which is critical for emergency medical care situations.

As a consequence of that, in the case of network disconnections, RHM applications could completely fail without providing the requested decision support to the users.

In order to face such issues with continuous supporting services, the architectural paradigm used to build RHM applications has been sensibly changed to retain part (or all) of intelligent data processing over mobile devices, which, meanwhile, have dramatically improved their communications and processing capabilities. Embedding knowledge-based decision support components over mobile devices can allow locally performing an accurate and continuous analysis of the patient's health status, by minimizing network transmission, avoiding communication delays or interruptions and maintaining appropriate levels of security and privacy. As a consequence of that, decision support could be timely supplied, even in cases where connectivity is lacking, so lessening non- (or too late) raised alerts, which could negatively impact the patient's health (Woensel et al., 2014).

Anyhow, the requirements for realizing these knowledge-based decision support components are likely to exceed the simple storage and retrieval of information. Thus, the critical point, still pending to date, regards the definition of reasoning systems based on semantic technologies able to efficiently deal with as well as reason over sensed data directly on the mobile devices and identify the appropriate suggestions, reminders or alerts.

Two types of reasoning, namely *frequent reasoning* and *incremental reasoning,* can be adopted in the construction of knowledge-based decision support components for RHM applications (Woensel et al., 2014).

Frequent reasoning is typically adopted in RHM scenarios, hereafter referred as *computation intensive*, where significant quantities of varying data are first collected and, successively, at a certain time frequency, for instance, once a minute or even hourly or daily, they are processed jointly to infer a more complete indication about the patient's health status, implying a high computational complexity. For these RHM scenarios, reasoning systems operate periodically on the collections of sensed data and have to guarantee a response within acceptable, but not strict, time constraints. Thus, they can be dynamically deallocated and reinstantiated every time depending on their frequency of invocation in order to preserve the memory utilization.

Incremental reasoning is utilized in RHM scenarios, hereafter referred as *real-time*, where sensed data change very frequently (e.g. fresh readings are produced by sensors several times per second and assume different values from the previous ones) and are processed immediately to produce a real-time response. For these RHM scenarios, reasoning systems are continuously invoked every time new data are sensed and have to respond to data changes in a timely manner. Thus, they have to be kept in memory and dynamically reapplied in order to infer new suggestions, reminders or alerts within strict time constraints.

There are some semantic-enabled proposals supporting these two typologies of reasoning (Ali and Kiefer, 2009; Gu et al., 2007; Kim et al., 2010; Lorecarra, 2009; Jang and Sohn, 2004; Tai et al., 2011; Motik et al., 2012). Nonetheless, from a theoretical point of view, their performance on mobile devices and under a realistic RHM scenario is unlikely to achieve acceptable levels with reference to responsiveness and memory consumption, since strongly influenced by various factors, such as the reasoning algorithms used and the data scale. Indeed, most of them rely on reasoning algorithms that are too resource-intensive to be applied either to real-time and frequently varying RHM scenarios or to computation intensive RHM scenarios operating with considerable amounts of slowly changing data.

Moreover, it is not clear to what extent these existing solutions have been optimized for mobile devices to efficiently exploit their available resources, or simply ported from desktop-oriented implementations (Woensel et al., 2014). A study reported in (D'Aquin et al., 2010) has shown that semantic-enabled reasoning can take several tens or hundreds of KB of memory per elementary data. Hence, reasoning over a set of sensed data, even if not extremely large, but directly on a mobile device, can easily drain its memory resources, even on latest high-end models where the total memory size has reached 1 GB or more, because the quantity of memory available to the user applications is, however, limited.