Tricorder: Consumer Medical Device for Discovering Common Medical Conditions

Maja Somrak[1], Mitja Luštrek[1], Jakob Šušterič[2], Tomo Krivc[2], Ana Mlinar[2], Tilen Travnik[3], Luka Stepan[4], Mitja Mavsar[3] and Matjaž Gams[1]
[1] Jozef Stefan Institute, Jamova 39, Ljubljana, E-mail: maja.somrak@ijs.si
[2] MESI, razvoj medicinskih naprav, d.o.o., Leskoškova cesta 9d, Ljubljana, E-mail: ana.mlinar@mesi.si
[3] Domenca Labs, d.o.o., Celovška 32, Ljubljana, E-mail: tilen.travnik@dlabs.si
[4] Gigodesign, oblikovanje in komunikacije, d.o.o., Kamniška ulica 49, Ljubljana, luka.stepan@gigodesign.com

Technical Position Paper

Keywords: diagnostic device, Tricorder, medical application, intelligent assistant

Received: January 12, 2014

The Qualcomm Tricorder XPRIZE $10 million competition will open the doors for health and wireless technology. The aim is to design a portable, wireless device that monitors and diagnoses health conditions of residents without medical knowledge. The radical innovation in healthcare will give individuals far greater choices in when, where, and how they receive care. In this paper we present the competition itself and in particular the research prototype of the MESI Simplifying diagnostics team. The deadline of the first part of competition is approaching, after which the ten chosen teams will compete in the second, final round. Our approach builds upon the previous research and applications of several interdisciplinary partners constituting the team. The modular prototype fulfills the Tricorder demands by 24/7 monitoring the user and asking for medical help when needed. In the first stage it enables medical analyses of 15 prescribed medical conditions without demanding any particular medical knowledge. The additional novelties include: a modular structure enabling inclusion of any medical or ambiental sensor or device through web of things, additional previous medical applications not set out in the competition such as the detection of diseases manifesting in movement or activity recognition, and designing an intelligent medical assistant to take care of the user.

Povzetek: Opisano je tekmovanje Qualcomm Tricorder XPRIZE s skupnimi nagradami v višini 10 milijonov dolarjev in pristop slovenskih partnerjev.

1 Introduction

The Qualcomm Tricorder XPRIZE competition [1] is a global competition with the aim to stimulate the development of technologically advanced medical devices bringing accurate health diagnostics to the consumers. Currently, there are 30 international teams actively registered in the competition, which represents around one tenth of those that expressed initial interest. Teams will compete in terms of diagnostic accuracy and functionality, as well as user experience. The submission deadline for the Qualifying Round is May 2014, and from those that will submit the desired contributions, ten selected teams will advance in September to the Final Round scheduled to take place in the first half of 2015. Finally, up to three teams will be awarded a prize in total sum of $10 million.

The device envisioned for the competition will integrate innovative sensing hardware with advanced artificial-intelligence techniques. The convenient and portable design of the device will allow for anytime, anywhere, reliable health assessment, independent of medical professionals. The application of such devices could improve the utilization of healthcare resources by reducing unnecessary doctor's appointments, as well as improve personalized health care. This radical new approach will give individuals far greater choice in their own health-care by delivering health-care tools directly into their hands. A significant emphasis is put on user experience to ensure that the user will be able to use it correctly (no medical background needed) and that they will want to use it.

The target device is also based on a permanent and even continuously increasing demand for simple, easily accessible, and reliable diagnostic methods and systems. Currently, one of the most widely available tools for health assessment are online symptoms analysers, for instance the WebMD [2] or Mayo Clinic symptom checkers [3]. However, these are known to be unreliable, especially in the absence of additional diagnostics or healthcare professional consultation [4]. Recently, new innovative devices for remote diagnosis had been developed. One of such products for physical examination is the Tyto™ care, which combines multiple different technologies in their telemedical device [5]. Additionally, there is the need for continuous monitoring of vital signs, which is not provided by the aforementioned solutions. Various devices for monitoring vital signs are available in the market and range from wearable chest straps to more user friendly wristbands with integrated medical-grade sensors [6]. A novel, innovative solution should encompass both the continuous monitoring and reliable, advanced diagnostic methods in a single device. One of the most successful attempts in this direction represents the Scanadu Scout™ [7], one of the favourites in the Qualcomm Tricorder XPRIZE competition.

Moreover, there has been a persistent desire over several decades to design artificial intelligence (AI) assistants for various human tasks [8]. In recent years, humans use advanced assistants that might be perceived as somehow intelligent, on a regular basis: Google Now [9], question-answering systems like iOS Siri [10] or Android Assistant [11]. However, even though these systems are massively used, none of them exhibits true intelligence. True, they seem intelligent to naïve users and they are able to solve the tasks they are designed for reasonably well; however, they fail in a couple of sentences in particular when examined by those familiar with the Turing test. The idea of intelligent assistants emerges once again in the next generation of diagnostic devices – why not use these assistants and integrate them into a medical system? However, a viable, state-of-the-art solution that would successfully integrate assistants and other successful AI methods as quite diverse technologies in a single system, remains a challenge.

This paper describes our entry in the Tricorder competition, with the goals specified. Our approach currently mainly integrates machine learning techniques and multiple diagnostic hardware modules. It allows for a highly scalable solution with a possibility of subsequent extensions.

2 Our approach

Our approach is based on a combination of the concept of the Sci-Fi Tricorder – a multifunction hand-held device used for sensor scanning, data analysis, and recording data [8] – and the engineering and market approach, resulting in a prototype that will be able to perform well in the experimental tests. These two conditions seem to be orthogonal on each other and finding the compromise is the central issue of the competition.

The AI part of our approach integrates the following:

1. Design an AI-based medical assistant that will gather all information from the sensors on the user and in the environment and other sources of knowledge, with the goal to use this information for the benefit of the user.
2. The agent-based system will be highly modular and interdisciplinary.
3. The system will be able to communicate with the user through the graphical user interface (GUI) and in natural language.
4. The system will diagnose not only the diseases set out in the Tricorder competition, but also additional diseases and conditions, and if possible also predict probability of future complications.
5. The task is to infer short-, mid- and long-term conclusions about the user's medical situation, tuned to characteristics of each user.

(1) Concepts like the *Internet of things* [12][13] and *Sensor fusion* [14] enable the integration of several independent sources of data into meaningful information. Our approach combines the two concepts and also provides an integration at a novel level, as presented in [15] where each of the sensor was treated as a context and all the other sensors as input data from which the machine learning model of the domain was constructed. We call this approach *multiple-context sensor fusion*. For example, if the level of physical activity is taken as the context, and all the other sensor data are used for machine learning, we can reason as follows: in the context of low physical activity, a high heartbeat (and some other characteristics) indicates an alarming situation. Use of context is essential for quality performance in real-life circumstances and we expect the same will happen when it will be implemented in the second stage of the competition. Currently, we have implemented the algorithm for several tasks and expect no problem implementing it also in the Tricorder prototype. However, one should note that successful use of several AI modules mostly designed in previous projects with several tens of thousands lines of code each, far exceeds the capacities of current mobile devices. Therefore, services in the cloud are the only operational option for now.

(2) The ability to gather information from any devices that can be contacted through predefined standard communication protocols is already an indication of modularity and interdisciplinarity. The system needs to gather all available information from all sources and make use of it. Integrating data from devices included in the Tricorder, and learning from observation, may – for example – enable automatic learning that lower temperature in a room can result in a common cold for an individual user. If the temperature in a living room during regular monitoring significantly decreases for a substantial amount of time, a warning is thus issued that a low temperature previously caused a cold and that it is recommended that the temperature is increased. Currently, this functionality is not implemented yet.

In addition, the Tricorder system should be modular in a sense that if another device component is connected, it should be easily incorporated into the whole system. For example, connecting a heart-rate monitor or disconnecting it should not cause any errors in the system. The solution for this has been long known – agent systems enable the most flexible architecture. Our approach is based on JADE [16] designing a research prototype with the desired flexibility. We have previously designed MASDA system for analyses of soccer strategies [17] and designing cognitive and behavioural clones for teaching team commander how to deal with hostile crowds [18]. Currently, we are using agent architecture in smart houses and smart cities [19][20]. These systems already demonstrate large amount of elasticity, autonomy, interdisciplinarity and ability to deal with several heterogeneous sources of knowledge and integrating them into one functional system. To implement it in our Tricorder device, the current architecture will have to be upgraded in the second stage of the competition. However, several versions of the systems were already designed at a level of independent prototypes, including systems for activity recognition based on accelerometers of a mobile phone, in a bracelet, or separately attached sensors to various body parts. Some of these systems already used a sophisticated architecture [22]. Another of our systems [23] combines text and image marked and sent through a smart phone to estimate the probability of the Lyme disease. The current systems is already flexible in a way that adding or deleting one or more sensors is a rather simple task; however, no complex architecture or agent approach is currently implemented in our system.

(3) User communication can be performed through a classical text or highly visual interface. In addition, a natural language assistant is planned to be a part of our Tricorder. Two examples of our already implemented natural language assistants are Robi for the Jozef Stefan homepage [23] or Svizec for the union homepage [25]. These systems are implemented in a cloud and are available for major mobile platforms like Android or iOS. Adding a natural language interface into our Tricorder system would demand a couple of months of work. An example of practical use would be improved communication when the user does not understand the question or a message displayed on the screen. Similarly, the system is supposed to talk to the user not only through predefined speech sequences, but also from the dynamically generated text. We have designed such a platform for several man-machine projects [26].

(4) In previous experiments, diseases that manifest in the patient's motion, like the Parkinson's disease, were already detected with over 90% accuracy on simulated and real patients [27]. The methods used were based on integration of machine learning, dynamic time warping (DTW) and semantic attributes for each disease. It was demonstrated that each of these mechanisms improves the classification accuracy and that the most successful combination integrates them all. While only 5 diseases were tested so far, this research was an indication that several diseases can be successfully detected from movements only. In addition, the increase in these characteristic signs also enables the prediction of how fast the condition will deteriorate, thus enabling preventive actions. Several systems of this kind have been developed at the Jozef Stefan institute. It is feasible to implement these services in the cloud in several months each. The usefulness of these modules is clearly dependent on the success rate in practical tests – one of the future-work tasks. But noticing that a person limps for several hours or that the hands shake more than normally, or that a spasm is present in one arm for some time period are sufficient to call for help. The system will be able to communicate with the user or even bypass the user and call immediately for help in case of emergence or lack of user response – if set up in this way for example for an elderly living home alone.

(5) Several of our systems already deal with short-, mid- and long-term situations. They observe the performance of a particular user and learn his/her habits and performances [28][29]. For example, if a person already limps, then only differences to the common gait are looked for. Similarly, blood pressure of 140/90 can be an alarming news for one person and a good news for another. The team has designed several systems already that proved their performance, e.g. the Jozef Stefan Institute has won the live activity and fall detection EvAAL competition [30] and demonstrated another system at the European AI conference [31].

While the team has been designing all these subsystems for specific tasks, meaning that we have these prototypes developed and several of them in regular use, they are used for different tasks (see the relevant publications). Due to time constraints we were not able to introduce these functionalities into our existing prototype, but we estimate it is feasible to modify and integrate the already designed subsystems in half a year, at least at a prototype level with most of the computing performed in the cloud. Currently, we can only demonstrate each of the subsystem on its own.

However, an early version of a system focused on the practical part of the Tricorder competition has already been implemented and tested, and is presented in the next section. Partially, it already encompasses some of the AI sub-models presented in this section.

3 MESI Tricorder competition entry

In this section we describe our working prototype designed for prototype use at the Qualifying Round.

3.1 The system

Our Tricorder system was already presented at the CeBIT fair [32]. Its schema is presented in Figure 1. It consists of a bracelet for monitoring the vital signs, a mobile device with application for communication with the user and several applications including local computing methods and connection to the services in the cloud, and specialized modules for additional tests to determine specific diseases.

The bracelet measures vital signs designed for everyday use:

1. The ECG is the recording of the electrical activity of the heart. The signal is obtained by touching three electrodes, two of them with a wrist and the third with a finger of the other hand. Although only two electrodes are enough to measure single channel ECG, we added a third one, which is used for the noise reduction – similar to a right-leg drive (RLD) electrode in standard ECG devices. We improved and minimized the technology to fit our tight

housing and occupy only one square centimetre of printed circuit board (PCB).

2. The oxygen saturation (SpO2) is a measure of percentage of haemoglobin that has already bounded with oxygen. It is calculated by measuring the reflected red and infrared light that is sent to the fingertip. After post-processing the acquired data the device calculates respiratory rate.
3. The third sensor inside the bracelet is the temperature sensor. For best performance and quicker measurement, we use an infrared sensor.

For accurate and continuous measurements of vital sign user is required to use the so called *shield.* It consists of wireless cuff for measuring blood pressure using an automatic oscillometric method and a patch located on the ribs for measuring oxygen saturation, temperature, electrocardiogram (ECG), respiratory rate and activity tracking. Data obtained by vital signs and activity recognition help diagnose several diseases such as atrial fibrillation, hypertension and sleep apnea.

The second part of device consists of four in-depth modules for diagnosing 15 diseases:

1. The first module is so-called *To see*. A charge-coupled device (CCD) image camera with controlled standard and polarized white light is used for detecting melanoma and streptococcal pharyngitis. Special mount enables standard distance from the skin for obtaining real size of melanoma. For detecting otitis media wideband technology is used. High frequency sound impulses similar to Dirac impulse are transmitted to ear channel to move the eardrum. The reflected sound is recorded with high sensitivity digital microphone and processed with microcontroller. It calculates the fast Fourier transform (FFT) of the reflected signal and compares it with the pre-collected database.
2. The second module is *To hear*. Smart electronic stethoscope is designed for detecting pulmonary diseases. It has a microphone with high signal-to-noise ratio (SNR) attached to a specially designed bell, similar to the standard stethoscope chest piece, to record the low frequency sounds. We are developing software that is able to identify sounds and noises typical for each of pulmonary conditions. On the other side of the module, there is a second microphone for measuring air lung volume and speed or flow of inhaled and exhaled air. It is used to diagnose chronic obstructive pulmonary disease (COPD).
3. The *Urine module* analyses urine using test strips that are scanned by camera and automatically processed with computer vision. The mount on the camera enables standardization of light and distance for accurate results.
4. The fourth is the *Blood module*. To achieve the best possible user experience we are avoiding invasive methods and taking advantages of the spectroscopy technology. The blood module is able to diagnose anaemia and diabetes.

Figure 1. The CeBIT version of our Tricorder system consists of the bracelet, a mobile device with mobile applications and additional hardware devices for particular tests.

3.2 User experience

The proposed device aims to despecialize an aspect of primary health-care by giving end consumers the insight into their health on an instant basis. Through a user-centric design process we intend to understand the basic human needs and appropriate the technology at hand in such a way that it makes sense to a non-specialist. At the same time we want to engage the user more frequently, encouraging involvement and provided a detailed insight.

The system additionally tries to record vital signs in a greater extent that can be used as a detailed insight into the state before visiting a general practitioner (GP). Our aim was to solve as much of the diagnosing process only through a user-friendly questionnaire while the separate modules are intended for diagnosis confirmation. The questionnaire on its own already enables a more informed referral to a GP. Many specialist units outperformed a single multi-functional unit in our user testing for readability and ease-of-use. Modules are essentially multi-sensor units, packaged as “digital senses” for the end consumer. These units act together as an internet-of-things ecosystem and can be expanded as necessary in the future.

3.3 The diagnostic algorithm

We developed a novel algorithm for the initial assessment of the user’s medical condition, which could be either a healthy condition or one of the preselected 14 diseases. The algorithm combines two approaches: user interaction with a questionnaire and machine learning to make a tentative diagnosis.

The algorithm is implemented on a mobile device (see Figure 2) and focuses on the first of the two stages in the diagnostic procedure. Namely, during the first stage the user answers questions from a questionnaire to establish an initial diagnosis that is primarily informative. In the second stage, a special, additional device – a diagnostic module – is used for medical test and final diagnosis. The diagnosis can be additionally confirmed by an expert, if the user chooses to seek professional help in given circumstances. However, the questionnaire algorithm is one of the key components in this diagnostic process since it results in the first diagnosis to be later confirmed or rejected.

The algorithm consists of five steps:

1. Information such as the user's profile, vital-signs measurements and pain symptoms serve as initial inputs to the algorithm.
2. The initial inputs are used for the system to deduce a list of probable symptoms, from which the user selects main symptoms that he/she is experiencing.
3. The selected main symptoms, together with the initial inputs serve as the input to the algorithm for the medical-condition probability calculation. The medical conditions are shown in Table 1.
4. The user is asked for additional symptoms until the medical conditions emerge out of an uncertainty range that represents an undefined condition.
5. The medical condition with the highest probability is presented to the user as a tentative diagnosis, which should be further confirmed by a diagnostic hardware module.

The output of the first step of the algorithm is a list of probable symptoms, created by utilization of association rules and information-gain (IG) ranking. The input about the symptoms is used by the third and fourth step to provide relevant questions to the user. A new question is chosen according to the expected most informative symptom (symptom with highest IG). Each time the user answers a question, the probability for each medical condition is retrieved from the J48 classifier, dedicated to that condition. The classifiers were trained on data set of 15.000 virtual patients, generated using expert medical knowledge. The virtual patients were generated utilizing the table that relates symptoms to the diseases. For example, a patient with a single disease is generated with one or several symptoms for a randomly chosen disease. The probability of a patient having any of the medical conditions was uniformly distributed. Additionally, the probabilities of frequent combinations of the diseases had been set according to the medical experts. Only the last two percentages of all virtual patients were designed to have symptoms of a randomly selected combination of any two diseases.

The calculated probabilities of medical conditions may fall within the so called certainty or uncertainty range. The uncertainty range is a range of probability values where specific medical condition is neither very probable nor very improbable. The thresholds (low and high) define the uncertainty range.

The final output of the algorithm consists of the predicted medical condition, its probability and a selection of a further diagnostic module for the test. In this test, user is subdued to a single medical test, according to the initial diagnosis. A particular test is capable of recognizing any medical condition that belongs under the same diagnostic module (see Table 2) and is as such not necessarily limited to confirming only the initially predicted condition. This means that even an incorrect initial diagnosis and a correct diagnostic module selection can lead towards correct final diagnosis. Therefore, patients that are initially incorrectly diagnosed in the first stage, will have correct diagnosis in the second stage, if the diagnostic module will be able to induce the correct diagnose and will perform reliably.

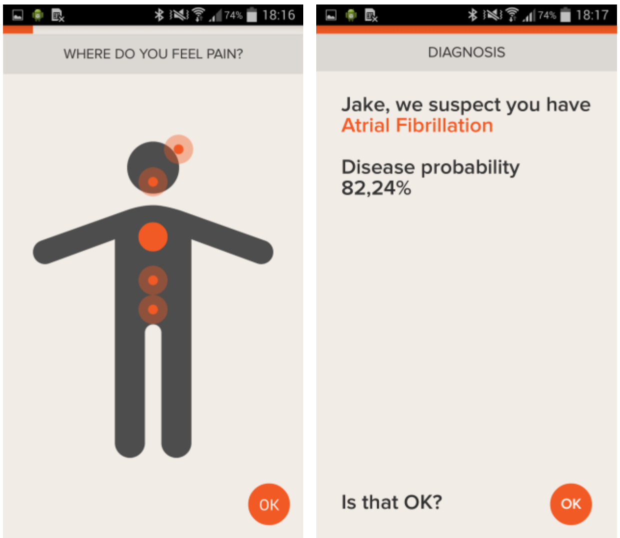

Figure 2. Examples of questionaries' menus and diagnostics.

4 Experiments

While learning was performed on 15.000 patients, testing was performed on a separate dataset of 1500 patients, where each medical condition was present in at least 100 patients. The test set was generated in a similar way as the original data set, in a way that the examples were different for each of the used sets.

Table 1. Initial prediction of the algorithm.

Initial medical condition prediction	Sensitivity	Specificity
Healthy	0.61	0.62
Hypertension	0.99	0.88
Atrial fibrillation	0.99	0.95
Acute haemorrhagic stroke	0.96	0.98
Obstructive sleep apnea	0.86	0.97
Hepatitis A	0.91	0.99
Otitis media	0.94	0.84
Streptococcal pharyngitis	0.95	0.97
Tuberculosis	0.96	0.99
COPD	0.99	0.95
Acute viral pneumonia	0.93	0.86
Lower urinary tract bacterial infection	0.99	0.99
Microcytic iron deficiency anaemia	0.83	0.95
Leucocytosis	0.59	0.6
Diabetes type 2	0.76	0.75

The results of initial medical condition prediction (initial diagnosis) are shown in Table 1.

Sensitivity is the probability that a person with a certain disease is correctly identified. Sensitivity is defined as a number of patients with the disease, who are correctly classified (true positives, TP), divided by the sum of both correctly classified and incorrectly classified patients with the disease (false negatives, FN). The formula for sensitivity is as follows:

$$\text{SENSITIVITY} = \text{TP} / (\text{TP} + \text{FN}).$$

Specificity is the probability that a person, identified as having the disease, is correctly identified. It is defined as a number of patients with the disease, who are correctly classified (TP), divided by the sum of both correctly classified patients with the disease and incorrectly classified patients without the disease (false positives, FP). The formula for specificity is as follows:

$$\text{SPECIFICITY} = \text{TP} / (\text{TP} + \text{FP}).$$

From Table 1 one can notice that healthy patients are least successfully identified. The reason for this is that both our diagnostics and the competition are primarily focused on correctly identifying the people in which a disease is present. However, on average, the obtained sensitivity of 0.884, specificity of 0.886 and accuracy of 0.883 seem quite acceptable for real-life trials.

Table 2. Corrected prediction of the algorithm.

Corrected medical condition prediction	Sensitivity	Specificity
Healthy	0.61	0.62
Hypertension	0.99	0.88
Atrial fibrillation	0.99	0.95
Acute haemorrhagic stroke	0.96	0.98
Obstructive sleep apnea	0.86	0.97
Hepatitis A	0.91	0.99
Camera module (*To see*)		
Otitis media	0.94	0.84
Streptococcal pharyngitis	0.95	0.97
Microphone module (*To hear*)		
Tuberculosis	0.96	0.99
COPD	0.99	0.99
Acute viral pneumonia	0.99	0.86
Urine module		
Lower urinary tract bacterial infection	0.99	0.99
Blood module		
Microcytic iron deficiency anaemia	0.92	0.95
Leucocytosis	0.63	0.77
Diabetes type 2	0.86	0.81

As mentioned in the previous chapter, there are some patient examples, which are incorrectly classified in the first stage, but their classification is corrected during the second stage of the diagnostic process, as a result of the correct module selection. The results that additionally account for these examples are presented in Table 2.

The average sensitivity is then 0.903, the average specificity 0.904 and average accuracy is 0.901.

The length of the list with probable symptoms was predefined to seven symptoms. From this list, on average two symptoms had been selected as present by the patient. The algorithm asked less than four additional questions, on average, to make the initial diagnosis.

The results suggest that such a questionnaire is both user friendly and efficient in terms of diagnostic accuracy.

5 Discussion

The Qualcomm Tricorder XPRIZE competition aims at revolutionizing home medical care through advances in hardware such as electronics and mobile devices, and software such as advanced applications and AI services. The last is related to the growing sense of optimism in the AI community. More and more established AI researchers believe that we are already in the transition period according to the "singularity theory" [33] According to Kurzweil [34] advances in electronics and artificial intelligence will enable the human civilization to jump ahead as use of metal enabled a jump from the stone- to metal age.

Our prototype already enables first analyses of vital signs and diagnoses of the selected 15 medical conditions. First tests were performed on virtual patients, since we were not able to perform the tests on live patients under the given circumstances. As a consequence, one might argue that the obtained results (sensitivity, specificity, etc.) were too optimistic. We agree that only clinical trials can represent efficiency in real life; however, the first experiments provide certain hope that this approach will be fruitful. One should bear in mind that according to many publications, internists achieve substantially lower performance in the first trial without specialized tests at other locations. Additionally, people without the necessary medical knowledge would benefit enormously if the device achieved comparable results than the one in our experiments.

Another cause why we have obtained such good results might be that number of the diagnoses was limited to 15 and each disease was very well characterized with its symptoms. Furthermore, each attribute was treated as correct without errors in judgment about the particular symptom. In real life, proper detecting of symptoms is a difficult task on its own. All these issues are to be addressed in future work.

In addition to correct tests with real-life patients and improvements of the existing device, another major issue is prevalent: to introduce the advanced AI and sustain high functionality in real-life circumstances we plan to modify and integrate the individual functional modules that partners had developed so far. That alone would enable design of a far more capable and intelligent system with agent structure, advanced multiple learning capabilities, sensor fusing, full internet of things, modular and interdisciplinary design, communication through both GUI and in natural language, user adaptability and adaptation. For the Final Round we intend to upgrade the system in that direction. The system is targeted to become a medical assistant making permanent observations and taking care of the user in a manner of intelligent agent assistants with certain degree of autonomy. The improved assistant observations would also play an important role in case the user wants analyses of his/hers medical conditions and assess probabilities of diseases and need of professional medical help.

In summary: as predicted by the organizers of the Qualcomm Tricorder XPRIZE competition, the device to revolutionize home care is on the brink of a break-through. We hope to contribute to these world-wide efforts.

References

[1] Qualcomm Tricorder XPRIZE, http://www.qualcommtricorderxprize.org/

[2] WebMD, http://symptoms.webmd.com/

[3] Mayo Clinic Symptom Checker, http://www.mayoclinic.org/symptom-checker/select-symptom/itt-20009075

[4] GUALTIERI, L.N. The doctor as the second opinion and the internet as the first Information. Technology Interfaces, 2009. ITI '09. Proceedings of the ITI 2009 31st International Conference on: June 22-25, 2009, Dubrovnik, Croatia.

[5] Tyto™ care, http://tytocare.com/

[6] W/Me medical-grade wristband, https://www.kickstarter.com/projects/723246920/finally-a-wearable-device-that-can-improve-your-li

[7] Scanadu Scout™, https://www.scanadu.com/

[8] Tricorder, http://en.wikipedia.org/wiki/Tricorder

[9] Google Now, http://bgr.com/2012/11/15/google-now-wins-popular-science-award/

[10] iOS Siri, http://www.apple.com/ios/siri/

[11] Android Assistant, https://play.google.com/store/apps/details?id=com.speaktoit.assistant

[12] IEEE Intelligent systems: Web of things, Vol. 28, No. 6, November/December 2013

[13] IEEE Computer: Internet of things, Vol. 46, No. 2, February 2013

[14] Information Fusion, An International Journal on Multi-Sensor, Multi-Source Information Fusion, Elsevier.

[15] Gjoreski, Hristijan, Kaluža, Boštjan, Gams, Matjaž, Milić, Radoje, Luštrek, Mitja. Ensembles of multiple sensors for human energy expenditure estimation. V: The 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, September 8-12, 2013, Zurich, Switzerland. UBICOMP 2013. Washington: Association for Computing Machinery = ACM, pp. 359-362.

[16] Bellifemine, Fabio Luigi, Caire, Giovanni Greenwood, Dominic. Developing Multi-Agent Systems with JADE, Wiley, 2007.

[17] Chong, Nak-Young, Mastrogiovanni, Fulvio. Handbook of Research on Ambient Intelligence and Smart Environments: Trends and Perspectives, IGI Global, 2011

[18] Tavčar, Gašper, Gams, Matjaž, Kvassay, Marcel, Laclavík, M., Hluchý, L., Schneider, B., Bracker, H. Graph-based analysis of data from human behaviour simulations. V: SAMI 2012: IEEE 10th Jubilee International Symposium on Applied Machine Intelligence and Informatics: proceedings: January 26-28, 2012, Herl'any, Slovakia. Piscataway: IEEE, cop. 2012, pp. 421-426.

[19] IEEE Internet Computing: Smart Cities, Vol. 17, No. 6, November/December 2013

[20] http://www.artemis-ia.eu/project/index/view?project=50

[21] OpUS, http://dis.ijs.si/Opus/

[22] Kozina, Simon, Gjoreski, Hristijan, Gams, Matjaž, Luštrek, Mitja. Efficient activity recognition and fall detection using accelerometers. V: BOTÍA, Juan A. (ur.). Evaluating AAL systems through competitive benchmarking: International Competitions and Final Workshop, EvAAL 2013, July and September 2013, [Madrid, Valencia]: proceedings, (Communications in Computer and Information Science, ISSN 1865-0929, 386). Heidelberg [etc.]: Springer, 2013, pp. 13-23.

[23] Čuk, Erik, Gams, Matjaž, Možek, Matej, Strle, Franc, Maraspin-Čarman, Vera, Tasič, Jurij F. Supervised visual system for recognition of erythema migrans, an early skin manifestation of lyme borreliosis. Strojniški vestnik, ISSN 0039-2480, Feb. 2014, vol. 60, no. 2, pp. 115-123, ilustr., doi: 10.5545/sv-jme.2013.1046.

[24] Robi, http://www.ijs.si/

[25] Svizec, http://www.sviz.si/svizec/

[26] Asistent, http://dis.ijs.si/projekt-asistent/?page_id=100

[27] Pogorelc, Bogdan, Gams, Matjaž. Detecting gait-related health problems of the elderly using multidimensional dynamic time warping approach with semantic attributes. Multimedia tools and applications, ISSN 1380-7501, 2013, vol. 66, no. 1, pp. 95-114, doi: 10.1007/s11042-013-1473-1.

[28] Kaluža, Boštjan, Cvetković, Božidara, Dovgan, Erik, Gjoreski, Hristijan, Mirchevska, Violeta, Gams, Matjaž, Luštrek, Mitja. A Multiagent care system to support independent living. International journal on artificial intelligence tools, ISSN 0218-2130, [in press] 2014, 10 pp., doi: 10.1142/S0218213014400016.

[29] Cvetković, Božidara, Kaluža, Boštjan, Luštrek, Mitja, Gams, Matjaž. Adapting Activity Recognition to a Person with Multi-Classifier Adaptive Training. Journal of Ambient Intelligence and Smart Environments, accepted for publication (2014).

[30] EvAAL competition, http://dis.ijs.si/?p=1320

[31] Luštrek, Mitja, Kaluža, Boštjan, Cvetković, Božidara, Dovgan, Erik, Gjoreski, Hristijan, Mirchevska, Violeta, Gams, Matjaž. Confidence:

ubiquitous care system to support independent living. V: RAEDT, Luc de (ur.). ECAI 2012: 20th European Conference on Artificial Intelligence, 27-31 August 2012, Montpellier, France including Prestigious Applications of Artificial Intelligence (PAIS-2012) systems demonstrations track, (Frontiers in artificial intelligence and applications, ISSN 0922-6389, v. 242). Amsterdam: IOS Press, cop. 2012, pp. 1013-1014.

[32] CeBIT, http://www.cebit.de/product/qualcomm-tricorder-xprize/470055/Q558791?source=pkl

[33] IJCAI 13, Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, Peking, China, ISBN 978-1-57735-633-2, (2013).

[34] Kurzweil, Ray. The Singularity is Near. New York: Viking Books, (2005).

[35] MESI team, http://simplifyingdiagnostics.com/xp/#team