



Human-in-the-loop Connectivity Management in Smartphones

David Nunes, Jorge Sá Silva, Carlos Herrera, Fernando Boavida

► To cite this version:

David Nunes, Jorge Sá Silva, Carlos Herrera, Fernando Boavida. Human-in-the-loop Connectivity Management in Smartphones. 14th International Conference on Wired/Wireless Internet Communication (WWIC), May 2016, Thessaloniki, Greece. pp.159-170, 10.1007/978-3-319-33936-8_13 . hal-01434875

HAL Id: hal-01434875

<https://inria.hal.science/hal-01434875>

Submitted on 13 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Human-in-the-loop Connectivity Management in Smartphones

David Nunes ^(✉)¹, Jorge Sá Silva¹, Carlos Herrera², and Fernando Boavida¹

¹ University of Coimbra, Department of Informatics Engineering, Coimbra, Portugal,
[dsnunes, sasilva, boavida]@dei.uc.pt

² Escuela Politécnica Nacional, Departamento de Electrónica y Telecomunicaciones,
Quito, Ecuador,
carlos.herrera@epn.edu.ec

Abstract. Cyber-Physical Systems which detect and use human psychological and physiological states as feedback to their control-loop are known as Human-in-the-Loop (HITL). The current understanding of the general requirements and theory behind this new class of systems is, however, still limited. This paper attempts to contribute towards their development by presenting a conceptual model that integrates some of the major ideas spread throughout the literature. In addition, it also discusses how the acquisition of a human's state requires reliable networking. In this respect, an architecture for HITL management of networking configurations is presented. This architecture is implemented in an experimental scenario where a smartphone system attempts to positively impact mood by dynamically adapting privacy settings and networking interfaces to the user's emotions. In this scenario, good hand-off performance is crucial for improving Quality of Experience and reliability. Several experiments regarding emotional classification and mobile hand-off are also presented and discussed.

Keywords: Emotion Detection, Human-in-the-Loop Cyber-Physical Systems, Network Management, Smartphones

1 Introduction

Recent years have brought a tremendous evolution in the areas of the Internet of Things (IoT), robotics and wireless sensor networks. It is now possible to develop intelligent Cyber-Physical Systems (CPSs) that easily sense and control our environment. Nevertheless, these technologies are still widely unaware of the human context, which is often considered an external and unpredictable element. Current research indicates that future CPSs will likely strive to become more "human-aware". This is the defining characteristic of "Human-in-the-loop" Cyber-Physical Systems (HiTLCPSs), where the human's emotions and actions are taken into consideration.

With this paper, our main goals are to contribute towards a logical and concise organization of HiTLCPS concepts and to provide an innovative HiTLCPS

case-study, which achieves emotion-aware smartphone connectivity management, such as the selection of interfaces and privacy levels. In short, this paper’s main technical contributions are as follows:

1. An overview of the area of HiTLCPSs, through a general model that organizes its major ideas.
2. A HiTL architecture directed towards smartphone connectivity management.
3. An implementation of our model targeting a mobile application for the improvement of user mood.
4. The results of several experiments that evaluate our system’s emotion classification and handoff performance, which have a direct impact on the Quality of Experience (QoE).

Our mobile application, HappyHour, is novel in the way that it is capable of “closing the loop”, that is, using an emotion inference result to actuate and provide suggestions that attempt to improve the user’s mood. Although there are previous works that attempt to perform emotion inference from smartphones, these are based on analyzing communication history and application usage patterns, instead of environmental and physical information, with the purpose of improving recommendation systems [1] or are primed for data collection, rather than feedback and user interaction [2]. HappyHour also innovates in the way that it is capable of adapting its and the smartphone’s functionality to the user’s condition, through the dynamic change of privacy settings and networking interfaces. As far as we know, this is the first attempt at using HiTL concepts to optimize networking aspects.

The rest of this paper is organized as follows: section II presents a general model for HiTLCPSs; section III presents a HiTL model for the management of smartphones connectivity aspects; section IV presents an implementation of this model targeting a mobile application; section V presents an evaluation of our implementation and its handover performance; section VI summarises the presented ideas and possible future work.

2 Human-in-the-loop Cyber-Physical Systems

Previous research in HiTLCPSs has proposed workstations that detect human distractions to save energy, systems that enable spatial control of the mobile wireless spectrum, semiautonomous wheelchairs controlled through electroencephalography (EEG), wheelchair-mounted robotic arms for disabled people, HiTL drug-delivery pumps and human-aware HVAC systems [3]. Each of these HiTLCPSs provides a tangible example on how the human context can be useful for control-loop decisions [4].

Our previous research [3] has allowed us to gain some insight onto HiTLCPSs. Thus, we present our own view of the HiTLCPS landscape through a conceptual model, shown in figure 1. This model represents, as far as we know, the first attempt at condensing the major processes of HiTLCPS control. Each HiTLCPS requires a **Human-in-the-Loop Intelligence** module, responsible for receiving

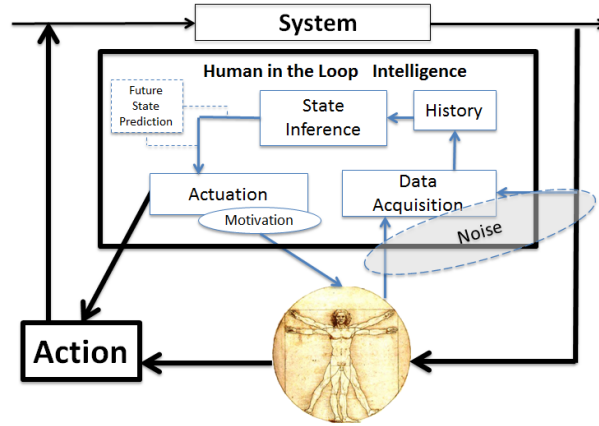


Fig. 1. The Processes of Human-in-the-Loop Control

input from the human sensors and for influencing the control-loop. On a first step, determining a human's state requires the **acquisition of data**, through the use of, for example, IoT devices that act as sensors and acquire data from both the human and the system.

History is also important for many HiTLCPSs, since humans are creatures of habit and previous data often offers insights that allow more accurate prediction of human context.

State inference techniques, such as advanced mathematical models or machine learning techniques, are usually applied to accurately detect human intents, psychological states and actions.

Actuation in HiTLCPSs can be classified into two major classes. Firstly, there is a direct actuation in the control-loop that results from the feedback provided by the system's current status and the inference of human state. The second class of actuation is related to the actions of humans, since they are far from passive elements and can, themselves, perform tasks.

Noise is a factor that affects the entire system. It is particularly important in data acquisition: for example, HiTLCPSs based on speech or video-captured gestures have to deal with ambient noise and moving background clutter; another example is the acquisition of vital signs, which is prone to interference from physiological functions that have little to do with what needs to be acquired.

In addition, each and every of these processes needs to be **reliable**. The inability to consistently infer a human's state in an accurate manner can have severe consequences on control-loop decisions and compromise the entire system. Additionally, **security and privacy** are equally important to protect industrial processes, medical data or sensitive personal information from external unauthorized exploitation. Networking is a crucial aspect to achieve this reliability and security, since HiTLCPSs are often large and distributed, with multiple sensing

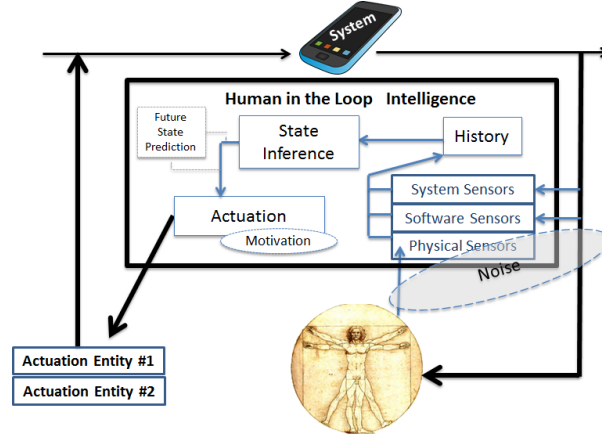


Fig. 2. HiTLON within the processes of Human-in-the-Loop Control

devices transmitting data between each other and to centralized remote system controllers.

3 Towards Human-in-the-loop management of smartphone connectivity

As far as we know, no solutions consider the human context in their control-loop decisions. We believe that this is limiting, in the sense that smartphones are personal devices that cater for their user preferences. In fact, in the context of HiTLCPSs, the humans' context must be considered, but human attention should not be required. Therefore, automated HiTL management of aspects such as privacy and network interface handoff can contribute to a more efficient networking distribution and system usability.

We propose a HiTL architecture directed towards intelligent network management of interfaces and privacy. We refer to it as "Human-in-the-loop over networks" (HiTLON), of which we are currently developing our own implementation. This HiTLON architecture fits directly in our general model, as shown in Figure 2. Our previous "System" box is now represented by a smartphone. The "data acquisition" is represented by physical, software and resource sensors. These "sensors" may have very disparate instantiations; for example, **Software sensors** are usually handled at the application layer and measure non-physical properties related with applications, such as QoS requirements, communication history of the user (e.g. number of SMSs, phone call durations), social networking data, among others. The **physical sensors** represent every sensor that is made available by the smartphone's hardware (e.g. accelerometer, microphone, camera, GPS, etc.) or even by physically wearable devices (e.g. smartshirts, smartwatches) that are not an integral part of, but can communicate with, the

smartphone. **System sensors** may implement calls to the system’s Kernel to return current battery level, CPU usage, wireless connection strength or available memory.

Encompassed by the ”Human-in-the-loop” control is the state inference. Sensory information is processed here, acquiring context from the raw data. For example, a state-inference task may be responsible for performing activity classification based on accelerometer data (e.g. standing, walking or running states).

The actuation process (decision-making), is then realized depending on the state-inference results. The decisions are carried out by **actuation entities**. Some examples of these entities might be handoff interfaces, security and encryption mechanisms, privacy managers, among others. This way, despite focusing on privacy and interface management, we intend to present our HiTLON architecture as a general approach for HiTL management of connection characteristics. In this particular scenario, direct human actuation does not apply, hence, the associated connections are not represented. Nevertheless, human motivation is represented by the high-level policies, which materialize the human intent.

The HiTLON concept can be applied to many different scenarios. For example, a health monitoring system may monitor the ECG of a patient to detect arrhythmias, sending this data to a remote doctor’s PC. Whenever critical data is detected, the reliability of the network connection could be increased, even at the expense of greater bandwidth or energy requirements which would be lower during normal circumstances.

Another example is a toddler monitoring application where sound or image processing is used to detect patterns of distress. When critical events are triggered, a HiTLON may prioritize connectivity and performance in the parents smartphone to receive a direct video-feed of the child.

Our model also accommodates other types of human-aware management, such as giving higher levels of QoS to messages exchanged with important people (e.g. delivery reports to close friends); sending sensitive and urgent data preferably through encrypted cellular connections to maximize connectivity and security; opportunistically using Bluetooth to save cellular traffic for non-critical and delay-tolerant data; and using multiple interfaces to maximize throughput if the smartphone is currently connected to an energy source.

HiTL control balances the decision between the three main pillars: the application’s QoS requirements, the device’s current status and the human context, including defined high-level goals and other restrictions (e.g. avoid cellular connection fees).

4 HiTL-over-networks in mobile BCI

In this section, we will explore a particular application of our HiTLON model targeting a mobile Behavior Change Intervention (BCI) scenario. Traditional BCIs involve therapy sessions where advice and support are provided to induce lifestyle changes that may help people coping with chronic diseases, smoking

addiction, diets or even depression. More recently, due to their sensing and processing capabilities smartphones have been used by behavior scientists for more directed interventions [2]. A HiTLON system can play a role in reducing anxiety; when a stressful event is detected, the system may attempt to guarantee the best networking performance in order to avoid further frustrating events. If battery levels are low, HiTLON could choose to preserve battery life instead, so that interventions can be delivered timely. To evaluate these ideas, we developed a mood-oriented BCI case-study over which we developed an implementation of a HiTLON system, based the architecture previously presented in section 3.

4.1 HappyHour App

Common knowledge and scientific research agree that moderate walking exercise and contact with natural environments can provide several cognitive benefits, such as improved memory, attention and mood [5, 6]. Other studies suggest that contact with natural environments not only makes people feel better but also makes them behave better, thus having both personal health benefits and broader social benefits [7]. HappyHour is a HiTLCPS based on this premise that takes a BCI approach to positively affect mood [8]. It unobtrusively senses emotions and presents timely walking suggestions when negative moods are inferred.

The system periodically processes data from the smartphone’s microphone and accelerometer, a smartshirt’s ECG³ and weather information from a web API⁴, feeding it to a neural network⁵ in order to infer the user’s emotional state. When negative emotional states are detected, the application motivates the user to go for a walk, showing a map with the position of nearby points-of-interest (POI).

The system also periodically collects its users’ GPS positioning, microphone sound and acceleration in an anonymous way and aggregates this information in a central server. This allows HappyHour to display the real-time attendance (based on user location) and overall movement (based on average accelerometer data) at each POI through heatmaps with different colors. This information allows users to pick livelier areas (greater attendance and movement) or calmer areas (less attendance and movement), which are more prone to soothing environments. The users’ microphone data is processed through a music recognition API⁶, allowing the application to display the background music at each POI. In a similar way, HappyHour is also capable of displaying the ”general mood” of a certain POI, by aggregating and averaging the individual moods of nearby users. Thus, users are able to select places to visit with the type of environment, mood and music that they feel are best for venting emotional stress.

Figure 3 shows HappyHour’s application flow. The core of our emotion-awareness lies in HappyHour’s machine learning ability to process different forms

³ VitalJacket: <http://www.vitaljacket.com/>

⁴ Open Weather Map: <http://openweathermap.org/>

⁵ Encog Framework: <http://www.heatonresearch.com/encogin>

⁶ The Echo Nest: <http://the.echonest.com/>

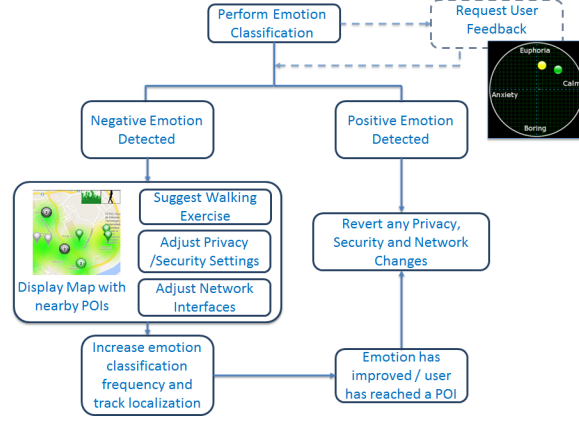


Fig. 3. HappyHour Application Flow

of sensing. With this work, we do not intend to propose robust methods for emotion detection, but instead, to provide a practical proof-of-concept that shows how emotional information can benefit HiTLCPSs. Thus, in order to determine the best machine learning technique for our application, we studied previous comparisons between the different possibilities. Previous work [9] has scored different classification algorithms in terms of correct classification rate and in terms of central processing unit (CPU) time needed for the classification. The latter is of particular importance for smartphone HiTLCPSs, since these are limited terms of available processing power and energy. This has led us to opt for an artificial neural network as our emotion inference tool, since it offers a reasonable correct classification rate while being one of the least time-consuming techniques.

As previously mentioned, HappyHour’s neural-network is fed with different types of data. The relationship between external factors and a person’s emotions is still unclear [10]. Nevertheless, we intended to consider at least three general sources of data: environmental clues, vital-signs information and meteorological information [8].

HappyHour also manages network interfaces depending on the human’s mood. By default, WiFi is used as the primary networking interface for positive emotions, while cellular communication is used as a backup due to its additional monetary cost. However, if the system detects a negative emotion state, the default network profile defines that both cellular connections and WiFi will be used at the same time, thus providing better QoE. We employ MPTCP [11] in our current implementation to achieve these multipath capabilities over a single TCP connection. Privacy management in HappyHour relates to the automatic sharing of location on a social network⁷. Since some people feel eased by social-interaction, the automatic sharing of location among friends may contribute towards positive socialization. Other people may prefer solitude to ease their

⁷ In the current prototype, the user can share his location through Facebook

minds. Thus, through the tailoring of privacy profiles, HappyHour can automatically adapt privacy settings to the current emotional context of the user.

While our current implementation is limited to changes in privacy and network interfaces, we can easily conceive more advanced HiTL mechanisms to control other aspects. For example, the redundancy tradeoff of multiple interfaces could be moderated by the device’s battery-level, since completely depleting the battery can result in even greater user frustration. Security measures could also be dynamically adapted depending on the human context, with additional levels of encryption being applied when his/her emotions, position or actions are of sensitive nature.

4.2 Emotion-Aware MPTCP

We believe that it is important to bring this human-awareness to the networking layer. In the particular case of MPTCP, due to the extreme importance of networking in HiTLCPS, HiTL functional parameters should be a part of the protocol’s interface with the application layer. These HiTL functional parameters could affect several aspects of the protocol.

MPTCP currently couples the congestion windows of each TCP subflow to achieve fairness at bottlenecks and resource pooling and to push most traffic to uncongested links. We propose that a different congestion controller for MPTCP, could aim at achieving different objectives for quality of service, reliability, and resilience. This congestion controlled resource of pooling/fairness/stability could, instead of focusing exclusively on maximizing throughput through all available paths, consider the possibility of choosing alternative paths base on the human’s needs. This choice could be motivated on the monetary cost of links, but could also consider other aspects such as connectivity, range, delay and jitter, since stability can be more important than throughput in HiTLCPSs.

Making effective choices based on the human context requires the network layer to have knowledge of the path ”cost” for the human user. This information would be provided by the application layer through an extension of MPTCP’s application interface. The application layer can configure send and receive buffer sizes via the sockets API (SO_SNDBUF, SO_RCVBUF). However, it would also be possible for a MPTCP implementation to set a human contextual value for dynamically modifying the send and receive buffers and treating this request as an implicit task of the networking layer. Other possibility is to restore TCP’s ability to send ”Urgent” data, which is not currently in use by MPTCP [12].

5 Experiments, Simulations and Evaluation

In this section, we will present some experiments and simulations which serve as an evaluation of our HiTL mobile BCI. Firstly, we will focus on the performance and accuracy of our emotion classification. Afterwards, we will discuss the handoffs in our HiTLON implementation.

Table 1. Testing training performance (150 emotions).

Hidden Layers	Epochs
One	100
Two	3000

Table 2. Testing neural network accuracy (41 emotions).

Hidden Layers	Sensitivity	Specificity
One	0,679	0,766
Two	0,720	0,830

5.1 Emotion classification and neural-networks

To test HappyHour’s emotion classification, we considered two major requirements: the amount of effort required for training the neural network (which is important in terms of processing power and battery drain) and accuracy of the network. In order to test the training effort, we began by generating simulated emotions. For each type of emotion, we empirically defined a probability value for different ranges of its input components (heart rate, cloudiness, movement, etc). This method allowed us to generate 150 simulated emotions, that, while not valid for testing accuracy (since they are not derived from actual human beings), are sufficient for testing the training performance. Thus, we counted the number of epochs necessary to successfully train the network for each configuration. The results, shown in table 1, show that using two hidden layers increases the training effort significantly. Therefore, we also needed to test if using more layers brought any benefits in terms of accuracy. Note that in a typical Artificial Neural Network (ANN) architecture, its neurons are usually grouped in layers; the first layer receives the input, while the last layer transmits the final output. In-between these layers are the hidden layers which allow the ANN to extract higher-order statistics from the data, by providing additional transformations and processing.

In order to assess the impact of the number of layers, we requested a test subject to use our application for a period of a week, during which his sensory data and emotional feedback were recorded for a total of 41 records. We then tested both neural network configurations using this data. Considering that negative emotions are the events of interest, we evaluated performance through two statistical measures known as sensitivity and specificity. In our case sensitivity is the proportion of negative emotions that were correctly identified as such; that is, it measures when our system was capable of detecting that it was necessary to adapt the smartphone’s settings. Specificity, on the other hand, measures the proportion of correctly identified positive emotions. A perfect emotional predictor would present the maximum value of 1 for each of these metrics.

The results shown in table 2 suggest that using a two layer configuration presents considerably better results. After pondering over the results, we decided that, despite being more demanding, a two layer configuration presented a better compromise in terms of training time and accuracy.

5.2 Evaluation of HiTLON handoff performance

Our mood-oriented HiTL implementation would not make sense if the handoffs between different network interfaces introduced even greater connectivity issues. The objective of these experiments was to determine if MPTCP could offer acceptable handoff performance while maintaining the connection flow. This served as preliminary work before trying to further modify the protocol in order to introduce human-awareness components in its congestion control mechanisms and application interface. While our current HappyHour implementation runs on a LG Nexus 5 smartphone with a MPTCP-enabled Android kernel, we performed our experiments on a laptop. This was done to achieve a finer control over the testing conditions. The laptop ran an MPTCP-enabled Linux kernel and was equipped with a WiFi antenna and a 3G USB Dongle. It used these interfaces to connect to a remote host and continuously sent ECG and accelerometer data. The connection was made through a WiFi network supported by a cable Internet backhaul, offering a bandwidth of around 1 Mbps upstream and 18 Mbps downstream, and a 3G network, offering a bandwidth between 0.2 and 0.4 Mbps. To




Fig. 4. WiFi being cut off




Fig. 5. 3G being cut off

evaluate our default negative emotion scenario, we used both interfaces simultaneously. One interface was disabled mid-test, in order to emulate a disconnection event. Figures 4 and 5 show the throughput of the TCP connection on the remote host. As we can see, a drop on the general throughput is noticeable, but the dropped packets are quickly resent and the connection effortlessly continued through the remaining interface. A second series of tests was performed, where one interface was configured as "backup". Figure 6 shows a typical positive emotional state situation, where WiFi is defined as the primary interface and a 3G cellular connection acts as a backup. As the user goes on with his walking exercise, he eventually loses connection to WiFi hotspots, which have limited range. MPTCP is then responsible for smoothly rerouting traffic through the 3G cellular connection. As the graphic shows, in our experiments the connection to the server took about 4 seconds to recover.

Figure 7 shows another situation where a user has defined a 3G cellular connection as the primary interface in HappyHour's network profile. This would be a configuration used whenever it is important to promote connectivity and avoid




Fig. 6. WiFi as primary interface and 3G acting as backup, with WiFi being cut off




Fig. 7. 3G as primary interface and WiFi acting as backup, with WiFi being cut off

handoffs unless they are absolutely necessary; e.g. the user could be riding a bus or a train, where the connections to WiFi networks can be very transient. WiFi, now being relegated to "backup", recovered the connection in approximately 3 seconds, which is faster than 3G. This was to be expected, since WiFi has a much greater downlink capacity and its TCP congestion window increases much faster.

6 Conclusion

In this work, we presented an overview on topic of HiTLCPS through a theoretical model of the processes associated with the control-loop. We then approached the problem of managing privacy and networking interfaces in smartphone-based HiTLCPSs. To do so, we devised an architecture for this connectivity management problem. We then applied this architecture to a real-world scenario based on a HiTLCPS for the improvement of mood. For evaluation, we performed some preliminary experiments which indicate that neural-networks might be an adequate machine-learning technique for these scenarios.

During our exposition, we also argued how reliability is an important factor in HiTLCPSs. Thus, another part of our experimental results focused on hand-off performance, since intermittent connectivity is a source of frustration with negative impact on the user QoE. The results showed acceptable handoff performance and a definite improvement over "brute-force" handoffs. Thus, MPTCP can be a promising protocol for supporting seamless handoff within our HiTLON architecture.

As future work, we will continue to develop our BCI application and implement some of the concepts that were considered in our model but that have been left out of our test scenario. In particular, we would like to focus on the management of the device's power-consumption based on the human's context.

Acknowledgments

The work presented in this paper was partially financed by Fundação para a Ciência e a Tecnologia, POPH/FSE and by SENESCYT - Secretaría Nacional de Educación Superior, Ciencia, Tecnología e Innovación de Ecuador.

References

1. LiKamWa, R., Liu, Y., Lane, N.D., Zhong, L.: Moodscope: Building a mood sensor from smartphone usage patterns. In: Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services. pp. 389–402. MobiSys '13, ACM, New York, NY, USA (2013)
2. Lathia, N., Pejovic, V., Rachuri, K.K., Mascolo, C., Musolesi, M., Rentfrow, P.J.: Smartphones for large-scale behavior change interventions. *IEEE Pervasive Computing* 12(3), 66–73 (2013)
3. Nunes, D., Zhang, P., Silva, J.S.: A survey on human-in-the-loop applications towards an internet of all. *IEEE Communications Surveys & Tutorials* 17(2), 944–965 (2015)
4. Liang, C.J.M., Stankovic, J., Lin, S.: Reducing energy waste for computers by human-in-the-loop control. *IEEE Transactions on Emerging Topics in Computing* 99(PrePrints), 1 (2013)
5. Berman, M.G., Kross, E., Krpan, K.M., Askren, M.K., Burson, A., Deldin, P.J., Kaplan, S., Sherdell, L., Gotlib, I.H., Jonides, J.: Interacting with nature improves cognition and affect for individuals with depression. *Journal of affective disorders* 140(3), 300–305 (2012)
6. Berman, M.G., Jonides, J., Kaplan, S.: The cognitive benefits of interacting with nature. *Psychological Science* 19(12), 1207–1212 (2008)
7. Weinstein, N., Przybylski, A.K., Ryan, R.M.: Can nature make us more caring? effects of immersion in nature on intrinsic aspirations and generosity. *Personality and Social Psychology Bulletin* 35(10), 1315–1329 (2009)
8. Carmona, P., Nunes, D., Raposo, D., Silva, D., Herrera, C., Silva, J.S.: Happy hour - improving mood with an emotionally aware application. In: 15th International Conference on Innovations for Community Services (I4CS). n/a (2015)
9. Guinness, R.: Beyond where to how: A machine learning approach for sensing mobility contexts using smartphone sensors. In: Proceedings of the 2013 International Technical Meeting of The Institute of Navigation (2013)
10. Gunes, H., Pantic, M.: Automatic, dimensional and continuous emotion recognition. *International Journal of Synthetic Emotions* 1(1), 68–99 (January 2010)
11. Ford, A., Raiciu, C., Handley, M., Barré, S., J.Iyengar: Architectural guidelines for multipath tcp development. Internet Engineering Task Force (IETF) , Request for Comments 6182 (March 2011)
12. Scharf, M., Ford, A.: Multipath tcp (mptcp) application interface considerations. Internet Engineering Task Force (IETF) , Request for Comments 6897 (March 2013)