Journal of Universal Computer Science, vol. 26, no. 1 (2020), 4-32
submitted: 30/12/18, accepted: 15/11/19, appeared: 28/1/20

Pompilos, a Model for Augmenting Health Assistant Applications with Social Media Content

Henrique Damasceno Vianna
(University of Vale do Rio do Sinos, São Leopoldo - Brazil
hdvianna@gmail.com)

Jorge Luis Victória Barbosa
(University of Vale do Rio do Sinos, São Leopoldo - Brazil
jbarbosa@unisinos.br)

Abstract: Caused by habits such as poor diets, lack of physical activity practice or smoking, non-communicable diseases were elected by the World Health Organization as one of the greatest challenges of the twenty-first century, despite a lot of information produced in social media focused on preventing this type of disease. This paper presents the Pompilos Model, which aims at improving computer-aided social support by suggesting beneficial health resources and revealing what influences other people's health, so to foster better health behaviors in social relations. In order to evaluate the model's feasibility, we performed a random experiment during one month and half with two groups to assess the influence of messages related to the prevention of chronic diseases. Those messages presented information on a healthier diet, the practice of physical activities, and ways to lose weight, from monitored Twitter profiles on the habits of health assistant web application's users. So it would be possible to manage food intake, the practice of physical activities, and weight control. Messages related to the prevention of chronic diseases, such as a healthier diet, the practice of physical activities, and weight loss from monitored Twitter profiles were directed to an intervention group as a way to re-engage users in their care activities. With this information, we found a correlation between message reading and the access to the application history feature among intervention users.
Key Words: Architecture for Distributed Systems, Social Network Analysis, Non-communicable Diseases Prevention
Category: H.3.4, H.3.5, J.4

1 Introduction

According to data from the World Health Organization, chronic non-communicable diseases accounted for 68% of global deaths in 2012. Care for this type of disease transcends the patient's engagement by extending to their family, friends, as well as acquaintances. Those people may influence treatment either in a positive or negative way. In that context, social support is understood as the ability of the social network to alleviate the harmful effects of stress and other health risks.

Social Cognitive Theory (SCT) states that behaviors are learned and reinforced through social interactions [Bandura, 2002]. Self-awareness of how one can influence other people's health might also lead not just to the improvement of one's well-being, but also to the improvement of the well-being of people in their social circle. For example, a person who knows their eating habits are influencing people close to them may start making better eating choices to help those people obtain better health.

Moreover, according to [Milani and Lavie, 2016], the current model of care can no longer efficiently cope with the challenges of this century. An aging population, the increasing incidence of chronic diseases, and the imbalance between the demand for doctors and the formation of new professionals, created the need to elaborate a new model of care. In this new model, technology plays an important role as it provides continuous monitoring, access to real-time information, communication with health professionals, support in disease management, and fostering social support (and beneficial social relationships).

Hence, computer aided healthcare must learn that sometimes patients do not have a partner who can aid them achieving beneficial health gains. Thus, the necessary aid can come from people in their social network. Direct connected nodes of an ego could be a match for certain health-related activities, sometimes these nodes could serve as bridges to aid in the formation of new beneficial social ties.

The influence of the social environment on health is not a new theme, having already been addressed by Emile Durkheim in the nineteenth century. Durkheim contributed significantly with his studies on the weight of the social effect on individuals' morbidity [Durkheim, 1897, Berkman et al., 2000]. In his study "Suicide", Durkheim analyzed particularly the influence society has on the decision of an individual to commit suicide.

Social network has an important role in health as it regulates access to resources and opportunities to its members, as well as models their behavior, which may be of higher or lower risk. Hence, *social support* is the ability of that social network to alleviate the harmful effects caused by stress and other health risks, through the provision of material, emotional, and informational resources; and in influencing behaviors such as healthier eating, practice of physical activities, medication use, and seeking medical follow-up [House et al., 1988].

In essence, social support is the beneficial influence of social relations in people's health [House et al., 1988, Berkman et al., 2000]. In a computing context, social support is successfully achieved through Internet applications, such as chats, blogs, forums, wikis, or video sharing [Vianna and Barbosa, 2017]. Notwithstanding, few studies approach the possibilities and opportunities of integrating the ever-growing amount of computational devices and data for improving social support.

Today we are immersed in a world of computing devices and data. Data generated by personal devices, such as smartphones, mimic the behaviors of their owners, and therefore could be applied to reveal aspects that influence health, or to find people the contact of whom can influence the improvement of health. We present Pompilos, a model for enhancing health care applications with social media content and for informing users about the impact of their behavior on the health of others. With that in mind, we organized this paper as follows: Section 2 presents some related works; Section 3 describes the proposed model; Section 4 the prototype created for evaluating our model; Section 5 exposes the results found in two performed experiments; and finally, Section 6 presents final remarks about this work.

2 Related Works

This section presents five works that focused on improving social support on the individual level or are techniques for identifying social awareness with social media data, which are the goals of the proposed Model. At the end of this section we show a comparison between these works in terms of *pervasiveness*, *adaptability with patients' behavior*, *social support availability* and *how they enhance awareness about social influence on health*. We grouped works belonging to larger projects for the sake of objectivity.

2.1 WANDA

WANDA [Alshurafa et al., 2014a, Sideris et al., 2015, Alshurafa et al., 2017] is a remote health monitoring system that uses patients' baseline data and contextual information (collected by the patients' smartphones) to improve NCDs care. According to [Lan et al., 2012], monitoring systems failed to achieve significant improvements in heart failure treatment. In general, they tend to be invasive and reactive. *WANDA* is a proactive, non-invasive platform integrated to smartphones, which has an analytics engine and a social platform. The engine analyzes patients' data collected through automatic sensing to provide custom monitoring and notifications, while the social platform adds cooperation and competition features among the system's users.

WANDA has received great improvements since its first publication. The WANDA remote monitoring system was already used to predict the success in care of certain risk factors related to cardiovascular diseases, adherence to the practice of physical activities, daily questionnaire fulfillment, blood pressure measurement of the Women's Heart Health Study participants, treatment adherence prediction after intervention by social support, as well as to identify patients who are most likely to be successful in their treatment, so they can receive custom care plans in advance.

2.2 Patient Journey Record System

Patient Journey is a metaphor that understands that the patient, as well as a traveler, is on a journey, and they will need assistance to deal with situations they are not used to in different stages of the disease [Martin et al., 2011, Martin et al., 2012]. *Patient Journey Record System (PAJR)* is a patient-cen–tered framework proposal that aims at integrating information systems, social networks, and digital democracy, so that different agents can construct a health support system collaboratively, taking into account that each patient has their own journey.

PAJR was already used to analyze the responses given by patients with chronic diseases to semi-structured questionnaire asked by "Care Guides". The analysis performed by the PAJR verifies the severity levels of the reports made by patients and classifies them accordingly, identifying in advance the need for intervention in the patients. We evaluated the system from November 2010 through December 2011 in a controlled experiment, which had 153 participants in the Intervention Group and 61 in the Control Group. Overall the Intervention Group had a 50% lower admission number than the Control Group. In addition, the model used by the system was able to classify 100% of cases of unplanned urgent events. According to authors, although the system has achieved good results, there is still a need for further evaluation.

2.3 Accessible telehealth - Leveraging consumer-level technologies and social networking functionalities for senior care

Accessible Telehealth [Dhillon et al., 2013]) is a conceptual framework for elderly care. The framework design was elaborated by reviewing the absences observed between different types of care platforms, namely: Telehealth, Health Record, Health Information Web Sites, and Serious Games, and through the collection of requirements made with patients interview. Within the functionalities identified are social and emotional support. A prototype was developed implementing features of social support in the form of social network, where users could create groups and search for friends. An evaluation of the prototype was carried out by 43 senior citizens. In the evaluation, 35% of patients agreed that the social functionalities motivated them to use the system. While 31% agreed that the involvement of friends helped them manage their health.

2.4 Towards chronic emergency response communities for anaphylaxis

[Schwartz et al., 2014] propose a workflow approach to help chronic patients facing emergencies. In this approach, patients who face emergency situations

and do not have the resources required to deal with it are able to trigger an alert request. This request is passed on to individuals (members of the system) who have the resources necessary to help and are within a distance range that allows for a timely intervention. Members may accept or reject the aid request. This type of mechanism, where patients request help from each other in health situations, was called by the authors “social medicine”.

2.5 IntelliCare

IntelliCare is a set of 14 mobile applications oriented to help users coping with depression and anxiety [Mohr et al., 2017]. Designed as Behavioral Intervention Technology (BIT), each IntelliCare application aims at developing user skills in different behavioral strategies, such as setting goals, relaxation, or behavioral activation. A total of 105 users enrolled in a eight weeks field trial of the IntelliCare and the results showed an improvement in depression and anxiety score tests.

2.6 You Tweet What You Eat: Studying Food Consumption Through Twitter

Food consumption screening is done with the use of costly questionnaires. Social networking users, such as Twitter, often share information about foods they consume. Social network data analysis, such as data from Twitter, makes it possible to identify diet trends at a geographical level (e.g., State or City level), as well as personal habits, and the relationship these habits have with the social network of individuals [Abbar et al., 2015]. Using a Naive Bayes model, the authors classified the type of foods consumed according to the text of messages shared on Twitter. Crossing the information about food consumption and caloric level of foods, it was possible to correlate obesity and diabetes at the State level and to validate them with the results published by the American Centers for Disease Control and Prevention (CDC).

The authors were also able to predict obesity and diabetes at individual levels. To do so, a regression model inferred the risk of obesity based on the type of food shared by users in their messages. Finally, the social network analysis performed by the authors found that the probability of being obese or having diabetes increases when one is connected to other obese or diabetic nodes. That is, nodes of the same network share messages containing similar foods.

2.7 Related Works Comparison

It seems that the studied works provide social support by two main approaches: technological support or behavior awareness. Technological support may refer to the provision of software tools for supporting or getting in touch with others. Yet, behavior awareness refers to the addition of features for cognitive behavioral treatment, that is, features that help patients adapt their behaviors by acquiring knowledge of the positive and negative effects of their actions [Gertz and Culbert, 2009, Moher, 2018] related to health. The previously presented works are compared in terms of pervasiveness, adaptation according to users profile, social support, and awareness of social situation.

Table 1 summarizes this comparison. The column *Pervasive* indicates whether the work proposes the provision of assistance to users anytime and anywhere, that is, whether it is pervasive [Satyanarayanan, 2001]. *Adaptive* indicates whether the work proposes the adaption of its operation according to the users' behaviors or situations they are in, in other words, whether it supports context awareness [Dey et al., 2001]. Column *Social Support Enabled* indicates whether the work proposes, in some manner, the participation of others in improving the health of its users. Finally, *Social Aware* indicates whether the work proposes a means for providing awareness on the social influence of others on the health of its users, or the influence of its user on health of others, in some way that it can recommend connections that might benefit the health of its users, or improve the understanding of how the users' behaviors influence the health of others.

WANDA and Towards chronic emergency response communities for anaphylaxis are understood as *pervasive* since they provide a platform that can be used anytime and anywhere by its users, providing continuous care. They are also *adaptive*, as is the Patient Journey Record System. WANDA is able to predict user engagement in its care according to their behaviors. The Patient Journey Record System checks users' answers to questionnaires to identify in advance the need for intervention. By its turn, Towards chronic emergency response communities for anaphylaxis is aware of user's location and whether they have the necessary resources to indicate the possible caregivers in emergency situations. With the exception of You Tweet What You Eat, all other works offer some functionality of *social support*. WANDA detects the need of social intervention to improve engagement, the *Patient Journey Record System* is based on contacts between patients and caregivers, Accessible Telehealth offers a social network to improve the collaboration among users, and in Towards chronic emergency response communities for anaphylaxis the users can request the help of others. While, IntelliCare has Social Force, an application that helps users identify supportive contacts and encourage users to be in touch with them. And lastly, just You Tweet What You Eat is *social aware* once its model considers the influence of others in the probability of having a health condition. Thus, different from the related works presented in this paper, Pompilos aims at integrating pervasiveness, adaptability, and social data to improve social support by making the user aware of their social context. That is, give users social recommendations for

Title	**Pervasive**	**Adaptive**	**Social Support Enabled**	**Social Aware**
WANDA [Lan et al., 2012, Alshurafa et al., 2014b, Alshurafa et al., 2014a, Sideris et al., 2015, Alshurafa et al., 2017]	Yes	Yes	Yes	No
PAJR [Martin et al., 2011, Martin et al., 2012]	No	Yes	Yes	No
Accessible Telehealth [Dhillon et al., 2013]	No	No	Yes	No
Towards chronic emergency response communities for anaphylaxis [Schwartz et al., 2014]	Yes	Yes	Yes	No
IntelliCare [Mohr et al., 2017]	Yes	Yes	Yes	No
You Tweet What You Eat [Abbar et al., 2015]	No	No	No	Yes

Table 1: Related Works Comparison

the improvement of their health, and also to show how user's behaviors impact health of people close to them.

3 Pompilos Model

The essence of the Pompilos Model is expressed in Figure 1 within an activity diagram [OMG, 2007]. This diagram exposes a black box model showing the hierarchy of activities that must be performed in order to accomplish the Pompilos' goals of improving social support by means of revealing influence on health of others and suggesting beneficial health resources.

Contacts and *Context History* data are used as input to *Generate social network* of users. Social networks and context information are then used to *Manage models of influence detection*, that is, create and train models that can identify the influence of someone's behaviors on the health of others. *Check new resources* and *Infer influence between nodes* are executed in the data of different

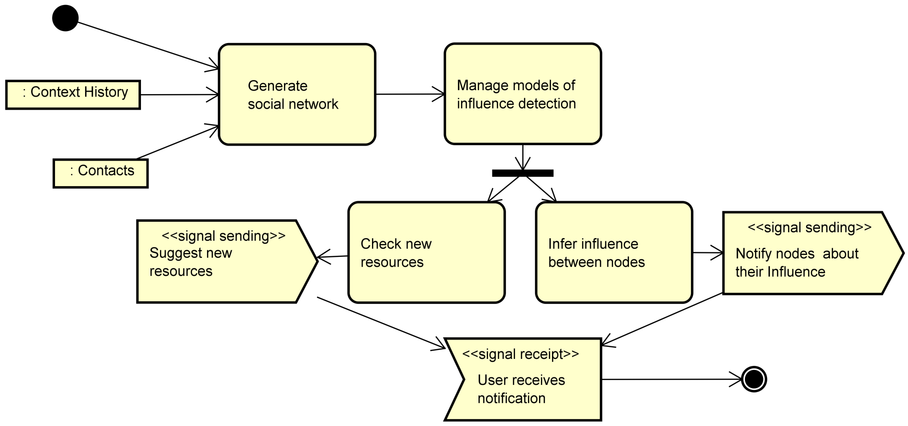

Figure 1: Pompilos General Model

nodes in order to Suggest new resources that may benefit users' health and *Notify nodes* about their Influence. Finally, *User receives notification* regarding the new resources and node influence information generated for them.

3.1 Social Network Generation and Models of Influence Management Dynamics

The participants and dynamics responsible for accomplishing the actions of *Generate social networks* and *Manage models of influence detection* and their interactions are expressed as a communication diagram [OMG, 2007] in Figure 2. In this communication diagram, and in the next, the actors represent the types of software agents and the objects represent the data storage components. Software agents types are supposed to have many instances that have autonomous behaviors, while data storage components are the proxy for the storage and retrieval of large amounts of data. Both the agents and the data storage components must interact with each other in order to perform the actions defined in the general model.

The *U'Ductor Node* [Vianna and Barbosa, 2014] represents instances of the U'Ductor middleware. These instances may represent people who generate contexts from smartphone activities, or places that generate contexts from user visitations, resource sharing or location related status (e.g. temperature or air humidity). *Context Acquisition Agent* instances may impersonate U'Ductor's users as a means of collecting context data of the user's activities on different applications. For example, a Context Acquisition Agent may be authorized to capture the Twitter activity generated by its user, which can be a like or a retweet on the tweets of others. The Personal Node and the Context Acquisition

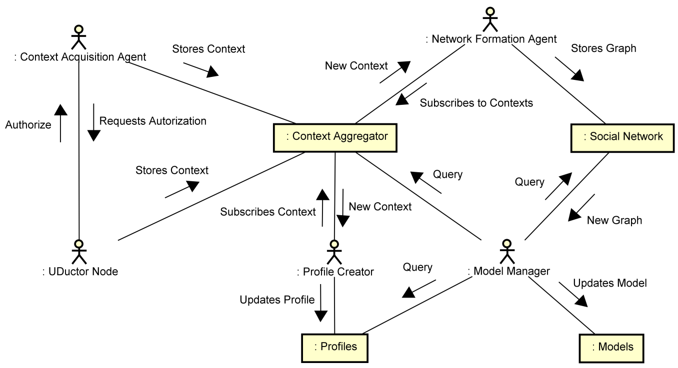

Figure 2: Social Network and Model Generation Dynamics

Agent store the collected context data in a *Context Aggregator*. *Network Formation Agent* instances use context data, such as contacts, phone calls, visited locations or any other relevant information received from the Context Aggregator to form social networks, which are stored by the *Social Network* storage component.

There are different models to predict the influence between nodes. So, distinct instances of the *Model Manager* will co-exist, aiming at discovering influences of nodes by testing different models with data from context history, profiles, and social network. The findings of the Model Manager are then stored in the *Models* data storage component to be used later to infer the influence on the nodes. Lastly, the *Profile Creator* instances are responsible for creating summaries of the users' context histories. For example, these summaries can indicate the engagement of users in a particular treatment or care plan. The profile information are stored in the *Profile* data storage component and may be useful for the creation of new models of causal influence.

3.2 Influence Notification and Resource Suggestion Dynamics

Figure 3 shows a communication diagram where the participants and the dynamics of the General Model's activities *Check New Resources*and *Infer Causal Influence between Nodes* (Figure 1) are expressed. The interaction of the *Influence Advisor* and the *Resource Recommender* with other components are very similar. Both select the most appropriate model for their tasks in the Models data storage. After the selection of the model, a request for execution is made to some instance of a *Model Executor*. These instances of the Model Executor run the requested model using data from the Context Aggregator, the Profiles, and the Social Network, according to the parameters of the request.

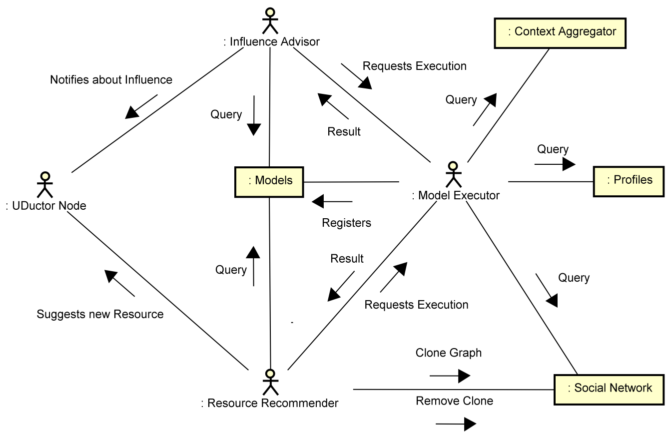

Figure 3: Influence Notification and Resource Suggestion Dynamics

After receiving the result from a Model Executor, the Influence Advisor and the Resource Recommender might interpret that result and execute further processing in order to send to Personal Nodes the influence that they exercise on others' nodes and to recommend new beneficial health resources. Specifically, the Resource Recommender may want to *clone social network graphs* in order to simulate new connections and infer their outcomes.

3.3 Architecture's Technical Aspects

The implementation of the model's components follows a proposal for the escalation of building context-aware applications for non-communicable diseases prevention presented in another work [Vianna and Barbosa, 2019]. As shown in Figure 4, the Pompilos' conceptual architecture has two categories of services, which use the aforementioned components, exposing them as interfaces: Pompilos Node (Figure 4a) and Storage (Figure 4b). Thus, these services are able to provide features for access control, resource and context sharing, notification, location, application, and node lookups. Both resource and context interfaces allow the query and storage of data, while the context interface also allows context subscription. Hence, these interfaces allow other services or Software Agents to, respectively, automatically manage access control, store, fetch and receive updates from the data maintained by the Pompilos Nodes.

The storage service is an extension of the Pompilos Node, being responsible for the management of data repositories, exposing implementations of the context or resource interfaces. Models and Social Network services extend the Storage service definition by exposing new interfaces. Models services expose the register interface that allows the registration of the Model Executor agents.

Nonetheless, the Social Network service provides interfaces for cloning and removing social network graphs that are used for training new knowledge discovery models. The registering, cloning, and uncloning interfaces are provided as node resources.

Pompilos Node services are also aware of their environment to actuate when necessary. This is achieved by exposing the notification interface used to receive data updates of subscribed interests from Storage services, or other Pompilos Nodes [Birman and Joseph, 1987, Dey et al., 2001]. The Resource Recommender and Influence Advisor extend the Pompilos Node service by adding a result interface. This interface is summoned by the Model Executor service to deliver results of knowledge discovery model executions requested with the execute interface. Both the result and the execute interface rely on the resource interface from the Pompilos Nodes, as they are offered as resource residing in nodes. Lastly, Context Acquisition Agents use the access control interface to negotiate authorizations for collecting contexts on behalf of the Pompilos nodes.

The application of the Pompilos can be illustrated by the following example:

> Considering a Twitter user that is willing to consume healthier food. The Pompilos model could offer them relevant healthy diet messages, for example, messages that identify which followed account influences the user in food consumption topics. Thus, a Network Formation Agent could use the followed account to form its network. Messages from the followed accounts would be used to create the message topic profile of each followed account. Moreover, messages with which the user interacted could be used to generate their user profile. Based on the user profile and the message topics profile it is possible to identify the profiles that influenced the user the most by message topic. Moreover, messages from healthy diets could be recommended to the user through the monitoring of accounts that publish these types of messages.

4 Evaluation

As a way to demonstrate the applicability of the proposed model in the prevention of non-communicable diseases, we developed MyUDuctor, an online mobile assistant for diets, weight management, and the practice of physical activities needed to prevent and control risk factors for chronic non-communicable diseases. Like any other assistant, it allowed the scheduling and recording of care activities, and allowed the user to visualize their care progression. However, unlike existing care applications, MyUDuctor collected messages related to the practice of physical activities and healthy eating from social networks. Such messages were presented to the users as a form of social support, which is recognized as a motivational factor for reinforcing user engagement to their care activities.

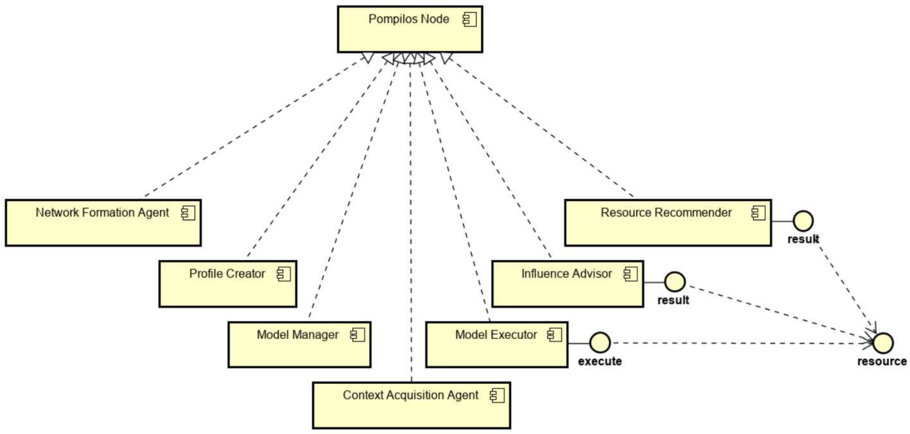

(a) Software Agents Services

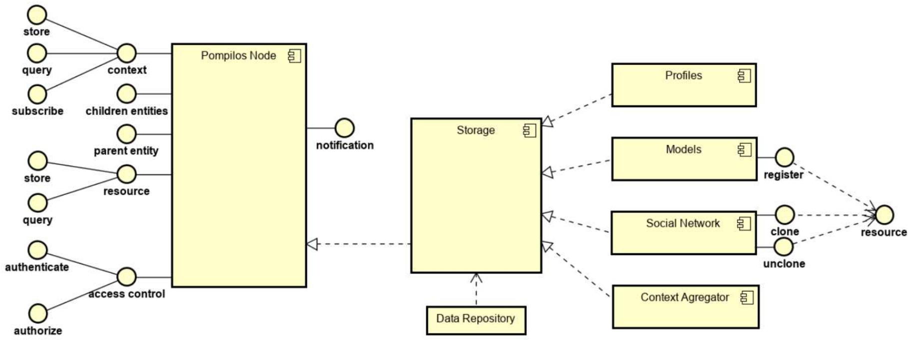

(b) Pompilos Storage Services

Figure 4: Conceptual Architecture Specialized Services

The next sections explain the existing features of MyUDuctor and the process for recommending messages related to the prevention of non-communicable diseases. The focus of this experiment is on evaluating the influence of receiving social media messages on application usage behavior. To record completed activities (weight, diet, and physical activity), forms were provided. In some instances of physical activity, such as walking and jogging, the form can be replaced by automatic data collection from wearable phone sensors. Nevertheless, to maintain the consistency of the application's use dynamics, we decided to use only forms, which allowed the addition of different categories of physical activity, such as bodybuilding, stretching, swimming, and martial arts. Those activities may have different energy expenditure, and are not always liable for data acquisition by sensors [Ainsworth et al., 2011].

4.1 MyUDuctor: a Health Assistant Application

MyUDuctor was publicly available at https://app1.uductor.com. To access the application, the user should have a Google account. This restriction was made to ensure the authenticity of the user's e-mail address as some of the communications from the application were made through electronic messages.

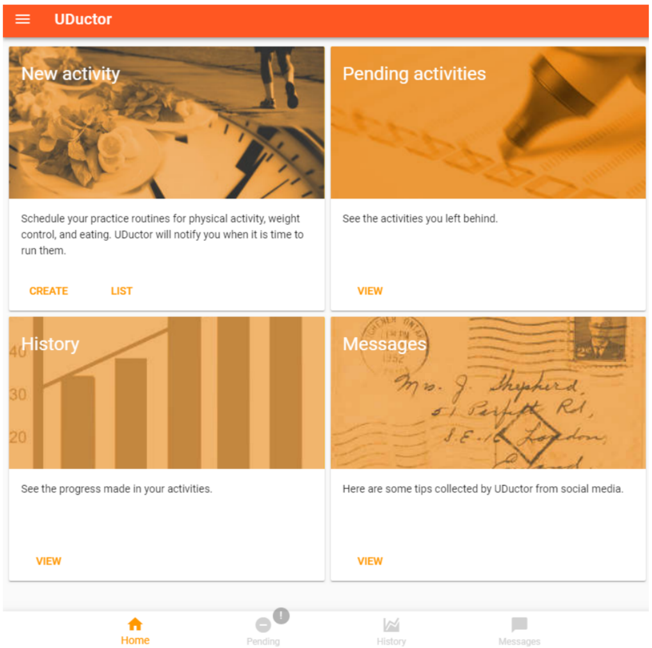

Figure 5: Application main screen

The first time the users authenticated, the application displayed the free and informed consent form. In this form, the user could also set the notifications receiving permissions and access to location data. The use of the application was only allowed after the user checked the option “I have read and I accept the terms and conditions set forth above” and pressed the “I accept” button. This way, the free and informed consent term was accepted electronically and online, allowing users from different locations to participate in the experiment. A copy of the informed consent form signed by the person in charge of the research was sent by e-mail to the users after the acceptance of the term in the application.

After accepting the informed consent form, users were directed to the main application screen (Figure 5). This screen described all the features in the application, as well as provided links so that the users could access these actions.

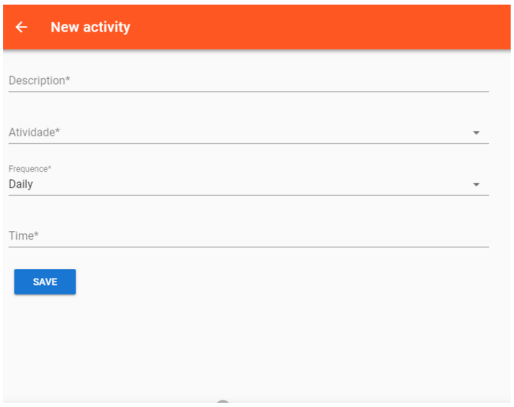

Figure 6: Daily activity

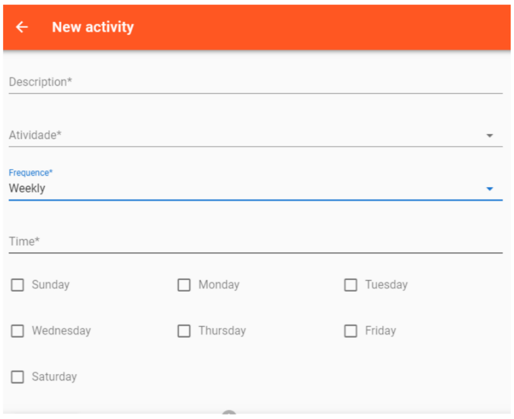

Figure 7: Weekly activity

Optionally, users could access these same features in the options menu or through the shortcut bar. The application provided five main features: *New Activity*, *Activity List*, *Pending Activities*, *History*, and *Messages*.

The "New activity" feature allowed the users to define the recurrence with which they would be advised about the need to carry out their activities. There were three types of possible activities, all related to preventing the risks of non-communicable diseases: Put their activities. There were three types of possible activities, all related to the prevention of risks of noncommunicable diseases: *Physical activity*, *Meal*, and *Weight*. Users could set activities to be performed daily in a certain hour (Figure 6); weekly, with a repetition rate per day of the week (Figure 7); or a monthly repetition, with a repetition rate per day of the month (Figure 8). In turn, the "Activity List" feature allowed the users to edit

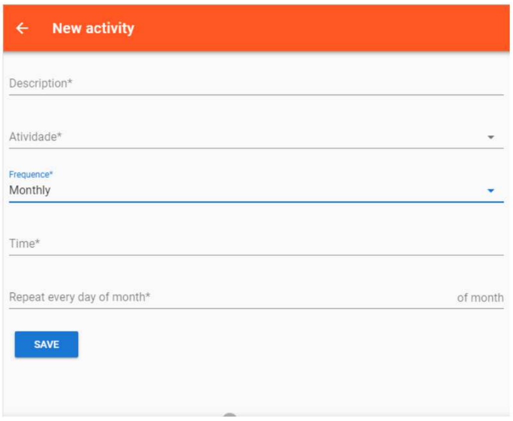

Figure 8: Monthly activity

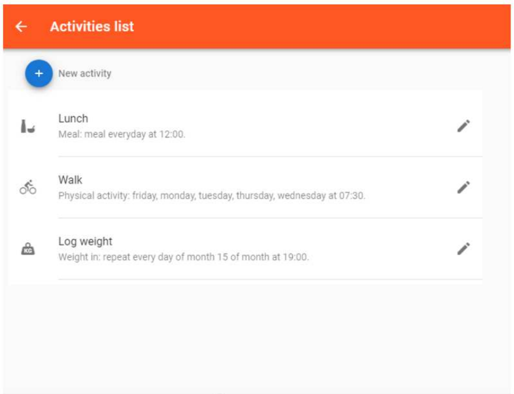

Figure 9: Activities list

the recurrence of previously registered activities (Figure 9).

The users received reminders for those activities that they had scheduled. Reminders for activities that the user was required to perform were emailed to them. Optionally, users were able to receive notifications via smartphone, as long as the permission to receive notifications had been accepted by the users.

Activities that the user did not register appeared in the "Pending Activities" action (Figure 9). When users clicked on any item on the pending activities list, the application displayed a log form according to the type of activity. The "Physical activity log" allowed the users to record what types of activity they performed, the pace of these activities, and for how long the activities were performed (Figure 10b. For example, in the same log the user could indicate that they walked for 20 minutes at a moderate pace and ran for 10 minutes at

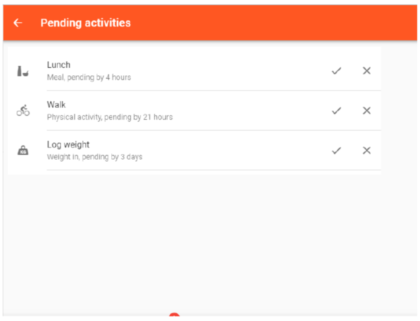

(a) Pending activities

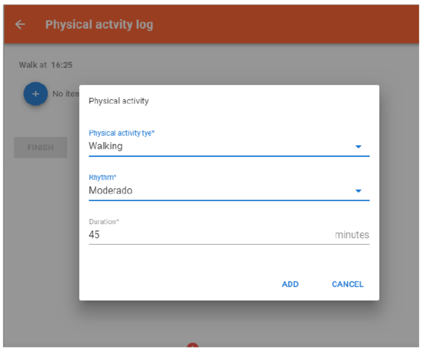

(b) Pending physical activity log

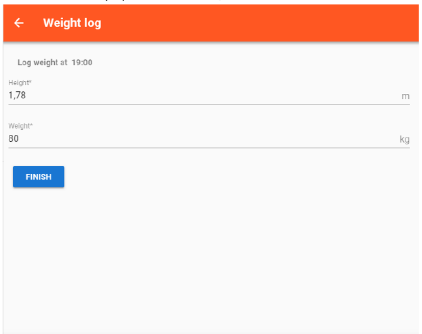

(c) Pending weight log

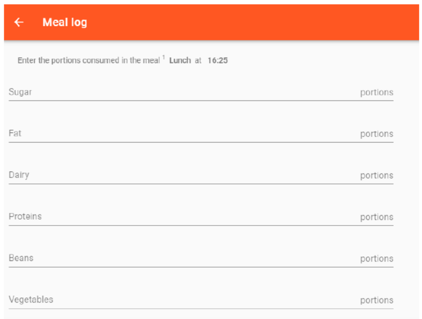

(d) Pending meal log

Figure 10: Pending activities

a low pace.

The "Weight Log" allowed users to enter their height and weight (Figure 10c), this way it was possible to calculate the user's minimum and maximum weight limits [WHO, 2018a]. The 'Meal Log" allowed the users to enter the portions they consumed in their meal for each food category (Figure 10d). When inserting a food portion, the application indicated the appropriate number of portions per day and presented some portion equivalences for some types of foods according to the one proposed by [Philippi et al., 1999].

The "History" feature allowed users to graphically view records made for each type of activity in a specific period. The physical activity history (Figure 11a) indicated the activity summary performed by the user and the green-time recommended activity time for adults by [WHO, 2018b]. The weight history (Figure 11b) presented, in addition to the user's weight records, the minimum and maximum limits according to their height [WHO, 2018a]. The Meal History (Figure 11c) presented the summary of portions consumed by the user by type and the recommended amount for the same period, as proposed by

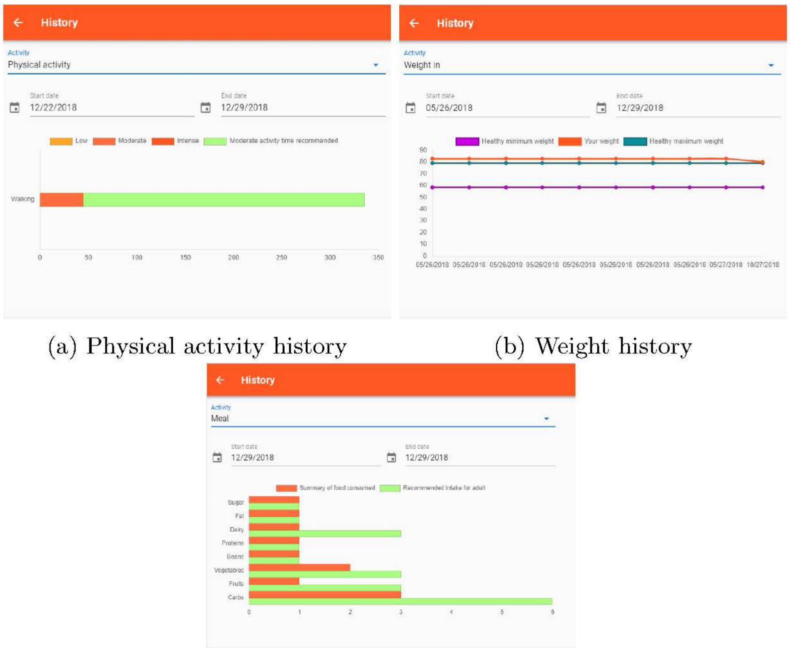

(a) Physical activity history (b) Weight history

(c) Meal history

Figure 11: History

[Philippi et al., 1999].

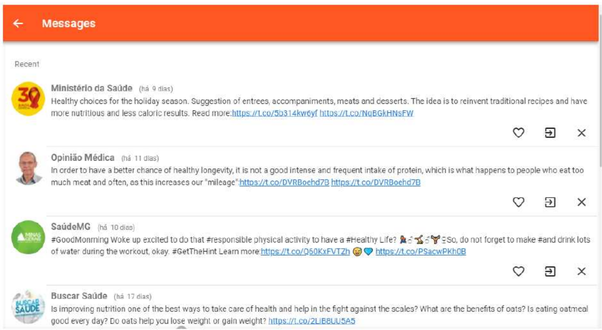

Figure 12: Messages

The application notified its users whenever a new message relevant to the prevention of chronic disease risk factors was identified. The list of recommended messages could be accessed by the user through the “Messages” action (Figure 12). Users had the ability to "like" or go to the source of each message received.

4.2 Messages Recommendation and Rank from Twitter's Health Profiles

To recommend the messages related to the practice of physical activity, healthy eating, and weight control, it was necessary first to identify profiles on the Twitter social network that publish these types of messages. First, a seed profile was defined to be examined. To do so, the profile of the Brazilian Ministry of Health (@minsaude) was chosen. Then a review of the messages sent by the profile related to healthy eating and physical activity was checked. If the related message had hashtags, these hashtags were used to search for Twitter messages to identify other related profiles. In the end, the following profiles were selected: @BuscarSaude, @saudavelebarato, @realfoodbrasil, @rederms, @saudavelcomida, @BlogSaudeBetter, @WorldBoaForma, @fitcomdaniele, @CamilaGramelick, @DetoxReceitas, @opiniaomedica, @juntoemisturadi, @belagil, @SES_RS, @SaudeMG, and @minsaude.

Figure 13a shows the activities performed in training the model for detection of NCDs preventive messages. First, all messages sent by the selected profiles with up to two years of publication were registered in a database. In the end, 24,971 messages were stored in the *tweets collection*. A *data classification* process was then started. From the tweets collection, 7,040 messages were classified as "Healthy Eating", "Physical Activity", "Weight Control" or "Unclassified". To perform the classification, a simple web page application was made available (Figure 13b). After the classification, the data were purged for impurity removal like stop-words, URLs, hash symbols (#) and duplicate messages. In the end, 4,747 remained for classification.

The model training consisted of two main activities: *topics generation* and *model training*. Given that each tweet can be understood as being composed of a mixture of topics, and that each word from that tweet has a probability of pertaining to a topic, the Latent Dirichlet allocation algorithm was used to map out the topical probability of each tweet from the tweets collection [Robinson and Silge, 2017]. The topical probability of each tweet was used as features to train a Support Vector Machine (SVM) model for classifying the tweets as NCDs prevention messages (healthy eating, the practice of physical activity, or weight control) or not related with NCDs prevention [Chang and Lin, 2011]. The number of generated topics was increased until an accuracy greater than 90% was obtained. This was reached when the SVM model was trained with 96 topics. The training performance by number of topics is summarized in 13c.

A profile monitor agent was notified whenever one of the profiles sent a new message. These messages were then forwarded to a specific agent, which verified

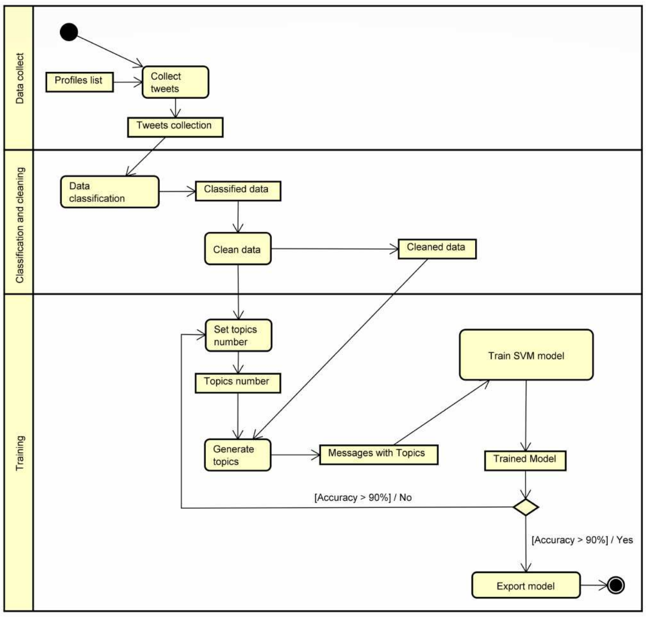

(a) Training Process

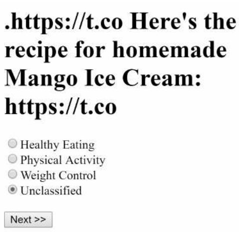

(b) Message Tagger

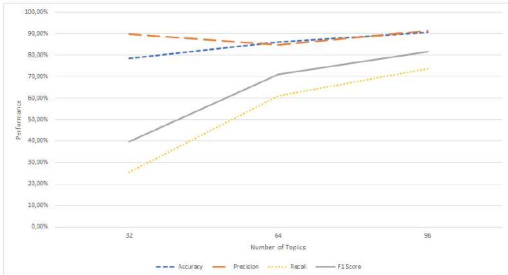

(c) Training Performance

Figure 13: Training Process and Performance

whether the messages were related to the practice of physical activity, healthy eating, or weight control using the trained model. If the messages were related to one of those topics, they were kept for further quality checks by the researcher in charge and only then forwarded to the users.

A rank of Twitter health care profiles was publicly available on the web at https://app1.uductor.com/rank. The monitored Twitter profiles were notified about the ranking in order to encourage them to improve the posting of non-communicable diseases prevention-related messages. The score of each profile took into account a Bernoulli distribution [Goyal et al., 2010] of users' engagement in terms of likes, followed links, and the use of the application's features in a 24 hour window by the intervention users after receiving the message.

5 Results

Two experiments were run concomitantly in order to assess the engagement of the application users and the behavior of the social network profiles in regards to the posting of messages related to non-communicable diseases prevention. The methods and findings of these two experiments are presented in the following sections.

5.1 Findings about MyUDuctor usage

In order to assess the user's engagement to the application, we designed a randomized experiment. For this, two variants of the application were made available at https://app1.uductor.com, they were *Control* and *Intervention*. The health message recommendations feature was enabled on the Intervention variation and disabled on the Control variation. Users were randomly assigned to one of the two as they registered to the application and were not informed about the group they were assigned to. The application was promoted by direct messages in instant message platforms to the author's contacts, through posts on the author's social platforms, and to four different university discussion lists at the dates 10/12/2018, 10/15/2018 and 10/16/2018. A total of 45 users registered to the application, 23 have been assigned to the Intervention version and 22 to the Control version.

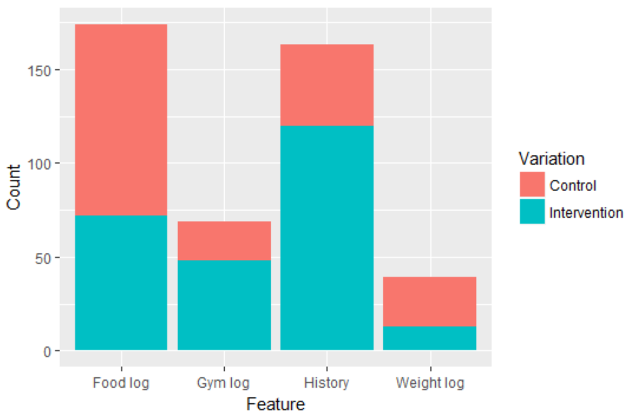

Figure 14: Total features usage

Users from the Control variation had a greater usage of meal and weight logging features, while the users from the Intervention group had a greater usage of the History and Physical Activity logging features. Figure 14 shows the number of times each type of feature was accessed by the users. Figure 15 shows the distribution of features usage through time. The Messages and Messages Interaction features were added to that chart, where its possible to observe that Intervention users have used the application for longer. Users of the Intervention variation accessed the application for more days, on average 10% more.

Figure 15: Distribution of features usage through time

A robust fitting of linear model was applied to check if History, Message Interactions, and Messages features had any correlation whatsoever. With that, a strong relation was detected as demonstrated in Figure 16 by the three lines that have the same trend over the distribution. Afterwards, a Wald test for multiple coefficients was applied in each feature distribution: History *p-value = 0.013, p-value <0.05*; Message *p-value = 0.022, p-value <0.05*; Message Interaction *p-value = 0.028, p-value <0.05.* Hence, it was possible to reject the null hypothesis and state that Messages and Message Interactions features have possibly influenced the use of History feature by the Intervention users.

5.2 Twitter Profiles Behavior's Analysis

The selected Twitter profiles were monitored from 07/08/2018 to 12/01/2018. In total, 7,011 messages were processed, from that, 470 were classified as non-communicable diseases prevention messages, and 420 passed the quality check.

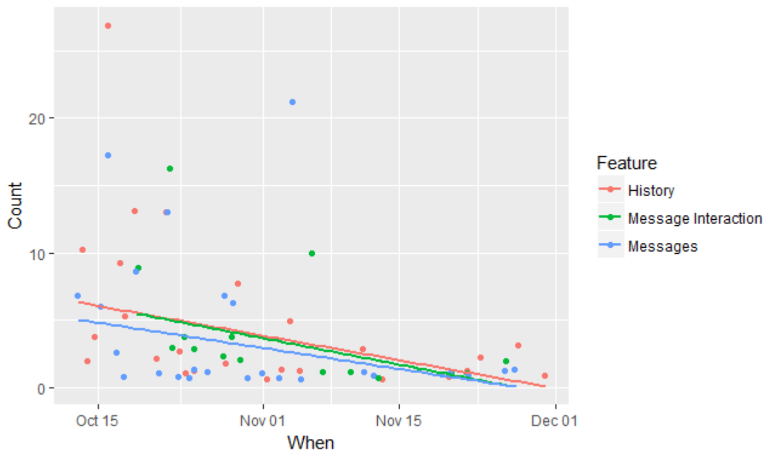

Figure 16: Distribution of History, Message Interactions, and Messages features

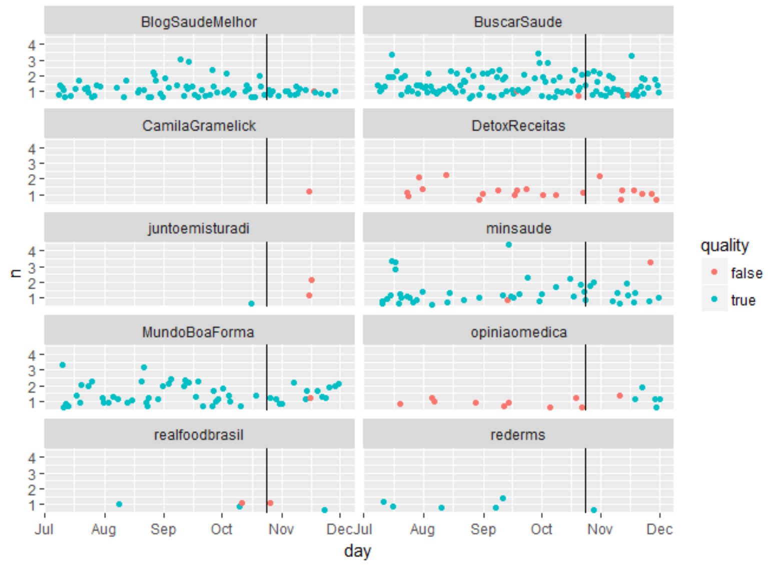

Figure 17: Distribution of the sending of non-communicable diseases prevention message by the monitored Twitter profiles

Quality check was done to ensure the delivery of messages that did not stimulate the cult of body [Boepple et al., 2016, Simpson and Mazzeo, 2017], so that the messages sent by the platform had, for the most part, educational biases [Smahel et al., 2017]. Moreover, messages should be self-contained, or at least, had a link pointing to more information.

Profile	Before Intervention			After Intervention			Causal Inference Test $\alpha = 0.05$
	Messages	Daily mean	σ	Messages	Daily mean	σ	p
BlogSaudeMelhor	72	0.67	0.71	17	0.44	0.55	0.06
BuscarSaude	121	1.13	0.80	43	1.10	0.72	0.42
CamilaGramelick	0	0.00	0.00	1	0.03	0.16	0.00
DetoxReceitas	16	0.15	0.41	8	0.21	0.47	0.25
juntoemisturadi	1	0.01	0.10	3	0.08	0.35	0.00
minsaude	48	0.45	0.83	18	0.46	0.85	0.48
MundoBoaForma	57	0.53	0.78	20	0.51	0.76	0.49
opiniaomedica	9	0.08	0.28	7	0.18	0.45	0.05
realfoodbrasil	3	0.03	0.17	2	0.05	0.22	0.30
rederms	5	0.05	0.21	1	0.03	0.16	0.36

Table 2: Causal inference test

By 10/23/2018, a contact with each monitored Twitter profile was made through direct messaging. The message presented the My UDuctor application and had a link pointing to the Twitter's Health Profile Rank. The profiles @opiniaomedica and @BuscarSaude accessed the rank once, as the others profiles did not access the rank at all. Figure 17 presents the distribution of non-communicable diseases prevention messages sent by the monitored Twitter profiles. The vertical line indicates the moment when the profiles were notified about the application, the red dots represent messages that did not pass the quality check, while the blue dots did.

A causal inference with Bayesian structural time-series model test was made to check if the contact made had an effect in the tweeting behavior of Twitters' profile [Brodersen et al., 2015]. *Table 2* summarizes the results of the test. In general, the test results were not statistically significant and could not have meaningful interpretations. However, two profiles were exceptions and are highlighted in Table 2. It seems that contact did a positive effect, in the profiles @juntoemisturadi e @opiniaomedica. It is also important to notice that the profile @opiniaomedica had an improvement in the quality of messages after the contact was made, as shown in Figure 17. Although the daily number of messages from some profiles was low, the observation of these profiles lasted approximately six months. This observation allowed us to evaluate some cases, such as the changing in text characteristics in the @opiniaomedica profile or the changing in the average of messages sent, such as the @BlogSaudeMelho profile.

6 Conclusions

Prevention of non-communicable diseases is a global concern as these types of diseases account for a significant part of global deaths. The current model of care can no longer cope efficiently with the challenges of this century; hence, technology might provide continuous monitoring, access to real-time information, communication with health professionals, support in disease management, and fostering social support. Moreover, models for detecting the dissemination of information in social networks, and also to identify influence that the nodes exert on others already exist [Chen et al., 2013, Christakis and Fowler, 2011, Garcia-Herranz et al., 2014, Christakis and Fowler, 2010, Goyal et al., 2010]. Social data is used for knowledge discovery, and was successfully applied in studies related to NCDs to detect diseases outbreaks [Lee et al., 2015, Ram et al., 2015, Zhang et al., 2016], self-disclosure discourses [Balani and De Choudhury, 2015] or health risks [Paul and Dredze, 2011, Culotta, 2014, Weber and Mejova, 2016]. The Internet is the default platform for offering social support for NCDs prevention and care. This support is generally offered through forums, social networks, wikis, blogs, chats, and shared videos. Some studies used data collected from patients to detect the need for interventions: [Martin et al., 2011, Lan et al., 2012, Schwartz et al., 2014, Alshurafa et al., 2014a, Sideris et al., 2015]. For example, [Alshurafa et al., 2014a] used data collected from patients answers in a smartphone app to send them notifications for behavior change. However, none of the studied works used social data or data automatically collected from patients to provide social support based on the influence received from their social network.

The Pompilos model is supported by the idea that social behaviors spread through people social networks. In this way, the model dynamics relies on certain activities, which are: collecting data from users, generating users' social networks, inferring users' profiles, and training models for detecting and computing social influence on spreading NCDs risk factors on people. By accomplishing these activities it is possible to recommend resources for engagement in health care prevention, and also to make people aware about their influence on the health of others. The model based on the achievement of these activities is the main contribution of this article as this concept was not explored in any of the studied works. The presented model is also able to integrate the ever-growing amount of computational devices and data for improving social support.

As shown in this article, pervasive software can be used to build solutions that positively enhance feelings of well-being and social participation in health care, which is improved when the different types of data are applied to enhance users' capabilities. The proposed architecture is technically feasible and the model is not complex to understand. Furthermore, it can be scaled up, guarantee privacy, and provide near real time updates, which ease the creation of context aware applications. The model was designed for accomplishing the requirements of preventive care of NCDs but can be extended to other domains. As the communication between the model's elements is based on the REST architecture style, a great part of its implementation can be achieved using existent web technologies. However, other architectures can be used to accomplish the requirements like, for example, the Enterprise Service Bus (ESB), which is an architectural framework that provides a common interface for integrating applications to different services [Bhadoria and Chaudhari, 2018, Bhadoria et al., 2018, Sharma et al., 2017].

For evaluating the model an online mobile assistant for diets, weight management, and the practice of physical activity was developed and tested by 45 users for one month and a half. Users' experience was enhanced when they received non-communicable diseases prevention messages from monitored Twitter profiles. This feature correlated with an improvement of the applications' History feature usage, suggesting that the users were more concerned in following their health behavior. Besides that, the monitored Twitter profile behaviors were analyzed and it seems that at least for one profile the awareness of social influence had a positive effect. However, this evidence is small and will be addressed in future works.

Finally, this study may not be understood as generally conclusive and is open to further improvements and following researches. The experiment presented was very open in the sense that users were not controlled. This has some advantages, as it elevates the users' perception of blindness, that is, their perception in forgiving being observed by the researcher [Dale et al., 2016], but hinders the collection of more cohesive data. Hence, a more controlled experiment with students of a Physical Education course is under way. By the end of this experiment we expect to acquire information about the model effectiveness, and also provide an interlink between the application acceptance, usage, and effectiveness. Additionally, the model could be adopted by Public Health Departments as a way to improve NCDs prevention policies and check their effectiveness. This last proposed work is in the design phase in partnership with the Department of Health of the State of Rio Grande do Sul, Brazil.

Acknowledgement

The authors wish to acknowledge that this work was supported by Fapergs/Brazil (Foundation for the Supporting of Research in the State of Rio Grande do Sul - http://www.fapergs.rs.gov.br), CNPq/Brazil (National Council for Scientific and Technological Development - http://www.cnpq.br) and Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code

001. We are also grateful to University of Vale do Rio dos Sinos - Unisinos (http://www.unisinos.br) for embracing this research.

References

[Abbar et al., 2015] Abbar, S., Mejova, Y., and Weber, I. (2015). You tweet what you eat: Studying food consumption through twitter. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 3197–3206, Seoul, Republic of Korea. ACM.

[Ainsworth et al., 2011] Ainsworth, B. E., Haskell, W. L., Herrmann, S. D., Meckes, N., Bassett, D. R. J., Tudor-Locke, C., Greer, J. L., Vezina, J., Whitt-Glover, M. C., and Leon, A. S. (2011). 2011 compendium of physical activities: A second update of codes and met values. *Medicine & Science in Sports & Exercise*, 43(8).

[Alshurafa et al., 2014a] Alshurafa, N., Eastwood, J., Pourhomayoun, M., Liu, J. J., Nyamathi, S., and Sarrafzadeh, M. (2014a). A framework for predicting adherence in remote health monitoring systems. In *Proceedings of the Wireless Health 2014 on National Institutes of Health*, WH '14, pages 8:1–8:8, Bethesda, MD, USA. ACM.

[Alshurafa et al., 2014b] Alshurafa, N., Eastwood, J. A., Pourhomayoun, M., Liu, J. J., and Sarrafzadeh, M. (2014b). Remote health monitoring: Predicting outcome success based on contextual features for cardiovascular disease. In *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 1777–1781, Chicago, IL, USA. IEEE.

[Alshurafa et al., 2017] Alshurafa, N., Sideris, C., Pourhomayoun, M., Kalantarian, H., Sarrafzadeh, M., and Eastwood, J. A. (2017). Remote health monitoring outcome success prediction using baseline and first month intervention data. *IEEE Journal of Biomedical and Health Informatics*, 21(2):507 – 514.

[Balani and De Choudhury, 2015] Balani, S. and De Choudhury, M. (2015). Detecting and characterizing mental health related self-disclosure in social media. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '15, pages 1373–1378, Seoul, Republic of Korea. ACM.

[Bandura, 2002] Bandura, A. (2002). *Social cognitive theory of mass communication.*, chapter Social cognitive theory of mass communication., pages 121–153. LEA's communication series. Lawrence Erlbaum Associates Publishers, Mahwah, NJ, US.

[Berkman et al., 2000] Berkman, L. F., Glass, T., Brissette, I., and Seeman, T. E. (2000). From social integration to health: Durkheim in the new millennium. *Social Science & Medicine*, 51(6):843–857.

[Bhadoria and Chaudhari, 2018] Bhadoria, R. S. and Chaudhari, N. S. (2018). Provisioning for Sensory Data Using Enterprise Service Bus: A Middleware Epitome. In *Enhancing CBRNE Safety & Security: Proceedings of the SICC 2017 Conference*, pages 197 – 203. Springer, Cham.

[Bhadoria et al., 2018] Bhadoria, R. S., Chaudhari, N. S., and Vidanagama, V. T. N. (2018). Analyzing the role of interfaces in enterprise service bus: A middleware epitome for service-oriented systems. *Computer Standards & Interfaces*, 55:146 – 155.

[Birman and Joseph, 1987] Birman, K. and Joseph, T. (1987). Exploiting virtual synchrony in distributed systems. *SIGOPS Oper. Syst. Rev.*, 21(5):123–138.

[Boepple et al., 2016] Boepple, L., Ata, R. N., Rum, R., and Thompson, J. K. (2016). Strong is the new skinny: A content analysis of fitspiration websites. *Body Image*, 17:132 – 135.

[Brodersen et al., 2015] Brodersen, K. H., Gallusser, F., Koehler, J., Remy, N., and Scott, S. L. (2015). Inferring causal impact using bayesian structural time-series models. *Annals of Applied Statistics*, 9(1):247–274.

[Chang and Lin, 2011] Chang, C.-C. and Lin, C.-J. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27. Software available at `http://www.csie.ntu.edu.tw/~cjlin/libsvm`.

[Chen et al., 2013] Chen, W., Lakshmanan, L. V. S., and Castillo, C. (2013). *Information and Influence Propagation in Social Networks*. Morgan & Claypool Publishers.

[Christakis and Fowler, 2010] Christakis, N. A. and Fowler, J. H. (2010). Social network sensors for early detection of contagious outbreaks. *PLOS ONE*, 5(9):1–8.

[Christakis and Fowler, 2011] Christakis, N. A. and Fowler, J. H. (2011). Social contagion theory: Examining dynamic social networks and human behavior. *CoRR*, abs/1109.5235.

[Culotta, 2014] Culotta, A. (2014). Estimating county health statistics with twitter. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 1335–1344, Toronto, Ontario, Canada. ACM.

[Dale et al., 2016] Dale, L. P., Dobson, R., Whittaker, R., and Maddison, R. (2016). The effectiveness of mobile-health behaviour change interventions for cardiovascular disease self-management: A systematic review. *European Journal of Preventive Cardiology*, 23(8):801–817. PMID: 26490093.

[Dey et al., 2001] Dey, A. K., Abowd, G. D., and Salber, D. (2001). A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human–Computer Interaction*, 16(2):97–166.

[Dhillon et al., 2013] Dhillon, J. S., Wünsche, B. C., and Lutteroth, C. (2013). Accessible telehealth - leveraging consumer-level technologies and social networking functionalities for senior care. In *2013 6th International Conference on Human System Interactions (HSI)*, pages 451–458, Sopot, Poland. IEEE.

[Durkheim, 1897] Durkheim, E. (1897). *Suicide: A study in sociology.* Taylor & Francis e-Library, 2005.

[Garcia-Herranz et al., 2014] Garcia-Herranz, M., Moro, E., Cebrian, M., Christakis, N. A., and Fowler, J. H. (2014). Using friends as sensors to detect global-scale contagious outbreaks. *PLOS ONE*, 9(4):1–7.

[Gertz and Culbert, 2009] Gertz, D. S. and Culbert, T. (2009). Chapter 91 - pediatric self-regulation. In Carey, W. B., Crocker, A. C., Coleman, W. L., Elias, E. R., and Feldman, H. M., editors, *Developmental-Behavioral Pediatrics (Fourth Edition)*, pages 911 – 922. W.B. Saunders, Philadelphia, fourth edition edition.

[Goyal et al., 2010] Goyal, A., Bonchi, F., and Lakshmanan, L. V. (2010). Learning influence probabilities in social networks. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, WSDM '10, pages 241–250, New York, NY, USA. ACM.

[House et al., 1988] House, J., Landis, K., and Umberson, D. (1988). Social relationships and health. *Science*, 241(4865):540–545.

[Lan et al., 2012] Lan, M., Samy, L., Alshurafa, N., Suh, M.-K., Ghasemzadeh, H., Macabasco-O'Connell, A., and Sarrafzadeh, M. (2012). WANDA: An End-to-end Remote Health Monitoring and Analytics System for Heart Failure Patients. In *Proceedings of the Conference on Wireless Health*, WH '12, pages 9:1–9:8, San Diego, California. ACM.

[Lee et al., 2015] Lee, K., Agrawal, A., and Choudhary, A. (2015). Mining social media streams to improve public health allergy surveillance. In *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015*, ASONAM '15, pages 815–822, Paris, France. ACM.

[Martin et al., 2011] Martin, C. M., Biswas, R., Joshi, A., and Sturmberg, J. (2011). *Patient Journey Record Systems (PaJR): The Development of a Conceptual Framework for a Patient Journey System*, pages 75–92. IGI Global.

[Martin et al., 2012] Martin, C. M., Vogel, C., Grady, D., Zarabzadeh, A., Hederman, L., Kellett, J., Smith, K., and O' Shea, B. (2012). Implementation of complex adaptive chronic care: the Patient Journey Record system (PaJR). *Journal of Evaluation in Clinical Practice*, 18(6):1226–1234.

[Milani and Lavie, 2016] Milani, R. V. and Lavie, C. J. (2016). Health Care 2020: Reengineering Health Care Delivery to Combat Chronic Disease. *The American Journal of Medicine*, 128(4):337–343.

[Moher, 2018] Moher, M. (2018). Chapter 33 - diabetes mellitus. In Rakel, D., editor, *Integrative Medicine (Fourth Edition)*, pages 334 – 346.e3. Elsevier, fourth edition edition.

[Mohr et al., 2017] Mohr, D. C., Tomasino, K. N., Lattie, E. G., Palac, H. L., Kwasny, M. J., Weingardt, K., Karr, C. J., Kaiser, S. M., Rossom, R. C., Bardsley, L. R., Caccamo, L., Stiles-Shields, C., and Schueller, S. M. (2017). Intellicare: An eclectic, skills-based app suite for the treatment of depression and anxiety. *J Med Internet Res*, 19(1):e10.

[OMG, 2007] OMG (2007). OMG unified modeling language (OMG UML), infrastructure, v2.1.2. Technical report, Object Management Group (OMG).

[Paul and Dredze, 2011] Paul, M. and Dredze, M. (2011). You are what you tweet: Analyzing twitter for public health.

[Philippi et al., 1999] Philippi, S., Andrea, L., Cruz, A. T., and Ribeiro, L. (1999). Adapted food pyramid: a guide for a right food choice. *Revista de Nutrição*, 12(1):65–80.

[Ram et al., 2015] Ram, S., Zhang, W., Williams, M., and Pengetnze, Y. (2015). Predicting asthma-related emergency department visits using big data. *IEEE Journal of Biomedical and Health Informatics*, 19(4):1216–1223.

[Robinson and Silge, 2017] Robinson, D. and Silge, J. (2017). *Text Mining with R - A Tidy Approach*. O'Reilly Media.

[Satyanarayanan, 2001] Satyanarayanan, M. (2001). Pervasive computing: vision and challenges. *IEEE Personal Communications*, 8(4):10–17.

[Schwartz et al., 2014] Schwartz, D. G., Bellou, A., Garcia-Castrillo, L., Muraro, A., and Papadopoulos, N. G. (2014). Towards chronic emergency response communities for anaphylaxis. In *Proceedings of the 2014 IEEE 15th International Conference on Information Reuse and Integration (IEEE IRI 2014)*, pages 98–103, Redwood City, CA, USA. IEEE.

[Sharma et al., 2017] Sharma, U., Bhadoria, R. S., and Dixit, M. (2017). Featured Analysis of Enterprise Service Bus. In *Exploring Enterprise Service Bus in the Service-Oriented Architecture Paradigm*, pages 14 – 25. IGI Global.

[Sideris et al., 2015] Sideris, C., Alshurafa, N., Kalantarian, H., Sarrafzadeh, M., and Eastwood, J.-A. (2015). Effects of coaching on adherence in remote health monitoring systems: Analysis and prediction of participant adherence. In *Proceedings of the Conference on Wireless Health*, WH '15, pages 10:1–10:8, Bethesda, Maryland. ACM.

[Simpson and Mazzeo, 2017] Simpson, C. C. and Mazzeo, S. E. (2017). Skinny is not enough: A content analysis of fitspiration on pinterest. *Health Communication*, 32(5):560–567. PMID: 27326747.

[Smahel et al., 2017] Smahel, D., Elavsky, S., and Machackova, H. (2017). Functions of mhealth applications: A user's perspective. *Health Informatics Journal*, 25(3):1065–1075. PMID: 29121831.

[Vianna and Barbosa, 2014] Vianna, H. D. and Barbosa, J. L. V. (2014). A model for ubiquitous care of noncommunicable diseases. *IEEE Journal of Biomedical and Health Informatics*, 18(5):1597–1606.

[Vianna and Barbosa, 2017] Vianna, H. D. and Barbosa, J. L. V. (2017). In search of computer-aided social support in non-communicable diseases care. *Telematics and Informatics*, 34(8). 1419-1432.

[Vianna and Barbosa, 2019] Vianna, H. D. and Barbosa, J. L. V. (2019). A scalable model for building context-aware applications for noncommunicable diseases prevention. *Information Processing Letters*, 148:1 – 6.

[Weber and Mejova, 2016] Weber, I. and Mejova, Y. (2016). Crowdsourcing health labels: Inferring body weight from profile pictures. In *Proceedings of the 6th*

International Conference on Digital Health Conference, DH '16, pages 105–109, Montréal, Québec, Canada. ACM.

[WHO, 2018a] WHO (2018a). Obesity and overweight. *Online: http://www.who.int/news-room/fact-sheets/detail/obesity-and-overweight.*

[WHO, 2018b] WHO (2018b). Physical activity and adults.

[Zhang et al., 2016] Zhang, W., Ram, S., Burkart, M., and Pengetnze, Y. (2016). Extracting signals from social media for chronic disease surveillance. In *Proceedings of the 6th International Conference on Digital Health Conference*, DH '16, pages 79–83, Montréal, Québec, Canada. ACM.