AMCIS 2022 Proceedings | SIG SEC - Information Security and Privacy

Aug 10th, 12:00 AM

Assessing the Impact of Virtual Health Communities and Environmental Characteristics of Chronic Pain Mobile Health Apps on Users' Privacy Decisions: A Multilevel Perspective

May Bantan
Saudi Electronic University, mb2627@mynsu.nova.edu

Doaa Alrefaei
WPI, Dalrefaei@wpi.edu

Raed Alharbi
Saudi electronic university, r.alharbi@ufl.edu

Assessing the Impact of Virtual Health Communities and Environmental Characteristics of Chronic Pain Mobile Health Apps on Users' Privacy Decisions: A Multilevel Perspective

Completed Research

May Bantan
Saudi Electronic University
m.bantan@seu.edu.sa

Doaa Alrefaei
Worcester Polytechnic Institute
Dalrefaei@wpi.edu

Raed Alharbi
Saudi Electronic University
r.alharbi@seu.edu.sa

Abstract

Chronic pain has been identified as one of the most widespread health-related problems. Potential chronic pain apps users seek health communities for current and previous reviews to assess the quality of the apps and make a decision regarding disclosing their information to these apps. In this study, we present a multilevel perspective on how virtual health communities and environmental characteristics of chronic pain mobile health apps impact users' privacy decisions. We used Exploratory Data Analysis and Machine Learning (ML) to operationalize the Theory of Multilevel Information Privacy. The results revealed that the most influential factors affecting users' cost-benefit analysis are Chronic Pain MHA's characteristics related to user's information privacy. The ML results indicate that the existence of information privacy policy can be predicted through the ways the apps use to Collect Data, App's Category, Country, and Store Type, which in turn affect users' decisions.

Keywords

Chronic Pain, Mobile Health App, Information Privacy, Machine Learning.

Introduction

Chronic pain has been identified as one of the most widespread health-related problems, which remains for at least three months (Merskey 1986). According to the Centers for Disease Control and Prevention (CDC), around 50 million adults of the United States population suffer from chronic pain. Providing Pain-related mobile health applications is critical to assess, evaluate, and manage users' pain. The CDC reported that the prevalence of chronic pain was 20.4%, and the prevalence of high-impact chronic pain was 7.4% (CDC 2020). Previous research pointed out the importance of pain mobile applications to help people with chronic pain in assessing and managing their pain within an out-clinic environment (Thurnheer et al., 2018). Therefore, potential chronic pain apps users seek health communities for current and previous reviewers' feedback to assess the quality of the apps and make information disclosure decision to these apps. One of the criteria they search for is the existence of privacy features and policies to protect their health information. However, many current chronic pain mobile health apps (MHA) fail to provide sufficient information to evaluate their quality in terms of the confidentiality of the information collected (Salazar et al., 2018). Current research pointed out the existence of privacy issues within chronic pain MHA caused by the absence of an information privacy policy (Terhorst et al., 2021). In addition, they mentioned that user ratings and app descriptions in the stores are insufficient to inform potential users. However, it is still unclear how virtual health communities inform chronic pain MHA users' privacy-related decisions. Therefore, this study aims to answer the following main questions: Q1: How the environmental characteristics, social identity, social norms affect the decisions of chronic pain MHA users? Q2: How the environmental characteristics and users' ratings affect the existence of privacy policy in chronic pain MHA?

The Information Systems (IS) literature highlights the need to conduct multilevel research (Burton-Jones and Gallivan 2007; Sun and Compeau 2016), especially within the context of information privacy (Bélanger et al., 2014). Therefore, we attempt to answer the research questions by following Bélanger and James's (2020) approach, which suggests that users' privacy-related decisions are processed across multiple levels. We use machine learning to operationalize the Theory of Multilevel Information Privacy (TMIP) in order to assess users' Multilevel Information Privacy Decisions (MIPD), taking into consideration virtual health communities. In the following sections, we explain TMIP within the context of chronic pain mobile health apps. Second, we discuss the methodology implemented using ML and descriptive analysis. Third, we present our empirical results, followed by the discussion. Finally, we conclude our results and present limitations and ideas for future research.

The Theory of Multilevel Information Privacy in Chronic Pain Apps

Information privacy has been mainly defined as "the desire of individuals to control or have some influence over data about themselves" (Bélanger and Crossler 2011, p. 1017). (Bélanger and James's (2020) study shows that individuals follow a multilevel approach to make their privacy-related decision, MIPD. Based on that, they introduced the Theory of Multilevel Information Privacy (TMIP), which includes five elements: the multilevel information privacy decision (MIPD) and behavior (MIPB), the offline and online environment, the salient social identity, privacy calculus, and information privacy norms (Bélanger and James 2020). In the following subsections, we describe how we implemented these five elements within the context of chronic pain MHA.

Element 1: The Multilevel Information Privacy Decision (a) and Behavior (b)

Bélanger and James (2020) defined a MIPD as the application of the salient social identity's Information Privacy Norms (IPNs) of users, which include implicit or explicit rules constructed and maintained by their individual or the group social unit. The MIPD includes two main components: (1) information to be disclosed and (2) interaction to receive this information. Consequently, MIPD will lead to MIPB, which is the action users take to share information with others influenced by the social unit's rules. In this study, we assume that users of chronic pain MHB build their MIPD and MIPB based on the social identity's IPNs of their virtual health communities (VHCs).

Element 2: The Offline and Online Environment

In their TMIP, Bélanger and James (2020) assumed that environmental characteristics have an impact on the user's social identity. In the context of pain-management MHA, we highlight the crucial role of the technological environment to guide users' privacy decisions characterized by the following components: (1) location (i.e., virtual space, country of users), (2) people (i.e., virtually present in the MHA), and (3) information characteristics: data type (e.g., health, finance); data format (e.g., text, image, audio, video); ownership (i.e., users, apps owner); (4) chronic pain MHA characteristics (i.e., store type (i.e., Google Play for Android users, IOS for Apple users), country of the App regulation, App category based on the service provided (e.g., medical, health and fitness, educational, utilities), language; (5) privacy-related issues (i.e., existence of information privacy policy, sharing data with third parties, unauthorized access to data, secondary use of data). Similarly, we expect these four key components to have the greatest impact on users’ social identity’s IPNs, subsequently influencing their MIPD to use a chronic pain MHA after running a cost-benefit analysis as shows in element 4.

Element 3: The Salient Social Identity

In the TMIP, Bélanger and James (2020) proposed that the salient social identity determines which social unit’s IPNs (personal or a group) are activated to determine the MIPD. In addition, they rely on social identity theory (SIT) and self-categorization theory (SCT) to explain how different rule sets may be enacted as individuals move between groups, where SIT explains how individuals can identify with multiple groups. Whereas SCT explains how individuals can have both personal and multiple social identities and move between these identities. Together, these theories allow us to understand how a person or group can act in terms of their salient social identity at a given point in time. The users of Chronic pain MHA can form their salient social identities based on their individual or group social units (e.g., chronic pain VHC, chronic pain App’s store, chronic pain App’s reviewers).

Element 4: Privacy Calculus

Bélanger and James (2020) indicated that individuals will possibly behave in a way that contradicts their normative MIPD (i.e., the MIPD that individual or group ought to make based on their salient social identity’s IPNs) if the estimated benefit of engaging in a behavior outweighs the estimated cost. This cost-benefit analysis (i.e., privacy calculus) has then resulted in a counter-normative MIPD. The estimation of the costs and benefits varies based on environmental characteristics. In this study, we suggest that users of chronic pain MHA will consider the components of environmental characteristics to calculate the potential risk of sharing their personal information with these apps. However, if the desired benefit of engaging with chronic pain MHA overrides the estimated cost of their information privacy risk, users will still use such an app despite their information privacy concerns. For instance, if the reviews of a chronic pain MHA indicate privacy-related issues with regards to data misuse or information privacy policy non-compliance, users are more likely to avoid using such apps based on the shared and explicit norms of VHCs. However, if users’ estimated benefit of managing their pain outweighs the subsequent cost of sharing their data, then they will decide to release this data (i.e., counter normative MIPD).

Element 5: Information Privacy Norms

It is key to understand the factors influencing the mechanism that helps individuals to have control over their personal data(Bélanger and James 2020). Norms (i.e., shared understanding of acceptable behavior; or written rules that denote acceptable behavior) have been identified as essential factors that guide individuals’ behaviors in particular circumstances. The IPNs demonstrate the shared understanding or written rulesets for how the social unit ought to manage information and interactions. Multilevel information privacy mentions that individuals or groups could formulate, regulate, and implement the norms as a way of managing their information with others. Therefore, the ruleset is the key factor that is used to control information and interaction with others. It is possible to have multiple social identities available to an individual or group, and each social identity has separate norms for information and interaction management. In this study, we assume that users’ information privacy norms are shaped by their social identity’s IPNs and the environmental characteristics, which dictate the norms for managing information and interactions in chronic pain MHAs (i.e., Information privacy norm development). We also propose that previous users’ Experiential Feedback (e.g., reviews, rating) has an impact on potential users’ disclosure decisions in chronic pain MHA, which can be changed over time due to Experiential Feedback evolvement.

Methodology

This section describes the dataset collection and its characteristics. We explain the two types of analysis adopted to answer the research questions: Exploratory Data Analysis and Machine Learning.

Dataset Collection and Characteristics

We collected user reviews for chronic pain mobile apps from the United States Google Play and Apple App Store. The search for the mobile apps was performed within the Appbot tool to extract the reviews that were posted from 2012 until January 2022. The search terms used were: "chronic pain", "chronic pain management", "back pain", "migraine", "pain scale", and "pain diary". Appopt tool automatically extracted available information about chronic pain apps from the selected app stores including the mobile app name, store information, store link, users' rating, version, user country, users' reviews, and reviews sentiment analysis. In this study, apps' categories were collected manually from the store of each app as follows: health & fitness, medical, education, social, and utilities.

To assess the eligibility of the selected apps and the reviews, the dataset was reviewed by the authors manually. Inclusion criteria were a) app's name and description are chronic pain related, b) apps available in the US Apple App Store and the US Google Play App Store, c) the reviews were written in English, and d) apps have one or more reviews. A total of n= 150 mobile apps were collected from both stores. After screening for eligibility, a total number of 85 chronic pain apps were eligible for inclusion (Apple App Store = 33, Google Play App Store = 52). The included pain apps targeted several chronic pain types. For example, some chronic pain apps focused on topics related to headache and migraine, back pain, fibromyalgia, general pain management, and assessments. The total reviews that were collected from all apps and used in the analysis are n=6881.

Data Analysis Approach

To answer the first research question, we performed an Exploratory Data Analysis (EDA) on users' reviews of Chronic pain MHA using a descriptive data analysis to explore the impact of elements 3, 4, and 5 on users' decision-making regarding their privacy. This primary investigation of the dataset is critical to discover patterns, identify anomalies, and have a full understanding of the dataset to better interpret the following Machine Learning method. In addition to the extracted reviews, an information privacy assessment was performed manually at both the app level and the store level. To start with, we manually checked the existence of the privacy policy on each app. If the app included a privacy policy, we then checked if the corresponding store's privacy policy and app permission provided enough information for users about how the app will collect, share, modify, and access their personal information (i.e., identifiers, location, app content data).

To answer the second research question, we implemented a Machine Learning framework to build a predictor of the existence (i.e., the absent or the present) of private policy in chronic pain MHA based on the features discussed in elements 2 and 5. The objective of this prediction model is to investigate the environmental characteristics and determine which environmental feature has more weight in predicting the existence of a privacy policy in chronic pain MHA.

Model Selections

We utilized three well-known classification techniques: eXtreme Gradient Boosting (XGB), Random Forest classifier (RFC), Support Vector Machine (SVM). The following is a brief description of these modeling methodologies.

- XGB produces an additive predictive model by integrating several weak predictors, typically Decision Trees (Chen and Guestrin, 2016).
- RFC is an ensemble of classifiers composed of decision trees constructed using different randomization sources(Scornet et al., 2015).

• SVM is distinguished from other classification algorithms. It determines the liner or non-linear separation in the feature space that maximizes the distance from the nearest data points of all the classes (Zhang et al., 2006).

Model Training and Hyperparameter Optimization

The dataset is divided into training and testing. The training portion is used to train prediction models, while the testing is utilized to assess the performance of each model. In our experiments, the training-testing split is 3750: 1250. By creating random permutations of the hyperparameters, we use a Random search approach. A Random with 10 rounds of 5-fold cross-validation was applied to determine the optimal hyperparameters for the model.

For the RFC model, we optimize the number of trees in the forest, the maximum depth of the tree, the minimum number of samples required to split an internal node. Table 1 contains the complete list of final hyperparameters for the three models. To evaluate if increasing the amount of the k-folds used in cross-validation during randomized search influences projected accuracy, the determined models are retrained using 10-fold cross-validation rather than 5-fold cross-validation. The predictive performance of the 5-fold and 10-fold cross-validation models are similar to each other. Thus, we conduct our final analyses on a 5-fold cross-validation model(Wong and Yeh 2020).

Model Testing and Predictive Performance

To evaluate our models, we utilize different evaluation metrics including accuracy, precision, recall, and F1 score. (1) The accuracy is a well-known metric that is used for classification models which is defined as the number of the correct prediction over the total number of predictions. (2) Recall: a classification model's ability to recognize all data sets corresponding to a certain class. (3) Precision: a classification model's ability to return exactly data points corresponding to a certain class. (4) F1 score: a single metric that utilizes the harmonic mean to combine recall and accuracy(Goutte and Gaussier 2005).

We first evaluate the performance of each prediction model on the described dataset to investigate the existence of privacy policy on the selected MHA. The evaluation of the selected models is shown in Table 1. The RFC model performs with the highest accuracy in comparison to the other approaches, as demonstrated in Table 1 below.

Model \ Evaluation Metric	Accuracy %	F1-Score %	Precision %	Recall %
XGB Classifier	93.5	94	92	97
RFC (label encoding + one-hot encoding)	**94**	**96**	**94**	**97**
RFC (one-hot encoding)	90	92	91	95
SVM Classifier	87.13	91	91	92

Table 1. Performance Comparison

Table 1 illustrates the performance comparison result for different machine learning models. RFC is the highly predictive model among the models in terms of all evaluation metrics. The second highly predictive model is the XGB classifier with an accuracy of 93.5. The gap between the top and second top models based on accuracy performance is about 5 percent. We can also observe that using the appropriate encoding schema is crucial to boost the model performance, as it is shown in Table 1, where RFC (using hybrid encoding) outperforms the other RFC (using only the one-hot encoding). The label encoding is used to encode ordinal data where the order of the feature values matters, whereas the rest is encoded using the one-hot encoding.

Additionally, the area under the receiver operating characteristic (ROC) Curve (AUC) is used to gain a better understanding of our model prediction performance, where ROC is a graph depicting a classification model's performance across various classification selected thresholds(Pepe, 2000). Figure 1 shows the plot for the ROC curve where AUC is 0.74 (Classifiers that give curves closer to the top-left corner indicate better performance.)

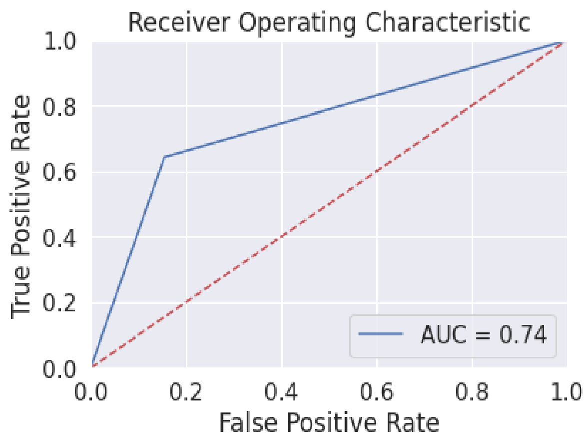

Figure 1: The ROC Curve

Explainer Selection

Machine learning algorithms explainers are used to infer the outcome behind the generated predictions(Shu et al., 2019), and the underlying reasons for privacy policy concerns, by a broad understanding of the model's behavior. It also improves model transparency, which results in a boost to model trust. Hence, we use a well-known machine learning explainer, Shapley Additive Explanations (SHAP)(Lundberg et al., n.d.). SHAP is a model-agnostic algorithm for unifying explanations for Convolutional Neural Networks (CNN) (Vu and Thai 2020). The SHAP is based on a game-theoretic approach for obtaining Shapley values that quantify the contribution of each feature to a prediction. An additive feature attribution approach is used to estimate the model's conditional expectation function's Shapley values.

Analysis and Result

Descriptive and Exploratory Approach

Our dataset contains a total of 20 out of 85 (23%) apps that don't provide a privacy policy, 48 apps out of 85 (56%) shares personal information with third parties, 17 (20%) Apps were categorized under "unknown" status where the developer either didn't provide a privacy policy or didn't specify if they share users' data with a third party. We also noticed that two apps provide different privacy policies for each app store. Individuals' privacy awareness has also increased based on the amount of posted reviews over time (period from 2013 to 2022).

Only 0.0040 out of the total reviews (n=6881) reflected on users' privacy concerns and decision-making with six major themes: privacy policy, fake and scam content, third-party engagement, users' consent, sharing personal information, and HIPAA regulations. We also extracted all records that mentioned the virtual health community and the users' engagements. Out of the collected reviews, only 10 reviews were concerned about the privacy policy and sharing personal information while using the chronic pain apps. These reviews clearly state users' unwillingness to use the app due to different factors. The following reviews present reviewers' privacy concerns in using the chronic pain MHA due to the lack of privacy settings, absence of privacy policy, secondary usage, or linkage to personal social-media accounts like Facebook or Twitter, which violates HIPAA:

""Immediately uninstalled. The only thing you get after installing is to create an account full of personal details. You can't give the app a try or even just look around at the options. No settings to ensure privacy. Can't try it this way. Pity."

"If the "privacy policy link would have actually been a real website that existed, I might have known it linked to fb!???? And they are privy to all your personal info and for what?"

"The privacy policy is troubling. I will quote from it directly: (We may use your information for market research and other marketing purposes, and for the purposes specified in this Privacy Policy. Please do not submit any personal information if you do not want it to be collected). Use the app and your private medical information is for sale. No thank you."

"Use the app and your private medical information is for sale. No thank you."

"Bugs are fixed with the weather, I am glad they got on that... Great program, I've been using it fir years... My only complaint is the lack of email support, their web page is very generic, they want you to use Facebook or twitter, I don't, my personal life is not any of this program or its creators business (not being mean, but...) when liking a Facebook page, you can see everything about me. one word, HIPPA! Other then that, excellent programming, keep it up.""

Other reviews indicate the issues regarding scam or fake content, but the users didn't reveal their decisions regarding the usage of the application. The scam reviews were mostly about the subscription cancellation process and unrealistic pain killer content as presented below:

"Basically another scam targeted at those who are in constant pain and desperate to find any kind of relief."

With regards to the virtual health community, 37 relevant records were extracted from the dataset. Users were extremely pleased with the apps' communities they are members of, which provided them with support, knowledge, sharing symptoms, and advice. The polarity of the reviews within the virtual health community section was positive, which indicates the direct influence of these communities on users' decisions, actions, and feedback. The following reviews present users' consideration of VHCs to make informed decisions about using chronic pain MHAs, with regards to the app's tracking of their data in order to provide the service.

""Overall the app is really great at tracking and giving you data/insights on your migraines. I appreciate the app trying to improve and help the community of migraine sufferers. Many thanks to Migraine Buddy team!"

"Even if you stick with the free version this app will give you a community of support, heaps of knowledge and advice as well as a superb way to keep track of your migraines. I honestly feel better prepared, both physically & psychologically, since downloading Migraine Buddy. I cannot say enough good things about this app.""

In various ways, users acknowledge the fact that chronic pain MHAs share, track, and collect data while providing the app services. The review below highlights the impact of the app's limited access to users' private information on their decisions to use the app, which reflects on users' cost and benefits decisions.

"Love this app! It is simple to use, detailed, and easy to understand my headache trends. Makes tracking my headaches less of a pain (pun intended). Plus it doesn't require access to parts of my phone info that I would rather keep private (like some other popular headache tracking apps)."

However, some users' may care more about managing the pain as an outcome than protecting their information from a secondary use by the app owner or the third party. The following review indicates that if the benefit of using the chronic pain MHA to help users in managing their pain outweighs their privacy concerns, then users might still decide to use this app.

"I have been very impressed with the information on this app so far. I am liking it a lot. Worth the money. However, be aware that they are collecting information about you and selling it to third parties. Don,t answer the personal questions and try to give them as little information as possible would be my advice. Other than that this app is filled with very useful information, and I am happy to see these kind of healing techniques becoming more well known and understood."

Machine Learning Approach

In this section, we first evaluate the models' performance to appreciate its outcome in determining the existence of a privacy policy. The SHAP explanations are presented to illustrate the most important features of the model predictions as explained below. The machine learning approach used in this paper is to detect the existence of a privacy policy, based on the features discussed in element 2 regarding the environmental characteristics including: Location (Country of the User), Information Characteristics (Data Type), MHA Characteristics (Store Type, Language), Privacy related issues (the Existence of Privacy Policy); in addition to features from element 5 regarding experimental feedback including: User Rating and its corresponding date, and reviews sentiment polarity. We selected these features based on the MIPD theory that assumed these components to have the greatest impact on users' social identities. We expect to gain more insights into the most important feature(s) determining the existence of a privacy policy in chronic pain MHAs.

Models Performance

Table 1 describes the model performance for three well-known classifiers including XGB, RFC, and SVM models. The RFC model outperforms the other two methods in terms of all evaluation measurements. The RFC performance is 94, followed by the XGB and SVM models with accuracy performance of 93.5, 87.13, accordingly. Also, we use different encoding techniques to boost the model performance as we can see for the RFC models. Using a hybrid technique is curial since one of the techniques may work better than the other based on certain features. In particular, we use label encoding to encode the data where order matters, such as the rating, whereas the one-hot encoding (a frequently used method to deal with categorical data), is used to demonstrate binary encoding for different categories when the order is not taken into consideration.

SHAP Interpretation Features

The SHAP values for each individual feature are presented in the summary plot (Figure 2). The features are sorted in order of their overall importance in making the final prediction (whether or not the chronic pain application has a privacy policy). Each line on the plot represents one specific feature, where each dot represents one user case with colors ranging from red to blue. The red and blue colors refer to high and low feature values, accordingly. The positive SHAP values (on the right side) show the impact of the prediction toward the existence of the privacy policy, whereas the negative SHAP values (on the left side) illustrate the prediction impact to the absence of the privacy policy. Figure 2 illustrates the top important factors (features) of our model based on the predefined dataset.

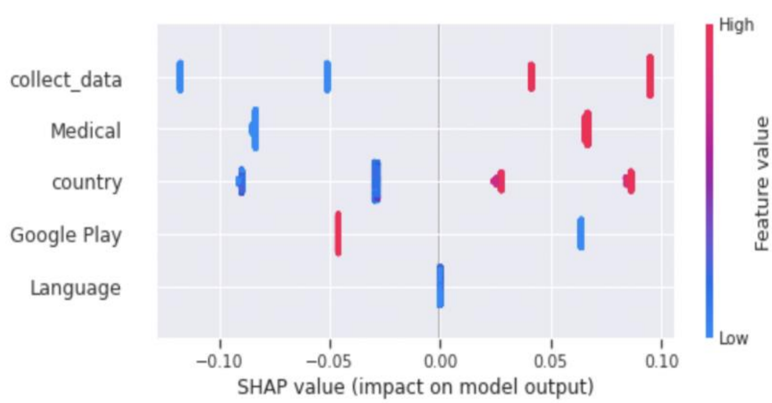

Figure 2: SHAP Values

The most influential factor to the model decision is the Data-Collection related information. The Application Category, Users' Country, and Store Type play a crucial part in influencing the final decision of the model. In particular, if the chronic pain MHA is categorized as "medical," there is a high chance that this app will have a privacy policy. However, if the category of the app is "health and fitness," "education," or "utility" (non-medical), there is a low chance that the app will have a privacy policy. Similarly, the Store Type is another important feature that has a major impact on the model. If the app is placed on the Google Play store, there is a high chance of privacy policy absence; whereas if the app is placed on the IOS store, there is a high chance of a privacy policy existence. On the other hand, Users' Language is found to have an equal influence on the model prediction.

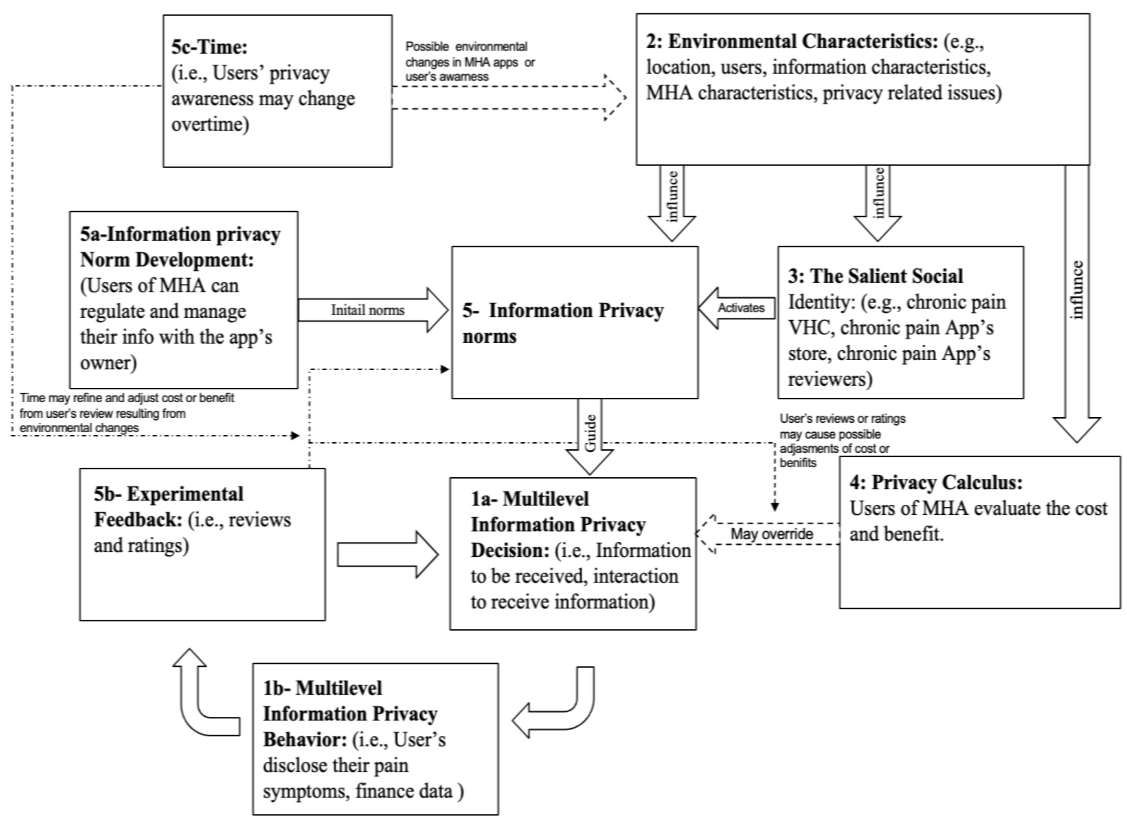

Figure 3: Theoretical Framework for the Theory of Multilevel Information Privacy in Chronic Pain Mobile Health Apps

Conclusion and Limitation

In this paper, we provided a multilevel perspective on how virtual health communities and environmental characteristics of chronic pain MHAs influence users' privacy decisions. The results of our EDA revealed that the most influential factors affecting users' cost-benefit analysis towards their privacy-related behavior and decisions are Chronic Pain MHA's characteristics related to user's information privacy. In addition, the ML results indicate that the most influencing environmental characteristics that predict the existence of information privacy policy are related to the app's ways to Collect Data, App Category, Country, and Store Type, which in turn affect users' decisions through their privacy calculus. We also found that VHC and previous users' feedback have an influence on users' decisions with regards to privacy settings, absence of privacy policy, secondary usage, linkage to personal social media accounts, which violates HIPAA, scam or fake content, and data access and tracking. This research study is limited to the uneven number of chronic pain MHAs extracted from both stores (i.e., IOS and Google Play) available in US stores only. In addition, our research focused on the sentiment analysis that was generated from the Appbot tool. To address these limitations, we suggest conducting further research taking into consideration apps operating in different countries for both app stores. We propose measuring users' reviews using Natural language processing (NLP).

REFERENCES

Bélanger, F., Cefaratti, M., Carte, T., & Markham, S. E. (2014). Multilevel Research in Information Systems: Concepts, Strategies, Problems, and Pitfalls. Journal of the Association for Information Systems, 15(9), 1.

Bélanger, F., & Crossler, R. E. (2011). Privacy in the Digital Age: A Review of Information Privacy Research in Information Systems. In Source: MIS Quarterly (Vol. 35, Issue 4).

Bélanger, F., & James, T. L. (2020). A theory of multilevel information privacy management for the digital era. Information Systems Research, 31(2), 510–536. https://doi.org/10.1287/ISRE.2019.0900.

Burton-Jones, A., and Gallivan, M. J. 2007. "Toward a Deeper Understanding of System Usage in Organizations: A Multilevel Perspective," MIS Quarterly (31:4), pp. 657-679.

CDC. (2020). Chronic Pain and High-impact Chronic Pain Among U.S. Adults, 2019. NCHS. https://www.cdc.gov/nchs/products/databriefs/db390.htm#ref1

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 13-17-August-2016, 785–794. https://doi.org/10.1145/2939672.2939785

Goutte, C., & Gaussier, E. (2005). A Probabilistic Interpretation of Precision, Recall and F-Score, with Implication for Evaluation. 345–359.

Lundberg, S. M., Allen, P. G., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. https://github.com/slundberg/shap

Merskey, H. (1986). Pain.- Classification of chronic pain- Descriptions of chronic pain syndromes and definitions of pain terms. Pain, Suppl3.

Pepe, M. S. (2000). Receiver Operating Characteristic Methodology. Journal of the American Statistical Association, 95(449), 308–311 .https://doi.org/10.1080/01621459.2000.10473930

Salazar, A., de Sola, H., Failde, I., & Moral-Munoz, J. A. (2018). Measuring the quality of mobile apps for the management of pain: Systematic search and evaluation using the mobile app rating scale. JMIR MHealth and UHealth, 6(10). https://doi.org/10.2196/10718

Scornet, E., Biau, G., & Vert, J. P. (2015). Consistency of random forests. Annals of Statistics, 43(4), 1716–1741. https://doi.org/10.1214/15-AOS1321

Shu, K., Cui, L., Wang, S., Lee, D., & Liu, H. (2019). Defend: Explainable fake news detection. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 395–405. https://doi.org/10.1145/3292500.3330935

Sun, Z., & Compeau, D. (2016). Multilevel Research in Information Systems: An Exploration of Causal Linkages (Issue 1).

Terhorst, Y., Messner, E. M., Schultchen, D., Paganini, S., Portenhauser, A., Eder, A. S., Bauer, M., Papenhoff, M., Baumeister, H., & Sander, L. B. (2021). Systematic evaluation of content and quality of English and German pain apps in European app stores. Internet Interventions, 24. https://doi.org/10.1016/J.INVENT.2021.100376

Thurnheer, S. E., Gravestock, I., Pichierri, G., Steurer, J., & Burgstaller, J. M. (2018). Benefits of mobile apps in pain management: Systematic review. In JMIR mHealth and uHealth (Vol. 6, Issue 10). JMIR Publications Inc. https://doi.org/10.2196/11231

Vu, M. N., & Thai, M. T. (2020). PGM-Explainer: Probabilistic Graphical Model Explanations for Graph Neural Networks.

Wong, T. T., & Yeh, P. Y. (2020). Reliable Accuracy Estimates from k-Fold Cross Validation. IEEE Transactions on Knowledge and Data Engineering, 32(8), 1586–1594. https://doi.org/10.1109/TKDE.2019.2912815

Zhang, H., Berg, A. C., Maire, M., & Malik, J. (2006). SVM-KNN: Discriminative nearest neighbor classification for visual category recognition. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2, 2129–2136. https://doi.org/10.1109/CVPR.2006.301