Scrutable & Persuasive Push-Notifications*

Kieran Fraser, Bilal Yousuf, and Owen Conlan

ADAPT Centre, School of Computer Science and Statistics, Trinity College Dublin, Dublin 2, Ireland
{kfraser,yousufbi,owen.conlan}@scss.tcd.ie

Abstract. Push-notifications have the potential to reinforce positive behaviours when applied in an intelligent manner. This paper explores a method of improving the delivery process of push-notifications by extracting scrutable persuasive features and refining prediction of notification outcomes. Additionally, a method is proposed for generating recommended notifications, based on the extracted persuasive features, to maximise potential engagement for scenarios such as behavioural interventions. The results illustrate that the persuasive features extracted contributed toward improved push-notification action prediction and that the personalised persuasive notifications recommended vastly increased the *Click Through Rate* (CTR) of notifications.

Keywords: Push-notifications · Synthetic data · Scrutable persuasion

1 Introduction & Related Work

In today's attention economy, an abundance of information is pushed from every direction and device, seeking to constantly engage regardless of the emotional state or health of those targeted. Push-notifications are an example of a powerful design tool used to persuade engagement [8]. The intent of a push-notification is to add timely value, however due to their inherent design and unintelligent management, notifications contribute toward adverse smartphone use and behaviours resulting in poor digital health (e.g. NoMoPhobia, FOMO). This paper investigates a means of improving the design and subsequent engagement behaviours associated with push-notifications by: 1) Extracting scrutable persuasive features from notifications; 2) Using the extracted persuasive features to improve prediction of open/dismissal of notifications; 3) Generating personalised persuasive notifications which can be used for positive behavioural change interventions.

Push-notification were a topic of recent study by Morrison et al. [7] whereby the impact of timing and frequency of notifications on user responses and subsequent use of a health-intervention app was explored. Similarly, Smith et al. [9] studied the impact of personality on choosing a persuasion type for personalised

* This research is supported by the ADAPT Centre for Digital Content Technology under the SFI Research Centres Program (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

reminders in melanoma patients. Cialidini's 6 principle's of behaviour [1] were used for identifying the suitability of a reminder in a given situation. Comparably, Thomas et al. [5], also use Cialdini's principles to craft personalised messages for encouraging healthy eating. This paper uses Cialdini's principle's in an identical manner, in the sense of generating personalised persuasive notifications, but also leverages the scrutable facets they offer for enabling transparency and explainability.

2 Method

2.1 Data Collection

A smartphone application was created to capture push-notifications in-the-wild for the purposes of identifying negative notification-engagement behavioural patterns and developing intelligent systems which could improve behaviour toward notifications. 15 participants (2 female, 13 male; all Android users) engaged with the WeAreUs app over a period of 4 months allowing for the collection of over 30,000 push-notifications as well as 291 questionnaires. The questionnaire was used to identify and verify features such as the sender of the notification (using the participants contact list) and the subject of the notification. For this study, the WeAreUs data set is limited to 11 users (male; aged 21-64) as 4 of the participants had under 100 notifications logged (due to notification settings of their device). During the study, the participants were not restricted to any particular smartphone activity (e.g. business or leisure).

2.2 Feature Engineering

The original notification features captured through the WeAreUs app are as follows:
$\{$ *app*, *category*, *priority*, *subject*, *time*, *day*, *updates*, *contactSignificantToContext*, *action (open/dismiss)* $\}$. *Dismiss* is defined as a user removing (swiping away) a notification without opening (clicking) it. On inspection, these features do not easily reveal whether a push-notification will be persuasive. An objective of this study is to enable end-users to identify persuasive facets of notifications they receive in order to promote improved self-awareness of notification-engagement behaviour and prevent addictive habits forming. Therefore, 6 features of persuasiveness (P1-P6) are derived with respect to push-notifications using Cialidini's 6 principle's as a guide and the original features as seed. Cialidini's principles are as follows: **Scarcity**: people will place higher value on something that is rare; **Authority**: people follow and respect requests made by an authority; **Reciprocity**: people feel obliged to return a favour; **Commitment and Consistency**: people tend to follow through on their word and uphold behaviours associated with their own self-image; **Liking**: people will follow what they like; **Social Proof**: people will do what they see their peers doing. These principle's are applied to push-notifications as follows (P1-P6 scores are weighted evenly, and sum to a max value of 6):

1. P1 (Authority) - a combined measure of a) priority; b) number of updates; and c) a contact's significance to a given context; indicates persuasiveness. Assumptions: *the app is considered an authority on knowing how important a notification is; an associated contact is an authority if found relevant to the context.*
2. P2 (Scarcity) - a measure of how rare a notification is, indicates persuasiveness. Assumption: *notifications which are rarely seen are more tempting to open.*
3. P3 (Liking) - a measure of previously liked feature content, taking the action 'opened' as an indicator of likeness, indicates persuasiveness. Assumption: *users are persuaded by notifications which contain content they like.*
4. P4 (Social Proof) - a measure of similar notifications opened by other users indicates persuasiveness. Assumption: *users tend to act similarly to their peers.*
5. P5 (Commitment and Consistency) - a measure of similar notifications opened by the user (essentially their habits), indicates persuasiveness. Assumption: *users tend to behave consistently with their notifications.*
6. P6 (Reciprocity) - a measure of how recently content was consumed in an app before the app sent a notification, indicates persuasiveness. Assumption: *if content was recently consumed in an app, the user acknowledges they received value and are more likely to be persuaded to open a notification from it.*

2.3 Prediction and Generation

Once extracted, the 6 persuasive principle's (P1-P6) derived with respect to push-notifications were evaluated using a selection of machine learning algorithms to ascertain the effectiveness of the facets toward predicting whether an incoming push-notification is opened or dismissed. Additionally the persuasive features were also evaluated using *Mean Decrease Impurity* (MDI) [6] for identifying the level of importance of each feature when predicting open or dismiss. The hypothesis being that the derived persuasive features should improve the performance of predicting a notification's *Click Through Rate* (CTR), as highly persuasive notifications should indicate a higher likelihood of opening a notification. Subsequently, the persuasive features should also be identified as of higher importance when attempting to predict opens/dismissals for this same reason, highlighting their value toward prediction performance.

Assuming the persuasive features are a good indication of CTR, synthetic notifications generated with a high combined sum of P1-P6 should be opened more frequently by the receiver. A conditional *Wasserstein Generative Adversarial Network with Gradient-Penalty* (WGAN-GP) [4] was used to synthetically generate push-notifications with combined P1-P6 values of between 5 and 6. The synthetic notification data was first evaluated using the *Train on Real, Test on Synthetic* (TRTS) [2] method to ensure convincing samples were being generated. Subsequently, the synthetic notifications were then classified as opened or dismissed by the selection of machine learning algorithms trained on the original notification data (simulating a real world scenario) and the results were compared against a random benchmark and the original notification data.

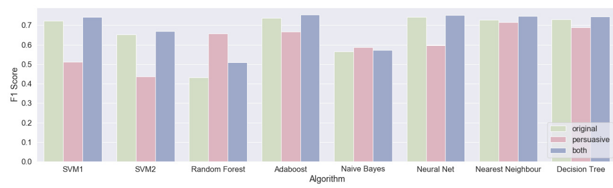

Fig. 1. F1 scores of selected algorithms when predicting notification action

3 Results and Discussion

Figure 1 illustrates the results obtained when a selection of algorithms were used to predict the open or dismissal of a notification. Three scenarios are shown for comparison. In the *original* scenario, each algorithm was trained and tested using only the original features of the notification. In the *persuasive* scenario, only the extracted persuasive features (P1-P6) were used for training and testing. The last scenario, *both*, is the union of all features used in the first two scenarios. The F1 score was taken as the metric for measuring performance due to the necessity for balancing *Precision* and *Recall*. Stratified 10-fold cross validation was used for each scenario. As can been seen from the results, the scenario in which both the original and persuasive features were used together yields best performance across all algorithms but for Naive Bayes and Random Forrest, in which the persuasive scenario performs best. While only a marginal increase in most cases over use of original features alone, this demonstrates that the persuasive features add value in performing predictions of open/dismissal of notifications.

Additionally, the persuasive features add a scrutable element as they were derived based on well defined principle's of persuasion. Therefore, a system implementing features P1-P6 could illustrate their values to end-user's when explaining automated decisions or facilitating self-reflection and steering behavioural change. For example, Figure 2 illustrates categories of notifications with associated P1, *Authority*, values. The chart highlights that notifications with categories *msg* and *reminder* generally have a high *Authority* persuasion factor. Armed with this information, user's could adjust their behaviour by ensuring they don't open notifications of this type simply because they feel they are authoritive.

Although the combination of original and persuasive features performed best, by ascertaining feature importance via MDI implemented through the Scikit-Learn *ExtraTreesClassifier*, select persuasive features were identified as most important. P3 and P6, *Liking* and *Reciprocity* respectively, were ranked highest for 8 of the users while original features such as the app and the day of the week were ranked most important in only 2.

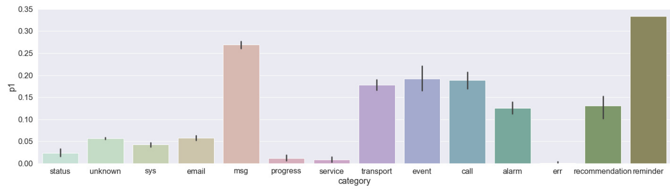

Fig. 2. P1 values (in the range 0-1) of notifications split by category

The WGAN-GP was chosen for synthetic generation as it shows enhanced training stability and enables categorical, as well as continuous, feature generation. The generated data was evaluated using TRTS, such that all algorithms were first trained on the real notification data and then tested using the generated data. By comparing the resultant F1 scores with those from Fig 1, the similarity between the synthetic and real data could be evaluated. The Root Mean Squared Error (RMSE) calculated across F1 scores of all algorithms identified the scores differing in the range of 0.02-0.07, illustrating that convincing samples were generated. Discrepancies in prediction performance using synthetic data could be attributed to a loss of nuance generated within the data, whereby the real distribution of all features was not mapped fully to the generator's latent space. For each user, 1000 notifications with a combined P1-P6 sum of between 5 and 6 were generated using the conditional WGAN-GP generator as an example of an intelligent system which can recommend highly personalised and persuasive notifications on demand. The recommended notifications were then tested using the selected algorithms which were trained on real notification data for each user. A data set of 1000 randomly generated notifications was used as a benchmark and the original notifications were also used for comparison. Figure 3 depicts the CTR of each scenario for all user's and illustrates that the synthetically generated persuasive recommendations typically have a much higher CTR than those randomly generated, or those originally sent to the user. The result of this means that persuasive notifications can be generated, based on Cialidini's principle's, such that users will open them. Notifications such as these could be used to motivate positive behavioural change at opportune moments while ensuring, through the scrutable persuasive features, that the user is aware of the type of persuasion they are subjected to.

4 Limitations, Future Work & Conclusion

Due to the intimate nature of notifications, encouraging participation in a study such as this is challenging [3], hence the dataset of 11 male users is a limiting factor of this study. Similarly, technological restrictions regarding notification monitoring set in place on iOS devices prevented those users being included in the study. However, this was used as motivation for generating synthetic notifications, a method of which is proposed in this paper. Future work will aim to

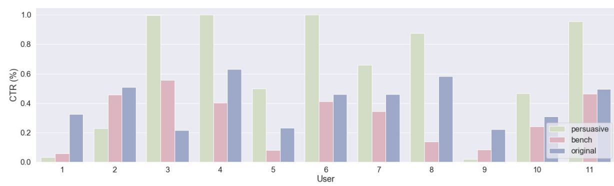

Fig. 3. Comparison of CTR of users using differing notification data sets.

improve synthetic notification generation with respect to persuasive notifications as well as extraction of additional features which indicate persuasive strategies. In addition, measuring extracted persuasiveness against that perceived by the user and identifying actions which follow is also proposed for future work.

In conclusion, the goal of this paper was to evaluate a method of extracting scrutable persuasive features from push-notifications for the purpose of improving predicted action outcomes and enabling users to reflect on the persuasive characteristics. The results illustrated that persuasive features could be extracted and visualised and that performance in predicting notification action outcomes could be improved using the persuasive features. Recommended persuasive notification could also be generated on demand and were also shown to increase the Click Through Rate (CTR) of users when simulated.

References

1. Cialdini, R.B.: Influence, vol. 3. A. Michel Port Harcourt (1987)
2. Esteban, C., Hyland, S.L., Rätsch, G.: Real-valued (medical) time series generation with recurrent conditional gans. arXiv preprint arXiv:1706.02633 (2017)
3. Fraser, K., Yousuf, B., Conlan, O.: Synthesis and evaluation of a mobile notification dataset. In: Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization (2017)
4. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of wasserstein gans. In: Advances in Neural Information Processing Systems. pp. 5767–5777 (2017)
5. Josekutty Thomas, R., Masthoff, J., Oren, N.: Personalising healthy eating messages to age, gender and personality: Using cialdini's principles and framing. In: Proceedings of the 22nd International Conference on Intelligent User Interfaces Companion
6. Louppe, G., Wehenkel, L., Sutera, A., Geurts, P.: Understanding variable importances in forests of randomized trees. In: Advances in neural information processing systems. pp. 431–439 (2013)
7. Morrison, L.G., Hargood, C., Pejovic, et al.: The effect of timing and frequency of push notifications on usage of a smartphone-based stress management intervention: An exploratory trial. PloS one (2017)
8. Oulasvirta, A., Rattenbury, T., Ma, L., Raita, E.: Habits make smartphone use more pervasive. Personal and Ubiquitous Computing **16**(1), 105–114 (2012)
9. Smith, K.A., Dennis, M., Masthoff, J.: Personalizing reminders to personality for melanoma self-checking. In: Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization. pp. 85–93. ACM (2016)