Enticing Notification Text & the Impact on Engagement

Kieran Fraser
kfraser@scss.tcd.ie
ADAPT Centre
Trinity College Dublin
Dublin, Ireland

Owen Conlan
owen.conlan@scss.tcd.ie
ADAPT Centre
Trinity College Dublin
Dublin, Ireland

ABSTRACT

Push-notifications are a design tool used by mobile and web apps to alert subscribers to new information. In recent years, due to widespread adoption of the technology and the shrinking level of user attention available, marketing techniques have been deployed to persuade subscribers to engage positively with notifications. One such technique, known as the curiosity gap, exploits Lowenstein's Information-Gap theory. This paper explores the impact of enticing notification text, instilled by the curiosity gap, on subsequent engagement actions. A classifier was defined to identify enticing language in notifications. Features commonly paired with enticing text were identified. Intelligent notification delivery agents, trained using data captured in-the-wild, were evaluated using enticing and non-enticing notifications to demonstrate the influence of enticing text. Additionally, a solution was proposed and briefly evaluated for limiting subscriber susceptibility to enticing notifications.

CCS CONCEPTS

• **Human-centered computing** → **Smartphones**; • **Computing methodologies** → **Natural language processing**; **Reinforcement learning**.

KEYWORDS

Push-notifications; Information-Gap; NLP; Reinforcement learning

ACM Reference Format:

Kieran Fraser and Owen Conlan. 2020. Enticing Notification Text & the Impact on Engagement. In *Adjunct Proceedings of the 2020 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2020 ACM International Symposium on Wearable Computers (UbiComp/ISWC '20 Adjunct)*, September 12–16, 2020, Virtual Event, Mexico. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3410530.3414430>

1 INTRODUCTION

Push-notifications were originally created to inform users they had received a new email[8]. However over the past decade, with the number of smartphone apps and ubiquitous devices capable of pushing notifications rising, the number of scenarios in which notifications are pushed has expanded. Notifications have progressed from being pushed via a single channel with a single intended purpose

(updates on email), to spanning multiple channels, each with a different purpose (e.g. social network interactions, meeting reminders, sensor readings, public transport updates etc.). The motivation behind push-notifications has also advanced. Push-notifications, once intended as a simple update, are now crafted with an intention to incite particular behaviours, such as increased platform engagement or purchasing. As such, notifications are now created with the goal of maximising positive engagement and, to achieve this, marketing techniques are implemented to purposefully influence subscribers to open notifications and improve marketing campaign metrics, such as Click-Through-Rates (CTRs). The CTR, in this work, is defined as the sum total of opened notifications divided by the sum total of notifications pushed. Whilst it is a commonly tracked metric of push success, it does lack nuance with respect to certain notification engagement scenarios (e.g. notifications positively received but dismissed as they fulfilled their intended purpose). One such marketing technique leveraged for improved notification engagement is the *Curiosity Gap*. The Curiosity Gap exploits Loewenstein's Information-Gap theory [5] which suggests that curiosity peaks when a small amount of information is revealed, but key parts are withheld. An example of a notification title exploiting the curiosity gap might be: *Business Schools That Will Dominate In The 21st Century*. In this text title, key information, the names of the Business Schools, are withheld to entice subscribers to open the notification and find the answer. Similarly, a common Facebook notification: *You have a new friend suggestion*, also exploits the curiosity gap, as key information, the persons name, is omitted from the notification tempting users to open and engage with the Facebook platform to find the answer. Indeed, Kang et al. [4] illustrate, through experimentation with trivia questions, that curiosity due to information-gaps can elicit an inclination to pay for missing information, such is the power of its influence.

This paper leverages existing data sets comprised of both clickbait and non-clickbait news article headlines to create a classifier for identifying enticing notification text content. Potthast et al. [12] define *Clickbait* as "a certain kind of web content advertisement that is designed to entice its readers into clicking an accompanying link. Typically it is spread on social media in the form of short teaser messages". Additionally, Chakraborty et al. [1] further define clickbait headlines as exploiting the curiosity gap, enabling attention distraction and contributing to readers cognitive overload. Clickbait implies the *teaser message* does not accurately represent the content within, whereas the curiosity gap can be applied while still accurately representing content, albeit with missing information in the headline. Therefore, for the purposes of identifying enticing notifications exploiting the curiosity gap, the corpus of clickbait titles were used for identifying notifications with text similar to those found in clickbait headlines.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UbiComp/ISWC '20 Adjunct, September 12–16, 2020, Virtual Event, Mexico

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8076-8/20/09...\$15.00

<https://doi.org/10.1145/3410530.3414430>

This paper is structured as follows: section 2 briefly discusses related work in the areas of notification management and identification of enticing text; section 3 analyses notification text enticement for a data set procured in-the-wild and evaluates a classifier for identifying enticement; section 4 further explores notification delivery using the same data set and evaluates a RL method of managing notifications; section 5 describes a second notification data set used for evaluating a summarisation method as a potential solution to enticement in notification content; section 6 briefly discusses limitations and future directions of the work and finally, section 7 presents a summary of the paper.

2 RELATED WORK

A Do Not Disturb Challenge was reported by Pielot et al. [11] whereby 30 volunteers were asked to go without all device notifications for 24 hours. The results found that notifications had a negative impact on work which required focus as, when disabled, participants were more productive and less distracted. However, it was also found that anxiety levels in participants increased throughout the study due to fear of missing out on important information. This illustrates a clear need for improving our current relationship with notifications but also highlights that jumping to complete abstinence also has negative effects. Reducing the enticement found in the text of notifications, which this work explores, is one method of potentially reducing the impulse to read and open all notifications. The creation of an intelligent agent which manages notifications on behalf of a user could also alleviate the anxiety they induce.

Similarly, in a design experiment, Hiniker et al. [3] created an Android app for encouraging improved time-limited behaviour with smartphones through monitoring app usage and alerting users when self-appointed limits were up. The objective being to aid users achieve goals for smartphone non-use and reclaim time spent. Again, the study found that participants wanted to change their unhealthy behaviour with technology with 58% admitting they wanted to spend a little less time on their smartphone. Through a survey, a small corpus of behavioural change desires was also formed which describe how users wish to change. The app itself employed various techniques to encourage the user to achieve their self-set usage goals such as showing the current amount of time they have spent in apps already. In comparison, this work attempts to instill improved behaviour in the intelligent agents which manage user notifications to reinforce better engagement behaviour with notifications, which would subsequently reduce time spent on devices. This is achieved by reducing the bias found toward opening enticing notifications found in habitual user data and allowing other more contextual features dictate whether a notification is valuable and healthy for a user at a given moment.

Mehrotra et. al [7] have also investigated mining an association rule set for notification preferences. PrefMiner focuses on interruptibility management and learns user preferences for receiving notifications using intelligible rules for user consumption. The solution has a very high performance in predicting notification dismissal, but most importantly, it also gives the end-user insight into how their notifications are being managed by the intelligent system. Transparency such as this instills trust between user and machine and is a necessity for widespread adoption. Solutions such as this

are needed now, more than ever, as notification creators employ increasingly intelligent tools for influencing behaviour. For example, Phrassee¹ is a tool which creates textual content, such as email subject lines, primed for maximising open rates. Other websites also advertise generative text capabilities for creating clickbait-like headlines². Indeed, OpenAI expressed reservations toward releasing their generative model for fear of misuse.

3 NOTIFICATION ENTICEMENT

This section briefly introduces a notification data set captured in-the-wild followed by the creation of a classifier for identifying enticing notifications. Enticing notifications are defined in this paper as *notifications with text similar to article headlines identified as clickbait, indicating potential exploitation of the curiosity gap*. On applying the classifier to the notification data set, notifications and subsequent user actions were evaluated and insights drawn.

In 2018, an in-situ study of mobile notifications in-the-wild was carried out with 18 participants between March and August. The WeAreUs app logged incoming notifications and actions taken on them. The resultant data set has over 35,314 notifications and includes participant engagement actions (open or dismiss) toward notifications as well as text extracted from the content of each notification. Of the 18 participants, 7 had little engagement with the study and/or incomplete notification text data available which could be used for this work. They have been excluded from the analysis. Of the remaining 11 participants, on average, each received ≈ 40 notifications per day, $\approx 52\%$ of which were opened, as illustrated by the median Click-Through-Rate (CTR) in Figure 1.

Figure 1: Median CTR of notifications per user.

There are many features which influence the action of a user toward opening or dismissing a push-notification - the app which posted the notification; the location of the user on receiving the message; the category of notification - previous work has explored these features in detail [2]. This paper focuses on the textual elements of notifications and, in particular, the potentially adversely enticing aspects of the text which manipulate participants to act positively (open) toward a notification, regardless of contextual value.

For the purpose of preserving the privacy of participants in the WeAreUs study, the full title, ticker and content texts of the notification were not recorded, but were instead used to create a privacy-compliant text feature as follows:

¹<https://phrassee.co/product/>

²<https://www.contentrow.com/tools/link-bait-title-generator>

Figure 2: Notification enticement classifier accuracy and loss

Table 1: Clickbait and Non-clickbait examples

Clickbait	Text
Yes	19 Hilarious Jokes About Being An Intellectual That Will Make You..
Yes	Business Schools That Will Dominate In The 21st Century..
No	Coldplay's new album hits stores worldwide this week
No	James Harden is Our NBA MVP at the All-Star Break..

- (1) The title, ticker and content text elements of each notification were combined
- (2) The topic of the notification was inferred and recorded
- (3) Named Entity Recognition was performed to tag tokens pertaining to places, people etc.
- (4) Personally Identifiable Information such as people's names, places, email addresses and phone numbers were removed
- (5) Keywords were extracted from the remaining text and duplicate bigrams were removed (as notification title, ticker and content texts regularly contain duplicate phrases)

The resultant notification text feature is made up of keywords representing the title, ticker and content text elements of the notification. The sentiment of notification text was inferred using the TextBlob [6] python package and an analysis indicated that, for the majority of notifications, the sentiment polarity of text was shifted positively.

3.1 Enticement Classifier Creation

This section details the creation of an enticement classifier for detecting notifications which are similar to clickbait article headlines exploiting the curiosity gap. Potthast et al. [12] hosted the *Clickbait Challenge* in 2017 to enable the comparable evaluation of clickbait detectors and released a data set comprised of annotated clickbait examples as part of their challenge. Chakraborty et al. [1] gathered a data set comprised of both clickbait and non-clickbait samples from sources such as *Wikinews* and *BuzzFeed*. For this work, the annotated clickbait data sets of Potthast et al. and Chakraborty et al. were chosen to train a classifier for identifying enticing notifications. The combined data set was split into a training set of 47,538 clickbait/non-clickbait samples and a test set of 4,000 samples. An example of clickbait and non-clickbait samples are illustrated in Table 1.

A Multilayer Bidirectional RNN with two stacked GRU layers was chosen as a model for classifying notification text as enticing. An embedding layer was also added and initialised with GloVe's [10] 100 dimension pre-trained word vectors³. The model was trained using a batch size of 128 for 5 epochs, reaching an accuracy of 89.2%. The training and validation accuracy and loss are illustrated in Figure 2.

Subsequently, the enticement classifier was then applied to the notification texts from the WeAreUs data set to investigate the distribution and impact of enticing text found in differing types of push-notification. Figure 3a illustrates that notifications with shorter character lengths are more enticing than notifications with larger amounts of text. With respect to sentiment, enticing notifications included more extreme positive and negative polarity values than non-enticing notifications (Figure 3d). With regard to the text content, Figures 3c and 3f illustrate the most frequent unigram and trigrams of notifications classified as enticing. Social media phrases seem common in text classified as having clickbait-like enticement: *highlights tweet look*, *message messages chats*. It is also worth considering other notification features with respect to enticement, such as the apps and topics from which the majority of enticing text, which exploits the curiosity gap, was generated. Figures 3e and 3b illustrate the distribution of notifications classified as enticing across apps and topics respectively. Facebook's *Messenger* app was the largest generator of enticing notifications followed by *YouTube* and *WhatsApp*. Topics which produced the highest number of enticing notifications were *comedy*, *entertainment* and *healthy living*.

The actions taken toward notifications by the 11 participants were also analysed. Notifications with text content classified as enticing had higher CTRs among participants (67.51%) compared with the CTR of notifications classified as non-enticing (48.44%). Of the 11 analysed, the enticing notifications elicited an increased CTR in 10 participants, with 1 outstanding participant having a lower CTR for enticing notifications. An exact sign test was used to compare the differences in CTR between enticing and non-enticing notifications. The enticing notifications elicited a statistically significant median increase in CTR (15.59%) compared to the non-enticing notifications, $p = .012$.

The utility of the enticement feature in contributing to predictions of participant engagement actions toward notifications was also assessed. Permutation feature importance was used to identify the contribution made by the enticement feature compared to other notification features commonly used for predicting notification actions (open/dismiss). The permutation feature importance algorithm implemented in the scikit-learn python library was used to identify importance values. Feature importance values were extracted for each of the 11 participants and the mean was calculated to create a holistic ranking of features. The AdaBoost classifier, implemented in scikit-learn, was the estimator of choice for predicting open/dismiss actions. Figure 4 illustrates the final ranking of notification feature importance, in which the enticement feature ranks 4th. Interestingly, the enticement feature ranks above others which might be considered more indicative of notification engagement behaviour such as *category* and *timeOfDay*. Understandably, features such as the *appPackage* and *dayOfWeek* rank highest.

³<https://nlp.stanford.edu/projects/glove/>

(48%) of notifications delivered in the WeAreUs study were dismissed (as illustrated in Figure 1). Using the notification text as input and the notification action as ground-truth, a RL agent was trained to predict whether a participant would open or dismiss a notification based on its textual content. A Deep Q-learning (DON)

Whilst the RL agents show potential for replicating participant actions taken toward notifications, this may not be a desired outcome. The agents could be learning poor engagement habits inherent in a participant’s previous behaviour with respect to notifications. For example, a participant may be susceptible to enticing notification text and therefore, an agent, trained using their past notification engagements, would also be susceptible, thus reinforcing poorer habits as opposed to improving and encouraging better notification-engagement behaviour. To test this hypothesis, the trained agents were tested using two data sets. The first was comprised only of notifications classified as enticing and the second only of non-enticing notifications. The agents predicted actions

of open or dismiss and the subsequent CTR for each agent was recorded. The results, illustrated in Figure 5, demonstrate that the median CTR for enticing notifications was much larger than that of non-enticing notifications. Therefore the intelligent agents, trained on real participant notification data, were just as susceptible to enticing notifications.

Figure 5: Intelligent notification delivery agent actions (open/dismiss) taken on enticing and non-enticing notifications

4 big tech trends for 2021
Hindustan Times
Do you know which tech trends you should be looking out for in 2021? We reveal our top 4 picks here.
technology, cloud computing, Artificial Intelligence ...

Figure 6: Example of Google Alert notification

5 NOTIFICATION SUMMARISATION

This section introduces a second notification data set which includes a richer source of text data to evaluate enticement. A summarisation method was also explored as a potential solution to reduce the influence of enticing notification text. Evaluations were performed in a simulated environment where agents, trained using the WeAreUs notification data captured in-the-wild, simulated user actions of opening or dismissing notifications.

In the previous sections, push-notification text was analysed for enticement and its subsequent impact on notification engagement actions were discussed. However, due to the privacy-preserving process executed when gathering data during the WeAreUs study, the quality of the notification text available was a limiting factor. This section details an alternative push-notification data set which includes higher quality textual features: the *Alert* data set.

Google Alerts is a content change detection and notification service enabling users to subscribe and receive notifications of new stories published to the web based on search terms. For this study, a number of generic search terms were subscribed to and the generated notifications, made up of titles and content (illustrated in Figure 6), were recorded in real-time. The set of terms subscribed to were: {*breaking*, *entertainment*, *health*, *beauty*, *style*, *science*, *sport*, *tech*, *world*} and the subsequent data set generated was comprised of 35,841 notifications. An extractive summariser [9], which leverages BERT token embeddings and implements a clustering algorithm, was used to extract summaries from the content feature of the notifications. The hypothesis was that unbiased summaries of the content could be used in place of the enticing titles which were exploiting the curiosity gap, in order to reduce the persuasive nature

of the notifications and provide subscribers with adequate detail to make informed actions on the notification. Values for enticement and sentiment were inferred for both notification titles and summaries using the same models described in previous sections.

Figures 7a and 7d illustrate the top 20 bigrams extracted from titles classified as enticing and their corresponding summaries respectively. The top bigrams for the titles included values such as: *need know*, *start day*, *love island*, *things need*, *know start* - which correspond to clickbait-like titles such as *5 Things You Need to Know* or *Start Your Day with*. In contrast, the top summary bigrams did not include these terms and were much less sensational in appearance. For example, the title bigrams included abbreviations (e.g. *man utd*) while the summaries did not (e.g. *manchester united*). Similar results were found in the extracted trigrams illustrated in Figures 7b and 7e. The enticing notification titles included sequences such as: *things need know*, *need know today*, *things tech need*, *10 things tech*, which were also similar to clickbait-like text exploiting the curiosity gap.

Figures 7c and 7f illustrate how enticement and sentiment were distributed among notification topics. Similar to the analysis of previous sections, *healthy living* and *comedy* were among the top topics generating enticing notifications. The sentiment across all topics was evenly spread, but for crime which, understandably, was negatively oriented.

To ascertain the effects of enticing *Alert* notifications on user engagement, the DQN agents trained in the previous section were used for evaluation. The *Alert* notification titles and summaries were converted to vectors using the pre-trained BERT-base model of the sentence-transformers library [13]. The agents were then given each notification text type (title and summary) in turn and tasked with predicting user actions (open or dismiss). Figure 8 illustrates the CTR of each participant's DQN agent when given only enticing notifications. The median CTR for notifications when the agents were given notification text summaries was 48%. In contrast, the median CTR was 57% when the agents were given the notification titles. Therefore, summarising the content of notifications, and using the summaries as a feature for managing notifications, could improve notification management and reinforce healthier engagements with notifications.

6 LIMITATIONS & FUTURE WORK

A limitation of this work included the granularity of notification text data available. The WeAreUs data set provided textual notification content, however due to disclosure restrictions the quality of the text was greatly reduced thus hindering the effectiveness of applying classification methods. The Alert data set provided a much richer source of text content, however, this data set did not include end-user notification engagement features indicating how a person would react (open or dismiss) to a notification. Simulated users, in the form of DQN agents, were used as proxy for evaluation purposes. Future work will include recruiting a larger user base to test the proposed methods of identifying and mediating the impact of enticing notifications. Additionally, further work will also address the autonomous creation of notification text and whether enticing notifications can be translated in-real-time so as to minimise their impact on the end-user.

Figure 7: Analysis of enticing notification text found in the Alert data set

Figure 8: DQN notification CTRs for enticing titles (red) and summarised enticing notifications (green)

7 CONCLUSION

The goal of this paper was to explore the prevalence of enticing text (exploiting the curiosity gap) in push-notifications and its impact on notification engagement. Through an analysis of notification data collected in-the-wild, the CTR was identified to be statistically significantly larger for notifications with text classified as enticing than for those classified as not enticing. Additionally, text enticement was identified as an important feature when predicting notification engagement action. Finally, a summarisation method was proposed and evaluated, in a simulated environment, as a solution for reducing the influence of enticing notification text on subscribers. Results showed that the CTR was reduced when enticing notification text content was summarised.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] Abhijnan Chakraborty, Bhargavi Paranjape, Sourya Kakarla, and Niloy Ganguly. 2016. Stop clickbait: Detecting and preventing clickbaits in online news media. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 9–16.
- [2] Kieran Fraser, Bilal Yousuf, and Owen Conlan. 2019. Scrutable and persuasive push-notifications. In *International Conference on Persuasive Technology*. Springer, 67–73.
- [3] Alexis Hiniker, Tadayoshi Kohno, Julie Kientz, and Ray Hong. 2016. MyTime: Designing and Evaluating an Intervention for Smartphone Non-Use. (05 2016).
- [4] Min Jeong Kang, Ming Hsu, Ian M Krajbich, George Loewenstein, Samuel M McClure, Joseph Tao-yi Wang, and Colin F Camerer. 2009. The tick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological science* 20, 8 (2009), 963–973.
- [5] George Loewenstein. 1994. The psychology of curiosity: A review and reinterpretation. *Psychological bulletin* 116, 1 (1994), 75.
- [6] Steven Loria. 2018. *textblob Documentation*. Technical Report.
- [7] Abhinav Mehrotra, Robert Hendley, and Mirco Musolesi. 2017. Interpretable machine learning for mobile notification management: An overview of preminer. *GetMobile: Mobile Computing and Communications* 21, 2 (2017), 35–38.
- [8] Catherine A Middleton and Wendy Cukier. 2006. Is mobile email functional or dysfunctional? Two perspectives on mobile email usage. *European Journal of Information Systems* 15, 3 (2006), 252–260.
- [9] Derek Miller. 2019. *bert-extractive-summarizer*. Retrieved July 6, 2020 from <https://github.com/dnmiller612/bert-extractive-summarizer>
- [10] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*. 1532–1543.
- [11] Martin Pielot and Luz Rello. 2017. Productive, anxious, lonely: 24 hours without push notifications. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*. 1–11.
- [12] Martin Potthast, Tim Gollub, Matthias Hagen, and Benno Stein. 2017. *The clickbait challenge 2017*. Retrieved July 6, 2020 from <https://www.clickbait-challenge.org/>
- [13] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- [14] Rowan Sutton, Kieran Fraser, and Owen Conlan. 2019. A Reinforcement Learning and Synthetic Data Approach to Mobile Notification Management. In *Proceedings of the 17th International Conference on Advances in Mobile Computing Multimedia (MoMM2019)*. Association for Computing Machinery.