January 23, 2018

Selective notifications based on message priority determination

Annika Matta

Hannah Keiler

Carmen Klett-Wilkinson

Jacob Bank

Timothy Sohn

See next page for additional authors

Follow this and additional works at: http://www.tdcommons.org/dpubs_series

Inventor(s)

Annika Matta, Hannah Keiler, Carmen Klett-Wilkinson, Jacob Bank, Timothy Sohn, Kaushik Roy, and Yi Zhou

Selective notifications based on message priority determination

ABSTRACT

Users of computing devices, e.g., smartphones, wearables, tablets, personal computers, etc. that choose to be notified about incoming messages, e.g., email are often overwhelmed with notifications. Currently, there is no differentiation between notifications for high priority and low priority messages. Too many notifications can overwhelm or distract users. The techniques of this disclosure selectively provide notifications about email messages that are determined to be of priority or importance. The techniques utilize a trained machine-learning model that determines notification thresholds based on user preferences and user-permitted factors. Implementation of these techniques can reduce interruptions and distractions.

KEYWORDS

- Notifications

- High priority

- Interruption

- Alert

- Incoming email

BACKGROUND

Users of computing devices, e.g., smartphones, wearables, tablets, personal computers, etc. that choose to be notified about incoming email are often overwhelmed with such notifications. Currently, such notifications do not differentiate between high-priority and lower priority emails. More notifications can lead to users checking the email application more and can sometimes overwhelm or distract users.

DESCRIPTION

This disclosure describes techniques to selectively provide notifications only when high-priority email messages are received, e.g., emails that are determined to be urgent, to require a reply, or to be deemed important to the user. User permission is obtained to use one or more factors to determine the priority of an incoming message. Such factors can include, e.g., contents of the email message, user data such as emails (sent/ received), meeting, and calendar data. Only such factors as permitted by the user are utilized, and use of such data is specifically for determination of the priority of an incoming message.
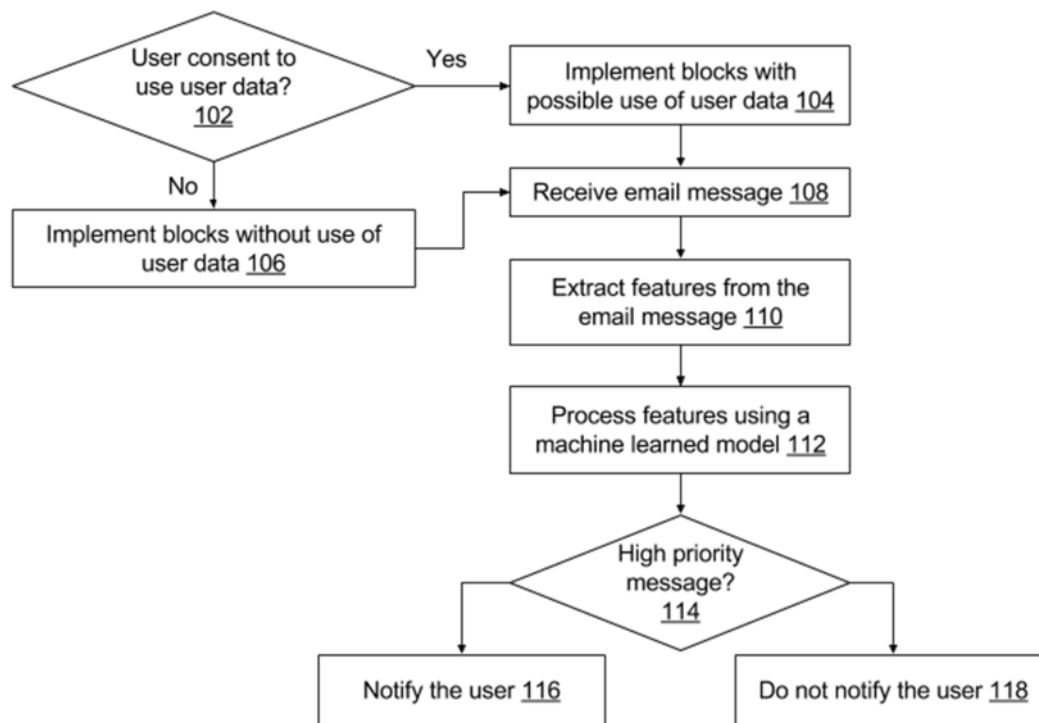

Fig. 1: Selective notification for email

Fig. 1 illustrates an example process to predict whether a notification of an incoming email message is to be provided to a user. A determination is made (102) as to whether user consent (e.g., user permission) has been obtained to use user data in the implementation of the process. For example, such user data can include user data from an email application, a calendar application, etc.

The described techniques may use such user data in some implementations. If user consent has been obtained from relevant users for which user data may be used in these techniques, then a determination is made that the techniques can be implemented with possible use of user data (104). If user consent has not been obtained, then a determination is made that the techniques are to be implemented without use of user data (106). If user consent has not been obtained, the techniques are implemented without use of user data.

Upon receiving an email message (108), features are extracted (110) from the email message. The extracted features are analyzed by a classifier (112) to determine the probability that the user will open or reply to the message. The classifier can be implemented using a machine learned model, e.g., a deep neural network, to determine the probability. If the probability meets a predetermined threshold, then the message is determined to be of high priority (114) and the user is notified (116). If the probability does not meet the threshold, the email is determined to be of low priority and notification is not provided to the user (118).

The set of features extracted for a message, determined with user permission, includes an affinity score that represents, a frequency and extent of interactions between the sender and recipient of the message. The affinity score is calculated using user-permitted data such as calendar events and meeting invites that include the sender and recipient of the email and previous communications between the sender and recipient. Other features extracted with user consent can include bigrams and trigrams from the email body and subject and information relating to the email thread. The email thread is analyzed to determine the number of messages in the thread and whether the user replied to or authored the messages. The determined features are inputs to the machine learned model.

The model is configured to determine if there is a high probability of a user opening or replying to an email message. A single model can be applied to categorize different messages across different groups of users. If users consent, the model is personalized based on permitted user-specific factors such as affinity scores, the frequency and number of notifications received, etc. Other factors can include the number of messages opened by users in response to notifications and the number of messages that the user replied to post-notification. The threshold probability for notification can be determined based on such user-permitted factors.

The techniques of this disclosure selectively provide notifications about email messages only upon determining that the messages are of priority or importance. The techniques do not require users to perform manual configuration and utilize a trained machine-learning model that determines e notification thresholds based on user preferences and user-permitted factors. Implementation of these techniques can reduce interruptions and distractions. Rule-based notifications where the rules are explicitly configured by users can be used as an alternative to or in addition to the techniques described herein.

Further to the descriptions above, a user may be provided with controls allowing the user to make an election as to both if and when systems, programs or features described herein may enable collection of user information (e.g., information about a user's social network, social actions or activities, profession, a user's preferences, or a user's current location), and if the user is sent content or communications from a server. In addition, certain data may be treated in one or more ways before it is stored or used, so that personally identifiable information is removed. For example, a user's identity may be treated so that no personally identifiable information can be determined for the user, or a user's geographic location may be generalized where location information is obtained (such as to a city, ZIP code, or state level), so that a particular location of

a user cannot be determined. Thus, the user may have control over what information is collected about the user, how that information is used, and what information is provided to the user.

CONCLUSION

The techniques of this disclosure selectively provide notifications about email messages that are determined to be of priority or importance. The techniques utilize a trained machine-learning model that determines notification thresholds based on user preferences and user-permitted factors. Implementation of these techniques can reduce interruptions and distractions. Rule-based notifications where the rules are explicitly configured by users can be used as an alternative to or in addition to the techniques described herein.