Technical Disclosure Commons

August 26, 2019

Triggering context-appropriate reminders dependent on user activities

Victor Cărbune

Andrii Maksai

Sandro Feuz

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

Triggering context-appropriate reminders dependent on user activities

ABSTRACT

Users often set reminders to be alerted for performing specific tasks. Such reminder functionality often deals with tasks connected to specific times and/or places. However, some tasks, such as those that require planning based on dependencies between multiple actions, require a more nuanced understanding of the context. This disclosure presents techniques to trigger a reminder for a given action that is dependent on another action that will be performed by the user. The techniques infer a user's future actions with permission from the user. Alternatively, or in addition, if the user permits, the user's likely future activities can be inferred from other relevant contextual information or specified manually by the user. Subsequently, the inferred and/or user-specified future user activities with a sufficiently high likelihood of occurring are used to trigger reminders for actions that are connected to and/or dependent on the inferred future activities.

KEYWORDS

- Reminder

- Personalization

- Context-based reminder

- User activities

- Voice assistant

- User feedback

- Interdependent activities

- Machine learning

BACKGROUND

Users often set reminders to be alerted for performing specific tasks. Such reminders can be set by explicit user input, e.g., via typed text, selection of user interface (UI) elements such as buttons and dropdowns, voice commands, etc. Voice assistants are often used to trigger reminders for various tasks. Such assistants typically include functionality to extract intended reminder triggers from the received voice input. Such reminder functionality often deals with tasks connected to specific times and/or places. For instance, a user may choose to be reminded to purchase milk at the grocery store, pay rent on the first of every month, clean the bedroom when at home on a weekend, etc.

However, some tasks, such as those that require planning based on dependencies between multiple actions, require a more nuanced understanding of the context beyond simple time and location. For instance, a user may wish to be reminded to take the gym bag when leaving home in the morning but only on the days on which the user plans to go to the gym. If the user's gym schedule is flexible and variable, setting reminders based on time and/or location can be ineffective and/or cumbersome. Reminders triggered in the morning at home may be too frequent with many being unnecessary if the user is not going to the gym that day. On the other hand, reminder triggered at the gym are too late as the user may have already left home without the gym bag. Moreover, reminders set manually once a gym visit is known can be cumbersome and inflexible due to their manual nature. Reminders based on time and/or location are similarly suboptimal for various other situations encountered commonly in everyday life such as taking recycling materials to a recycling center before going to a shop near the center, taking specific items to be returned to a store prior to leaving home for that store, etc.

DESCRIPTION

This disclosure presents techniques to trigger a reminder for a given action that is dependent on some other action that will be performed by the user. Specifically, the described techniques remind a user to perform an action X when the user is planning to do another action Y that is connected to action X in some way. For example, the user can be reminded to take the gym bag prior to leaving home on days when the user plans to go to the gym, take the recycling before leaving for a location near the recycling center, buy an online ticket before an upcoming train journey, etc.

To this end, the techniques rely on inferring a user's future actions with permission from the user. The action inference can be based on one or more of a variety of approaches. For instance, if the user permits, a trained machine learning model can be used to predict a user's likely future actions based on a variety of user-permitted factors such as the user's current location, places visited by the user in the past, events in the user's calendar, history of the user's digital activities such as searches on the web or in maps, use of specific application, etc. Based on such factors, the model can learn and predict the probability of the user performing a specific activity in the near future. For instance, entries for a gym visit in the user's calendar can be used to detect the days on which the user is planning to visit the gym. Similarly, if the user permits, the user's intended activities can be predicted based on the user's online activities such as checking a store's opening hours on the web, looking up directions to a location on the map, etc.

Alternatively, or in addition, if the user permits, the user's likely future activities can be inferred from other relevant contextual information such as the time of the day, the day of the week, the user's physical activity and/or movement, etc. Further, a user can manually specify upcoming future actions via appropriate mechanisms which can include keyboard or voice input. For instance, a user can declare: "I am about to leave for the gym."

The inferred and/or user-specified future user activities with a sufficiently high likelihood of occurring are used to trigger reminders for actions that are connected to and/or dependent on the inferred future activities. For instance, if it is inferred that the user is likely to visit the gym on a given day, the user can be reminded to take the gym back prior to leaving home in the morning that day. Whether the likelihood of an inferred activity occurring is sufficiently high is based on comparing it with a threshold value that must be met. The threshold value can be specified by the developers and/or adjusted by the users and/or dynamically determined depending on the context.
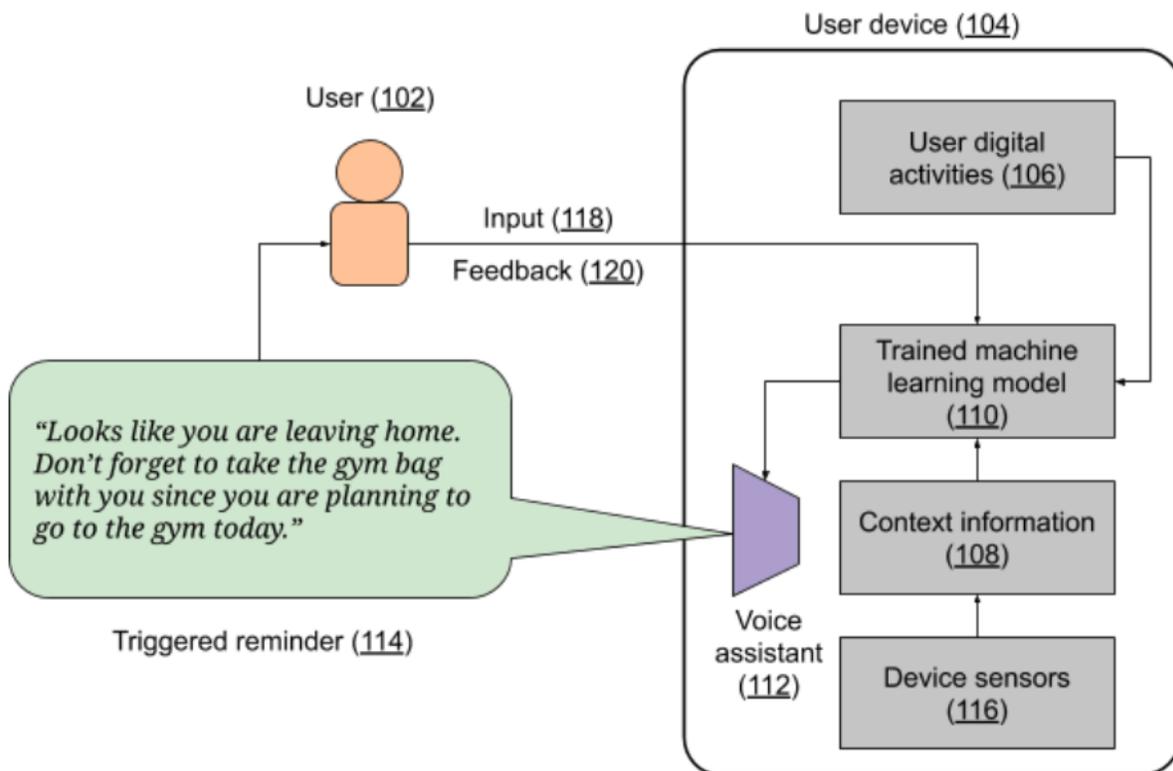

Fig. 1: Triggering reminders based on inferred future activities

Fig. 1 shows an operational implementation of the techniques described in this disclosure. A user (102) uses a device (104) for various digital activities (106). While the techniques are described with reference to a single device, it is possible that user information is obtained from

multiple devices (when permitted by the users) and reminders are triggered on particular devices as appropriate. For example, user devices such as smartphone, tablet, wearable device, laptop, smart speaker, etc. can be used to obtain user information and to provide appropriate reminders.

With the user's permission, the user's digital activities as well as other context information (108) are provided as input to a trained machine learning model (110). The context information can be obtained from the device sensors (116) and/or other external sources. Further, the user can provide explicit manual input (118) to the trained machine learning model about the user's activities and/or desired reminders. The output of the machine learning model indicates whether a reminder connected to the user's inferred upcoming activities is to be triggered.

The triggered reminder (114) is then delivered to the user, e.g., via a voice assistant (112) or other mechanisms, e.g., a notification, a text message, etc. For instance, Fig. 1 shows that the user is being reminded to take the gym back prior to leaving home on the morning of a day on which the user plans to visit the gym. The voice assistant hardware and/or software can be separate from the user device and/or be included within the user device.

The model can be further adjusted and personalized by incorporating user feedback as part of the training data as appropriate. The user can be provided with mechanisms for responding to the triggered reminders in various ways. For instance, the user can confirm the usefulness of a reminder via mechanisms such as a button for indicating approval, voice input suggesting appreciation, etc. Alternatively, the user can use similar mechanisms to indicate that the reminder was irrelevant or unneeded with a corresponding reason, if applicable. For instance, a user can respond to an inappropriate gym-related reminder triggered on a weekend by informing the system: "I never go to the gym on Saturdays."

The described techniques can be extended by taking into account dependencies between various user activities and/or reminders. For instance, a reminder to "learn German tomorrow" can result in the automatic creation of a dependent reminder to "take German books" when leaving home in the morning the next day. In such cases, the original reminder that triggers the other dependent reminders can be manually specified by the user or can itself be generated automatically via the techniques generated above.

Implementation of the described techniques can significantly enhance the user experience (UX) of contextual reminders beyond time and/or location-based reminders. Further, automated triggering of reminders based on user actions predicted with permission can reduce the time and effort of manual reminder specification. As a result of the increased efficiency and effectiveness of the reminders, the user can avoid wasting time, effort, or money due to forgetting a task because of the lack of an appropriately timed reminder.

The described techniques are implemented with specific user permission to access user data used to automatically trigger reminders. Users are provided with options to select specific factors that are permitted for such use, or to deny permission entirely. Users can select parameters related to the reminders, e.g., where and when they'd like to receive reminders, devices and modalities used to provide recommendations, etc. The user can turn off recommendations at any time.

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

This disclosure presents techniques for triggering a reminder for a given action that is dependent on another action that will be performed by the user. The techniques rely on inferring a user's future actions with permission from the user. Alternatively, or in addition, if the user permits, the user's likely future activities can be inferred from other relevant contextual information or specified manually by the user. Subsequently, the inferred and/or user-specified future user activities with a sufficiently high likelihood of occurring are used to trigger reminders for actions that are connected to and/or dependent on the inferred future activities. The model for reminder triggering can be further adjusted and personalized by incorporating user feedback as part of the training data as appropriate. Implementation of the described techniques can significantly enhance the user experience (UX) of contextual reminders, reduce the time and effort of manual reminder specification, and avoid wastage of resources due to the user forgetting a task.