March 2020

User Activity Prediction for Device Screen and Power Management

Jon Napper

Omri Amarilio

Dan Erat

Jia Meng

User Activity Prediction for Device Screen and Power Management

ABSTRACT

To save battery, consumer devices, e.g., laptops, smartphones, etc., enter low-power states based on user inactivity. Typically, a device enters a low-power state after a fixed time has elapsed since the last user activity, e.g., keyboard or touchscreen activity. However, a fixed timeout does not work evenly for all users; for example, it is found that a substantial fraction of users reactivate the device immediately after the device screen is dimmed or the device has entered sleep state. This disclosure describes machine learning techniques to predict the transition to a low-power state based on user activity patterns and the state of the device. The techniques result in improved user experience due to better prediction of the start of user inactivity and increased battery life due to accurate power management.

KEYWORDS

- Power management
- Battery management
- Sleep state
- Screen dim
- User activity prediction
- Operating system
- Machine learning

BACKGROUND

To save battery, consumer devices, e.g., laptops, smartphones, etc., enter low-power states based on user inactivity. For example, the screen of a laptop can be turned off, the laptop can be put in a sleep state, or can be powered off. Typically, a device enters a low-power state after a fixed time has elapsed since the last user activity, e.g., keyboard or touchscreen activity. However, a fixed timeout does not work evenly for all users; for example, it is found that a substantial fraction of users reactivate the device immediately after the device screen is dimmed or the device has entered sleep state.

DESCRIPTION

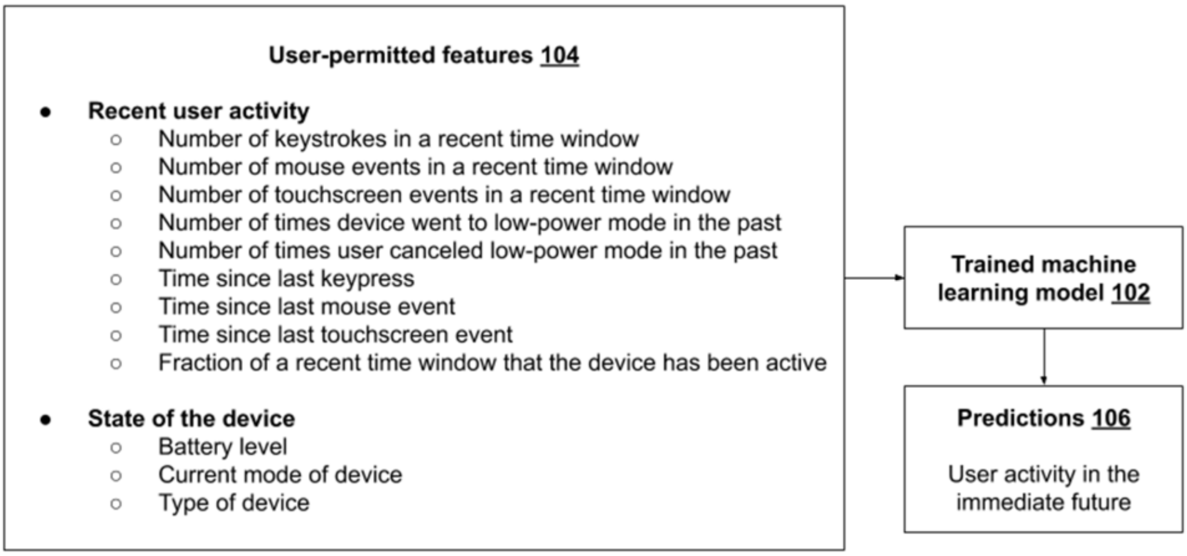

Fig. 1: Using machine learning to predict user activity to enable power management

As illustrated in Fig. 1, this disclosure describes machine learning techniques that predict, after a short period of user inactivity, user activity in the immediate future, e.g., if the user is expected to continue remaining inactive for a longer period or not. A trained machine-learning model (102) accepts as input user-permitted features (104) such as recent user activity and the current state of the device in order to predict user activity in the immediate future (106). The user is provided with options to choose features that are used for such prediction. The user can deny permission for individual features of user activity, can choose a granularity of data for the feature, and can entirely disable the use of user activity data for user state prediction. Further, the user can selectively enable or disable user activity prediction for power management.

The prediction of user activity is used to decide if the device should transition to a lower power state or not. One or more features related to recent user-activity can be used, as permitted by the user. Example features related to user activity that can be used for user state prediction include:

- The number of keystrokes in a recent time window, e.g., the last hour.
- The number of mouse events in a recent time window, e.g., the last hour.
- The number of touchscreen events in a recent time window, e.g., the last hour.
- The number of times the device went to a low-power mode in the past.
- The number of times the user canceled low-power mode in the past.
- The time since the last keypress.
- The time since the last mouse event.
- The time since the last touchscreen event.
- The fraction of a recent time window, e.g., the last hour, that the device has been active.

With user permission, recent device-state features can also be used, and can include:

- Device battery level.
- The current mode of the device.
- The type of the device.

The machine learning model can be, e.g., a regression learning model, a neural network, etc. Example types of neural networks that can be used include long short-term memory (LSTM) neural networks, recurrent neural networks, convolutional neural networks, etc. Other machine learning models, e.g., support vector machines, random forests, boosted decision trees, etc., can also be used.

The described techniques result in accurate prediction of the onset of user inactivity, and can therefore enable better power management. For example, the techniques increase the probability that the device does not go into a low-power state while the user is active, and also reduce the duration of time that the device is in a normally-powered state while the user is inactive.

Further to the descriptions above, a user is 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 actions or activities of providing input to a device, a user's preferences, or a user's current location). In addition, certain data are 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 is treated so that no personally identifiable information can be determined for the user. Thus, the user has 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 describes machine learning techniques to predict the transition to a low-power state based on user activity patterns and the state of the device. The techniques result in improved user experience due to better prediction of the start of user inactivity and increased battery life due to accurate power management.