April 01, 2019

Dynamic selection of assistive features based on utility scores

Victor Cărbune

Sandro Feuz

Dynamic selection of assistive features based on utility scores

ABSTRACT

Assistive features are generally enabled for all features for which an assistance interface is available or implemented. While the number of assistance-capable features continues to grow, meaningful quantification of their utility to the user is not taken into account when deciding which of the features ought to be prioritized. This disclosure describes techniques to prioritize and select from available assistive features, and enable the selected features for a user in a given context. With permission from the user, a centralized model is provided with the set of available assistance-capable features. The model is applied to prioritize and select the assistive features that are enabled.

KEYWORDS

- Assistive feature
- Smart assistant
- Virtual assistant
- User context
- Multiarmed bandit
- Reinforcement learning
- Utility score

BACKGROUND

Various apps and devices, such as smart speakers, smartphones and other devices, include features that provide automated assistance to the user. Such assistive features are typically enabled for all features for which an assistance interface is available or implemented. However, such an assistive layer for a feature is not always relevant or useful for all users in all contexts. In addition to the context, the utility of a feature is also dependent on a specific user's expertise. While the number of assistance-capable features continues to grow, meaningful quantification of their utility to the user is not taken into account when deciding which of the features ought to be prioritized. Selecting the assistive features relevant and useful for a given user in a given context is a challenging problem.

DESCRIPTION

This disclosure describes techniques to prioritize and select from available assistive features, and enable the selected features for a user in a given context. With permission from the user, a centralized model is provided with the set of available assistance-capable features. The model is applied to prioritize and select the assistive features that are then enabled.

If the user permits, prioritization and selection of the assistive features is based on a utility score that indicates the utility derived by the user based on the assistive action connected to the feature. Since the absolute utility is specific to each feature, the utility score serves as an independent unit that expresses utility such that multiple utility scores can be combined and/or compared.

Utility scores can be derived from a variety of factors such as, e.g., time spent in achieving a goal with the assistive feature relative to the time taken without assistance, reduction in cost due to the use of the assistive feature, etc. A utility score can be a normalized score and/or weighted combination of multiple factors. The factors used to determine the absolute and relative utility of each context, with or without assistive features, can be measured through experiments that record user task performance variables. Such measurement is performed with user permission.

Once a utility score is determined for each individual assistive feature, the model is applied to maximize a total utility score based on selective enabling and disabling or the set of available assistive features. The goal of the techniques is to select assistive features that raise user effectiveness, and to avoid automation and assistance for situations in which the user does not require help.

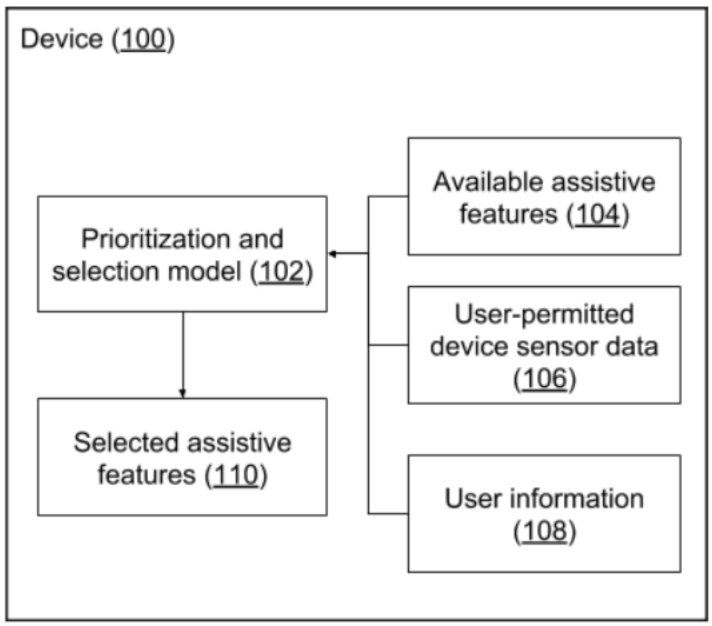

Fig. 1: Selection of assistive features using selection model

Fig. 1 shows an operational diagram of the techniques of this disclosure. With the user’s permission, available assistive feature (104), user-permitted data from device sensors (106), and other permitted user information (108) are made available to a prioritization and selection model (102). The prioritization and selection approaches described in this disclosure are applied by the model to generate a subset of the assistive features (110) that are enabled for the device user for help for the user’s current task.

The model is trained using a typical online learning setup in which outcomes are observed only based on the decisions that have been taken, e.g., whether the user chose to use an available assistive feature, or selected a particular feature from multiple available features. In this instance, enabling an assistive feature yields only the utility achieved by having thea feature enabled, but no indication is obtained of the utility derived in case the feature was unavailable. Formalizing the selective enabling or disabling of assistive features in such a manner results in a typical contextual bandit problem that can be expressed through a multi-arm setup. In a given context, there are N arms, each arm representing an assistive feature that can be enabled, and one of the arms is chosen. Alternatively, the operation can be setup as a simple two-arm scenario with an effective enable/disable ranking for each available assistive feature.

In complex scenarios, different assistive features can influence each other. For example, enabling a particular assistive feature can affect the features that are available for the next decision step. In such cases, reinforcement learning models with additional complexity can be used instead of the contextual bandit setup described above. With user permission, the models can use relevant contextual data, such as the app being used, the information displayed on the screen, measurements obtained by device sensors, metadata from the operating system, etc. Alternatively or in addition, if the user permits, the models can incorporate user information obtained through learned user-embeddings or other appropriate aggregation mechanism.

As typical in bandits or reinforcement learning approaches, learning such models involves an exploration cost. The exploration cost is incurred because the model operation requires disabling certain assistive features in a subset of context. Such disabling can impose a large cost on the user, e.g., if the features are central for the user’s context and task. Therefore, practical implementations of the techniques of this disclosure can benefit from research approaches that help bound the exploration cost of the strategy implemented by the model.

Application of the described techniques can surface assistive features that are the most relevant and needed for a user's current task and context. Further, if the user permits, the chosen assistive features can be customized to take into account the user's expertise pertaining to the task at hand. As a result, with permission from the users, an app or a device can provide personalized combinations of assistive features for different users as suited for the needs, expertise, and context of each user.

If users permit, an alternative approach can include classifying users into buckets based on various aspects of their expertise and enabling assistive features based on predetermined categorization that maps these features to the user buckets. However, unlike the model- based solution, such an approach may not scale with the growing number of assistive features.

Selective enabling of assistive features as described in this disclosure can result in the beneficial side-effect of reduced use of device power because assistive features deemed as not necessary can be stopped and prevented from consuming power. In the case of battery-powered devices, these power savings can provide improved battery life. These power related benefits can also be incorporated within the model, e.g., by incorporating those with appropriate weights when calculating utility scores.

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 describes techniques to prioritize and select from available assistive features, and enable the selected features for a user in a given context. With permission from the user, a centralized model is provided with the set of available assistance-capable features. The model is applied to prioritize and select the assistive features that are then enabled. Utility scores from different assistive features are determined and the model is applied to maximize a total utility score based on selective enabling and disabling of the set of available assistive features. The techniques enable selective provision of assistive features that raise user effectiveness and avoid automation and assistance in situations in which the user does not benefit from assistive features. Application of the described techniques can surface assistive features that are the most relevant for a user's current task and context.