July 17, 2018

Recommending User Action Sequences to Reach an Inferred User Goal

Victor Carbune

Alexandru Damian

Recommending user action sequences to reach an inferred user goal

ABSTRACT

Users vary in their knowledge and skill pertaining to the use of technology. Less experienced users or users that are not technically savvy may not notice or utilize the most efficient functionality suitable for their goals even when it is available within the user interface. With user permission, this disclosure utilizes a machine learning model to observe sequences of user actions across a large number of users. The sequences are utilized to infer the user goal and provide corresponding guidance via the user interface to assist the user in achieving the goal more effectively.

KEYWORDS

- Action sequence
- User goal
- Machine learning
- User action path
- User interface state
- Assistive user interfaces
- User guidance
- Training

BACKGROUND

Users vary in their knowledge and skill pertaining to the use of technology. Less experienced users or users that are not technically savvy may not notice or utilize the most efficient functionality suitable for their goals even when it is available within the user interface.

For instance, a less technically savvy user may not notice or know about utilizing various filters available in a shopping application to narrow shopping choices. Therefore, when searching for an item of a specific color, the user may click through multiple listed items to manually inspect the color, instead of utilizing filters provided within the user interface to narrow the set of items to those of a specific color. Not knowing about, noticing, or understanding various possible actions within the UI can be inefficient, waste user's time and effort, and contributes to user dissatisfaction and frustration.

DESCRIPTION

The techniques of this disclosure utilize a machine learning model that infers user goals. With user permission, user actions performed utilizing a user interface of an application are logged along with the state of the application that corresponds to the respective user action. For example, within the context of an application pertaining to photos, user actions such as tapping or swiping on specific elements within the user interface are recorded along with the UI screen displayed as a result of that action. Completing a task that achieves a goal often results in a screen that is the same or similar, e.g., a page confirming order submission or email delivery. Using this insight, the recorded user actions are analyzed to infer the user goal and accordingly cluster a sequence of actions into a corresponding user session.

The application state is represented as a screenshot of the corresponding displayed screen along with the history of actions and screens within the user session that led the user to arrive at that screen. Data regarding actions and screens is collected with permission from a large set of users that provide consent for collection and use of such data. The data is used to train a machine learning model based on the analyzed states and user sessions within the data. Once trained, the

model can be applied continuously and may run within the application or the device. Alternatively, the model may run external to the application.

At any current state within the application, the model is utilized to assign probabilities to each of the possible user goals. When the probability for a plausible user goal is higher than a confidence threshold value, the shortest path of sequences of user actions needed to reach the goal state from the current state is determined by the trained model. The user interface is updated to provide an indication of the goal and the recommended action sequence to the user, e.g., with corresponding step-by-step assistance. The user is free to dismiss the suggested goal, select a goal other than the ones offered by the model, or deviate from the suggested action at any step within the suggested sequence. Alternative to being implemented within an application/ device, the techniques of this disclosure can also be invoked via an application programming interface (API) that enables application developers to offer customized recommendations and assistance.

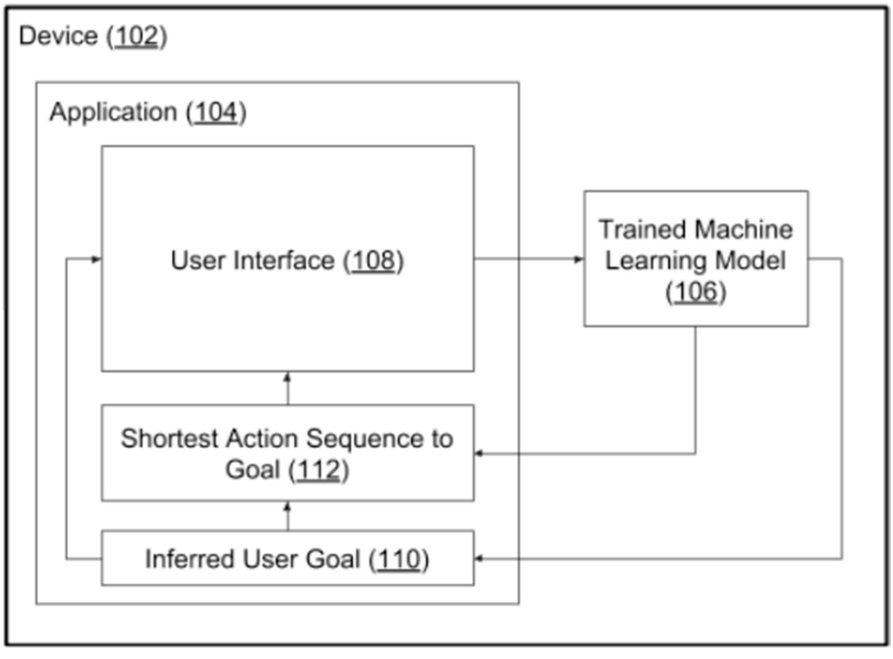

Fig. 1: Recommending effective action sequence for user goal

Fig. 1 shows an example implementation of the techniques of this disclosure. A user interacts with an application (104) running on a device (102). In the example shown in Fig. 1, a trained machine learning mode (106) is utilized to provide recommendations. The model is trained based on user-permitted training data of user actions and corresponding states, e.g., collected from prior user sessions of interaction with the application using a log server,.

During operation, a current user interface (108) of the application is provided as an input to the trained model. The model assigns probabilities to different user goal states that may be reached from the input screen. The calculated probabilities are compared to a specified confidence threshold value. If none of the calculated probabilities are higher than the threshold value, no action is taken. If one or more of the calculated probabilities is higher than the threshold value, the goal (110) with the highest calculated probability is inferred as the likely goal of the user.

The goal is then presented to the user via the user interface of the application along with the shortest sequence (112) of actions that a user can perform using the user interface to reach the goal state from the current state of the user interface. At each step within the action sequence, the user is provided with guidance for the corresponding action. The goal and sequence recommendations and corresponding guidance are presented via user interface elements that may be dismissed by the user at any time, e.g., if the goal was identified incorrectly.

The techniques of this disclosure can be extended to offer multiple recommendations of goals and corresponding action sequences in cases where the calculated probabilities of multiple goals are higher than the specified threshold. Alternatively, a single goal and corresponding action sequence may be chosen, e.g., randomly, when multiple calculated probabilities above the threshold value are the same. Similarly, if multiple action sequences are equally efficient to

achieve a given user goal, all such sequences can be suggested for that goal or one or more sequences may be chosen at random. The value of the confidence threshold can be specified by the model developer, application developer, the user, the device maker, or obtained in another manner as suitable for achieving effective operation.

The techniques of this disclosure facilitate more effective use of an application, especially ones that provide elaborate features that can overwhelm novice users. Moreover, the techniques provide the scaffolding that enables such users to transition to skilled users in a shorter amount of time. The techniques can be provided to different parties via an API.

With user permission, application developers may utilize the API to access and analyze anonymized logs of user interactions within their applications. Such analyses can enable user research that sheds light on typical user action sequences and can uncover confusing, inefficient, or problematic user interaction and navigation patterns, screens, and features. Additionally, the API can reduce the time and effort needed to implement the user interface delivery of recommendations and guidance desired by the application developers. For instance, the API can be used to implement features like shortcuts or deliver tutorials.

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

With user permission, this disclosure utilizes a machine learning model to observe sequences of user actions across a large number of users. The sequences are utilized to infer the user goal and provide corresponding guidance via the user interface to assist the user in achieving the goal more effectively. The techniques of this disclosure facilitate more effective use of a software application, especially one that provides elaborate features that may overwhelm novices. Moreover, the techniques provide the scaffolding that enables novices to transition to knowledgeable users within a short amount of time. The techniques of this disclosure can be provided to various parties, e.g., via an API. With user permission, application developers can utilize the API to conduct user research and to deliver custom recommendations and/or guidance.

REFERENCES

1. Chughtai, Muhammad Waseem, Ali Bin Selamat, and Imran Ghani. "Goal-based hybrid filtering for user-to-user Personalized Recommendation." International Journal of Electrical and Computer Engineering (IJECE) 3, no. 3 (2013): 329-336.