July 01, 2019

Persona Profiles for Prioritization of Content Output

Victor Carbune

Pedro Gonnet

Persona Profiles for Prioritization of Content Output

Abstract:

A persona-specific ranking model creates a persona-specific profile layer in an operating system of a user equipment device that helps users to prioritize their ongoing activities by explicitly enabling a specific persona profile. A profile (*e.g.,* user profile) can be used to manage a visual display of personal data associated with a specific user or a customized desktop environment. Much like today's work profile, each persona profile is specific to a user context or activity, such as runner, reader, traveler, studier, house cleaner, philosopher, software engineer, rower, manager, skydiver, parent, and so on. The persona-specific ranking model creates a list of persona profiles suitable to the user based on user features (*e.g.,* user interaction, activities, and priorities) and recognizes when a particular persona is active. A system layer is implemented to enable individual apps to expose information through the persona-specific ranking model, such that the most relevant information to the user's current activity is prioritized accordingly within any application, including third-party applications. The persona profiles are manually, or autonomously, selectable.

Keywords:

Profile, work profile, persona, ranking model, persona-specific ranking model, persona-specific profile layer, system-level trained ranking model,

Background:

Today, securely managing a work environment without restricting users from using their work-provided devices for personal applications and data can be challenging. One solution to this

problem is for users to use separate devices, for example, one mobile phone for work that is managed by the organization's IT department, and a second mobile device that is solely for personal use. Another solution is to create a "work" profile on the mobile device that allows administrators to manage business data and applications they care about, but leave everything else on a device under the user's control. The work profile is controlled by administrators and kept separate from personal accounts, applications, and data. Corporate applications, data, and management policies are restricted to the work profile, keeping them secure and separate from personal data while maintaining user privacy. While the "work" profile is a useful feature for effectively focusing on tasks and activities related to the work environment, "work" and "personal" are not necessarily the only suitable classifications for most users. Further, prioritizing and filtering information flow based on current on-going activity is not only relevant in numerous other scenarios but is also non-trivial.

Description:

Generally, a profile (*e.g.,* user profile) can be used to manage a visual display of personal data associated with a specific user or a customized desktop environment. A profile may refer to an explicit digital representation of a person's identity. Implementing a "work" profile typically requires a hard separation of information flow through a different account than that of a "personal" profile. The techniques described herein introduce a "persona" layer in the mobile operating system that, unlike the work profile, aims to help the user prioritize their ongoing activities by explicitly enabling a specific profile. Accordingly, the techniques described herein generalize the concept of the "work" profile to a set of "persona" profiles, as is often the case that an individual wears multiple "hats" both in their personal and professional lives. For example, a user on vacation

could use a "traveler" profile. For a user sitting down to read a book, a "reader" profile may be more applicable. A user out jogging could use a "runner" profile. Each of these activities may handle notifications, messages, and incoming streams of information differently such that the most relevant information to the user's current activity is prioritized within an application to present a smart-filtered view.

For manual selection, a persona may be presented as a selectable user interface (UI) element (*e.g.,* icon). In one example, the UI element can be located proximate to a work profile UI element on the display of the mobile device, such as in a settings menu. When a user selects a particular persona, such as by tapping the corresponding UI element, the mobile device responsively personalizes all context-specific system and application functionality according to the selected persona. For example, when a reader persona is active and the user opens a news feed, the mobile device exposes articles highly relevant to a reading activity. The exposed articles may include articles related to a book being read or previously read in an e-reader application on the mobile device. Further, from a social aspect, a book review application can expose book reviews from friends.

In one implementation, when the user checks their email, the "priority inbox" may re-rank emails to prioritize those that correspond to the currently active persona, without hiding emails that do not correspond to the currently active persona. For the reader persona, for instance, emails typically considered a priority, such as "please RSVP to my dinner invite next week," which have nothing to do with the user's current activity (*e.g.,* reading), may be re-ranked below other emails less actionable but related to books. This is because book-related emails may be more interesting to the user during the user's reading activity.

In aspects, a map application can expose locations related to the current persona. For instance, the map application can expose nearby reading locations, *e.g.,* a library nearby that the user can visit to read a book. In addition, selection and location of the UI elements on the home screen, for example, can change based on which persona is active. For example, for the reader persona, news or e-reader applications may be displayed more prominently than other apps unrelated to the reader persona.

The mobile device can also automatically change settings based on the selected persona. Changeable settings may include volume, screen brightness, notification policies, and so on. Changing settings can reduce distractions during the user's current activity, such as to make reading easier for the user.

In some aspects, the mobile device can detect the user's current activity and in response, autonomously switch to the appropriate persona without requiring the user to manually select the persona. This detection can be based on currently open and active applications, time of day, day of week, observed patterns of user behavior, a calendar event, and so on. The mobile device can detect the user's current activity based on sensor data (*e.g.,* amount and/or speed of movement (running, walking, cycling, etc.), ambient temperature, location (global navigation satellite system (GNSS)), and so on). The mobile device may also allow the user to opt out of this automated persona switching.

To enable these persona profiles, several utility components and system changes can be implemented. For example, the operating system of the mobile device can implement a machine-learned model for classifying, based on user features, a list of persona profiles suitable to the user and recognizing when a particular persona is active. The model analyzes data generated by the user, *e.g.,* content consumed previously, general user features, location(s), time of day (session

based), how much movement (*e.g.,* the user might not be at work), lack of movement (*e.g.,* the user might be at a desk), etc. Combining this ancillary information provides a clearer idea of the user's current activity, and enables the mobile device to detect the user's activity and output a most-likely class or persona for the user. These classifications can also be manually refined by the user. The list of persona profiles can be presented to allow the user to choose which personas to implement. Then, the chosen persona profiles are made available for selection, *e.g.,* presented next to a work profile toggle.

An infrastructure component can be implemented that allows third-party applications to make use of this smart-filtered view for the user. The infrastructure component can be invisible for the application for certain types of system functionality (*e.g.,* notifications). However, for content-ranking, it could also expose to the application a particular persona mode (if the user grants this type of permission). The utility components also include machine-learned models for ranking content for relevancy to a particular persona.

Challenges may arise, however, in that some personas may not have customizations in place for some users. Typically, emails, text messages, social media feeds, or other media that a user consumes that are relevant to personality preferences and interests of the user (*e.g.,* in a runner-specific context) in a particular situation are mixed together with all the user's other activity interests.

In aspects, a two-stage training procedure can be utilized. In the first stage, a server trains a persona model by analyzing a list of possible personas (*e.g.,* reader, rower, runner, philosopher, software engineer, manager, paraglider, skydiver, parent, etc.) using a variety of information for associating relevant content to each persona on the list. This content includes both generally available information and (anonymized) personal information, such as relevant articles, materials,

trips, places, and so forth. Then, the user device trains a clustering model, based on the personas and the relevant content in the persona model, to classify the user in one or more of the personas on the list. Existing technology, *e.g.,* activity and sentiment analysis, can also be leveraged when training the clustering model.

To refine the flow of information for the user, a system-wide component is selectable (manually or autonomously) to toggle a mode, similar to the work profile today but it would set a specific persona. When a persona profile is selected, the mobile device informs one or more (including all) third-party applications associated with the selected persona profile (if the user provided a corresponding permission, which would also be introduced). If such permission is granted by the user, the persona profile can be used within the applications at their own convenience. Further responsive to selection of a persona profile, a machine-learned system-level ranking model is applied. This ranking model is implemented to prioritize different notifications, settings, messages, and any other functionality exposing content. In some instances, one or more applications may be entirely de-prioritized while others are further emphasized to the user. For example, a display of launcher icons of applications that are related to the current persona profile can be brightened, highlighted, and/or repositioned on the home screen to be more easily accessible to the user during the current persona mode. Other launcher icons for applications not related to the current persona profile can be de-emphasized (*e.g.,* faded or greyed out) and/or repositioned so as to be less prominent on the display.

For persona-specific content filtering and ranking, a generic machine-learned persona-level ranking model is implemented that scores how relevant a piece of content is for a persona. Then, the ranking model is further refined as the user interacts with the mobile device during a specific

persona profile to continuously adapt. Additional refinement can be applied based on federated learning.

The persona information can also be provided as an additional signal to other system services and applications, which can use the signal to enhance their own machine-learned models. In one example, if using a "runner" persona and the mobile device detects that the user is running, the mobile device can determine to switch off the mobile data communications and keep the location function on to reduce potential distractions from incoming communications.

In some aspects, a machine learning model can be used to identify the top "personas" (*e.g.,* top five) of the user. These modeling approaches may leverage neural network architecture from numerous related applications for analyzing content (*e.g.,* text, sound, images, and videos). Recurrent neural networks, attention-networks (with transformers), convolution networks, and so forth can also be used to train for a relevancy score. An example learning algorithm usable by a neural network includes triplet loss.

In contrast to today's work profile that requires a hard separation of information flow through a different account than that of a "personal" profile, personas have a soft separation of information such that some overlap may occur between personas and the information flow is re-ranked rather than simply being displayed or not displayed in a binary manner. Fig. 1 below depicts a diagram illustrating a soft separation of some example persona profiles each with particular prioritization settings. For example, Fig. 1 illustrates three example personas: a runner persona, a traveler persona, and a reader persona.

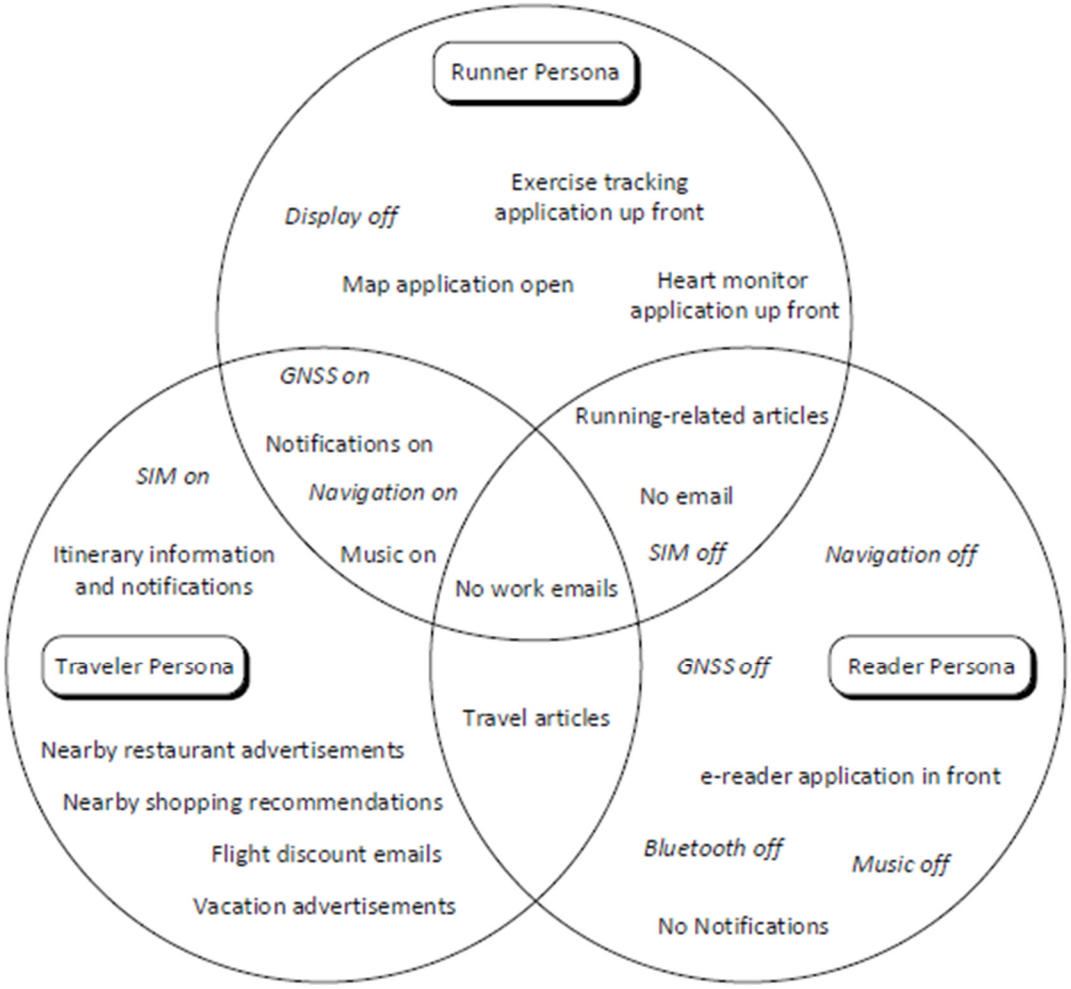

Fig. 1

In the "runner" persona, the mobile device prioritizes certain applications to be open or in focus, such as an exercise tracking application, a heart monitor application, or a map application. In addition, the user may wish to not be distracted by emails, particularly work emails. However, the user may want to receive notifications (including filtered notifications or all notifications), such as messages from a spouse or loved one. Further, while the user is running, the display may be turned off even though the mobile device is being jostled or certain buttons are pressed. The user may also wish to not be disturbed by incoming calls while running, so the subscriber identity module (SIM) may be disabled. Other functions may be activated or prioritized during the running

persona, including GNSS (*e.g.,* Global Navigation System (GPS)), navigation, music application, and so on. If the user is not running while the "runner" persona is active, running-related articles may be exposed for the user to read.

Continuing with the example illustrated in Fig. 1, the mobile device, when using the "reader" persona, may disable certain functionalities such as GNSS, Bluetooth™, music, navigation, and SIM that may not be useful to the user during a reading activity. The user may not want to be distracted by emails while sitting to read for a period of time. Also, the mobile device may automatically open an e-reader application. In a news feed, the various articles may be prioritized over others based on the content being related to other personas applicable to the user, such as running-related articles or travel articles.

In the traveler persona shown in Fig. 1, some settings may overlap with the runner persona, such as "GNSS on," "notifications on," "navigation on," "music on," and "no work emails." However, the SIM may be on to allow incoming calls. Information relevant to a user planning a trip while using the "traveler" persona may include travel articles, flight discount emails, or vacation advertisements. For a user currently traveling, relevant information may include itinerary information and notifications (including changes), nearby restaurant advertisements, or nearby shopping recommendations.

As depicted in Fig. 1, some information relevant to a particular persona may also, or only, be relevant because of one or more additional personas applicable to the user. For example, although the user is a reader and a runner, the news feed may de-prioritize articles on cycling if the user is not also a cyclist. If a user is a reader and not a traveler, then the news feed or email application may de-prioritize articles and advertisements related to traveling.

These techniques can help enable users to immerse in their mobile experience, deepen their knowledge in their preferred subject matter, and avoid being distracted by ongoing information flow, notifications, etc. These features long-term can also benefit career goals. For example, a "manager" persona may constantly help by recommending articles adapted to the level of experience of the user. Additionally, these techniques can customize the personas such that the personas are specific to the user's interests and activities.

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.

In conclusion, the techniques described herein abstract today's "work" profile to enable a mobile device to prioritize information relevant to a user's interests and current activity. Various persona profiles are customized to the user and further adapted as the user interacts during a specific persona mode. System-level and application-level prioritization of information and functionality can be implemented, based on the selected persona profile, to customize the information flow to the user, such as by modifying the display, activating or deactivating an

application or operation, enabling or disabling a component of the mobile device, changing settings, and filtering incoming information for relevancy specific to the user and the user's current context. In some implementations, the mobile device can detect the user's current activity and autonomously switch to a corresponding persona profile.

References:

[1] Laaden, Oren, Omer Yehiel Eiferman, Ranit R. Fink-Isaacs, and Amir Goldstein. Systems and Methods for Customizing Services to Users on Mobile Technology Platforms Using Personas. US Pub. 20150149588, filed February 6, 2015, published May 28, 2015.