

BROWSER PERSONAS: TOWARDS A REASONABLE MIDDLE GROUND

Ben Livshits \ Microsoft Research

Today's web users demand a more personalized browsing experience. This often requires a degree of *personalization*. Common examples of personalization include web page personalization, such as a personalized news page provided by Google News; search personalization, based on the past search history, location, etc. as commonly performed by many of today's popular search engines; and behavioral ad targeting.

These approaches to creating a more personalized web often run afoul of the user's desire to keep information about themselves – both sensitive PII as well as somewhat sensitive personal preference information – private.

CLIENT-SIDE INTEREST GATHERING

Recent research efforts such as RePriv, Adnostic, and Privad have advocated mining information about the user within the browser, so, as the user visits pages, the profile of her interests is continuously refined. Several representations of such interest profiles exist, a common one being a set of Open Directory Project (ODP) categories. For instance, the profile of a particular user may be

- *Hobbies > Sports > Hockey*
- *Arts > Cinema > Movies > Film noir*
- *Activities > Outdoors > Hiking*

This – fairly generic – *interest profile* enables multiple interesting applications such as augmenting a page to better match the interests, selecting the most relevant in a set of potential ads to present to the user, reshuffling search results to better match the interest profile, etc.

CLIENT- AND SERVER-SIDE PERSONALIZATION

Augmentation of a particular page can occur on both the client and the server; each of these approaches has tradeoffs.

- Server-side personalization
 - **Pros:** personalized versions of pages can be pre-computed to save time.
 - **Cons:** The client needs to share (a portion of) its interest profile to get personalized results; a common categorization needs to be shared between the server and the client.
- Client-side personalization
 - **Pros:** no sharing with the server is necessary
 - **Cons:** can increase page delivery latency

In practice, we can see both the client- and server-side approaches being adopted. For server-side personalization to work and work efficiently a common, preferably small, fixed set of profiles needs to be developed.

PERSONAS

Here we advocate the notion of *browser personas*, which correspond to roughly defined sets of interests. Here are some examples of such personas:

- College student
- Working mother
- Retired person
- Professional
- Business executive

While far from any notion of completeness, a simple list of personal such as this may be a good start towards aiding with simple personalization tasks.

NEWS SITE EXAMPLE

A news site such as Google news can be personalized to match the persona of the current browser user. The news relevant to a *college student* might be different from the news relevant for a *business executive*. The latter might emphasize business and political news; the former – entertainment news.

Optionally, the browser persona *may* be proactively shared with the Google news server by embedding it in a header. Of course, the ultimate control about whether to share this information should be left to the end-user.

ADDITIONAL CONSIDERATIONS

Several other considerations and observations emerge from the example above.

- Multiple personas supported: similar to the notion of a browser profile, multiple personas can be associated with the same browser. These could be a *work* and a *home* persona (i.e. business executive/working mother). They could also correspond to different members of a household, sharing the same browser. How to properly expose controls for switching between these personas is an interesting research question in the area of usable security.
- What happens if the user is deliberately deceptive about their persona, pretending to be someone they are not?
- Should there be a canonical set of all personas? Who would define this as a standard?

CONCLUSIONS

In this paper, we advocate a very weak form of identity, browser personas. The main idea is to allow for a degree of personalization, while browsing the web in an unauthenticated mode. We feel that personas provide a legitimate alternative to the current approach of user tracking and behavioral profiling, giving the user control over what data about themselves to share and when.