



ELIZABETH F. CHURCHILL
eBay Research Labs | churchill@acm.org

Elizabeth Churchill is director of human computer interaction at eBay Research Labs, San Jose, CA. She also serves as vice president of ACM SIGCHI.

Putting the Person Back into Personalization

Personalization. It's one of the words of our times. I heard it five times yesterday in different conversations. I think I also uttered it once or twice, pushing the count up further.

As I reflected on the conversations of the day, it became clear there are many different denotations and connotations of the word *personalization*. This reminded me of my ruminations on the word *social*. So, I decided to spend some time researching the concept, trying to understand what personalization means and what happens if you do it.

To be a person implies you are an individual, human-like character, an entity with a specific identity, with specific nuances of expression and emotion, and with quirks of character and taste. To personalize is to tailor for a specific person. According to the dictionary, if one were to personalize something, one might mark it with initials, a name, or a monogram, as in "personalizing stationery" [1]. Or one might "make something personal," specializing a general state of

affairs by adding flavor and nuance that reflect the individual's traits. One might also design or tailor something to meet an individual's specifications, needs, or preferences. Thus, personalization is about individuation and specialization.

Arguably, personalization is central to the discipline of human-computer interaction. As my friend Gordon Baxter puts it, "HCI is about particular people doing particular things in particular contexts." Only when we have understood that particularity across a number of cases can we be confident our generalizations and abstractions are truly useful and/or appropriate. This viewpoint exemplifies the case-study approach that is one of the mainstay methods in our practice.

In the world of designing Internet applications and services, personalizing is largely about filtering content to satisfy an individual's particular tastes. In this construction, personalization is a marketing technique, the logical extension of market segmentation. Here, we see the word used more or less

synonymously with (and regularly co-occurring with) concepts such as customization, individualization, segmentation, targeting, profiling, one-to-one marketing, and adaptive services. Influenced by the work of Edward Chamberlain [2] and H.W. Robinson [3] in the 1930s, Wendell Smith in 1956 defined market segmentation as "viewing a heterogeneous market as a number of smaller homogeneous markets, in response to differing preferences, attributable to the desires of consumers for more precise satisfaction of their varying wants" [4]. Smith's vision places the consumer's uniqueness or particularity over and above their status as an anonymous service recipient.

How Is Personalization Achieved?

Someone once said the Internet is the human-activity observer's Hubble telescope—ideal for collecting data to illuminate heretofore unseen nuances of human behavior.

It is certainly true that mountains of behavioral and transaction data are collected every second, and that data is the grist for the mill of online personalization. Captured, managed, and mined/munged data has revolutionized the marketing world, leading to *specialized market segmentation* (hyper-segmentation into narrow clusters), *progressive*

profiling (incremental data collection across sessions and interaction points), and *addressable marketing methods* (gathering information about online behaviors including site visitation, site and content engagement, and advertising exposure). Coupled with site structure, domain knowledge, user demographics, user profiles, and social networking data, a powerful foundation is created for predictive models that aspire to reflect what people need, want, desire—or fear losing. Every click, every download, and every purchase is useful information in the process of building your personal “taste graph.” Based on this, businesses and marketers can select targeted content that is more relevant, more desirable, and therefore more likely to be consumed.

The construct of personalization has been operationalized in multiple ways, with some variation being attributable to domain-related requirements. Usually, however, there are three main phases of personalization: *learning* (explicit from profiles and intentional user-generated signals [e.g., “likes” on Facebook] and implicit from transaction and activity trace data), *matching*, and *recommendation*. At the risk of oversimplifying, and acknowledging that new methods are being created every day, it is worth noting

that the matching process breaks down into:

- content-based matching, which filters data based on someone’s previously stored data;
- collaborative filtering, matching, and predicting based on collaborative information from multiple users to offer options for newcomers;
- rule-based filtering, where business rules or expert policies are followed to make recommendations; and
- hybrid methods, where a soupçon of all the approaches above are used.

Most contemporary services use hybrid methods. Emerging techniques for real-time data analytics promise swift adaptation, although it is the case that Web-based personalization is an iterative process that takes time to get right.

What Is Missing, and Where Do We Need Some Deeper Thinking?

Leaving important privacy and data issues on the side for now, in my view there are two exciting challenges for personalization where human-centered design thinking can combine with research into algorithm design to take us to the next stage of personalization: hyper-specialization of outcome personalization and improvements in process personalization.

Before diving in, let’s distinguish between two interrelated concepts: *outcome personalization* and *process personalization* [5].

Outcome personalization is about options. A salad bar gives me options for what to pile on my plate. I can personalize my meal. Internet personalization algorithms operationalize outcome personalization: A selection of items is offered from a larger set of options based on what is known about me, the consumer. I see books that are somehow related to other books I have purchased, movies that by some criteria correspond to others I have watched, shoes by a designer I have in some way indicated appeals to my taste, and so on. Sometimes these recommendations are creepily accurate, but more often they are pitifully off-beam. I suspect most of us like playing the game of trying to work out exactly *why* something was recommended. In any case, the objective, put in the best possible light, is to counter “option glut” by reducing the consideration set so that it is manageable, and thus more likely to result in a person finding something that is actually desirable. The alternative is information overload and decision fatigue, which leads to being overwhelmed, confused, and exhausted from reviewing the overabundant options. There is a lot

Person-Adaptive Systems and Current/Potential Input Data for Personalization

Person (“user”) data types

- demographic data
- knowledge-, expertise-, and experience-dependent personalization (used for most learning systems)
- skills and capabilities (used for most help systems and assistive technologies)
- interests and preferences (used for most recommender systems, driver of “taste graphs”)
- intent, goals, and plans (used for targeted browsing, plan and intent recognition)
- emotions (recent work in affective recognition and adaptive interactive interfaces and agents)
- physiological state

Behavioral and usage data

- movements (increasingly used for adaptive fitness applications)
- selective actions (information path analysis, link selection)
- ratings (now commonly used by sites like Netflix; early roots of this work derive from research at MIT and the work of the GroupLens researchers in Minnesota)
- purchases and purchase-related actions (e.g., suggestions of similar goods)
- broader task-related actions (bookmarking, printing, saving, sharing)
- temporal aspects of viewing behaviors (adaptation based on viewing time, analysis of micro-actions)

Usage regularities available from data

- usage frequency
- situation-action correlations, meeting requests, calendar entries
- action sequences, recommendations based on frequently used action sequences and those of others

Environmental data

- software environment, adaptation based on browser version, platform, device, configuration (e.g., location services on or off)
- hardware environment, adaptation based on device, bandwidth, processor or download speed, input devices, display devices
- locale (adaptation based on current geo location, social location, etc.)
- current conditions (hot, cold, ambient noise, light, dark)

of evidence that giving people too many options is deadly if you actually want them to make a decision.

Process personalization is about the quality of the interaction. Process personalization acknowledges that delivery, as well as content, needs to be personalized. Personalized service encounters are not just about offering content options; managing the presentation of the options through a sensitively choreographed, ongoing interaction is paramount.

In face-to-face encounters, this means reading signals carefully from body posture, from eye gaze, from words uttered. It is a careful dance of control between the customer and the service agent, smoothed with small talk.

Process personalization can be further broken down. There is a scale of engagement that goes from efficient recommendation, to friendliness enacted through social niceties (“Have a nice day!”), to active collaboration. Process personalization can be low effort and programmed. Think of the interaction you have every time you go to a chain restaurant—the same greeting, the same uniform, the same conversational exchange. The Walt Disney Company has perfected this kind of programmed personalization at its theme parks. Process personalization can also be customized. Customized personalization is about personalizing to the consumer’s interactional style and needs in the moment, as well as more stable or longer-term facets such as their demographic profile and/or manifest tastes. High-touch services (e.g., haircuts, therapy, medical encounters) tend to offer customized personalization, while low-touch service encounters (e.g., grabbing a pint of milk at a supermarket, booking a routine flight, paying a bill) tend to be programmed, with potentially minimal personalization. Notably, it is not the case that customized personalization is always better. Obsequiousness has its place. If I expect minimal interaction and I get someone paying me too much attention and working too hard to please me, it can be intrusive and annoying. However, if I expect customized personalization—if I want to feel special, and anticipate some care and attentiveness—and I get programmatic interaction, I feel

stonewalled, neglected, disturbed, and affronted.

To get back to interesting opportunities for research: The first area where we can really bring to bear a deeper understanding of human passions and proclivities is in course correcting the *hyper-specialization* inherent in outcome personalization today. Hyper-specialization refers to the fact that personalization algorithms based on taste graphs, like terriers, tend to keep digging deeper down the same hole. Most of us are familiar with Eli Pariser’s concept of the *filter bubble*, which has been much discussed in the context of filtered news, the balkanization of political opinions, and concerns about specialized search results [6]. The filter bubble refers to the idea that each of us is increasingly living in our own personal ecosystem of information. This happens through our own attitude-reinforcing actions (e.g., seeking out sources that reinforce what we already think, such as following like-minded people on Twitter and reading the same news sites every day). But, more importantly, according to Pariser, the process is exacerbated by algorithmic filtering in service of market forces. To sidestep the important activist concerns and focus on the algorithmic issues for now, it tends to be the case that as personalization algorithms get smarter, they also get narrower. Success in this world is more accurate prediction of what you like based on what you (appear to) have liked in the past (or possibly what your friends appear to have liked). Algorithms today are not good at working out that you really don’t like something anymore, that your tastes have shifted. For example, on one shopping site, I am still being shown items from a line of designer clothing I have not considered in over two years. And

algorithms are not good at making sense of what is missing, what might be appealing but could not be predicted from your past. Thus, we are increasingly seeing services injecting “serendipity” by suggesting content that is not what the individual’s taste graph would have predicted. This is what human editors are paid to do: to tell us what is cool or interesting that we may otherwise not have seen and to show us macro-trends and correct some of the micro-trend-reinforcing personalization algorithms. So, the design of effective human *plus* algorithmic recommendation systems is the way of the future. Human curation is back and curators need better tools.

To address the second area of possible innovation, process personalization, Web personalization is largely information centric, and has not focused on process personalization. In physical-world settings, the concierge, the salesperson, and the receptionist leave a lasting impression. The point of contact between the business or service and the customer is critical. Thus, full personalization includes consideration of how and when to recommend something in the light of the recipient’s physical and psychological circumstances at the time. It’s about managing the recommendation process, about sensitivity to personal circumstances as well as wants and needs. Compelling, persuasive service is about judging *when* someone is most likely to be receptive to an idea. In the digital world, the points of contact between a service and the consumer are interfaces and interactive scripts. These ought to be sensitive to the recipient’s current physiological and psychological circumstances. These need to be adaptive, contextualized, and personalized, just as the back-end algorithms produce personalized

content adaptively. Achieving process personalization at the interface requires some smart rethinking of what the role of an interface is.

Process personalization is also about expectation setting, transparency, and a good set of repair strategies when expectations are misaligned. When I go to a vibrant market, I expect merchants to potentially invade my auditory and visual space and quite possibly my physical space. When I go to high-end department stores, my expectations are different. Such expectations frame the kind of interactional choreography I expect. If I don’t get what I expect, it jars me into reassessing the situation—I have to engage my brain to figure out what is going on. Like most people, I don’t want to be forced to think when I was not expecting to have to. (Steve Krug points this out in his Web-usability text *Don’t Make Me Think* [7]). This is perhaps why people are change averse: because it forces them to have to think. This is a facet of the “innovator’s dilemma,” the situation where existing customers’ established practices and service expectations tug against change, where a customer base resists rethinking and relearning new ways of interacting with services and applications—even if the changes will more effectively serve latent or unstated needs [8].

So, What Is Next for Personalization?

Personalization is a multidimensional concept and it is important to unpack it. If we are really talking about personalization, rather than simply co-opting a word to make the technically feasible seem imaginative and socially sophisticated, we need to think more deeply. We can’t bucket test [9] our way incrementally through algorithm and

interface tweaks to develop the next phase of truly personalized services. We have to think more clearly and imaginatively about adaptive and sensitive outcome and process personalization. For the latter, we will need better models of social transaction and interaction. We will need a better grasp on the effects of context, not just informational but also physical, social, and psychological. We will need to develop better models that are sensitive to time—human time, experienced time, not only clock time. We will need more socially adept computational agents. We will need to put the person, or perhaps the people, back into personalization.

ENDNOTES:

1. <http://dictionary.reference.com/browse/personalization>
2. Chamberlin, E. *Theory of Monopolistic Competition*. Harvard University Press, Cambridge, MA, 1933.
3. Robinson, H.W. The equilibrium price in a perfect inter-temporal market. *Econometrica* 6, 1 (1938), 48-62.
4. Smith, W. Product differentiation and market segmentation as alternative marketing strategies. *Journal of Marketing* 21, 1 (1956), 3-8.
5. I have borrowed and extended the distinction between outcome and process personalization from some very nice work on face-to-face service encounters: Surprenant, C.F. and Solomon, M.R. Personalization in the service encounter. *Journal of Marketing* 51, 2 (Apr. 1987), 86-96; <http://www.ida.liu.se/~steho/und/htdd01/4999772.pdf>
6. Pariser, Eli. *The Filter Bubble: What the Internet Is Hiding from You*. Viking/Penguin Press, London, 2011. See also Bosker, B. Tim Berners-Lee: Facebook threatens Web, beware. *The Guardian* (Nov. 22, 2010); <http://www.guardian.co.uk/technology/2010/nov/22/tim-berners-lee-facebook>
7. Krug, S. *Don't Make Me Think. A Common Sense Approach to Web Usability*. New Riders, 2005.
8. Christensen, C. *The Innovator's Dilemma: The Revolutionary Book That Will Change the Way You Do Business*. Harper Business, 2011.
9. Bucket testing or A/B testing is “a methodology in advertising of using randomized experiments with two variants, A and B, which are the Control and Treatment in the controlled experiment. Such experiments are commonly used in web development and marketing, as well as in more traditional forms of advertising. Other names include randomized controlled experiments, online controlled experiments, and split testing.” See http://en.wikipedia.org/wiki/A/B_testing