Personalized Digital Services:
Power, Equity and Transparency in "Digital Familiars"
An exploration of the issues of personalization and customization

Jason Black, Kieran Downes, Frank Field, Aleksandra Mozdzanowska

(a working draft with links, etc. is at https://msl1.mit.edu/twiki/bin/view/Scratch/PersCustWorkingDraft3)

> ARIEL
> All hail, great master! grave sir, hail! I come
> To answer thy best pleasure; be't to fly,
> To swim, to dive into the fire, to ride
> On the curl'd clouds, to thy strong bidding task
> Ariel and all his quality.
> > *The Tempest*; Act I, Scene II; William Shakespeare

As computers and computerized services have become ubiquitous, there has been a concomitant increase in the mechanisms and modalities of personal interaction with these devices. However, the accessibility and understandability of the services being offered has continued to outstrip the public's grasp of the possibilities of these technologies.

One strategy that has been employed to ease the human-machine service interaction has been to shift the burden of understanding the operation and/or capabilities of a machine service away from the user and onto the programmer. By devising software interfaces that adaptively respond to signals from the user, the programmer can embed program facilities that can "learn," identifying limitations in the user's appreciation of the operation or the features of the machine service being provided and attempting to anticipate the user's needs and wants without the user having to master the machine or system.

These efforts to create adaptive interfaces have gone through many iterations, and have seen varying degrees of success. As the computer has become more and more ubiquitous, these adaptive interfaces have become a part of the daily lives of the public.

One of the main streams of this kind of system design goes under the names of "personalization" and "customization," largely distinguished from one another by the extent to which the user actively participates in shaping the performance and behavior of the software intermediary. While these tools are descendents of the original efforts to simplify the user's experience with complex software services, as their application has moved from the specialized to the mundane, a host of important issues has begun to emerge.

A brief historical exploration of a handful of the current exponents of this set of technologies can serve to illustrate the nature of these issues. Personalization's preeminent exemplar is Amazon.com, where a web-based software agent acts as the customer's personal shopper. Based upon the customer's purchasing history, and the purchasing history of all other Amazon.com shoppers, the software makes purchasing suggestions, leading the customer to what it expects are products that are likely to satisfy his wants. The user has the option to actively participate in "teaching" the software by ranking past purchases and commenting upon products purchased elsewhere, but the fundamental effort has been to mimic the behavior of an attentive salesman, who has a perfect memory and a firm grasp on the interests and buying habits of all of Amazon.com's customers.

At one level, there is nothing new about this kind of marketing, of course. While the field has gone through a host of incarnations, at its heart marketing has been about identifying what a customer has bought in the past, and drawing inferences about that to promote new purchasing options. Probably the key development in marketing has been the discovery that, sometimes, it is more effective to promote the seller/producer of the product than the product itself – leading to the rise of the notion of "brand" and its marketing. And the focus of that effort is the development of "brand trust" – the sense that, even if the customer does not know everything about the product offered, the fact that it is being offered by a company that the consumer has good feelings about is enough to close the sale.

Customization, where the user takes an active role in shaping the interaction with the machine service, has also moved into areas beyond software system interaction. For example, the delivery of news and other time-sensitive information is increasingly being offered through the agent of computer-driven site that sift the news data stream according to criteria set by the user. Whether delivered by pull (e.g., web-based) or push (e.g., email newsletters) technologies, the user can instruct the server, within a range of choices set by the programming, to classify and deliver information according to those classifications.

Again, in certain respects, nothing new here. News gathering and delivery strategies have evolved over hundreds of years, including coffee shops & taverns that specialized as collection points of certain kinds of information, diverse newspapers & news magazines focusing upon specific content or editorial positions, and various kinds of broadcast services whose content and delivery might commonly be shaped by the technologies employed and the strictures set by public policy.

The software industry has continued, of course, to work to refine the agents that act to improve the usability of their own products, seeking to increase the utility of (and, thus, the demand for) their products. Yet, it appears that there are limits to their abilities to accomplish this on the scale of the individual. The ubiquitous Microsoft Office family of tools is rife with instruments that actively seek to help the user. Auto-correction of typographic errors in Microsoft Word has been a boon to many – except when the software insists on making a change that the user does not want. Microsoft's Office Assistant ("Clippy," almost certainly an application of the ill-fated "Microsoft Bob" technology) has seen massive investment, only to be side-stepped owing to the virulent hatred that it has engendered in a sizable fraction of the population the tool was expected to service.

The potential of this family of technological developments is huge. As our tools for collecting, manipulating and acting upon information become more capable, they have also tended to become more complex, limiting their accessibility to those who might benefit most from their use. The introduction of the notion of a software-driven mediator between the user and the service that is programmed to adapt to the user's strengths, limitations and revealed preferences is a clever approach to the problem, and proponents claim it has led to substantially more penetration of these services than might otherwise be expected. However, as these tools have migrated from the land of pure software and into the realm of information services more generally, there are reasons to inquire into how these "digital familiars" change the conventional into something less so.

Rather than "agent," the notion of a "familiar" seems a more apt term in this context. As with Prospero's Ariel, the familiar is an agent of the magician/user, and its talents in channeling magical forces helps the user to achieve his ends. However, the familiar does not slavishly hew to the user's will, but is instead an independent agent whose motives are not necessarily aligned with the user's.

For both of the non-software examples, sales/marketing and information/news gathering, the notion of "trust" emerges as a key element of the activity. In the case of sales, the goals of the sales agent

is to develop a trust relationship with the consumer, so that she will be willing to make a purchase that will leave both the salesman and the consumer better off. In the case of news/information gathering, the issue of trust arises through the notion of the authority of the information source – can the information be "trusted?"

For both sales/marketing and news/information, society has seen the need to erect institutions to ensure equitable and transparent relationships in these transactions, particularly as the scale of the services has increased. Fair trade, truth in advertising, "fair and balanced" news — these are public policy goals that have emerged in the face of increasing concentrations of power on the side of the service providers, particularly as these interactions have moved from the personal to the corporate/industrial.

The "digital familiar" is presented to the consumer as an electronic servant. The attentive salesman, the easily-directed news clipping service and the host of complements being deployed today extend the abilities of the user/consumer in a host of exciting and innovative ways. But their introduction into daily life also raises a set of issues that, if addressed at all, is being handled without a complete consideration of their scope.

Fundamentally, the key issue is that, although the "digital familiar" poses as the servant of the consumer, the consumer is not the master of the familiar – either personalized or customized. While the familiar can mimic the development of a relationship, no such relationship is actually forming. Rather, the familiar is seeking to engender trust, without the reciprocal responsibilities that are a part of normal relationships.

As a consequence, the familiar is perfectly capable of sharing queries, information, analyses, and assumptions about the user that would be considered gross invasions and betrayals in the real world. Information divulged to a friend is constrained by the relationship, trust and social mores; the "digital familiar" may build upon those cultural assumptions, but it is not constrained by them. Moreover, the familiar, armed with an appreciation of the user that is developed through inhumanly attentive collection and analysis of user behavior, is positioned not merely to serve, but also to shape (if not control) the actions of a user whose whole relationship with the familiar founded upon his/her relative ignorance of the familiar's true capabilities and inner workings.

Already the information collection, organization and reselling business has grown to a scale that has raised concerns among public interest groups and activists. As "digital familiars" become an ever-increasing part of the way in which users interact with the world around them, it is going to be vitally important that there be an exploration of the ways in which the gaps that these tools exploit can be filled, either through the expansion of current institutions of management and control, or the development of new ones. And, increasing attention will have to be paid to exposing the interests and enforcing the responsibilities of the men "behind the curtain" who are financing the programming of these familiars.

Otherwise, when the public finds out the extent to which the "servants" have been talking behind their backs to characters they deem unsavory, their reaction may lead to the crippling not only of this area of innovation, but many others in digital communications.

Open Issues/Questions

1) Relationships
   a) Is forming a relationship with a computer program the right working metaphor for constructing a familiar? A healthy one?
   b) What are the pitfalls of constructing a "cutout" in a relationship, particularly when the cutout is under the (complete) control of one party? What are the implications of increasingly substituting digital, software-based intermediaries for traditional interactions?
2) Trust
   a) A chain of trust -- The development of digital familiars creates a complex chain of relationships and dependencies, with many actors working to refine their own piece of the problem. Does this increase or decrease the stability of the trust relationships? The design objectives? What about other social objectives – responsibility, liability, monitoring?
   b) How might liability and other forms of obligation and responsibility be assigned along this "chain of trust?" What institutions exist already to service these requirements (contracts, product liability, etc.)? What failures need to be addressed?
   c) What might be learned from other domains where trust/reputation are key elements, yet complexly derived through indirect and direct interactions – e.g., academia, scientific research? (How to interpret the list of authors on an academic paper, for example)
3) "Where is the brain?" (From *Harry Potter and the Chamber of Secrets* – "Never trust anything that can think for itself if you can't see where it keeps its brain")
   a) Who's really in charge – (a) the algorithm or (b) the data collected to drive the algorithm, or (c) someone else?
   b) The locus of power in expert/client relationships has evolved over time. What are the underlying assumptions, and how has the introduction of these sorts of technologies changed them? Have the institutions that have grown up around those assumptions changed along with them?
4) What's the nature of the "harm" that worries those who consider these technologies now?
   a) Direct "harms" might include price discrimination, sales of one's information (privacy), manipulation into doing things one might otherwise not do
   b) Indirect "harms" may include "cocooning," the elimination of diversity through the creation of an "echo chamber," a world shaped to reflect the individual's view rather than reflecting reality
5) What does concentration in the ownership and application of these technologies do to shape the concerns? Would ubiquity and universal access change the problems, or just change the emphasis?
6) Does the fact that most creators of these technologies are commercial, rather than governmental, moderate these concerns? Should it? Do pernicious uses by commercial or governmental entities present the greatest long term threats?
7) Familiars
   a) Fundamental technical question – how far can one go with this technology?
      i) What limits its effectiveness today? Are there efforts to tackle these limits now, or is the field focusing on other issues?
      ii) What are the limits on data mining? What are the boundaries of the field today?
      iii) What about our understanding of human cognition in this domain? How sophisticated is our ability to direct or induce conscious and unconscious responses? To condition?
   b) What constitutes effectiveness in this application? How do the alternative strategies (active vs. passive, visible vs. invisible) influence this?
   c) How to balance the benefits of mimicking real world relationships against the issues raised in the paper. Is it worth it?