

Personas' Progress

Reflections on a Group User Research Project

Michael Wojcik

November 5, 2009

On Personas and Rhetoric

In the classic rhetorical triad of author, work, and audience, personas are most obviously related to the last, since they serve as an abstraction of the intended or ideal audience.¹ Of course this is true to some extent for all user research products—which are after all aimed at discovering users—but some forms, such as usability testing, exist at the interface between audience and work. Personas on the other hand may take the nature of the work into account only in determining the contours of the intended audience. Viewed as a rhetorical tool, then, personas are in one sense narrow (because they focus almost exclusively on audience) though also broadly applicable (because they can potentially be used in refining a wide range of works). This was one reason why we extended our report with some examples of how our personas might comment on the OurMichiganAve.org <http://OurMichiganAve.org> website: we wanted to connect them to the work, in the results illustrated by the examples, and to the authors, in our capsule narratives of how website designers and developers might go about using the personas to derive those results.

It's also interesting to consider personas in terms of the other key structure of classic European rhetoric, the five canons. Personas can be used in various invention strategies; for example, considering each of a set of personas in turn as an audience member (a user of the website, in this case) can reveal missing features users of that type would want, or obstacles they would encounter. Style and delivery, which are difficult to separate in non-oral media, can also benefit from the insights gained by reviewing personas. Style matters are sometimes left to visual designers, for example, but personas can remind them that some users have different preferences and requirements (particularly special-needs users); and users with limited access, such as dialup, will be particularly penalized by bandwidth-heavy sites, which is clearly a question of delivery. On the other hand arrangement, in the website context, is often primarily a question of usability and market research; here personas can be used to offer suggestions, but as often a persona-based analysis produces multiple "ideal" arrangements with no guidance as to which is preferable. (Sometimes this motivates sites to offer different arrangements based on user preferences.)

Besides arrangement, though, the canon where personas may be most useful is memory. (Ironically, in the age of mechanical reproduction, and particularly in the modern era of

¹Under other rhetorical models, such as Burke's pentad, personas might be positioned somewhat differently. Further investigation along these lines might yield useful insights about personas, but I don't have the time or space to consider it here.

ubiquitous information technology, questions of memory are often neglected; some rhetors refer to it as the “forgotten canon”.² If you have TelePrompTers and PowerPoint, why memorize your speech?) This was a point we tried to emphasize in the introductory part of our presentation, and at a number of places in the final report. Remembering the diversity of your users is difficult, particularly when that information is present only as demographic statistics or market analyses (or not at all, as is all too common in corporate practice). Personas leverage the innate social skills of the human mind to help designers and developers keep a wide range of audience members at the forefront of their considerations.

So personas clearly have a fairly natural fit in the schemas of classical rhetoric, even when they are developed for non-oral, non-print media. What about visual and digital rhetorics? Certainly personas make use themselves of visual rhetoric, by incorporating visuals such as the portrait photograph (one of those social “hooks”, and also shorthand for the persona’s dominant attitude) and graphs. Delivering personas primarily as single sheets—often also supplied in poster form (as we did with our report) for group use, or card form for individuals—makes for easy reference and allows arranging them in various ways, eg for ad hoc side-by-side comparisons or tacked to a wall as row or column headers. Single-page deliverables with plenty of visual content also assure the reader that this is a tool which makes only moderate demands on the user, in terms of opportunity costs; there isn’t a lot to read. And the visually attractive presentation encourages reading the persona sheets and makes them more memorable.

The relationship of personas to digital rhetoric is a bit less clear, partly because the boundary between traditional and digital rhetorics is of necessity not well defined. As artifacts, personas can be delivered in print or electronic forms (we did both), and for this application both have their uses: it’s easy to distribute digital persona sheets, but the print versions are handy for spreading on a table or tacking to the wall. (This duality is not surprising; as Elizabeth Losh points out, print rhetorics and digital rhetorics continually inform one another.) It’s probably more useful to look at some specific examples of digital rhetorical theory to see how it applies to personas. In the report we cited Phillip Agre’s article, for example, because his conception of a user “community” as a set of people with shared positions and goals—and not necessarily other common attributes, in particular interpersonal connections among community members—is particularly relevant to websites, which may be used from all over the world by people who have no other common contact. Here one of the strengths of personas is that they seek to represent clusters of user needs, attitudes, and goals, without necessarily trying to also account for other dimensions of social existence (as a demographic survey might, for example). That lets a persona capture the essence of an Agreist community without the “noise” of other axes of connection or social difference. Another example: Bill Hart-Davidson’s well-known “Texts that Transform” essay argues for digital texts that present one of a number of forms, based on user needs, and the conflicting—and, crucially, unresolved—requirements of different personas can both motivate and lay out a roadmap for such transforming websites.

²I don’t have a citation for this. I’m not thinking of any specific instance—I just recall having heard this comment more than once in passing, for example in an audience comment to my 2009 CCCC presentation on color and typography in websites.

Corporate Approaches to Personas

I found most of the persona readings from corporate sources useful mostly as background information, particularly as I'd already studied personas in WRA410, where we read about them in Jesse James Garrett's web design book and discussed how they're used in practice with a TechSmith employee. Still, there were useful new bits in the readings. From Quesenbury, I picked up a list of usability aspects that can be investigated with personas, such as role, context, etc. Calabria's article had some useful tips for researching for personas, and the connection it made to Alan Cooper's work generated some ideas for me. Ford's description of "attitudinal" or "behavioral" personas was a useful concept, with its emphasis on defining personas in terms of how people approach a task. I was less enthusiastic about the Masten & Plowman "Digital Ethnography" piece, which I found weak on technical accuracy and a bit rich in technophilia excess. In any event, their digital ethnography method did not end up influencing our methodology in any significant way.

Certainly the most influential of the corporate readings, for our group, was the Maxwell & Schulman "Personified Segmentation" presentation. Their explanation of how to combine quantitative data with conventional persona research techniques inspired us to look more closely at the quantitative data already available for the Corridor (in the various reports we had available), to gather more quantitative data (notably the OMA profile statistics), and to do quantitative analysis as part of our persona development process. This led to some of our secondary deliverables, such as the tag clouds, as well as influencing our creative choices.

On Labor and the Division Thereof

In some group efforts, it's easy to determine (during or after) how, when, and what each member contributed. Certainly that's been my experience in the workplace, where people are acculturated to thinking of themselves in specific roles, and management imposes ubiquitous work assignment and output tracking practices. Even in an agile environment, where people may participate in multiple feature teams and work assignments are short-lived, we do fairly rigorous tracking.

That wasn't the case with this project. Much of our work was intellectual and happened in group sessions through brainstorming and other forms of discussion. Outside group meetings, team members typically self-selected tasks based on lists the group made of remaining work. (For the final report, we broke tasks down into a Google Docs spreadsheet; group members would put their names next to individual items as they started working on them, and update the progress column as they went.) There was no sense in the group, as far as I can tell, that work distribution was significantly unbalanced, so no one kept a formal history of who contributed what. Everyone in the group expressed regret that the groups wouldn't be maintained for the whole semester; while we understand the pedagogical and practical advantages of organizing into new groups, I think that demonstrates a shared sense that the group functioned well as a unit, and everyone pulled their weight. That said, I believe I can describe some of my individual contributions.

First, a survey of what the group produced: We had four major work products that we delivered to class: the initial presentation on persona theory and its accompanying class activity; the checkpoint presentation; the formal presentation at the CCED; and the report (with its accompanying materials). Everyone in the group worked on all of

these, and everyone participated in all three presentations, except for Andrea's unfortunate medical-related absence from the final presentation (but she had prepared a section to deliver, and contributed major portions of the presentation itself, such as the final designs for the persona sheets). Besides those, we had a variety of minor deliverables, such as in-class writing exercises.

Internal work products were more varied. At each of our group meetings, except those where we were working directly on a work product, someone took notes. I did this once or twice; I believe everyone else did so as well, at different times. Those notes were sometimes distributed as emails and sometimes as Google Docs. Since group face-to-face meetings were difficult to arrange outside class,³ we typically worked through email and Google Docs. The group produced over 150 email messages that I'm aware of (having been either the originator or a recipient), and 19 Google Docs with numerous revisions. Intermediate and partial drafts of work products were frequently sent around the group—for example, for the checkpoint presentation each member contributed a few slides, and we emailed those around so everyone could see what other members were doing. Obviously, the major internal work products were gathering and analyzing data that went into the construction of the four personas.

Of the four major work products:

- The initial presentation on persona development: I created the initial three slides on persona types, and edited together everyone's slides into the final presentation (mostly because we were sitting around my laptop at the time). I also created the prototype persona slide at the end.
- The checkpoint presentation: again, I dropped in the example persona slide (taken largely from the previous presentation). My main contribution was the final four slides of examples of other kinds of deliverables, such as the vignettes.
- The final presentation: I created the cover slide and the following three slides ("What's a Persona?", "Research Methods", "Analysis"), which I also presented. I edited the slides created by each participant into the final presentation. I made any number of other minor contributions here too, such as downloading and cropping the persona portrait photos (which other group members mostly found—I think I may have been the one to find the photo we ended up using for our "James" persona).
- The report: I created the initial skeleton and outline, with input from (and revisions suggested by) other team members. I wrote the abstract and introduction, the section "Gathering Data from Other Sources", some of the text in "Overview of Personas", and some fragments of text elsewhere in the report to smooth transitions or explain some detail. I did most of the collecting of content supplied by other group members, editing it into the final document, correcting style inconsistencies (always a problem with Microsoft Word), and proofing, though the last set of changes was done by Andrea. I also did all of the material related to the OMA profiles.

During the research and analysis phase I was not able to collect as much data as some other group members, partly because I was out of the country for a week, and was working

³In most weeks, the only time during reasonable hours when everyone in the group was available, outside class, was a small window on Wednesday evening. This is an inherent problem with groups that mix undergraduate and graduate members, with their very different class schedules.

intensively on the Feminisms and Rhetorics Conference prior to that. I did do one interview (which was useful because it added a datapoint that was demographically distinct from most of our other data). I also did the collection and analysis of contributor data from the OMA site, which gave us some interesting insights for current site users.

Again, though, I'd like to emphasize that the majority of the work the group did was joint and difficult to disentangle into individual efforts. I'm proud of my personal output for this project, but much more so of the results which can only be properly credited to the group as a whole.

On Directions and Destinations

It's universally acknowledged that, at least for intellectual labor, group productivity does not scale linearly—two people are not twice as productive as one, four not twice as productive as two, and so on.⁴ There is inevitably a fair degree of overhead in group projects for coordination and communication, particularly in a context like this, where (unlike the typical workplace) we have very limited time together during the week and highly conflicting schedules. Then too, with a creative project like this one, much time has to be spent sharing ideas and confirming the direction in which to proceed; each group member is going to arrive at useful ideas, only a small fraction of which can be employed in the final product. And, finally, our loose, ad hoc system for managing task assignment—essentially “take what you like from what needs to be done”—didn't lend itself to efficient coordination, though I'd be hard-pressed to give it up, since it led to such a productive and pleasant work process. So there were many flurries of moderately anxious emails as we tried to find out who was doing what, and what still needed to be done. In short, there are certainly things that could have been done faster and more consistently had this been an individual project or a significantly smaller team.

On the other hand, since everyone in the group made major, creative contributions, of course the products benefited in many ways from the group structure. Our four different personas were created by four different members, using input from the whole group; that led to important divergences in the imaginary personalities and demographics behind each persona. (This is of course good, since the set of personas needs to represent a diverse user base.) The group members who conducted interviews had contacts with different populations, so we were able to widen our target population. Members had a variety of skills and knowledges; only Laura would have been able to do the kind of data analysis she performed, for example, and Franny's idea of mapping personas onto orthogonal axes was brilliant. I personally was particularly glad to use my writing and editing skills and presentation experience to help the group.

Now to the process itself. Ultimately, I'm not sure just how valuable our direct user research—the interviewing in particular—was to the construction of our personas; I think we could have arrived at a similarly valid and useful persona set without it. However, it was a useful exercise in itself, and it had useful side consequences, such as the discussion with Penny Gardner and others at the Community Networking Forum. The data analysis, and particularly the keyword analysis and the OMA profile analysis, gave us a lot of important data. So did reviewing the two Corridor reports.

The creative process was terrific, even with the (sometimes a bit overwhelmingly) tight

⁴In my professional field, software development, the classic exposition on this theme is Fred Brooks' *The Mythical Man-Month*.

schedule and workload. From our initial brainstorming sessions to the final push to finalize our persona sheets and create first the final presentation and then the report, everyone seemed to sincerely enjoy the creative work we did.

I know some group members are still nervous about presenting, but I enjoy it, and working with interesting presentation materials like the ones we created is always good. In some ways, the three presentations were as useful for us in reflecting on the site as the final four personas will be, I think. Developing the additional deliverables (vignettes, persona suggestions for the site, etc) was a way to explore how developers might use the personas.

The final report was something we're all proud of, I think. (I have to give special credit here to Andrea for producing the persona posters and getting the report printed and bound.) At fifty-odd pages, including a variety of genres, and attachments, it's a substantial piece of work, and I think the degree of quality would be appropriate for corporate or academic use.

Persona research is inherently a softer form of user research than usability testing or experience mapping, and captures less real-world detail than ethnographic surveys. But it has the advantage of encouraging creative exploration of the problem space, and its final products are pleasant and easy to use. I'm not ordinarily inclined to perform user research,⁵ but even I found this experience thoroughly engaging and informative.

⁵I'm a software developer. We *hate* users.