

CHAPTER 11

Personas

I often speak with people who tell me they have a good understanding of their users, but are puzzled because they can't seem to agree on what the product should be or do. In order for your efforts as a designer to succeed, you must ensure that every member of the product team understands at least the fundamental characteristics and needs of your customers and users, or you'll spend a lot of time talking in circles. The research methods described earlier in the book will help you gather the right information; personas will help you determine what that data means, convey that meaning to product team members in a compelling and memorable way, make better design decisions, and build consensus around a direction.

Definition and Uses

Personas are archetypes that describe the various goals and observed behavior patterns among your potential users and customers. In the terms introduced in Chapter 10, personas are a specialized type of composite model resulting from cross-case analysis, using primarily inductive reasoning.

A persona encapsulates and explains the most critical behavioral data in a way that designers and stakeholders can understand, remember, and relate to. Unlike simple lists of findings or other types of models, personas use storytelling to engage the social and emotional aspects of our brains, which helps each team member either visualize the best product behavior or see why the recommended design is good. The very name "persona"—from the Latin term for a character in a play¹—emphasizes their roles as storytelling devices.

Regardless of the sort of product or service you're designing, you will almost always find two or more distinct types of thought or behavior among potential users. The persona description that represents each user type includes a name, a photo, and a set of goals, as well as a narrative that covers mental model, environment, skills, frustrations, attitudes, typical tasks, and any other factors that seem critical to understanding the behavior pattern. The following narrative, Katie Bennett, is an example of a simple persona description for a digital camera user.

1. Yes, the correct Latin plural is *personae*, but popular usage made the argument moot long ago.



Katie Bennett

Thirty-two-year-old Katie would have gone into fine art if she felt she could have made a living at it; now she runs the business side of her husband's small landscaping firm and saves her creative ambitions for the weekend.

A couple of years ago, Katie bought a pocket digital camera so she could post photos of completed jobs on the company's Web site, which she put together using iWeb on her Mac. As she started experimenting with getting the best images, Katie realized that photography offered many of the creative opportunities she enjoyed in painting. She was hooked. Looking for a more capable camera that wouldn't break the bank, Katie went to CNET.com for advice. After looking at a few comparisons but not reading detailed reviews, she went to the nearest Best Buy and bought a Nikon D70 with its kit lens and an inexpensive tripod, relegating her compact camera to snapshots at family events. She also considered Canon's Digital Rebel, but chose the Nikon because it "felt more like a professional camera."

Katie got home and sat down with her new camera and its somewhat intimidating manual. After half an hour of fiddling, she was overwhelmed by the options and decided to give the auto mode a try. Katie started hiking about on weekends to shoot landscapes, from sweeping skylines to dew-covered flowers. She was pleased with some of her shots, but wondered why some weren't much better than what she could do with the pocket camera; many did not meet her expectations. After reading a few issues of *Outdoor Photographer*, she decided she might do better with different lenses. Confused by all the letters, numbers, and lens specifications, Katie went to the local specialty camera shop for advice on which macro and wide-angle lenses to buy; she did not expect the staff at Best Buy to provide good advice. She was reluctant to buy the cheaper lenses made by other manufacturers because surely Nikon would make the best lenses for their own cameras.

Katie is thrilled with her new ability to capture images of the local flora as she would have composed them on canvas. Though Katie enjoys it when people admire her photos, she's more motivated by the satisfaction of achieving her own creative vision. She can now capture the compositions she wants, but still isn't quite happy with some of her photos.

Katie gets up early on Saturdays to catch dramatic sunrises, frequents every park and beach in the area, and takes the occasional day trip. She loves the excuse to get out into nature. She goes out equipped with her camera, lenses, tripod, and a couple of 4 GB memory cards. Katie takes 100 to 300 shots on the average outing. She can often take her time composing a shot because plants and scenery don't move much, but sometimes needs to move quickly to capture a butterfly perched on a flower, or a shaft of light coming through the clouds just so. She usually takes a photo on the auto settings first, pointing the auto focus at the area where she wants to capture detail in the hope that this will set the correct exposure. She then dials the aperture up and down and takes a couple of shots to bracket the exposure; she read about this technique in her magazine. She still gets overly dark areas or blown-out highlights in many photos; she's increasingly frustrated by the intricacies of correct exposure. She deletes the worst photos from the camera on the spot.

Katie brings her camera home and plugs it into her Mac using the USB cable. She dumps the images into iPhoto and sees what she can learn from the bad ones before deleting them. She makes a few minor adjustments, but is generally reluctant to manipulate her photos, believing she should be able to get the right image in the camera to begin with. She posts her favorites on her personal Web site, uses them on her computer desktop, and occasionally orders large prints of especially good images via iPhoto. Katie feels a bit limited by iPhoto's organization options, but appreciates its ease of use and integration with other tools.

Katie is considering upgrading to a higher resolution camera, but is reluctant to spend the money unless she knows she can get the results she wants.

KATIE'S GOALS:

- **Be able to capture what she sees in her "mind's eye."** Katie knows she has an eye for composition, but is frustrated when her inability to master difficult lighting makes for a lackluster photo.
- **Enjoy the scenery.** Katie takes photos of nature as a way to enjoy its beauty. She doesn't want to be so focused on the mechanics of using her camera that she forgets to enjoy what she sees.
- **Feel like a "real" photographer.** Katie is proud of some of her images, but hesitates to think of herself as a photographer because she feels she hasn't mastered some of the fundamentals.

If you read carefully, you can see that nearly everything in the description conveys or reinforces something important about Katie's attitudes, goals, and behaviors. Even the mention of specific store, Web site, and product names has a purpose: They reinforce that Katie has ambitions but isn't sure where to get the best advice. Although the fact that she runs a landscape business is clearly fictitious and not directly relevant, one or two such details can make a persona seem like a real human being instead of a sterile set of characteristics.

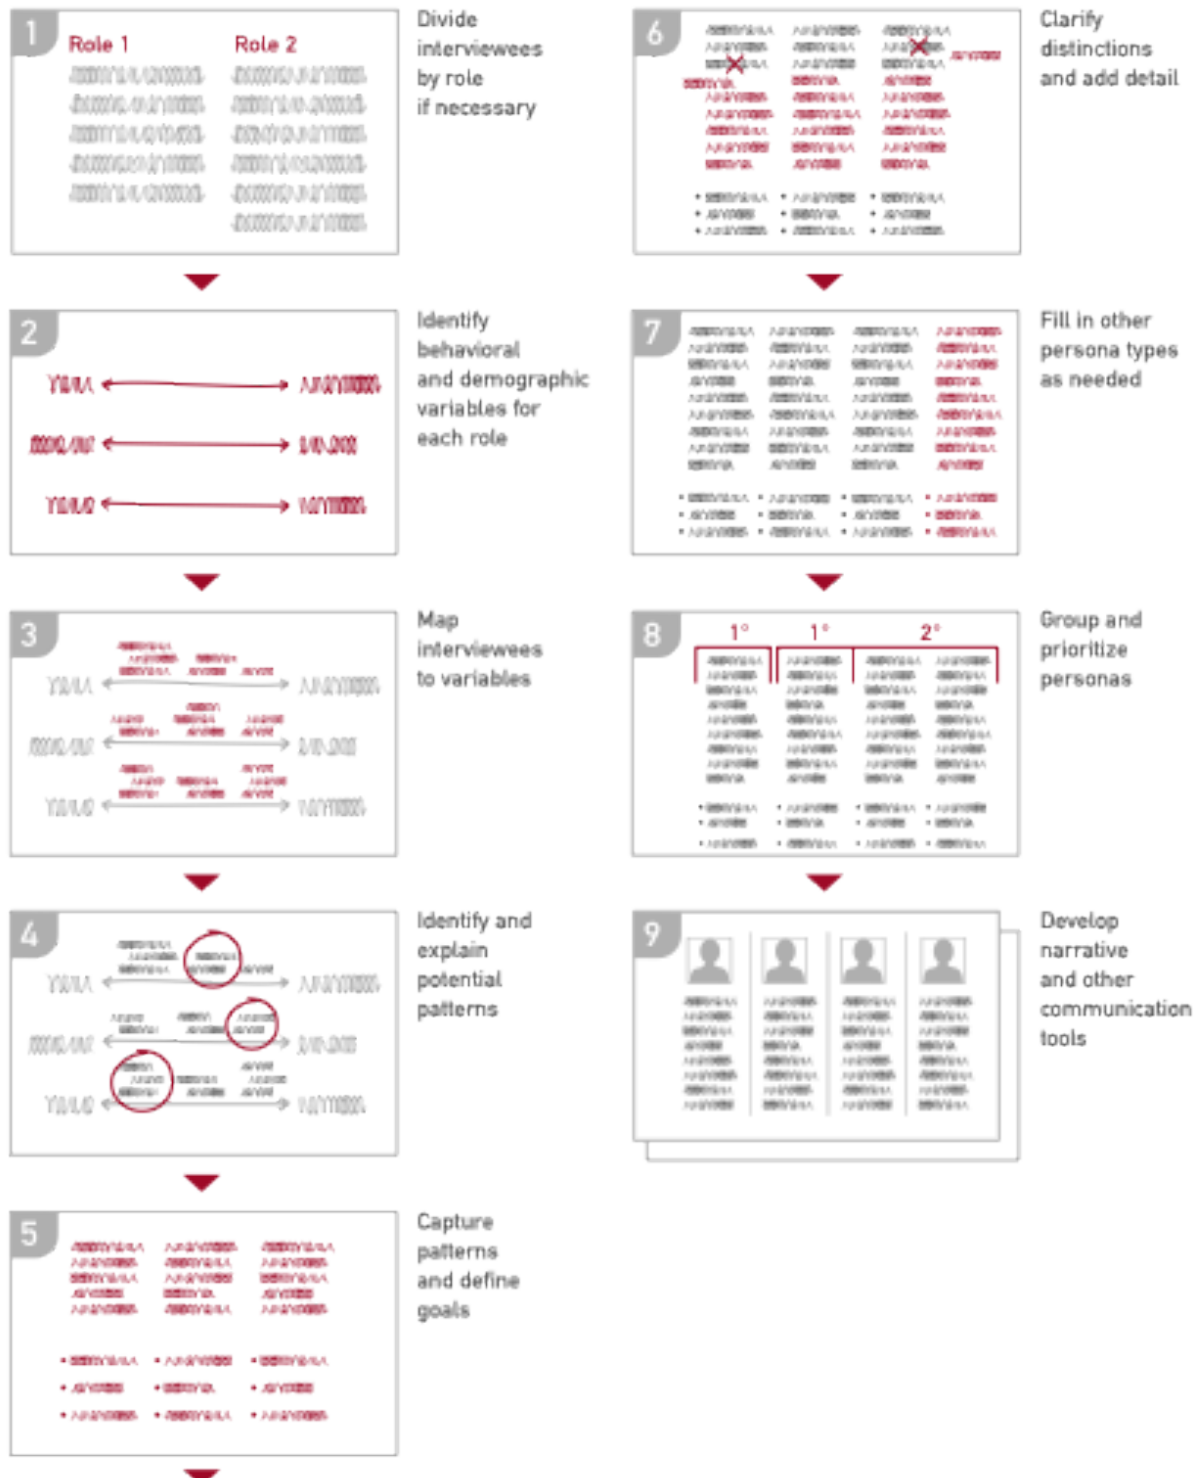
Personas are helpful in creating and iterating a design, building consensus, marketing the product, and even prioritizing bug fixes.

**Creating
personas involves
identifying the
critical behavior
patterns and
turning them into
a set of useful
characterizations.**

Creating Personas

Creating personas involves identifying the critical behavior patterns and turning them into a set of useful characterizations. A methodical analysis of the data is essential when those patterns are difficult to identify, and is worthwhile even if those patterns seem to leap, fully formed, from the data. It's entirely possible that your initial impression of the pattern is correct, but in many cases the obvious pattern is based on demographics rather than behavior, or is otherwise missing some critical factor. Even the most experienced persona creators can benefit from a rigorous approach, which usually requires only a day or two of effort.

Like most aspects of the design process, persona creation is not strictly linear, but does follow an approximate sequence, which is detailed in the following sections and illustrated at a high level in Figure 11.2. If there are clearly defined roles among your respondents, begin the process by comparing the interviewees in only one role at a time. From your data, identify behavioral variables—ways in which user behavior differed—and any demographic variables that seemed to affect behavior. Map the interviewees against the variables, then



Modeling

Figure 11.2 Overview of the persona creation process.



look for people who cluster together across multiple variables. Formulate explanations for that clustering to see if it really is a valid behavior pattern, then keep looking for any other patterns. Once you've exhausted the patterns within a given role, do the same for the other roles. Turn each behavior pattern into a persona by articulating goals and adding detail from the data. Fine-tune the personas as a set by clarifying the distinctions among them. Consider whether you need any other personas for political reasons before reviewing your rough drafts with stakeholders. Finally, prioritize the personas and develop the narrative and any other artifacts needed to describe them.

Simple as this all sounds, there are a hundred things you can do along the way that will make your personas either indispensable or a waste of time. That's not to say you can't experiment with how to create personas—after all, that's how they evolved as design tools in the first place—but when you're creating a real product or service with stakeholders who are skeptical, hurried, or otherwise not disposed toward design, you may only have one chance for your personas to succeed. The advice in this chapter should help you get them right the first time.

Step 1. Divide interviewees by role, if appropriate

As discussed in Chapter 6, roles are largely defined by *what* tasks people perform, rather than by *how* people perform those tasks. Roles usually seem clear in enterprise settings, such as end users versus system administrators. You can expect them to be less clear outside of a work environment, but they may still emerge; in some families, for instance, one family member primarily pays the bills, keeps in touch with relatives, or takes most of the photos.

It's often necessary to do your initial research recruiting based on job titles, which sound like roles but may differ widely among organizations. Don't be surprised, however, if it turns out that your research participants can't immediately be slotted into neat categories based on tasks. Job titles can hide specialization; for instance, a human resources manager might do a little of a lot of things, or might be focused on recruiting, staff development, or compensation and benefits. Sometimes the specialization is even less clear than this and may not be crisp enough to define separate roles. On a couple of projects involving financial analysts, for example, my team saw that some analysts were focused on routine tasks, such as monthly and quarterly roll-ups across various lines of business, while others did less predictable work, such as responding to unexpected changes or modeling various future scenarios. Many analysts did some mixture of all three, so treating these overlapping responsibilities as separate roles would have been problematic.

When the division between roles is very clear, such as between a surgeon and anesthesiologist or a loan officer and bank teller, it's best to treat the research participants in each role as a separate group for the purpose of identifying patterns. This is because large differences tend to obscure smaller differences; when you're comparing apples and oranges, for instance, it's harder to see that the apples are two different varieties, because the difference between the two types of fruit is so glaring. When you take the oranges away, you can see that one variety of apple differs slightly from the other in size, shape, and color. (See Figure 11.3.) For the same reason, customers and served populations should always be treated separately from potential users, as well. When the distinctions are muddy, as with the financial analyst example above, treat those interviewees as a single group. Table 11.1 illustrates this idea using the examples from Chapter 6, as if we've just completed the interviews.






Figure 11.3. Glaring differences, such as between the apples and oranges, tend to obscure more subtle differences, like those between the two varieties of apples.

Table 11.1. Examples of how to divide interviewees to begin persona creation.

Product	Expected roles	Observed behavior and next steps
E-mail system	<ul style="list-style-type: none"> – System administrator – E-mail account holder – External e-mail recipient – Purchase decision-maker 	System administrators, purchase decision-makers, and typical account holders were clearly distinct. All account holders were also external e-mail recipients, so they should be treated as one group.
Consumer digital camera	<ul style="list-style-type: none"> – Photographer – Buyer (usually the photographer) 	The people who bought cameras for others also used cameras themselves, so they're not distinct enough to separate.
Camera company Web site	<ul style="list-style-type: none"> – Potential buyer – Current camera owner – Camera dealer – Investor – Press – Job seeker 	Distinctions between buyers and owners were fuzzy, so they should be treated as one group. Dealers, investors, press, and job seekers were all distinct enough to treat as separate roles.
Inbound call center software	<ul style="list-style-type: none"> – Call center agent – Escalation agent – Call center supervisor – System administrator – Customer on the phone – Purchase decision-maker 	Most roles were distinct, but agents and escalation agents seemed to overlap quite a bit, so it's safest to treat them as one role.

Continued



Product	Expected roles	Observed behavior and next steps
Complex purchasing application	<ul style="list-style-type: none"> – Person who requests things – Person who processes requests – Person who authorizes purchases – Person who receives shipments – Person who pays bills – System administrator – Purchase decision-maker 	Many interviewees were clearly in one of these roles, while others overlapped two or more roles. You could include the multi-role interviewees in each of the distinct roles that applied.
Family calendaring system	<ul style="list-style-type: none"> – Family member – Perhaps one adult manages calendar 	Some adults clearly managed the calendar, but most shared management to some extent. Children and teens did no management, so it might be best to treat them as a distinct role.
Device used to deliver intravenous medications in a hospital	<ul style="list-style-type: none"> – Person who administers medication – Person who prescribes medication – Person who dispenses medication – Person who monitors patients – Patient – Purchase decision maker 	Patients and purchase decision makers are clearly distinct roles, as are the pharmacy staff who dispense medications to physicians and nurses. Nurses monitor patients and administer medications. Doctors may do those things once in a while, and only they can prescribe. There's a lot of overlap between doctors and nurses, but enough distinction that it's probably best to treat them separately.
Clothing store targeting women aged 25 to 40	<ul style="list-style-type: none"> – Woman shopping for her own clothing – Someone buying a gift 	Many gift shoppers also wound up shopping for themselves once they were in the store, so it's probably best to treat all interviewees as one group.

As you can see from the examples, whether to separate interviewees at this stage is really a judgment call. When in doubt about whether a role distinction makes sense, it's safest *not* to separate interviewees by role. It's easy enough to try a different approach if your grouping doesn't work well for subsequent analysis.

When roles are so specialized and consistent that every interviewee seems like a clone of the others, you might be able to skip steps two through four, though I recommend walking through them anyway. Often, you will find two or more important differences within each role. When the roles are very broad, as with consumer products and services, I strongly recommend not skipping any steps. Identifying good personas in large, diverse



audiences is sometimes more difficult than for most enterprise products because roles don't provide a convenient starting point.

WHAT YOU SHOULD HAVE AT THE END OF THIS STEP

Before proceeding to the following step, you should have a set of roles (if you have more than one) and a determination of which interviewees fit which roles.

Exercise

Determine how to separate interviewees by role, if at all, for the RoomFinder or LocalGuide. Use your own interview data or the data on the Web site.

Step 2. Identify behavioral and demographic variables

For each role, begin the cross-case analysis by identifying **behavioral variables**, which are aspects of behavior and attitude that seemed to differ across interviewees. Task frequency, mental models, and goals are all common types of behavioral variables. (If every interviewee demonstrated an identical concern or behavior, that's important information, but not useful at this stage.)

Once you have covered the behavioral variables, add any **demographic variables** or other facets of your interview data, such as environment, that seemed to affect behavior. Using the biggest writing surface you can find, lay each of these out as a spectrum or, occasionally, a set of multiple-choice options. This serves as an organizing structure for comparing individual interviewees in the next step, so leave plenty of room to write respondent names. Figure 11.4 is an example of such a variable set for people interacting with the health care system.

Almost any behavioral variable can be expressed as a continuum, whether this is a range from low to high or a contrasting pair, such as liberal to conservative. When the variable is quantifiable, label one end something like "many x" or "often does x" and the other as "few x" or "seldom does x." Don't worry about specific quantities while brainstorming variables; it slows the process. You can quantify later when mapping the data to the variables. To avoid confusion, always put the low end of any quantifiable variables on one side of the spectrum (usually the left) and the high end on the other. Table 11.2 shows some examples of how you might express behavioral variables related to shopping.

Begin the cross-case analysis by identifying *behavioral variables*, which are aspects of behavior and attitude that seemed to differ across interviewees.



Figure 11.4. A typical display of continuum and multiple-choice variables.

Table 11.2. Continuum behavioral variables.

Variable	Ends of spectrum	
Frequency of shopping	Seldom	Often
Price and brand sensitivity	More concerned with price	More concerned with brand
Price and quality sensitivity	Always chooses cheapest	Always chooses best quality
Attitude toward shopping	It's fun	It's a chore

There are sometimes a few variables, such as goals, that are difficult to express as a continuum, or for which a continuum doesn't make sense because the values are mutually exclusive. Express these as multiple-choice variables. For example, if you heard that people shop primarily to buy something they need, to spoil themselves, or to enjoy the challenge of finding a bargain, each of these reasons would become a choice. Table 11.3 shows some other examples of multiple-choice variables.

It's unusual to have more than a couple of multiple-choice variables in the whole set. If you find yourself wanting to express many variables as multiple choice, or if your multiple-choice variables have more than three or four choices, you're probably focused on the tools people use to accomplish the behavior rather than on the behavior itself. For example, it might be tempting to list many methods for sharing photos, such as sending e-mail, posting to personal Web sites, posting to online photo-sharing services, mailing hard copy prints, showing off physical albums, showing people photos on a mobile phone, and running slideshows from a laptop or set-top box. People use such a wide variety of tools that the

choices quickly become bewildering. Instead, examine what people are trying to accomplish with those tools. Some are sharing photos in a non-social way with people who aren't nearby, while others turn sharing into an interactive or social event; these could be characterized as remote sharing and local/social sharing.

In addition to focusing on behaviors rather than tools, good variables minimize subjective judgment by focusing on a single aspect of behavior. For example, whether someone is "very organized" or "not very organized" is somewhat subjective. I work with someone whose desk always looks like a disaster to me, but who can always put her hands on the desired information in about ten seconds or less, which makes it hard to argue that she's not organized. It's better to break the complex idea of "organization" down into several more objective variables, such as how many different systems someone has, how consistently he applies those systems, how much time he spends on staying organized, how often he misses deadlines, what the basic structures of his organizing system(s) are, and how quickly he can find a specific bit of information he needs.

Table 11.3. Multiple-choice behavioral variables.

Variable	Options			
Reasons for taking photos	Artistic expression	To share events with friends and family	To remember people and events	For business and insurance purposes
Most important criteria for choosing a car	Overall cost	Features	Emotional appeal	Environmental responsibility
Organizes messages by	Date and time	Sender or recipient	Topic	
Surgery starting point	Femur	Tibia		



By the time you're done identifying behavioral and demographic variables for a role, you'll probably have somewhere in the neighborhood of 20 variables, perhaps a few more or less. If you have considerably fewer, you may be missing important behavior. If you have a lot more, you're probably focused on mechanisms or trivial behaviors rather than fundamentals, which is problematic because it will make the critical behavior patterns harder to identify among all the noise. However, it's safer to err on the side of including a variable rather than excluding it. You'll see in subsequent steps that not every variable will be critical to identifying the behavior patterns.

TYPICAL BEHAVIORAL VARIABLES

Behavioral variables differ considerably from project to project and from role to role, so you'll have to identify many of them inductively from your data. One good way to do that if you haven't done a thorough job of coding is to pull out the notes for a couple of interviewees who struck you as very different, list the ways in which they differed, and then expand the list using other interviewees who seemed much different from the first two. In a sense, listing variables is like summarizing the answers to the questions you asked, along with the answers to questions you didn't ask. You can also get started by listing the following variables, which are almost always useful.

Mental models

As discussed in Chapter 7, understanding users' mental models of data and processes is critical to designing a system they can understand and use. Mental models are often one of the key distinguishing factors among personas. Such distinctions are usually most evident in how people organize their tasks or information. For example, an audiophile might think of albums as the fundamental organizing unit for music, since an album may tell a particular story or capture a point in an

artist's career. Others might not think of albums at all, except as the inconvenient way they had to access music when their collections were stored in cabinets rather than on hard drives. Chances are this mental model difference is related to other distinctive behaviors and attitudes.

Motivations and goals

The reason someone bothers to perform a task in the first place often leads to tremendous differences in behavior, so goals and motivations are almost always important variables for identifying patterns. Self-image and relationship to brands tend to fall in this category; someone's goal for buying a certain brand might be related to status, perceived quality, or other factors. If you find only two primary motivators, whether you represent this as a spectrum or multiple choice depends on whether some people are motivated by a little of each; people who are a little of both can be represented in the middle of the spectrum. Goals are usually best expressed as multiple choice unless you saw only two.

Frequency and duration of key tasks

Identify the common and important tasks in a broad sense. If you're designing a music library, for instance, these key tasks might be listening, organizing, sharing, and acquiring new music. Consider how often your various interviewees performed each task and for how long at one sitting. If the respondents differed much with respect to any of these behaviors, turn it into one or more variables, being sure to distinguish frequency from duration, such as "listens seldom/listens often" and "listens for short periods/listens for long periods."

Quantity of data objects

Users who deal with high volumes of data often have skills, behaviors, or needs not shared by



people who handle smaller quantities. It's often helpful to consider total quantity of data as well as quantity handled at one time. For example, someone importing music from a big pile of compact discs is handling a lot of data at once but may have a small library, while someone who only downloads one or two tracks a day may have an enormous library overall. Like frequency, quantity may or may not be critical to identifying the behavior patterns, but can often make those patterns more clear because of its relationship to other variables.

Attitude toward tasks

Does someone perform a task because she enjoys it, because it's her job to do it, or because it serves some higher goal? This variable is often revealing; without considering attitude, it would be easy to assume that two people who both spend a lot of time on a task are similar, even if they perform the task for entirely different reasons.

Technology and domain skill

Skill and experience may be primary drivers of behavior, but sometimes they're incidental to it. Be sure to separate skill or comfort with technology from skill or comfort with the domain; a nurse, for example, might not be comfortable with a computer, but she is expert at being a nurse, whereas a call center agent—who is unlikely to have much training—might need help doing his job.

Tasks people perform

While a role is defined by similar tasks, not everyone within a role will perform exactly the same tasks. In a work setting, some users may focus more on particular stages of a process, have more freedom in how they execute their tasks, or be more likely to do ad hoc tasks rather than the strictly routine. These factors tend to be affected by environment, skill, or seniority. In a consumer setting, tasks people perform may be more related to goals; someone who isn't worried about bouncing a check or motivated to save money, for instance, may be less likely to track every expenditure.

DEMOGRAPHIC AND OTHER VARIABLES

Once you think you have the important behavioral variables covered, consider whether there are any demographic or environmental factors that seemed to affect behavior. It's rare for more than a few demographic factors to matter, so including too many will add distracting

Behavioral variables differ considerably from project to project and from role to role.



clutter to your variable list, but feel free to include any that your gut tells you might be important. If you're listing multiple demographic variables, it might help to put them together at the bottom of your workspace.

The demographic variables that most often affect behavior are the ones described as recruiting considerations in Chapter 6: age, family structure, and geography for consumer products, and seniority, company size or industry, and geography for business environments.

Your user's height and physical strength seldom make much difference in the design of a Web site, but physical characteristics and abilities matter a great deal in designing physical products that are ergonomically appropriate. Physical characteristics should be included as a variable if you have a broad audience, but need not be considered a variable if your target population is within a narrow range of size and strength.

Physical environment is likewise a common variable for many physical products; does the respondent have a lot of space or a little, a cluttered space or a neat one? Again, use these factors as variables only when you saw distinctions.

WHAT YOU SHOULD HAVE AT THE END OF THIS STEP

Before proceeding, you should have a list of 20 or so variables that cover most of the codes in your notes, with most expressed as continua and perhaps a few as multiple choice. If you have any codes in the interviews that aren't reflected in your variable list, consider whether they should be included; you might have some codes that simply aren't useful. Compare a few pairs of respondents to one another and see if they have any differences, other than trivial ones, that aren't reflected on your list. If you have someone available to you who is experienced with this technique and at least somewhat familiar with your research, ask them to review your work for possible miss-

ing variables, variables that are too complex or subjective, variables that are really tools, and any other issues.

Exercise

Using either your own data for the Roomfinder or LocalGuide or the interview data from the Web site, develop the set of behavioral and demographic variables.

Step 3. Map interviewees to variables

The next step is to map the individual respondents to your variables as in Figure 11.5. This is easiest if you've coded your data, but not terribly difficult as long as you have a manageable data set and at least two people on the design team who've attended every interview. Team members who have attended a subset of the interviews can contribute to this discussion, but since it involves comparing each interviewee to every other, the most informed team members are ultimately responsible for the mapping.

The idea is to place each interviewee relative to the others along each spectrum (and in the appropriate multiple-choice categories, if applicable). Your placement need not be precise; this is not an exact science, so it really doesn't matter if respondent A is five or ten percent further along the scale than respondent B. What does matter is if respondent A is at one end and respondent B is toward the middle. This technique is similar to using a Likert scale, in that trying to break each continuum into more than about five zones is unlikely to be helpful. You'll find that some variables are polar, with people clustering near one end or the other and no one in the middle. Others may show clusters across the whole spectrum. If every interviewee is clustered around the same part of the spectrum, it means that variable won't be useful in distinguishing the patterns, so you can erase it from the board.



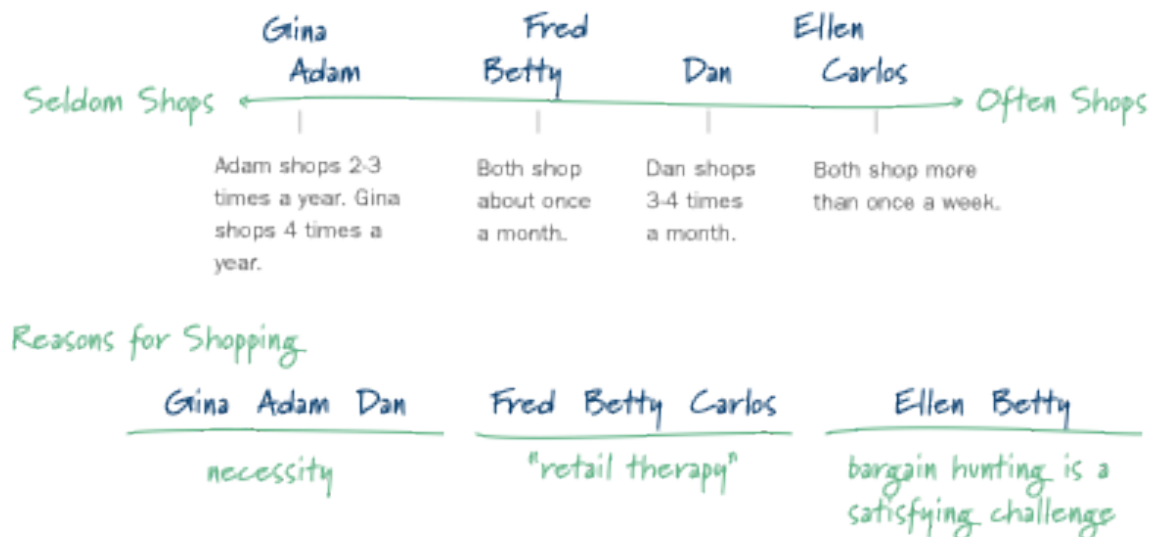


Figure 11.5. Example of interviewees mapped to continuum and multiple-choice variables.

APPROACHES TO MAPPING

There are two common approaches to mapping interviewees, both illustrated in Figure 11.6. Some people prefer to focus on one case at a time, placing a single interviewee on every variable before moving on to the next interviewee. This prevents digging through the notes for every interviewee as you consider each variable, but may mean you have to move a few interviewees you placed incorrectly. For example, you might think respondent A was close to the extreme end of a spectrum, but realize later on that he's closer to the middle when compared to respondent G. This need to erase and reposition is one of many reasons a giant dry-erase board makes a good tool for this activity. This approach can be challenging if the whole team doesn't have practice with making quick mental comparisons to other interviewees; if you're erasing and moving most people around as the process goes on, try the other technique.

You can also place every interviewee on a single variable at a time, so every variable is filled in

before you move on to the next. This may prevent a lot of erasing names and moving them around, but can require more flipping through interview notes for detail on each interviewee. This approach may be slower for experienced persona creators, but seems to be an easier starting point for most people. Consider trying both to see which works better for your team.

In either case, place interviewees on each scale relative to one another, not to what you believe of the population at large. Ensure that the placement is based on the data rather than your instincts or memory alone. An interview participant who felt strongly that she did something often might also leave you with that impression, whereas an objective comparison might show that she was about average in your sample. Look for some objective observation or statement to help you judge each quantifiable variable. Look for verification of any self-reported numbers; rely on your observation over the interviewee's statement unless you heard an explanation for the mismatch during the interview.

Place interviewees on each scale relative to one another, not to what you believe of the population at large.




Figure 11.6. You can place every interview on one variable at a time or place one interviewee at a time across every variable.

For variables that are less quantifiable, such as skill or attitude, consider any work product you saw, the terminology someone used, body language, and any other clues. For example, you can tell from an interview whether a camera user has no idea what aperture means, knows what it means but struggles with getting it right, or lives and breathes exposure without thinking hard about it. Ignore the interviewee's own assessment of his skill or comfort relative to other people; his points of comparison differ from the ones in your interview set.

When one designer proposes a placement for an interviewee, refer to your notes and see if you agree with the proposed position relative to other respondents. The discussion should continue until all team members agree. At least two of you presumably sat through all of the same interviews, so if there's much controversy, this is usually a sign that either someone is relying on memory rather than referring to the data, or something is wrong with the choice or expression of the variable. Sometimes there's disagreement about the meaning of the words you've used to label the variable; you can solve this by asking what each team member thinks is meant and not meant by the term. It's also possible that your variable is too complex or subjective, as discussed earlier in Step 2; break it down into more granular behaviors.

If you feel it might be helpful to go through this process partway through your interviews, give it a try. Because each part of the analysis and design process builds on the earlier ones, it can shed light on what you could have done better in the previous step. If you find you're unable to map every interviewee to nearly all of the variables, it's a good indication that your interviewing or data-capture skills need work.

TRICKY MAPPING SITUATIONS

Not every mapping situation is straightforward because there's no such thing as a perfect data set. The following sections outline how to deal with some tricky (but common) situations.

Behavior that varied by circumstance

It can be difficult to place an interviewee on a spectrum if her behavior varied by circumstance. What to do about that depends on whether one behavior was clearly dominant. If someone almost always behaves one way, but there are very rare exceptions, it's usually best to ignore the exceptions. For example, imagine that your multiple-choice variable is "motivation for taking photos" and the options are "artistic expression," "memories," and "sharing." Someone who sends photos of the kids to relatives once a month is clearly motivated by sharing. Someone who never shares, except for giving a framed landscape photo as a gift every few years, really doesn't belong in the sharing category; yes, she does share once in a great while, but it's not the reason she takes photos.

If the interviewee's behavior is consistently varied by circumstance, though, it's sometimes useful to note that interviewee on the scale in two positions, with a note about the circumstance that drives that behavior. Figure 11.7 is an example of this. If your spectrum ranged from "organizes photos by content" to "organizes photos by time or event," for instance, you could put me in the middle because I do some of both. However, it would be more informative to put me in two

places because I always assign attributes (such as species, location, and the pose or behavior) to my wildlife photos, but I always organize photos from family events by time. Thankfully, this sort of thing isn't that common because it can clutter up the display. If it happens with more than one or two interviewees, consider dividing the problematic variable into some finer categories.

Speculation about behavior

At some point, you probably heard interviewees describing what they would like to do if they had more time, if their tools were better, or if they themselves were better in some way. The emphasis in mapping is on each interviewee's real, current behavior, not on speculation or wishful thinking. People who are strongly motivated to do a certain thing find ways to do it, regardless of their tools or the time they have available. Someone who only wishes your Web site made it easy to compare prices with other sites doesn't belong in the same category with someone who spends the time to look up multiple sites, bookmark them all, and make a spreadsheet comparing prices and shipping costs across all of them. Mind you, she also doesn't belong in the same category as someone who didn't seem to care about price at all; words carry some weight, but not nearly as much as behavior.

Third-party behavior

As a general rule, you should exclude the behavior of third parties from consideration. If someone



Figure 11.7. You can note behavior that differs by circumstance this way.

tells you his wife performs a certain task, that's not helpful in profiling his behavior. The exception is when multiple people, such as the members of a family, will share a single device or application. Don't record this behavior on the spectrum as if it were your interviewee's, but consider how you might incorporate it if it seems to have a significant impact. One way might be to add another variable or two dealing with whether a device or application is shared, and the extent to which a third party's data or preferences interfere with your interviewee's tasks.

Missing data

It's possible that you will encounter missing data, especially in your earliest interviews or as your interviewing skills improve. If your notes don't contain the right information to help you place an interviewee on one of the scales, the best option is to avoid introducing potentially incorrect assumptions by not placing that interviewee on that spectrum at all. You might feel the need to fudge this a bit if your interview set is thin. Only do so if you're very confident in your conjecture and know that no major business decisions are dependent on the personas; otherwise, consult with your stakeholders to decide whether to get a little more data or go with your best guess. If you feel you must speculate about where to put an interviewee, indicate this visually, as shown in Figure 11.8; this will help you see what parts of your conclusions are built on shaky data.




Figure 11.8. If you must incorporate questionable data, indicate it visually so you can see where your analysis may be weak.

WHAT YOU SHOULD HAVE AT THE END OF THIS STEP

Before going on to identify patterns, make sure your display captures every interviewee in this role in relation to every variable you can fill out. If your data has a lot of holes in it, consider whether you need to do more research. Be sure that the whole team agrees with the placement of each interviewee. Capture the current state of your diagram; a digital camera is a convenient way to capture the contents of a whiteboard without doing a lot of tedious copying.

Exercise

For the Roomfinder or LocalGuide, map your data (or the data supplied on the Web site) against your set of variables.

Step 4. Identify patterns

The whole point of developing a set of variables and mapping your interview data to them is to facilitate the identification and verification of potential behavior patterns. What comes out of this step is a set of proto-personas: two or more behavior patterns defined by the correlation of multiple variables. Patterns are usually easy to spot in narrow roles (especially with a small set of interviewees) and more difficult to identify in unspecialized ones.

SPOTTING POTENTIAL PATTERNS

You might have a hunch about what the patterns are, but you might also have to look hard. Look at your entire set of variables at once. Rather than trying to find large clusters of people, start by looking for two or more people who frequently occur together. As in Figure 11.9, circle them with a colored marker so it's easy to see where their behavior and attitudes coincide. They almost certainly won't occur together on every variable, but if

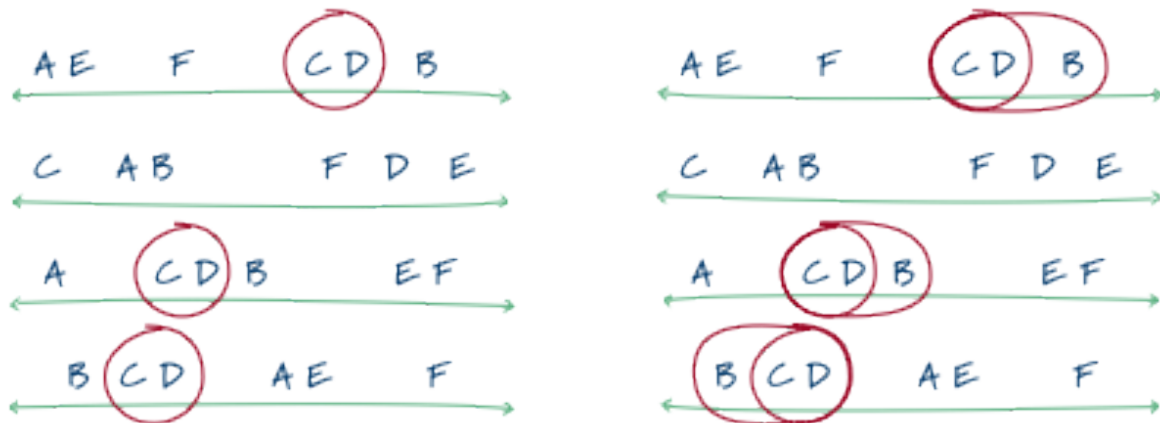


Figure 11.9. Start by circling the two people who show up together most often. Next, expand your circles to include anyone else who seems to appear with them most of the time.

they show up together on more than a third of the variables, they might represent a pattern. Step back and see if anyone else seems to be occurring frequently with your first pair. If so, expand your initial circles to include them. You might then find that these people appear with one of your original pair, but not the other, on some variables; circle these, too.

If the people in your possible pattern don't show up exactly together, but aren't all that far apart on the scale, these variables might still provide insight. Indicate these more tenuous relationships in a way that lets you see them, but with less emphasis. A dotted line is one good way to do this, as shown in Figure 11.10.




Figure 11.10. Use a dotted line to indicate more tenuous relationships.

Once you've circled everything that looks like it might be part of this first pattern, look at all of the variables where these people occur together. Try to explain how the variables are related (see Chapter 10, the section "Explanations and relationships," for more on explaining relationships in your data). Do you see one or more of the variables that seem like they might be the root cause for the other behaviors you observed? How is each variable related to at least one other? Is there a plausible explanation for that relationship, or does it seem like a coincidence due to your small sample? For example, it makes sense that a shopper's budget would affect how much time she spends comparing car models, but it's harder to believe that the amount of time she spends would be affected by the city she lives in.

After you've identified the variables that are legitimately related, consider the variables where these people did not appear together. Do any of them seem like they should be related, even though the results seem inconsistent with the rest of the pattern? If so, does that mean your potential pattern isn't really valid, or is there some reasonable explanation for the apparent mismatch?