

6 User Research

Get to Know the Guests You're Inviting to the Party

There are many user research techniques that can be used throughout the project lifecycle, either to better understand your users or to test out their behavior on versions of a site. This chapter will focus on some of the methods that are most commonly used in the beginning stages of the project.

These techniques will help you define the user groups that should be of highest priority during the project, put their needs and frustrations in context, and assess the performance of the current site (if one exists) using best practices in the field of user experience design.

Carolyn Chandler



Basic Steps of User Research

- 1. Define your primary user groups.** This involves creating a framework that describes the main types of users you're designing for—allowing you to focus your efforts in recruiting users for research.
- 2. Plan for user involvement.** This includes choosing one or more techniques for involving user groups in research, based on the needs of your project.
- 3. Conduct the research.** We'll cover the basic techniques here, such as interviews and surveys, and provide tips on how to go about them.
- 4. Validate your user group definitions.** Using what you learned from the research, you can solidify your user groups model. This model will then serve as a platform for the development of more detailed tools, such as personas (discussed in Chapter 7).
- 5. Generate user requirements.** These are statements of the features and functions that the site may include. You'll add these statements to your business requirements (discussed in Chapter 5) and prioritize them to become project requirements (discussed in Chapter 9).

This chapter will cover the first three steps, starting with the first: defining your user groups.

Define Your User Groups

Planning for user research at the beginning of a project can feel like a chicken-or-egg dilemma (which comes first?). How do you make sure you're talking to the right people, if you don't know yet who those people need to be?

One way to get started is to create an initial or provisional definition of the users you'll be designing for. This describes your site's primary user groups, which can help you focus your research efforts for the right roles, demographics, or other variables that may have an impact on how users will experience your site. User group definitions can be high level (a list defining each of your target user groups) or detailed and visual (a diagram showing multiple types of users, as well as how they interact with each other).

A high-level definition for a company's primary .com site might include the following user groups: potential purchasers, current purchasers, partners, and job seekers.

As you begin defining groups for user research, you'll start prioritizing user groups in more detail.

Your initial definitions are based on the collective knowledge of business stakeholders and project team members regarding the potential types of users who may be interacting with the site. Those definitions can be built by collecting some of the goals and attributes that different user groups may have. Here are the basic steps for defining your user groups:

- 1.** Create a list of attributes that will help you define the different users of your site (the next section will cover some of the most common)
- 2.** Discuss the attributes with those at the company who have contact with relevant types of users (for example, customers)
- 3.** Prioritize the attributes that seem to have the largest impact on why and how a potential user would use your site or application
- 4.** Define the user groups that you will focus on in research and design

The next sections take a closer look at some brainstorming techniques to help you collect these attributes and how to prioritize and model them (creating representations of the different user types that will help you focus your research efforts).

Create a List of Attributes

A good start for your attribute list is to gather and absorb any research or other documentation the organization has that could provide direction with regard to users. Here are some potential sources:

- ▶ Documents explaining company strategies, such as company goals, competitive information, marketing strategies, and business plans
- ▶ Market segmentations of current customers and other demographic data gathered by the marketing department
- ▶ Previously conducted user research (see Table 6.1 later in this chapter for some examples)

- ▶ Surveys, such as user satisfaction surveys and feedback forms
- ▶ Customer service reports covering frequently occurring issues

Next, identify people within the organization who have some insight into current or prospective users. The number and variety of people you should include depends on the type of project and its scope and timeline. If you know the initial definition of your user groups may have a short lifespan (for example, it's in use for only a month or two while user research is being planned), you may include just two or three participants. If you think the initial definition may need to hold you through a good portion of the design process (for example, if you only have this one to work with until you conduct usability testing, after some design has been done), include more participants and ensure you have a cross section of perspectives.

Some possible participants include marketing staff who are responsible for brand representation, segmentation, and campaigns; sales staff; product managers; customer service or support representatives; and trainers.

It's also good to include project team leadership and other business stakeholders in this exercise.

Ask the group to think of the different types of potential users they tend to interact with. Then ask them to list some of the common attributes they've encountered. Here are some examples of what could vary:

- ▶ **Primary goals**, as they relate to the subject matter of your site. Why are users coming to it and what are they trying to accomplish? For example, purchasing an item, trading a stock, or getting a specific question answered are common goals.
- ▶ **Roles**. This can be defined in many ways, but one way is to tie roles to the user's primary goal: job seeker, support seeker, potential client, and so on. Once you have more user information, roles can be subdivided by different needs or styles; for example, on an e-commerce site shoppers could include bargain-hunters and connoisseurs.
- ▶ **Demographics**, including age, sex, family (single, married, children), income level, and region
- ▶ **Experience** including level of education, level of familiarity with relevant technologies (often referred to as *technical savvy*), level of subject matter expertise, and frequency of usage (one-off, occasional, frequent)

- ▶ **Organizational attributes**, including the size of the company users work for, their department, type of job (entry level, freelancer, middle management, executive), tenure (long-term or high turnover?), and work patterns (remote work, amount of travel)

Once you have a list of some of the attributes that come up most often when stakeholders are describing potential users, you can start to prioritize them by their level of importance and then use that hierarchy to begin defining and modeling user groups.

Prioritize and Define

Which of the attributes listed above do you think have the greatest influence on how and why different user groups might use the site? Focus on the ones that you think will have the greatest impact on a user's goals or behavior. Prioritize those attributes, and remember the objectives you created in Chapter 4—they will help drive your choices as well.

An example best illustrates how to prioritize attributes. Say you're working with a company that provides tools for online trading of stocks, options, and futures. This particular company has determined that part of its strategy will be to engage *nonprofessionals* who are trading stocks on their own, online, and to encourage them to try trading new types of products such as options and futures. The company plans to do this by providing trading tools that are easy to use and targeted to those who want hands-on learning in a safe environment.

In discussing attributes with business stakeholders, you may find that the following ones seem to have the biggest impact on how individuals might use these tools:

- ▶ **Current frequency of trading**, specifically, frequency of direct online trading (for example, once a quarter, once a day, several times a day). Those who just dabble in trading (say, once a month) may not be serious about trying something new, while those who are already trading full time may not find much value in tools targeted to newer traders. But those who are active part-time traders could have a strong interest in the company's tools.
- ▶ **Number of product types traded**: just stocks or stocks, options, and futures. Those who are already trading all types of products may already have a preference for their own tools, but those who only trade one type may be ready to branch out to others.

- ▶ **Level of subject-matter expertise** (for example, familiarity with trading terms). This will help determine how much help they'll need along the way, with tutorials and glossaries.
- ▶ **Level of technical savvy** (for example, familiarity with making purchases online and online banking and trading). This will influence how much reassurance they'll need about information privacy and how advanced or simple the online interface needs to be.

You can prioritize these attributes because they may affect the user types you'll be targeting for research. If where traders live doesn't seem to have a real impact on how or why they trade, the Region attribute can drop off the list as a consideration for research participants. On the other hand, if the importance of a particular attribute sparks a lot of discussion it may be a good subject for a survey question or interview question (we'll be discussing surveys later in this chapter).

Comparing two or more attributes can help you prioritize as well. For example, if you make a chart using two attributes for online traders, you can start to see how groups fall within some of the ranges. **Figure 6.1** is an example of a rough user model you could make using the two attributes of Frequency of Direct Trading and Number of Product Types Traded; it also shows the resulting user groups that might form out of the discussion.

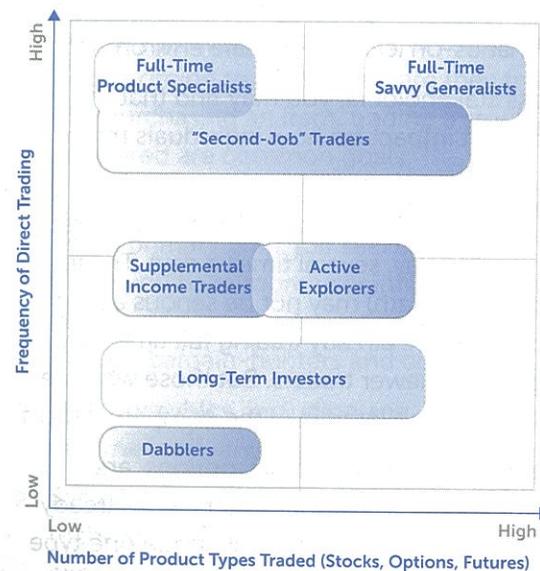


Figure 6.1 A chart of two attributes, representing a rough user model. Creating this model collaboratively can facilitate discussion about potential differences in user motivations and experience.

This user model provides a high-level way to discuss different user types. It's not meant to be the final model, and it doesn't label user groups exclusively (a user could be a long-term investor in stocks and also be actively exploring other possibilities in options or futures). But it does begin to express your understanding of different user groups and how they may be motivated to use your site.

This discussion concerning important attributes also helps you discover which ones you'll want to focus on when recruiting users for research. If you determine that Frequency of Trading is important, and that the priority will be to engage those who currently have a medium level of frequency, you'll want to define what *medium frequency* means (one to three times a week, for example) and recruit your research participants accordingly.

Speaking of research, let's talk about techniques you can use to involve users in your project.

Can You Design from User Models Alone?

There's debate within the user experience field about creating user models before research is conducted, because doing so can color your thinking before you have real user data, and because your project team or project sponsor may see the model as a replacement for user research. Using an unvalidated model does introduce more risk that your assumptions will be incorrect. In projects where you'll have no contact with users at all, however, a well-thought-out model (verified with sources outside the project team, such as a customer service group or training group) is better than having no model to use during design.

Choosing Research Techniques

Now that you have a rough idea of the user groups you want to include, it's time to plan the next step: your recommendations for the amount and type of user research activities to conduct during the project.

Table 6.1 presents some information on the most commonly used research techniques and when they are often most useful. Use this table as a reference to help you choose which ones best apply to your project. The next section describes each technique in more detail.

TABLE 6.1 Common User-Research Techniques

ACTIVITY	WHAT IT IS	WHEN IT'S USEFUL	CHALLENGES	TYPICAL TIME FRAME *
User Interviews	A one-on-one conversation with a participant who belongs to one of the site's primary user groups.	There is access to users, but type of access (in person, by phone, etc.) varies. You want to gain context but can't go to the user.	Getting straightforward opinions. It can be hard to gather information about attitudes and context, especially if interviews are conducted remotely.	2-4 weeks for 12 interviews: Up to a week to plan, 1-2 weeks to interview, and up to a week to compile results.
Contextual Inquiry	An on-site visit with participants to observe and learn about how they work in their normal, everyday environment.	The project team has little information on target users. Users work in a unique environment (e.g., a hospital). Users are working with fairly complex tasks or workflows.	Gaining access to participants. Going to users' environment may raise concerns about security, intellectual property, and intrusiveness. For business applications, it can be easier to visit on a workday.	3-4 weeks for 12 inquiries: 1 week to plan, 1-2 weeks to observe, 1 week to analyze and report results.
Surveys	A series of questions consisting of mainly closed-end answers (multiple choice) used to identify patterns among a large number of people.	You want to state results in more quantitative terms (e.g., "80% of the target user group said they never purchase cars online"). You're more interested in gathering information about preference than actual performance.	Getting an appropriate sample. Making sure questions are well-written so that you get accurate answers without leading respondents to a particular answer.	3-4 weeks for a short-run survey: 1 week to plan and write the survey, 1-2 weeks to run the survey, 1 week to analyze and report results.

TABLE 6.1 Common User-Research Techniques (continued)

ACTIVITY	WHAT IT IS	WHEN IT'S USEFUL	CHALLENGES	TYPICAL TIME FRAME *
Focus Groups	A group discussion where a moderator leads participants through questions on a specific topic. Focuses on uncovering participants' feelings, attitudes, and ideas about the topic.	The team believes that users' attitudes will strongly influence their use of the solution (e.g., if there have been problems with it historically).	Understanding how to target your questions to get the right information out. Facilitating the group effectively.	3-4 weeks: 1 week to plan and write questions, 1-2 weeks to conduct focus groups, 1-2 weeks to analyze and report results.
Card Sorting	Participants are given items (such as topics) on cards and are asked to sort them into groups that are meaningful to them.	You're working on a content source site with many items and want an effective structure for your user groups.	Determining which topics would be best to include.	3-4 weeks: 1 week to plan and prepare, 1 week to conduct research, 1-2 weeks to analyze and report results.
Usability Testing	Users try to perform typical tasks on a site or application while a facilitator observes and, in some cases, asks questions to understand users' behavior.	An existing solution is being improved. Competitive solutions are available to test. You have a prototype that lets users complete (or simulate) tasks.	Choosing the appropriate tasks to focus on. Determining how formal to make the test.	3-4 weeks for 10 users and medium formality: 1 week to plan and write the tasks, 1 week to run the tests, 1-2 weeks to analyze and report results.

* Typical Time Frame represents the time often needed from the point users are scheduled. Two groups of six to eight users are assumed (except for surveys, where the number of users should be larger). This does not include time for recruiting, which can take one to two weeks after creation of the screening questionnaire.

How Many Research Activities Can I Include?

Before you choose among the activities, ask yourself how much money and time the team can dedicate to user research. Consider the following situations to understand how much appetite your client company has for user research.

If project leadership and the project sponsors are comfortable with user research and are interested in using it for known goals, such as ensuring the site meets specific project objectives, then you're likely to have more leeway in planning for two or more activities, or for one activity that you conduct multiple times (for example, testing a design, changing it based on your results, and retesting the new design).

Considerations When Planning Research

When planning for any research techniques, consider the following:

- ▶ Why you're conducting the research: what you want to learn from it
- ▶ Who you're including: the primary user groups you outlined above
- ▶ How you'll get participants: recruiting people to participate and screening them (that is, asking questions to make sure they fall into the user groups you're targeting)
- ▶ How you'll compensate participants
- ▶ What space, equipment, and software you'll need, based on whether you'll be in person or you will conduct research remotely (see Chapter 14 for more on remote research)
- ▶ What you're covering: the primary topics
- ▶ How you're capturing information: the number of people involved and the tools they're using

Chapter 14 will cover each of these considerations as part of a detailed look at one of the most common techniques used by UX designers: usability testing.



Surfing

Steve Baty wrote an article describing different methods and how to choose among them based on the phase of development, your information needs, and the flexibility you have to incorporate user research. It's titled "Bite-Sized UX Research," by Steve Baty, UXmatters: <http://uxmatters.com/MT/archives/000287.php>.

If no one at the organization is familiar with user research and there's some resistance to it altogether, you may be better off proposing one round of research and picking the technique that you think will bring the most value to you, the project team, and the business stakeholders. Once you have the results of the research, the project team will have a better idea of what's involved and how the project can benefit. You'll then have a strong case for including more research later, if needed.

If you have room for at least two rounds of research, a good approach is to include one round during the Define phase, or early in the Design phase, to better understand the users. Then include one more round before development starts, to validate the design. For example, for a task-based application you might conduct user interviews before designing and then perform usability testing on a prototype later in the process. Or for a content source you might start with contextual inquiry and then include a card sorting exercise.

Note See Chapter 2 for more on task-based applications and content sources.

Let's take a closer look at each of these techniques and the ways they're commonly used.

User Interviews

User interviews are structured conversations with current or potential users of your site. These can be conducted over the phone, via video webconferencing tools (such as GoToMeeting or Skype), in person in a neutral location (such as a conference room), or, ideally, in the environment in which the user is likely to use the site. (This last situation is also great for conducting a contextual inquiry, covered later in this chapter.)

Interviews help you understand participant preferences and attitudes, but they should not be used to make formal statements about actual performance. If you're looking for specific information on how people interact with a site, it's better to observe them using it (for example, in a contextual inquiry) or ask them to perform tasks on the site (during usability testing). Site analytics can also give you some insights into some performance information that can be particularly strong when paired with interviews or inquiries that provide context for the data.

The Basic Process

For user interviews, the UX designer creates a list of questions aimed at eliciting information such as the following:

- ▶ Relevant experience with the site or with the subject matter
- ▶ The company's brand, as experienced by the participant
- ▶ Attitudes, for example, toward the subject categories covered (for a content source), the process being designed (for a task-based application), or methods of marketing (for a marketing campaign)
- ▶ Common goals or needs that drive users to your site or that of a competitor
- ▶ Common next steps users take after visiting the company's site
- ▶ Other people who are involved in the experience. For example, does a user tend to collaborate with someone else as part of the larger goal they're trying to achieve? Are they likely to share information or ask opinions of others along the way?
- ▶ Any other information that will help you validate the assumptions you've made about user groups up to this point—for example, whether the variables you discussed when creating a provisional user model really seem to influence the way users are experiencing your site

If more than one person is conducting interviews, it's a good idea to have a set list of questions and a scripted introduction that can be used to maintain consistency across interviews.

Choose ahead of time how structured you want the interviews to be. If you're going for a formal report, you'll probably want a high degree of structure,

where question order doesn't vary much and every question is asked, with few additions. If richness of data is more important than consistency, you may decide to opt for semistructured interviews, where you start with a list of questions but allow the conversation to follow a natural path, with the interviewer asking questions to further explore interesting comments (called *probing*).

The length of your interview can vary; 45 to 60 minutes is often the best range to shoot for. It gives you enough time to build a rapport and cover a wide range of questions without fatiguing your participant.

User interviews provide a rich set of data that you can use to write personas, which are covered in Chapter 7.

Interviewing Tips

The quality of the information you get out of an interview has a lot to do with the quality of the questions you ask.

Focus on participants' personal experiences. Don't ask them to speculate on what they may do in the future or on what others may do. This kind of information rarely predicts what they actually will do.

Don't ask *leading questions*—questions that imply a specific answer is the correct response or that influence participants to answer in a certain way. Ideally, questions are simple, neutral, and open ended. Some examples of leading questions are:

- ▶ What do you like about PseudoCorporation.com?

This assumes the user likes using the site. Use this question only if you also ask what they dislike about it.

- ▶ Does PseudoCorporation.com meet your expectations?

This can be answered with a simple yes or no, which doesn't give you much detail to help with your design efforts.

- ▶ Would you rather use PseudoCorporation.com or CompetitorVille.com and, if the latter, why do you think they are better than Pseudo?

This has a couple of problems: It's asking two questions in one statement, and it forces an implied opinion on the participant.

Better questions to ask are these:

- ▶ Tell me about your last visit to PseudoCorporation.com. Why did you go there?
- ▶ What do you remember about your visit?

If you're doing a large-scale, more formal set of interviews, you may want to include some multiple-choice questions. For the most part though, these don't give you very rich information. They can be hard for participants to follow when asked verbally, and they don't allow users to elaborate. In general, save that type of question for screeners or for surveys.

Perform a test run with someone, perhaps someone within the organization who isn't a member of the core team. This will help you discover questions that may not be clear and will also help you refine the timing and flow.

If it's possible, and the participant consents to it, record the interview so that others can benefit from hearing answers straight from the participant's mouth.

Contextual Inquiry

Contextual inquiry combines user observation with interviewing techniques. The UX designer goes to participants, ideally to the environments in which they're likely to use the site. For example, for an office application contextual inquiry would involve sitting at the participant's desk.

This method gives you rich information about the context a participant works within, including these:

- ▶ The real-life problems users are facing
- ▶ The kind of equipment they're working with
- ▶ The space they're working within—in particular, the amount of space they have, how much (or little) privacy, how often they are interrupted, and how they use the phone and paper (pay special attention to printouts they've posted or notes they keep handy).
- ▶ Their preference in using a mouse versus keyboard. This can greatly affect your design choices, especially if you're designing a tool that requires a lot of data entry.
- ▶ How they're working with others, in terms of both collaboration and sharing resources. If more than one person is using the same computer, for example, it will affect how you design login and security features.

- ▶ Other tools they're using, both online and off. How people use paper is especially interesting—for some tasks, it can be hard to design an online solution that competes with paper!

Inquiries combine observation time and interviewing time. They can last anywhere from a few hours to several days.

If participants can't dedicate at least 2 hours, you should consider just performing an interview. During an observation, it takes some time for the participant to adjust to your presence and act somewhat naturally, and this doesn't happen after just 15 minutes.

The Basic Process

Prepare a 10- to 15-minute introduction you can use with each participant. It should include the purpose of the inquiry, a high-level description of what you'll be doing together (the observation and interview), and how the information will be used. This is also a good time to get signatures on consent forms and to assure participants that what they share will be kept confidential.

Begin with some high-level questions about the participant's typical processes, especially ones that are relevant to the design of the site.

Let the participant know when you're ready to stop talking and start observing. Observation can range from active to passive. With *active observation*, a common approach is to have the participant take the role of the master while you take on the role of the apprentice. The master explains what he is doing as if teaching you his process. Active observing often gives you more background on the reasons for the participant's behavior, but it may affect how the participant works.

In *passive observation*, you encourage the participant to act as if you're not even there. Your goal is to observe behavior that is as natural as possible. For example, if a participant is talking to you, she may be less likely to take a call or go ask someone a question on a problem she's trying to solve, but if you're observing passively, you're more likely to see this happen. You can then follow up during the interview portion to ask about the reasons behind some of the behaviors you observed.

Either approach can work well. Generally, if you don't have a lot of time with participants (let's say, only 2 to 4 hours each) you may decide to use active

observation to ensure you get the depth of information you need. If you have a full day or more, passive observation offers a good balance of natural behavior and discussion.

Once you have information from your inquiries, you'll have a lot of rich data to sort through! So how do you identify patterns or trends in your results?

One way that's helpful is a technique called *affinity diagramming*. There are many great resources available on this topic, but here's a short description.

A Quick Guide to Affinity Diagramming

Affinity diagramming is the technique of taking a number of distinct and separate items (like statements made by users or observations made by a researcher) and grouping them together to form patterns and trends. Here are the steps involved in a simple affinity diagramming session:

1. Gather the team that performed the inquiries, with their notes
2. Give each person a pack of sticky notes and ask them to write a statement on each one, along with a short code that will allow you to track that statement back to a participant, such as their initials. Focus on statements that seem to have relevance to the site design, either specifically (a feature statement) or in a more general way (a statement that represents a participant's attitude to the company or subject matter).
3. Have everyone put their sticky notes up on the wall. You'll need a big blank wall if you're working on a large study; try to get one that you'll have access to for at least a few days.
4. Once all the notes are up, start grouping similar statements next to each other. This portion of the exercise can include the larger team. It's a great way to start sharing results.
5. Once groups start to form naturally, start labeling the groups to provide further structure. If some sticky notes belong to more than one group, you can write duplicates and place them in each appropriate group.

Note This method works well for contextual inquiry but can be applied to many other situations. For example, it's a great way to collaboratively create categories for unsorted topics, so it can help you move card-sorting results into additional levels of structure.

Patterns can emerge in many ways, so it's best to let them form on their own. However, here are examples of the kinds of categories that you might see, including the kind of statement you'd find in them:

- ▶ Goals: "I try to clear off all the open items here before I leave for the day."
- ▶ Mental models (includes statements that demonstrate how users are mapping external experiences to internal thinking): "I use this online tool as my briefcase, for things I reference a lot but don't want to carry around with me."
- ▶ Ideas and feature requests: "I wish this would allow me to undo. I keep moving the whole folder accidentally and it takes me forever to cancel out of it."
- ▶ Frustrations: "I'd ask the help desk about this, but half the time they don't know what the problem is either."
- ▶ Workarounds: "This takes so long to do here that I just end up printing out the list and working with it throughout the day. Then at the end of the day I enter in the results."
- ▶ Value statements: "This tool here saves me a lot of time, so if you're making changes don't take it away!"



Deep Diving

The quintessential resource on contextual inquiry is *Contextual Design*, by Hugh Beyer and Karen Holtzblatt (Morgan Kaufmann, 1997). The book also includes detailed information on interpreting results through techniques such as affinity diagramming.

For more information on mental models and how to understand them, take a look at *Mental Models: Aligning Design Strategy with Human Behavior*, by Indi Young (Rosenfeld Media, 2008). This is especially helpful when you're working on the information architecture for a content source.

Surveys

Surveys involve a set collection of well-defined questions distributed to a large audience. They most often consist of closed-ended questions (such as multiple choice questions) that can be easily collected with a tool that can display patterns among responses.

Surveys are good tools when you want to be able to state results in more quantitative ways (for example, "Of those surveyed, 82 percent of those who work from home state they have some form of high-speed Internet connection") than you would get with the kinds of open-ended questions that are used in interviews. However, you can gather qualitative information from them as well, about user perceptions and attitudes.

In the user experience field, surveys are often used to measure user satisfaction (with existing sites or applications) or to build or validate user models like segmentations or personas.

The Basic Process

As with user interviews, you don't want to ask questions that require users to speculate. Don't ask "If you got Feature X, would you use it?"

Unlike with interviews, in surveys multiple choice or Yes/No, True/False questions are best and easiest to analyze afterwards. They're also quicker for participants to answer.

Use surveys when you have questions that are factual requests for demographic data, such as these:

Of the devices listed below, which do you personally own? Choose all that apply.

- Computer
- Mobile phone
- Game system, such as Xbox, Playstation, or Wii

Or use surveys when you need to ask questions that are attitudinal with a set range of distinct choices. For example:

Read the following statements and select the degree to which you agree or disagree with each of them.

The Customer Service at Pseudo Corporation is responsive to my needs.

- Strongly Agree
- Agree
- Neither Agree nor Disagree
- Disagree
- Strongly Disagree

In particular, questions like the one shown in the second example are often used to supplement usability testing tasks. You can use this type as a follow-up question to find out if participants were frustrated when completing a task. Participants don't always like to state a negative opinion out loud, but they are often willing to express one when faced with a ranking system.

This brings out another point: Surveys are an excellent supplement to other forms of research you may be doing. For example, you can gain quantitative data from surveys to supplement qualitative data gained from user interviews or contextual inquiry (see Chapter 14 for more on quantitative and qualitative data). Combining two research methods provides a richer picture of the user than one method can provide on its own.



Surfing

If you want a high degree of confidence in your results and have the budget for it, there are formal tools available for measuring user satisfaction with regard to ease of use. These tools include questions that have been tested to ensure they are not leading or confusing to a broad audience. Some of the most commonly used are

ACSI (American Customer Satisfaction Index): www.theacsi.org/

WAMMI (Website Analysis and MeasureMent Inventory): www.wammi.com

SUMI (Software Usability Measurement Inventory): <http://sumi.ucc.ie>

When planning a survey, consider the following:

▶ Who are you targeting?

Use your provisional model to determine this. It'll make a difference in how you answer the rest of the questions here.

▶ What method for distributing the survey will give you the best results?

If your primary user groups tend to congregate in a particular location, you may get more results if you go there and set up a table for people to fill out the survey on paper. If your user groups are active Internet users, having an online survey could be the best choice for a large number of participants. Or you may decide your user group will be best found with phone surveys using a list of current customers.

▶ How much time will participants probably be willing to spend filling out the survey?

If you're providing some kind of compensation or they get some other benefit from filling it out, you can usually create a longer questionnaire—one that takes maybe a half-hour to complete. If not, you'll need to keep it short to help ensure people complete it—think 5 to 10 minutes. Either way, make sure participants are given an estimate of how long it will take and update them on their progress as they go through it (use page numbers like "2 of 4," for example, or show the percentage completed).

▶ How will you know when to start analyzing the data?

You may choose to run the survey until you reach a certain number of participants or until you hit a certain deadline, whichever takes priority.

▶ What tool will you use to collect and analyze the data?

If you're running the survey online, the tool you use to collect the data may have options for viewing and analyzing the results. If not, you'll need a method to enter the data into your tool of choice. For paper-based surveys this means a lot of data entry, so be sure you're planning for that time.

Focus Groups

Focus groups involve bringing together a variety of people within a target audience and facilitating a discussion with them. Common goals are to elicit opinions on topics relevant to the organization or its brand, such as past experiences, related needs, feelings, attitudes, and ideas for improvement.

A focus group is a good technique for several purposes:

- ▶ **Hearing a variety of user stories.** Open discussion is a great way to bring out the storyteller in all of us. When a focus group is going well, individuals build off each other's stories and ideas and remember situations they might not in a more structured one-on-one interview. The group format and energy can give people the time they need to recall these stories and share them.
- ▶ **Understanding relevant differences in experiences.** Most people are natural information sharers and want to compare favorite tools with others in their interest group. Often you can learn of competitive sites or services, or you'll hear tips for workarounds, resources, and support.
- ▶ **Generating ideas.** Although you don't want to make the group itself the designer, you often get some excellent ideas for new features or designs either directly from the group or from hearing about their work processes or frustrations. As with stakeholder ideas, be sure to trace these back to the core need (see Chapter 4) so you can be sure it's being addressed.
- ▶ **Understanding multiple points of a collaborative process.** If you're designing a process that involves multiple related roles and collaboration, groups can be a great way for you to fill in the gaps in your understanding of how people are interacting. For example, if you're working with a content source like an intranet, it can be helpful to gather a mix of those generating the content, editing the content, and consuming the content to identify the points where the process could be improved.

There's a lot of debate about the use of focus groups in UX research. It's not a good technique for testing usability (since users most often work individually, rather than in groups), and sometimes the group setting can unduly influence participants' statements. If planned and facilitated well, however, focus groups can bring out many insights that will be valuable to you as you're designing. Chapter 14 discusses this further in the context of concept testing.

The Basic Process

When writing questions for focus groups, consider the same tips you would use for writing user interview questions (covered previously in this chapter).

Begin with some of the easier questions, such as "Tell me about your last visit to PseudoCorporation.com. Why did you go there?" Save any questions focused on idea generation to the middle part of the group, when participants are feeling comfortable with you, each other, and the topic.

Assign time blocks to each topic and keep to them; it's easy for discussions to really get going and for time to slip by! If you're worried about time, put your most important questions in the middle of the topics list, after people have warmed to the activity but before any potential time crunch that could occur near the end.

Many of the logistics for focus groups will be the same as those for usability testing. (Chapter 14 offers suggestions on screening, recruiting, and scheduling.) The primary difference with focus groups is that you'll need a larger room with a table allowing participants to interact with each other easily. Shoot for six to eight people per 1- to 2-hour group session.

Give each person a nametag or a place card at their seat, so everyone can address each other by name.

The format of the discussion itself should include an introduction, which often hits these key points:

- ▶ Your role as moderator, and what you're expecting to get out of the discussion (for example, some of the points above)
- ▶ Why attendees were chosen to participate (for example, "You are all current users of the Pseudo Corporation site, and we've brought you together to find out about your experiences")
- ▶ How this information will be used—both in the design and from the standpoint of confidentiality
- ▶ That as the moderator, you're there to hear about their opinions and experiences. You want them to feel they can share honestly, so ask individuals to be straightforward but also respectful of others in the group.
- ▶ That there are many topics to cover, so at some point you will end a discussion on one topic to be sure you can cover all of them

This can then go into a round of introductions for group members, often including some kind of icebreaker question.

Your goal is to get everyone to talk on the first question, even if they just tell a short story. You can either start with one person and work around the table or let people answer naturally and then call on the people who haven't answered yet by name. Often you'll end up going around the table for the first few questions and then, when you feel the group is ready, with body language you can open up the questions to everyone.



Snorkeling: Body Language

A good understanding of body language can be an amazing tool when moderating focus groups or any user research conducted in person. It can help you understand when someone is feeling frustrated, excited, angry, or threatened, so you can identify when you should try to make someone more comfortable or probe on a particular comment.

The following book on the subject may take more than a weekend to read completely, but it's designed to be easy to flip through: *The Definitive Book of Body Language*, by Allan Pease and Barbara Pease (Bantam, 2006).

When you call on someone who hasn't answered yet, be sure to repeat the question in case they didn't understand it or weren't listening to the last few statements in the discussion. Also, avoid making a difference in opinion seem like a disagreement between two individuals.

Don't say, "Bob, we haven't heard from you yet. What do you think about what Chris just said?" but rather (looking at Bob), "How about you, Bob? What kinds of experiences have you had with Pseudo Corporation's customer service?"

As moderator, you control the flow of the discussion and you pass the virtual microphone around. You keep control using eye contact, volume of speech, arm movements, and orientation of your body. Most people will be very aware of your body language, and these cues can be useful signals if someone is dominating the conversation. If an overly vocal participant doesn't

get those hints, use a gentle but firm statement such as "OK, great, I'd like to open that thought up to others. Has anyone else encountered some of the same issues that Theresa has?"

When moving on to a new larger topic, give verbal notice that the previous discussion has finished and that a new one is beginning, so that people can clear their minds for the next topic.

Finally, when the activity is nearing its end, a simple look at your watch and shift in your body orientation can signal that the conversation should be wrapped up. As with any other activity, be sure to thank the group for their time.

Sharing results with your team typically takes one of two forms: findings are either shared according to the main topics being covered or are grouped into relevant categories much as they are for contextual inquiry. Affinity diagramming can be another effective way to bring together various trends and attitudes for illustration to the project team.

Card Sorting

In a *card sorting* activity, participants (working either individually or in small groups) are given items printed on cards and are asked to put them into groups that make sense to them. Either they group them into categories that are provided beforehand (called a *closed sort*) or they make their own groups and title each group themselves (called an *open sort*). At the end of the round of card sorting you should begin to see common patterns emerge in how people are sorting the items, as well as common areas of confusion or disagreement.

A common reason for doing this is to create a site map for a website or to create a hierarchy of content, categories, and subcategories containing items such as articles, documents, videos, or photos. This makes card sorting an excellent technique if you're working on a content source.

Note See Chapter 2 for more on content sources.

Say you're working on a common type of content source: the company intranet. Many intranets tend to categorize their information by the department that owns it, with navigation to human resources, operations, legal, marketing, and so on. For longtime employees this may not present an obvious

problem, because they have probably learned the lines of responsibility of each department and built an understanding of where to find information.

But for new employees, or for those who need information that they don't usually reference, it can be difficult to locate information that could fall within more than one department (or doesn't seem to fall into any). For example, where would you go to find a policy on signing of contracts with newly hired employees? It could fall under legal, or it could fall under human resources.

With card sorting, you can find common patterns in how potential users would categorize information, regardless of departmental lines.

The Basic Process

Collect the items you'd like to include in the card sort; 40 to 60 is usually a good range. You need enough to allow for a potentially large number of card groups to be created, but not so many as to overwhelm the participants with options (or to overwhelm *you* when you need to analyze the results).

Choose items that you think will be easy to understand and free from unnecessary jargon. You can include some subject-matter terms that you believe your user groups are likely to know, but avoid including too many "insider" terms. If you include too many company-specific terms or acronyms (such as "the SUCCEED campaign" for growing sales), you'll be testing the effectiveness of the company's marketing and communications, rather than building a common information hierarchy.

For the intranet example, you might include the vacation policy, 401(k) plan information, new-hire contract, vendor contract, nondisclosure agreement, new-employee orientation, health insurance information, and computer security policy.

This list represents a mix of clearly worded items that could be categorized in multiple ways. You could have one participant who groups new employee orientation and vacation policy together under human resources, and you could have another who groups new employee orientation and new hire contract together and names it "employee onboarding."

Once you have your list of items, put them onto cards that can be easily grouped and ungrouped. You can print labels and stick them onto index cards or print directly onto sheets of card stock that are perforated to separate into individual cards.

Perform a test run by asking someone to sort the cards into groups and give the groups names—for example, by putting a sticky note on the stack and writing the name on it with a pen. Ideally, your test participant is someone unfamiliar with the items and the activity. This will help you get a rough idea of how long the activity might take. If the test run takes over an hour, you may need to cut out some cards!

Once you have a finalized deck, you can bring in a real participant and give these basic instructions:

1. Arrange these cards in whatever groups make sense to you
2. Try to have at least two cards in a group. If a card seems to belong to no group, you can place it to the side
3. At any time as you're sorting, you can name a group. By the end of the activity, please name as many groups as you can.

Some trends will become obvious simply by observing the sessions. Others may take a little more analysis to bring out. There are several tools for entering and analyzing the results of card sorts; many of them come with tools that allow you to run card sorts remotely (see the following "Variations on the Card Sort" section for more on this).

In particular, OptimalSort (www.optimalsort.com/pages/default.html) and WebSort (<http://websort.net>) provide both remote sorting capabilities and helpful analysis tools. Or, if you want to do your own sorting in a more manual fashion, take a look at Donna Spencer's excellent spreadsheet, complete with instructions, available at www.rosenfeldmedia.com/books/cardsorting/blog/card_sort_analysis_spreadsheet.

Variations on the Card Sort

The discussion so far has focused on a card sort carried out with an individual, in person, where the participant is asked to name the categories he created. This is an open sort, meaning that the main categories have not been given to the participant—instead they are *open* to being named. This is a good approach when you're determining a new navigational structure or making significant changes to an existing one. For other situations, you might consider these common card sorting variations:

- ▶ **Closed sorts.** In a closed sort, you provide the high-level categories and participants add to them. The results are relatively easy to analyze, because you have a small set of possible categories and can focus on understanding which items fell most often into which categories. If you're adding large amounts of content to an existing information architecture or you're validating an existing site map, a closed sort can provide quick and actionable information to help with your categorization decisions.
- ▶ **Group sorts.** Rather than having an individual sort items into groups, you can have card sorting be a part of a focus group activity, where participants work together to sort items. Although the results don't necessarily reflect how any one individual would group the items, you can get a lot of insight into how people think about the items and their organization by hearing them work through the activity together, debating the rationale for each placement.
- ▶ **Remote sorts.** Sorting with physical cards can be a fun activity, especially for group sorts. But there are many great tools for performing sorts online with individuals. This also allows you to reach a greater number of participants or particular participants that may be difficult to meet with physically. OptimalSort and WebSort, previously mentioned, are two of the tools that make this type of online sorting easy.

Usability Testing

Usability testing involves asking participants to perform specific tests on a site or application (or a prototype of it) to uncover potential usability issues and gather ideas to address them.

You can perform usability testing during the Define phase if you want to gather information on how the current site can be improved. Or you can perform it on similar sites (such as competitive sites) to understand some of the potential opportunities for a more user-friendly solution.

Most often, usability testing is conducted as part of the Design phase, ideally in iterative rounds (where a design is created, tested, refined, and tested again). We'll discuss usability testing again in full detail in Chapter 14, "Design Testing with Users." That chapter includes tips for recruiting and planning that can help you conduct the activities discussed earlier in this chapter as well.

After the Research

Once you've completed one or more of these user research activities, it's time to revisit the assumptions you originally made about your user groups. Put those assumptions away for a moment, and ask yourself what user groups you would create now that you have more information. If some of your earlier assumptions weren't valid, consider any gaps you may have in your user research because a key group wasn't included. If this gap is identified early enough in your research activity, you may have time to adjust and add another set of participants to research in progress, to ensure you're getting a full picture.

With your new knowledge, you can revise your user definitions to more accurately reflect the groups that should be the focus. This will help you create more detailed tools like personas (discussed in Chapter 7) and will help you create user requirements for the list we began in Chapter 5.

In that chapter, we discussed the process of taking statements from business stakeholders and refining them into requirements. You'll follow a similar process with users—your work doesn't stop when you capture the idea or request. Dig down to the roots of needs and goals to make sure you understand them. This will ultimately help you design a solution that best meets those needs for all relevant user groups.

In the next chapter, you'll learn how to use the insight you gain in conducting user research to create tools that can bring focus to your user groups throughout design and development: personas.

