The Art of Building Great User Experience in Software

Free Sampler

Effective UI

Jonathan Anderson,
John McRee, Robb Wilson
& the EffectiveUI Team

O'REILLY®

Effective UI

Jonathan Anderson, John McRee, Robb Wilson,
and the EffectiveUI Team

O'REILLY®
Beijing · Cambridge · Farnham · Köln · Sebastopol · Taipei · Tokyo

Effective UI
by Jonathan Anderson, John McRee, Robb Wilson, and the EffectiveUI Team

Printed in Canada.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*http://my.safaribooksonline.com*). For more information, contact our corporate/institutional sales department: (800) 998-9938 or *corporate@oreilly.com*.

Editor: Steve Weiss
Development Editor: Jeff Riley
Production Editor: Rachel Monaghan
Copyeditor: Genevieve d'Entremont
Proofreader: Nancy Kotary
Indexer: Julie Hawks
Cover Designer: Karen Montgomery
Illustration and Interior Design:
The EffectiveUI Team

Printing History:

February 2010: First Edition.

ISBN: 978-0-596-15478-3
[F]

Contents

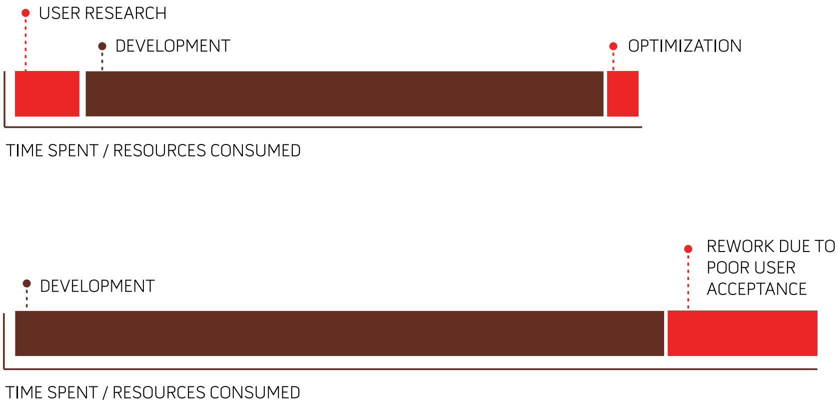

The time and cost of user research can make the overall project more efficient by reducing churn and avoiding poor user acceptance.

Key Concepts in User Research

User research is more than just an exercise in collecting statistics, trivia, and demographic information about users. Although that type of information is useful, it provides an incomplete and overly flat picture of users. Your project team needs a deep, rich understanding of users to be able to make correct decisions from the user's perspective. Since most members of the project team won't have the opportunity to get to know actual users, user research needs to provide the team with a means of seeing things from the user's perspective.

Empathy

By "empathy," we don't mean, of course, a sense of teary-eyed compassion for the plight of the user (although bad software UX definitely has the capacity to reduce users to tears). Rather, we mean it in the sense of the ability to understand, on both an intellectual and intuitive level, the users' needs and to see things from their perspective.

> *Sometimes when we suggest that the experience of interacting with software can be emotionally pleasurable and engaging, people (usually engineers) act puzzled at the idea of software having anything to do with emotion. But emotion underlies every human behavior and interaction, and software is no exception. It's often easier for people to imagine negative emotional experiences with software than positive ones, which indicates not the impossibility of positive experiences but the relative dearth of them so far in the history of software.*

Just as with every other component of the framework requirements, the goal in user research isn't to answer every possible question about the user's needs and perspective in advance, but rather to provide a framework for the team to answer those questions as they arise. Empathy on intellectual and intuitive levels allows them to take thoughtfully focused but relatively sparse information about users and use it to judge things and answer questions from their perspectives.

User Goals Versus Product Features and Tasks

The concept of "goals" is particularly significant in user research and UX design. To stakeholders and members of the project team, a product can seem to be a collection of features that gives users access to a set of capabilities. But to users, a product is a tool used to accomplish some higher-level goal. Those goals are accomplished using the features and capabilities available in the product, but the user's intention in using the product is to achieve his goals, not to simply employ the product's features and capabilities. Even with games, which are a type of software that has no strictly utilitarian purpose, the user's primary goal is to be entertained. He accomplishes that goal by operating the controller and progressing through the storyline, which are features of the game software. The use of those features isn't his goal, per se; they're the means by which he accomplishes his goal of being entertained.

Goals are also different from tasks. Tasks are the steps a user goes through when using the product to accomplish his goals, but the tasks themselves, like the employment of the features, aren't the end goal. The desire to achieve the goals is what motivates the user to employ a given set of features of a given product to perform a series of tasks that ultimately accomplish the goal. The meaningfulness of these distinctions is best demonstrated through the following example.

Essentially every official business activity is at some point discussed and formalized in some written form of communication: interoffice memos, sales proposals, contracts, legal demands, letters of introduction, tickets for business travel, and so on. Time is of the essence in almost every type of correspondence and documentation, so the goal of quickly exchanging this documentation has long been an important goal in business. The means by which that exchange has been facilitated—the features and capabilities of services and products that have been employed to accomplish the tasks required to achieve the goal—and the standards for quickness have changed dramatically over the last couple of hundred years, although the goal itself has not.

As an example, consider that a business in St. Joseph, Missouri, wants to send an agreement for the sale of a plot of land to its purchaser in Sacramento, California. This is a business scenario that was as normal and plausible in 1860 as it is today. The two high-level tasks required to accomplish this goal are first to prepare the agreement, and then to convey it to the purchaser. The means by which these tasks are accomplished have changed dramatically over the past century and a half.

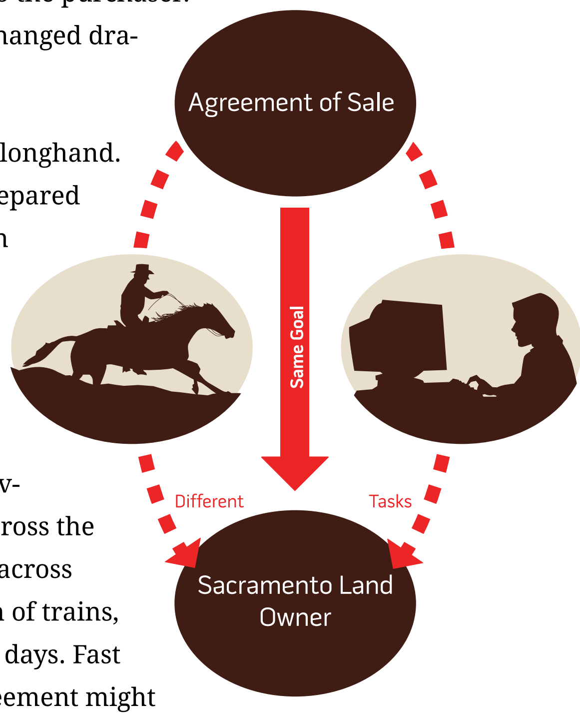

In 1860, the agreement would likely have been written in longhand. A couple of decades later, it would probably have been prepared using a typewriter. Many decades later, it might have been prepared on an early computer and printed out on a dot-matrix printer. In slightly more recent times, it would be prepared and printed with more sophisticated computer equipment and software.

In 1860, the handwritten agreement might have been delivered by the famed Pony Express on a 10-day horse ride across the country. Later, the agreement might have been conveyed across the country by the U.S. Postal Service using a combination of trains, horse carts, or early automobiles in a matter of about five days. Fast forwarding past the era of the dot-matrix printer, the agreement might never have been physically delivered; it might have been faxed for near-instantaneous delivery. There was—honestly—a time when 14,400 bits (yes, bits) per second was considered basically instantaneous. More recently, that agreement might have been delivered in digital format as an email attachment with no physical copy ever trading hands.

Though the nature of the tasks involved in achieving the goal of conveying a legal document from one party to another has changed tremendously over the years, the goal has remained the same. The tasks have required typewriters, fax machines, PCs, the Pony Express, the post office, and email. Each has its own peculiar features and tasks required to successfully operate them, but the goal has remained the same.

Because user goals are constant and knowable while everything else is in flux, they are the ideal conception of the user perspective for the framework requirements. They exist at a high level and stay consistent through a long history of changes in technology and product capabilities. Each new product development effort entails tailoring the current technological capabilities and business considerations to user goals. It also requires a fresh examination of user behaviors, which do change over time.

Qualitative Versus Quantitative Research Methods

It's difficult to describe a person's conscious experience in the form of quantitative data. Demographic research can yield quantitative results about the target market segment or user base (62 percent are men, 14 percent spend two or more hours a day watching television, and so on), and certain types of ultra-clinical usability research can result in quantitative or quasi-quantitative results. But there are limits to the value of this type of approach in user research and design for UX.

Some researchers try to dress up all information collected through user research in the clothing of quantitative results (for example, 47 percent of people thought our product was 30–70 percent better than our competitor's). The appeal of this approach is that these points of data make a beautiful line on a graph and hold greater credibility because of their apparently solid, scientific quality. But representing qualitative feedback from users in numeric form doesn't change the nature of that feedback from qualitative to quantitative.

Another reason why it's appealing to portray qualitative feedback in a numeric form is that it seems to suggest that statistical analyses could be done on the data to expose new information, but that's generally not the case. If you discover that 18 percent of users rate your product's UX at 2 or worse and 45 percent rate it at a 3 or worse on a scale of 10, you've learned something vaguely interesting but not terribly useful. You can't extrapolate any new information from these findings; they just let you know that improvements are needed…somewhere. This kind of information can remind you that you're on a damaged ship in a darkened sea full of rocks, but offers no map or illumination to help you navigate safely to shore.

The principal goal of user research is to communicate findings about a limited sample of users to give the project team to the ability to understand the broader user base and see things from the user's perspective. As a result, the information that user research produces must allow the project team to make inferences and extrapolate from it. This would seem to be an argument for more numeric data, but in fact, the opposite is better.

Qualitative research on a sample of users, even if converted to numeric form, cannot be extrapolated, interpolated, or inferred from using statistical methods. Most of it can't even be plotted on a graph in any truly meaningful way. Although there are formal methods of analysis (Grounded Theory, for example) that are useful in building from qualitative research, in practical reality anything that rigorous is seldom in the budget. Rather, the most effective mechanisms for making inferences from user research are human intuition and empathy.

Members of the project team can look at stories about users' lives and their experiences using the product and, using intuition and empathy, make accurate projections about the user's point of view regarding things not specifically addressed in the research. The empathetic and intuitive capacity of the team is what allows research on a limited sample of users to be used as an intelligent, expansible, and flexible part of the framework requirements.

So, although at first blush, qualitative user research might seem to be less scientific and therefore less credible, it's actually much more useful to the team. It gets closer to the core of the user's subjective experience in a way that allows the team to make the best intuitive and empathetic inferences. This makes it possible for every member of the team to represent the user's interests, ensuring that the user's goals—and, therefore, UX quality—are kept at the center of decision making.

The apparent precision of quantitative data can also give the false impression that it's the most simple and direct route to answers. Back in 2000, when tools to gather website analytics were relatively new and expensive, people were just beginning to understand the effect of "banner blindness"—the tendency of website visitors to unconsciously ignore anything that looked like a banner ad. A large online travel company conducted a study to find out how people were using their site. Through quantitative analysis of the site's analytics, they discovered that a significant percentage of users were not getting beyond the home page. But the quantitative data offered no explanation of why this might be.

The qualitative approach was cheap, fast, and offered an answer to the question, "Why?"

Usability experts (as they were called back then) were able to convince the travel company's executives that a qualitative study would shed better light on the situation. The first research participant who attempted to book travel online didn't make it past the first page. The "call-to-action" (the all-important button that advances you to the next stage in the purchase or conversion process) looked so much like a banner ad that the user simply did not see it. After the second and third consecutive users had this same problem, it was obvious that banner blindness was interfering with the site's success. This made both the problem and its solution quickly and plainly obvious. It was a win for the usability team, but more importantly the user struggling to book travel won relief and the business won increased revenue through radically improved conversion. The qualitative approach was cheap, fast, and offered an immediate answer to the question, "Why?" There are many cases like this in UX design where quantitative data might indicate that a problem exists, but a qualitative study is needed to know how to respond.

Who Should Be Involved in the Research

So much of user research is a combination of art and science that, as we've said, you'd be hard pressed to get through it successfully without the assistance of a professional user researcher. If you must press on without the aid of a professional, spend some time reading books that go into significantly greater depth on the process and discipline of user research than the high-level view we present in this book. If this is the route you take, you'll need to acknowledge that the results of your user research will be less robust, since it will be missing everything that the deep experience of a professional brings. Ensure that the project team is aware of this as they make use of the research results so they'll be prepared to identify and respond to weaknesses in the research. For the rest of this chapter, though, we'll assume that you'll have professional assistance.

You should be present for much or all of the user research process. This helps you get acquainted with the users directly and develop an early empathy for user needs. During the research process, keep quiet and simply observe. This prevents you from affecting the course of the discussion and the discoveries that result, and allows the one-on-one, comfortable interaction between the researcher and the user to proceed undisturbed.

It can also be helpful to bring stakeholders along for some of the interviews. It gives them an opportunity to see the process, listen to the questions, and understand the approach taken in user research. This reassures them that the research is valuable and is representing their customers' interests effectively. Stakeholders often become very engaged during the process, helping to ask questions and explore tasks that shed light on the problem. It's also a good opportunity for team building between the project leader, the researcher, and the stakeholders. It typically has the effect of getting everyone energized about the project and very considerate of the user's needs. This has obvious value as you continue through the project, and can also produce engaged stakeholders who will help you keep the other stakeholders on board.

There are risks, however, in involving stakeholders in user research. It's helpful when they constructively engage in the interview process, but some stakeholders might use it as an opportunity to push a particular agenda by asking leading questions or trying to direct the course of the discussion. Also, stakeholders who were involved with the previous version of the product can get very frustrated when they see users having trouble using the old system. We've seen some instances where a stakeholder intervened to try to educate the user on how to use the old system rather than simply observing the problems. There's also a risk that a stakeholder might seize on one particular challenge or observation of one particular user that might be minor or atypical and use it to try to push some personal agenda. But stakeholder support is such an essential element of a successful project that any risks to the research tend to be worth it. Good user research professionals know how to accommodate stakeholders in their research processes while maintaining its quality. You can also avoid many of these problems by preparing the stakeholders properly about the goals and practice of research ahead of time.

Finding Research Participants

In the business planning stage, you worked with stakeholders to produce a restrained set of key user attributes. The task of assembling the sample user set involves finding individuals who are good representatives of the key user attributes and who are currently using or are likely to use your product or one similar to it. The key user attributes will be useful in constraining the possible field of sample users into a manageable set, but within those constraints it's important to find a relatively diverse range of users. If your key user attributes describe, for example, marketing executives who manage budgets of $5–15M, then you might seek out:

- *Users working for companies in different industries*
- *Users managing budgets scattered throughout that $5–15M range*
- *Users representing a mixture of men and women*
- *Users of varying ages*

This practice leads to more roundly representative research. It's not always the case, by the way, that industry, gender, or age are the key variables of diversity. This will vary by project.

Stakeholders, salespeople, marketing managers, and other customer-focused colleagues help you identify potential research subjects. If the product is meant for internal use, then the job of finding sample users can be as easy as working through your company staff directory. Although it might seem like an imposition to ask someone to participate as a subject in user research, people often feel honored that you're interested in their perspective. Most participants are eager to help shape a product to better suit their own needs and to get a peek behind the curtain of software development. You can also promise a free copy of the final product to customers who participate in research, and people are usually more than happy to help out if there's a free lunch involved.

Determining the Research Sample Size

While quantitative research tends to require large sample sizes (usually in the thousands), qualitative research usually succeeds with surprisingly small sample sizes. This has been noted by the usability and UX experts Jakob Nielsen, Jared Spool, and Allen Cooper in some of their books and articles, and is something we encounter consistently in the user research we do. During business planning on a recent project, we determined that the following types of people would be the primary users of the product:

- *Network administrators in large businesses*
- *Billing analysts in large businesses*
- *Business owners in mid-size and smaller businesses*
- *Help desk agents from the product company's staff*
- *Sales representatives from the product company's staff*

We spoke to about 8–10 people in each group, for a total of 40–50 users in the sample. Working with a sample of this size, the user research stage took about three weeks.

There's no rule of thumb for determining how many people you need to work with; the size of the sample is dependent on project-specific considerations. The number of key user attributes identified through your work with stakeholders is a major factor. If only a small number of attributes were identified, then you need only find a representative diversity within the narrow constraints of those attributes. More user attributes will mean, of course, a larger sample. You needn't be overly meticulous when it comes to finding a diverse group internal to a given user attribute. Looking back to the example in the previous section, you needn't ensure that you have one male and one female marketing executive from every possible combination of industry, budget size, and age bracket. Often just a few people representing a given user attribute will suffice.

User research needs to balance expedience with thoroughness, and there's a rapidly diminishing return of value as you start to work with more and more people. This is particularly the case with user research for enterprise products, where we've found the user feedback is typically surprisingly homogeneous. The jobs found in large businesses—administrative assistant, help desk agent, network administrator, and so on—are often very similar across many companies. The reasons enterprise product users use the product, the demands imposed on them by their jobs, and the environments they're using the product in are relatively consistent.

For consumer products where the user's experience is guided less by standardized roles within organizations and more by personal considerations, more (but not dramatically more) research subjects are valuable. The target user attributes for consumer projects are typically strongly based on market segments, which are broad categories of people. You may have as a key user attribute, for example, the market segment "housewife." Housewives aren't a homogeneous group of people. Their backgrounds, home lives, motivations, brand loyalties, budget constraints, technical savvy, routines, and other wide-ranging factors will influence their relationship with a software product. Since "housewives" is such a broad user attribute, you'd need to plan to work with more research subjects to get a full understanding of the group. User researchers need to keep researching users until they notice and then confirm consistent patterns in the feedback.

The point where researchers are able to establish and confirm patterns is called saturation. It's the point where researchers start hearing the same things over and over, and stop identifying anything new or eye-opening. As we've said, researchers will usually hit saturation faster in enterprise products than in consumer ones, but in either case they may hit saturation sooner than they had originally expected. In these cases, it can be sensible to end the research early.

When choosing the size of the sample, be careful to consider your stakeholders' expectations and the effect on the credibility of your research. Your stakeholders' trust in and deference to the user's perspective as discovered and represented through user research will be critical to the project's success, so their expectations must be taken seriously. Sometimes talking to just eight users will be entirely sufficient from a research perspective, but won't seem thorough and compelling enough to garner the trust of your stakeholders. Small samples might also inadvertently exclude specific user or customer types that one of your stakeholders cares deeply about. Before starting the research, share your user researchers' plans with your stakeholders, whether individually or as a group. You need to ensure that they understand how user research works, that they support the research approach, and that they think the research sample is sufficiently representative. It's critical that the stakeholders support and respect the results of user research. You'll be pressing them to think about the product from the user's perspective and to be deferential to user needs, which are represented by the research results. If they're going to support and respect the research results, they first need to support and respect the way the research is conducted.

We recently went through an extensive user research phase for a huge enterprise product where we'd scheduled dozens of user interviews over a period of several days. We started hitting saturation very early, but we still continued on to complete over 60 hours of interviews. In this case, our client was very sensitive to the size of the research sample and keen to ensure all possible customers were included, so we continued on. Though the research itself yielded rapidly diminishing value for the project team, the value of our client's confidence in the results of the research and our diligence in respecting their interests trumped simple expedience.

Making Recordings

It will be useful all throughout the research process to make recordings of the time spent with users. The recordings are useful as reference material during the subsequent analysis and documentation of the research. Recordings are also a good way to share the user research process with stakeholders. They're more effective in building emotional buy-in and deference to the user's needs than documentation of the research findings alone. In the research stage of a product redesign effort we were involved in, we captured video of a user who got so frustrated and angry trying to use the product that he actually screamed and smashed the keyboard with his fist. When this video was shown to stakeholders, the product redesign suddenly acquired a much higher priority and budget.

How you record the sessions will depend on the situation. For the purposes of reference and stakeholder buy-in, video is preferable because it's more engaging and records more information, including the user's actual activities and workspace. You don't need a person in the room operating the camera; this is overkill and distracting. Just put a camcorder on a tripod in a position that lets it capture as much of the user's behavior and reactions as possible. One downside to video cameras is they can affect the way people behave and their comfort level in the conversation. Getting authentic information and putting users at ease are more important goals than capturing video. In some cases, placing the camera where it's easy to forget about it can mitigate this effect. It also helps to spend a little time with users and just have a casual conversation with them while the camera is running to get them acclimated to the environment and the presence of the camera.

If a video camera is too disruptive or intrusive, an audio recorder can work just fine. Audio recorders are much less conspicuous and are easily forgotten, even if they're sitting on the desk next to the user. People don't feel the need to give a performance in quite the same way as they do when they're on camera.

Whether you're recording video and audio or just audio alone, the quality of the audio capture is tremendously important. It's very aggravating to try to work your way through noisy, unclear recordings to try to find that one nugget of insight you're looking for. It also detracts tremendously from the emotional buy-in value to your stakeholders and makes the project look less professional. When setting up the video or audio recorder, consider the audio dynamics of the space and use decent quality equipment. When using the equipment in a new environment, make sure to test it to ensure that the recording will be usable. If positioning the video camera inconspicuously leads to bad audio capture, consider using a separate audio recording device as well. Also, be careful to keep audio recorders away from keyboards, air conditioning vents, and other sources of ambient noise, and also away from cell phones, as they cause interference in microphones even when they're not on a call.