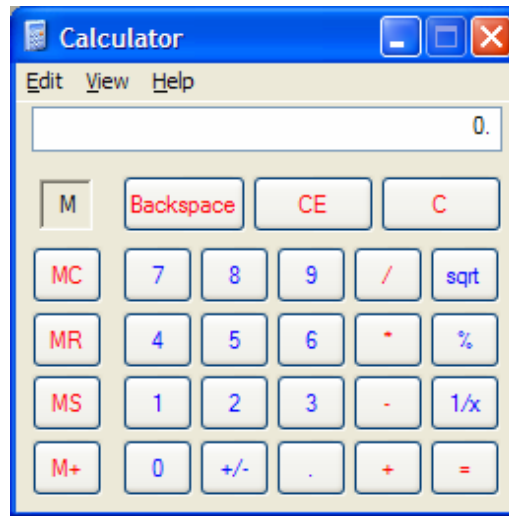


Lecture 6: Prototyping

UI Hall of Fame or Shame?



Fall 2003

6.893 UI Design and Implementation

2

Today's candidate for the Halls of Fame and Shame is the Windows calculator.

It looks and works just like a familiar desk calculator, a stable interface that many people are familiar with. It's a familiar metaphor, and trivial for calculator users to pick up and use.

Unfortunately it deviates from the metaphor in some small ways, largely because the buttons are limited to text labels. The square root button is labeled "sqrt" rather than the root symbol. The multiplication operator is * instead of X.

This interface adheres to its metaphor so carefully that it passes up some tremendous opportunities to *improve* on the desk calculator interface. Why only one line of display? A history, analogous to the paper tape printed by some desk calculators, would cost almost nothing. Why only one memory slot? Why display "M" instead of the actual number stored in memory? All these issues violate the **visibility of system state** heuristic. A more serious violation of the same heuristic: the interface actually has invisible modes. When I'm entering a number, pressing a digit appends it to the number. But after I press an operator button, the next digit I press starts a new number. There's no visible feedback about what low-level mode I'm in. Nor can I tell, once it's time to push the = button, what computation will actually be made.

Most of the buttons are cryptically worded (**recognition, not recall**). MC, MR, MS, and M+? What's the difference between CE and C? My first guess was that CE meant "Clear Error" (for divide-by-zero errors and the like); some people in class suggested that it means "Clear Everything". In fact, it means "Clear Entry", which just deletes the last number you entered without erasing the previous part of the computation. "C" actually clears everything.

It turns out that this interface also lets you type numbers on the keyboard, but the interface doesn't give a hint (**affordance**) about that possibility. In fact, in a study of experienced GUI users who were given an onscreen calculator like this one to use, 13 of 24 never realized that they could use the keyboard instead of the mouse (Nielsen, *Usability Engineering*, p. 61-62). One possible solution to this problem would be to make the display look more like a text field, with a blinking cursor in it, implying "type here". Text field appearance would also help the Edit menu, which offers Copy and Paste commands without any obvious selection (**external consistency**).

Finally, we might also question the use of small blue text to label the buttons, which is hard to read, and the use of both red and blue labels in the same interface, since chromatic aberration forces red and blue to be focused differently. Both decisions tend to cause eyestrain over periods of long use.

Today's Topics

- Paper prototypes
- Wizard of Oz prototypes

Today we're going to talk about prototyping: producing cheaper, less accurate renditions of your target interface. Prototyping is essential in the early iterations of a spiral design process, and it's useful in later iterations too.

Why Prototype?

- Get feedback earlier, cheaper
- Experiment with alternatives
- Easier to change or throw away

We build prototypes for several reasons, all of which largely boil down to cost.

First, prototypes are much faster to build than finished implementations, so we can evaluate them sooner and get early feedback about the good and bad points of a design.

Second, if we have a design decision that is hard to resolve, we can build multiple prototypes embodying the different alternatives of the decision.

Third, if we discover problems in the design, a prototype can be changed more easily, for the same reasons it could be built faster. Prototypes are more malleable. Most important, if the design flaws are serious, a prototype can be **thrown away**. It's important not to commit strongly to design ideas in the early stages of design. Unfortunately, writing and debugging a lot of code creates a psychological sense of commitment which is hard to break. You don't want to throw away something you've worked hard on, so you're tempted to keep some of the code around, even if it really should be scrapped.

The prototyping techniques we'll see in this lecture actually force you to throw the prototype away. For example, a paper mockup won't form any part of a finished software implementation. This is a good mindset to have in early iterations, since it maximizes your creative freedom.


Prototype Fidelity

- Low fidelity: omits details
- High fidelity: more like finished product

An essential property of a prototyping technique is its **fidelity**, which is simply how similar it is to the finished interface. Low-fidelity prototypes omit details, use cheaper materials, or use different interaction techniques. High-fidelity prototypes are very similar to the finished product.

Fidelity is Multidimensional

- **Breadth:** % of features covered
 - Only enough features for certain tasks
- **Depth:** degree of functionality
 - Limited choices, canned responses, no error handling



Fall 2003

6.893 UI Design and Implementation

6

Fidelity is not just one-dimensional, however. Prototypes can be low- or high-fidelity in various different ways (Carolyn Snyder, *Paper Prototyping*, 2003).

Breadth refers to the fraction of the feature set represented by the prototype. A prototype that is low-fidelity in breadth might be missing many features, having only enough to accomplish certain specific tasks. A word processor prototype might omit printing and spell-checking, for example.

Depth refers to how deeply each feature is actually implemented. Is there a backend behind the prototype that's actually implementing the feature? Low-fidelity in depth may mean limited choices (e.g., you can't print double-sided), canned responses (always prints the same text, not what you actually typed), or lack of robustness and error handling (crashes if the printer is offline).

A diagrammatic way to visualize breadth and depth is shown (following Nielsen, *Usability Engineering*, p. 94). A **horizontal prototype** is all breadth, and little depth; it's basically a frontend with no backend. A **vertical prototype** is the converse: one area of the interface is implemented deeply. The question of whether to build a horizontal or vertical prototype depends on what risks you're trying to mitigate. In user interface design, horizontal prototypes are more common, since they address usability risk. But if some aspect of the application is a risky implementation – you're not sure if it can be implemented to meet the requirements – then you may want to build a vertical prototype to test that.

A special case lies at the intersection of a horizontal and a vertical prototype. A **scenario** shows how the frontend would look for a single concrete task. Scenarios are great for visualizing a design, but they're hard to evaluate with users.

More Dimensions of Fidelity

- Look: appearance, graphic design
 - Sketchy, hand-drawn
- Feel: input method
 - Pointing & writing feels very different from mouse & keyboard

Two more crucial dimensions of a prototype's fidelity are, loosely, its look and its feel. **Look** is the appearance of the prototype. A hand-sketched prototype is low-fidelity in look, compared to a prototype that uses the same widget set as the finished implementation. **Feel** refers to the physical methods by which the user interacts with the prototype. A user interacts with a paper mockup by pointing at things to represent mouse clicks, and writing on the paper to represent keyboard input. This is a low-fidelity feel for a desktop application (but it may not be far off for a tablet PC application).

Paper Prototype

- Interactive paper mockup
 - Sketches of screen appearance
 - Paper pieces show windows, menus, dialog boxes
- Interaction is natural
 - Pointing with a finger = mouse click
 - Writing = typing
- A person simulates the computer's operation
 - Putting down & picking up pieces
 - Writing responses on the "screen"
 - Describing effects that are hard to show on paper
- Low fidelity in look & feel
- High fidelity in depth (person simulates the backend)

Paper prototypes are an excellent choice for early design iterations. A paper prototype is a physical mockup of the interface, mostly made of paper. It's usually hand-sketched on multiple pieces, with different pieces showing different menus, dialog boxes, or window elements.

The key difference between mere sketches and a paper prototype is **interactivity**. A paper prototype is brought to life by a design team member who simulates what the computer would do in response to the user's "clicks" and "keystrokes", by rearranging pieces, writing custom responses, and occasionally announcing some effects verbally that are too hard to show on paper. Because a paper prototype is actually interactive, you can actually user-test it: give users a task to do and watch how they do it.

A paper prototype is clearly low fidelity in both look and feel. But it can be arbitrarily high fidelity in breadth at very little cost (just sketching, which is part of design anyway). Best of all, paper prototypes can be **high-fidelity in depth** at little cost, since a human being is simulating the backend.

Why Paper Prototyping?

- **Faster to build**
 - Sketching is faster than programming
- **Easier to change**
 - Easy to make changes between user tests, or even *during* a user test
 - No code investment– everything will be thrown away (except the design)
- **Focuses attention on big picture**
 - Designer doesn't waste time on details
 - Customer makes more creative suggestions, not nitpicking
- **Nonprogrammers can help**
 - Only kindergarten skills are required

But why use paper? And why hand sketching rather than a clean drawing from a drawing program?

Hand-sketching on paper is faster. You can draw many sketches in the same time it would take to draw one user interface with code. For most people, hand-sketching is also faster than using a drawing program to create the sketch.

Paper is easy to change. You can even change it during user testing. If part of the prototype was a problem for one user, you can scratch it out or replace it before the next user arrives. Surprisingly, paper is more malleable than digital bits in many ways.

Hand-sketched prototypes in particular are valuable because they focus attention on the issues that matter in early design without distracting anybody with details. When you're sketching by hand, you aren't bothered with details like font, color, alignment, whitespace, etc. In a drawing program, you would be faced with all these decisions, and you might spend a lot of time on them – time that would clearly be wasted if you have to throw away this design. Hand sketching also improves the feedback you get from users. They're less likely to nitpick about details that aren't relevant at this stage. They won't complain about the color scheme if there isn't one. More important, however, a hand-sketch design seems less finished, less set in stone, and more open to suggestions and improvements. Architects have known about this phenomenon for many years. If they show clean CAD drawings to their clients in the early design discussions, the clients are less able to discuss needs and requirements that may require radical changes in the design. In fact, many CAD tools have an option for rendering drawings with a "sketchy" look for precisely this reason.

A final advantage of paper prototyping: no special skills are required. So graphic designers, usability specialists, and even users can help create prototypes and operate them.

Tools for Paper Prototyping

- White poster board (11"x14")
 - For background, window frame
- Big (unlined) index cards (4"x6", 5"x8")
 - For menus, window contents, and dialog boxes
- Restickable glue
 - For keeping pieces fixed
- White correction tape
 - For text fields, checkboxes, short messages
- Overhead transparencies
 - For highlighting, user "typing"
- Photocopier
 - For making multiple blanks
- Pens & markers, scissors, tape

Here are the elements of a paper prototyping toolkit.

Although standard (unlined) paper works fine, you'll get better results from sturdier products like **poster board** and **index cards**. Use poster board to draw a static background, usually a window frame. Then use index cards for the pieces you'll place on top of this background. You can cut the index cards down to size for menus and window internals.

Restickable Post-it Note glue, which comes in a roll-on stick, is a must. This glue lets you make all of your pieces sticky, so they stay where you put them. You can find this glue at Pearl Arts in Central Square; it's not found in the Coop.

Post-it correction tape is another essential element. It's a roll of white tape with Post-it glue on one side. Correction tape is used for text fields, so that users can write on the prototype without changing it permanently. You peel off a length of tape, stick it on your prototype, let the user write into it, and then peel it off and throw it away. Correction tape comes in two widths, "2 line" and "6 line". The 2-line width is good for single-line text fields, and the 6-line width for text areas. You can get correction tape at the Office Max in East Cambridge.

Overhead transparencies are useful for two purposes. First, you can make a selection highlight by cutting a piece of transparency to size and coloring it with a transparency marker. Second, when you have a form with several text fields in it, it's easier to just lay a transparency over the form and let the users write on that, rather than sticking a piece of correction tape in every field.

If you have many similar elements in your prototype, a **photocopier** can save you time. And, of course, the usual kindergarten equipment: pens, markers, scissors, tape.

Tips for Good Paper Prototypes

- Make it larger than life
- Make it monochrome
- Replace tricky visual feedback with audible descriptions
 - Tooltips, drag & drop, animation, progress bar
- Keep pieces organized
 - Use folders & open envelopes

A paper prototype should be larger than life-size. Remember that fingers are bigger than a mouse pointer, and people usually write bigger than 12 point. So it'll be easier to use your paper prototype if you scale it up a bit. It will also be easier to see from a distance, which is important because the prototype lies on the table, and because when you're testing users, there may be several observers taking notes who need to see what's going on. **Big is good.**

Don't worry too much about color in your prototype. Use a single color. It's simpler, and it won't distract attention from the important issues. Needless to say, don't use yellow.

You don't have to render every visual effect in paper. Some things are just easier to say aloud: "the basketball is spinning." "A progress bar pops up: 20%, 50%, 75%, done." If your design supports tooltips, you can tell your users just to point at something and ask "What's this?", and you'll tell them what the tooltip would say. If you actually want to test the tooltip messages, however, you should prototype them on paper.

Figure out a good scheme for organizing the little pieces of your prototype. One approach is a three-ring binder, with different screens on different pages. Most interfaces are not sequential, however, so a linear organization may be too simple. Two-pocket folders are good for storing big pieces, and letter envelopes (with the flap open) are quite handy for keeping menus.

How to Test a Paper Prototype

- Roles for design team
 - Computer
 - Simulates prototype
 - Doesn't give any feedback that the computer wouldn't
 - Facilitator
 - Presents interface and tasks to the user
 - Encourages user to “think aloud” by asking questions
 - Keeps user test from getting off track
 - Observer
 - Keeps mouth shut, sits on hands if necessary
 - Takes copious notes

Once you've built your prototype, you can put it in front of users and watch how they use it. We'll see much more about user testing in a later lecture, including ethical issues. But here's a quick discussion of user testing in the paper prototyping domain.

There are three roles for your design team to fill:

The **computer** is the person responsible for making the prototype come alive. This person moves around the pieces, writes down responses, and generally does everything that a real computer would do. In particular, the computer should *not* do anything that a real computer wouldn't. Think mechanically, and respond mechanically.

The **facilitator** is the human voice of the design team and the director of the testing session. The facilitator explains the purpose and process of the user study, obtains the user's informed consent, and presents the user study tasks one by one. While the user is working on a task, the facilitator tries to elicit verbal feedback from the user, particularly encouraging the user to “think aloud” by asking probing (but not leading) questions. The facilitator is responsible for keeping everybody disciplined and the user test on the right track.

Everybody else in the room (aside from the user) is an **observer**. The most important rule about being an observer is to keep your mouth shut and watch. Don't offer help to the user, even if they're missing something obvious. Bite your tongue, sit on your hands, and just watch. The observers are the primary note takers, since the computer and the facilitator are usually too busy with their duties.

What You Can Learn from a Paper Prototype

- Conceptual model
 - Do users understand it?
- Functionality
 - Does it do what's needed? Missing features?
- Navigation & task flow
 - Can users find their way around?
 - Are information preconditions met?
- Terminology
 - Do users understand labels?
- Screen contents
 - What needs to go on the screen?

Paper prototypes can reveal many usability problems that are important to find in early stages of design. Fixing some of these problems require large changes in design. If users don't understand the metaphor or conceptual model of the interface, for example, the entire interface may need to be scrapped.

What You Can't Learn

- Look: color, font, whitespace, etc
- Feel: Fitts's Law issues
- Response time
- Are small changes noticed?
 - Even the tiniest change to a paper prototype is clearly visible to user
- Exploration vs. deliberation
 - Users are more deliberate with a paper prototype; they don't explore or thrash as much

But paper prototypes don't reveal every usability problem, because they are low-fidelity in several dimensions. Obviously, graphic design issues that depend on a high-fidelity **look** will not be discovered. Similarly, interaction issues that depend on a high-fidelity **feel** will also be missed. For example, Fitts's Law problems like buttons that are too small, too close together, or too far away will not be detected in a paper prototype. The human computer of a paper prototype rarely reflects the speed of an implemented backend, so issues of **response time** – whether feedback appears quickly enough, or whether an entire task can be completed within a certain time constraint -- can't be tested either.

Paper prototypes don't help answer questions about whether subtle feedback will even be *noticed*. Will users notice that message down in the status bar, or the cursor change, or the highlight change? In the paper prototype, even the tiniest change is grossly visible, because a person's arm has to reach over the prototype and make the change. (If many changes happen at once, of course, then some of them may be overlooked even in a paper prototype, a clearly discernible. This is related to an interesting cognitive phenomenon called **change blindness**.)

There's an interesting qualitative distinction between the way users use paper prototypes and the way they use real interfaces. Experienced paper prototypers report that users are more deliberate with a paper prototype, apparently thinking more carefully about their actions. This may be partly due to the simulated computer's slow response; it may also be partly a social response, conscientiously trying to save the person doing the simulating from a lot of tedious and unnecessary paper shuffling. More deliberate users make fewer mistakes, which is bad, because you want to see the mistakes. Users are also less likely to randomly explore a paper prototype.

These drawbacks don't invalidate paper prototyping as a technique, but you should be aware of them. Several studies have shown that low-fidelity prototypes identify substantially the same usability problems as high-fidelity prototypes (Virzi, Sokolov, & Karis, "Usability problem identification using both low- and hi-fidelity prototypes", CHI '96; Catani & Biers, "Usability evaluation and prototype fidelity", Human Factors & Ergonomics 1998).

Wizard of Oz Prototype

- Software simulation with a human in the loop to help
- “Wizard of Oz” = “man behind the curtain”
 - Wizard is usually but not always hidden
- Often used to simulate future technology
 - Speech recognition
 - Learning
- Issues
 - Two UIs to worry about: user’s and wizard’s
 - Wizard has to be mechanical

Part of the power of paper prototypes is the depth you can achieve by having a human simulate the backend. A **Wizard of Oz prototype** also uses a human in the backend, but the frontend is an actual computer system instead of a paper mockup. The term Wizard of Oz comes from the movie of the same name, in which the wizard was a man hiding behind a curtain, controlling a massive and impressive display.

In a Wizard of Oz prototype, the “wizard” is usually but not always hidden from the user. Wizard of Oz prototypes are often used to simulate future technology that isn’t available yet, particularly artificial intelligence. A famous example was the listening typewriter (Gould, Conti, & Hovanyecz, “Composing letters with a simulated listening typewriter,” *CACM* v26 n4, April 1983). This study sought to compare the effectiveness and acceptability of isolated-word speech recognition, which was the state of the art in the early 80’s, with continuous speech recognition, which wasn’t possible yet. The interface was a speech-operated text editor. Users looked at a screen and dictated into a microphone, which was connected to a typist (the wizard) in another room. Using a keyboard, the wizard operated the editor showing on the user’s screen.

The wizard’s skill was critical in this experiment. She could type 80 wpm, she practiced with the simulation for several weeks (with some iterative design on the simulator to improve her interface), and she was careful to type *exactly* what the user said, even exclamations and parenthetical comments or asides. The computer helped make her responses a more accurate simulation of computer speech recognition. It looked up every word she typed in a fixed dictionary, and any words that were not present were replaced with X’s, to simulate misrecognition. Furthermore, in order to simulate the computer’s ignorance of context, homophones were replaced with the most common spelling, so “done” replaced “dun”, and “in” replaced “inn”. The result was an extremely effective illusion. Most users were surprised when told (midway through the experiment) that a human was listening to them and doing the typing.

Thinking and acting mechanically is harder for a wizard than it is for a paper prototype simulator, because the tasks for which Wizard of Oz testing is used tend to be more “intelligent”. It helps if the wizard is personally familiar with the capabilities of similar interfaces, so that a realistic simulation can be provided. (See Maulsby, 1993) It also helps if the wizard’s interface can intentionally dumb down the responses, as was done in the Gould study.

A key challenge in designing a Wizard of Oz prototype is that you actually have two interfaces to worry about: the user’s interface, which is presumably the one you’re testing, and the wizard’s.

Reading for Next Time

- Optional: Mullet & Sano, *Designing Visual Interfaces*