

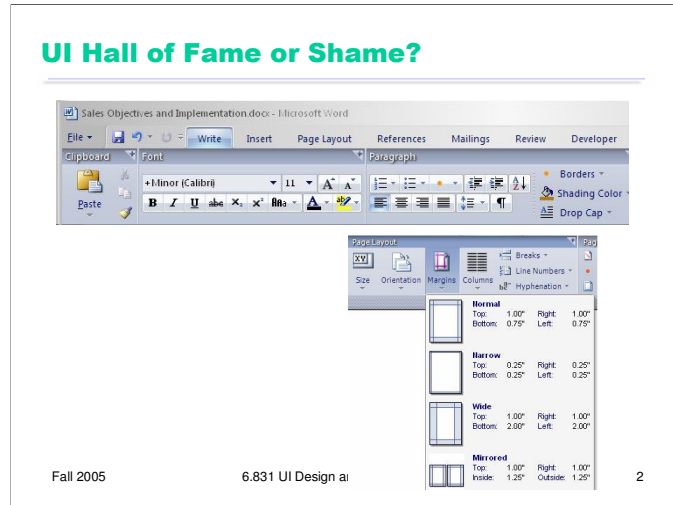
## Lecture 14: Heuristic Evaluation

---

Fall 2005

6.831 UI Design and Implementation

1



The next version of Microsoft Office (version 12) will have a radically different user interface. The menubar and toolbars will be merged into a single widget, called **command tabs**. Clicking on one of the tabs (“Write”, “Insert”, “Page Layout”, etc) switches to a different ribbon of widgets underneath. The metaphor is a mix of menubar, toolbar, and tabbed pane. (Notice how UIs have evolved to the point where new metaphorical designs are riffing on existing *GUI* objects, rather than *physical* objects. Expect to see more of that in the future.)

Needless to say, slavish **consistency** has been thrown out the window – Office 12 won’t have a menubar or toolbar anymore. But if we were slavishly consistent, we’d never make any progress in user interface design. If you look carefully at the interface, they *are* consistent in some important ways: (1) a File menu is still there, in the predictable place; (2) critical toolbar buttons are still there, like the Save and Undo buttons right next to File; (3) the command tab buttons resemble menubar menus, in both location and naming; (4) the ribbons look and act like rich toolbars, with familiar widgets and familiar affordances. So even though some new learning has to happen for Office 12 users, the knowledge transfer from other apps or previous Office is likely to be substantial.

One thing Office 12’s developers did very effectively is **task analysis**. In fact, they signed up thousands of Office users to a special program that collected statistics on how frequently they used Office commands and in which order – huge amounts of data that directly drove how commands were grouped into the command tabs, and which commands appear on command tabs as opposed to being buried in deeper dialogs. When a user interface designer can get this kind of data, you can do a lot to improve the usability for an average user. Web site designers are lucky, in this sense, because server logs give it to them for free! Microsoft had to do a lot more work to get it.

Some of the most complex parts of Office, such as page layout, will be packaged up as predefined templates displayed in **galleries** (like the dropdown menu shown here). This provides a **shortcut** to common choices for frequent users, while at the same time **simplifying** the interface for new users. Task analysis, again, is what drove the selection of these templates.

Office 12 will also provide more **feedback** about what a command will do, by showing a preview of its effect right in the document while you’re mousing over the command. So if you hover over the Wide Margins option in that dropdown list, your document will reformat to show you what the wide margins would look like. As long as your computer is fast enough to do it within 100ms, this would be a tremendous improvement to the visibility and feedback of the interface.

It’s hard to thoroughly evaluate Office 12 right now, since beta versions aren’t scheduled to be released until early 2006. Until then, you can find out more here:

<http://www.microsoft.com/office/preview/uioverview.msp>

<http://channel9.msdn.com/showpost.aspx?postid=114720>

## Nielsen's Heuristics

---

- **Meet expectations**
  1. Match the real world
  2. Consistency & standards
  3. Help & documentation
- **User is boss**
  4. User control & freedom
  5. Visibility of system status
  6. Flexibility & efficiency
- **Errors**
  7. Error prevention
  8. Recognition, not recall
  9. Error reporting, diagnosis, and recovery
- **Keep it simple**
  10. Aesthetic & minimalist design

Recall these 10 heuristics we discussed in an earlier lecture.

## Heuristic Evaluation

- Performed by an expert
- Steps
  - Inspect UI thoroughly
  - Compare UI against heuristics
  - List usability problems
    - Explain & justify each problem with heuristics

Fall 2005

6.831 UI Design and Implementation

4

One application of these 10 heuristics is a usability inspection process called **heuristic evaluation**. Heuristic evaluation was originally invented by Jakob Nielsen, and you can learn more about it on his web site. Nielsen has done a number of studies to evaluate the effectiveness of heuristic evaluation. Those studies have shown that heuristic evaluation's cost-benefit ratio is quite favorable; the cost per problem of finding usability problems in an interface is generally cheaper than alternative methods.

Heuristic evaluation is an inspection method. It is performed by a usability expert – someone who knows and understands the heuristics we've just discussed, and has used and thought about lots of interfaces.

The basic steps are simple: the evaluator inspects the user interface thoroughly, judges the interface on the basis of the heuristics we've just discussed, and makes a list of the usability problems found – the ways in which individual elements of the interface deviate from the usability heuristics.

The Hall of Fame and Hall of Shame discussions we have at the beginning of each class are informal heuristic evaluations. In particular, if you look back at previous lecture notes, you'll see that most of the usability problems are justified by appealing to a heuristic.

## How To Do Heuristic Evaluation

- Justify every problem with a heuristic
  - “Too many choices on the home page (Aesthetic & Minimalist Design)”
  - Can’t just say “I don’t like the colors”
- List every problem
  - Even if an interface element has multiple problems
- Go through the interface at least twice
  - Once to get the feel of the system
  - Again to focus on particular interface elements
- Don’t limit yourself to the 10 heuristics
  - We’ve seen others: affordances, visibility, Fitts’s Law, perceptual fusion, color principles
  - But the 10 heuristics are easier to compare against

Fall 2005

6.831 UI Design and Implementation

5

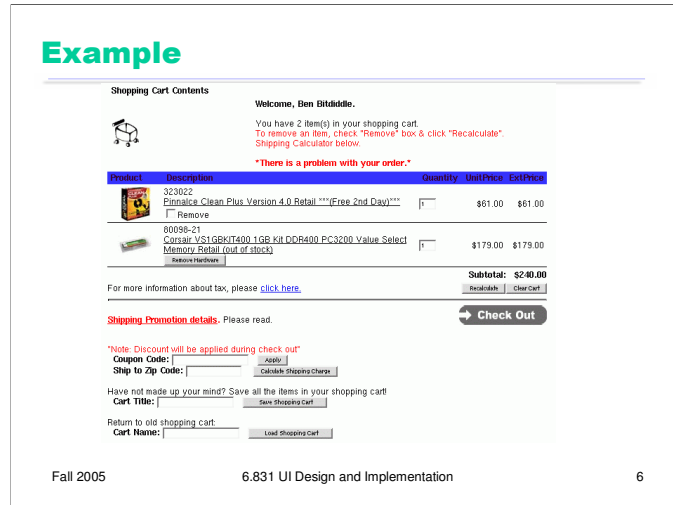
Let’s look at heuristic evaluation from the evaluator’s perspective. That’s the role you’ll be adopting in the next homework, when you’ll serve as heuristic evaluators for each others’ computer prototypes.

Here are some tips for doing a good heuristic evaluation. First, your evaluation should be grounded in known usability guidelines. You should justify each problem you list by appealing to a heuristic, and explaining how the heuristic is violated. This practice helps remove most of the (inevitable) subjectivity involved in inspections: You can’t just say “that’s an ugly yellow color.” (If it’s really yucky, you *should* pass that subjective opinion back to the design team, but you’ll be forced to identify it as subjective if you can’t find a heuristic to justify it.)

List every problem you find. If a button has several problems with it – inconsistent placement, bad color combination, confusing label – then each of those problems should be listed separately. Some of the problems may be more severe than others, and some may be easier to fix than others. It’s best to get all the problems on the table in order to make these tradeoffs.

Inspect the interface at least twice. The first time you’ll get an overview and a feel for the system. The second time, you should focus carefully on individual elements of the interface, one at a time.

Finally, although you have to justify every problem with a guideline, you don’t have to limit yourself to the Nielsen 10. We’ve seen a number of specific usability principles that can serve equally well: affordances, visibility, Fitts’s Law, perceptual fusion, color guidelines, graphic design rules are a few. The Nielsen 10 are helpful in that they’re a short list that covers a wide spectrum of usability problems. For each element of the interface, you can quickly look down the Nielsen list to guide your thinking.



Let's try it on an example. Here's a partial heuristic evaluation of the screen shown above. Can you find any other usability issues?

1. Shopping cart icon is not balanced with its background whitespace (Aesthetic & minimalist design)
2. **Good:** user is greeted by name (Visibility of system status)
3. Red is used both for help messages and for error messages (Consistency, Match real world)
4. "There is a problem with your order", but no explanation or suggestions for resolution (Error reporting)
5. ExtPrice and UnitPrice are strange labels (Match real world)
6. Remove Hardware button inconsistent with Remove checkbox (Consistency)
7. "Click here" is unnecessary (Aesthetic & minimalist design)
8. No "Continue shopping" button (User control & freedom)
9. Recalculate is very close to Clear Cart (Error prevention)
10. "Check Out" button doesn't look like other buttons (Consistency, both internal & external)
11. Uses "Cart Title" and "Cart Name" for the same concept (Consistency)
12. Must recall and type in cart title to load (Recognition not recall, Error prevention, Flexibility & efficiency)

## Heuristic Evaluation Is Not User Testing

- Evaluator is not the user either
  - Maybe closer to being a typical user than you are, though
- Analogy: code inspection vs. testing
- HE finds problems that UT often misses
  - Inconsistent fonts
  - Fitts's Law problems
- But UT is the gold standard for usability

Fall 2005

6.831 UI Design and Implementation

7

Heuristic evaluation is only one way to evaluate a user interface. User testing -- watching users interact with the interface -- is another. User testing is really the gold standard for usability evaluation. An interface has usability problems only if real users have real problems with it, and the only sure way to know is to watch and see.

A key reason why heuristic evaluation is different is that an evaluator is not a typical user either! They may be closer to a typical user, however, in the sense that they don't know the system model to the same degree that its designers do. And a good heuristic evaluator tries to think like a typical user. But an evaluator knows too much about user interfaces, and too much about usability, to respond like a typical user.

So heuristic evaluation is not the same as user testing. A useful analogy from software engineering is the difference between code inspection and testing.

Heuristic evaluation may find problems that user testing would miss (unless the user testing was extremely expensive and comprehensive). For example, heuristic evaluators can easily detect problems like inconsistent font styles, e.g. a sans-serif font in one part of the interface, and a serif font in another. Adapting to the inconsistency slows down users slightly, but only extensive user testing would reveal it. Similarly, a heuristic evaluation might notice that buttons along the edge of the screen are not taking proper advantage of the Fitts's Law benefits of the screen boundaries, but this problem might be hard to detect in user testing.

## Hints for Better Heuristic Evaluation

- Use multiple evaluators
  - Different evaluators find different problems
  - The more the better, but diminishing returns
  - Nielsen recommends 3-5 evaluators
- Alternate heuristic evaluation with user testing
  - Each method finds different problems
  - Heuristic evaluation is cheaper
- It's OK for observer to help evaluator
  - As long as the problem has already been noted
  - This wouldn't be OK in a user test

Fall 2005

6.831 UI Design and Implementation

8

Now let's look at heuristic evaluation from the designer's perspective. Assuming I've decided to use this technique to evaluate my interface, how do I get the most mileage out of it?

First, use more than one evaluator. Studies of heuristic evaluation have shown that no single evaluator can find all the usability problems, and some of the hardest usability problems are found by evaluators who find few problems overall (Nielsen, "Finding usability problems through heuristic evaluation", CHI '92). The more evaluators the better, but with diminishing returns: each additional evaluator finds fewer new problems. The sweet spot for cost-benefit, recommended by Nielsen based on his studies, is 3-5 evaluators.

One way to get the most out of heuristic evaluation is to alternate it with user testing in subsequent trips around the iterative design cycle. Each method finds different problems in an interface, and heuristic evaluation is almost always cheaper than user testing. Heuristic evaluation is particularly useful in the tight inner loops of the iterative design cycle, when prototypes are raw and low-fidelity, and cheap, fast iteration is a must.

In heuristic evaluation, it's OK to help the evaluator when they get stuck in a confusing interface. As long as the usability problems that led to the confusion have already been noted, an observer can help the evaluator get unstuck and proceed with evaluating the rest of the interface, saving valuable time. In user testing, this kind of personal help is totally inappropriate, because you want to see how a user would really behave if confronted with the interface in the real world, without the designer of the system present to guide them. In a user test, when the user gets stuck and can't figure out how to complete a task, you usually have to abandon the task and move on to another one.



## Formal Evaluation Process

1. Training
  - Meeting for design team & evaluators
  - Introduce application
  - Explain user population, domain, scenarios
2. Evaluation
  - Evaluators work separately
  - Generate written report, or oral comments recorded by an observer
  - Focus on generating problems, not on ranking their severity yet
  - 1-2 hours per evaluator
3. Severity Rating
  - Evaluators prioritize all problems found (not just their own)
  - Take the mean of the evaluators' ratings
4. Debriefing
  - Evaluators & design team discuss results, brainstorm solutions

Fall 2005

6.831 UI Design and Implementation

9

Here's a formal process for performing heuristic evaluation.

The training meeting brings together the design team with all the evaluators, and brings the evaluators up to speed on what they need to know about the application, its domain, its target users, and scenarios of use.

The evaluators then go off and evaluate the interface separately. They may work alone, writing down their own observations, or they may be observed by a member of the design team, who records their observations (and helps them through difficult parts of the interface, as we discussed earlier). In this stage, the evaluators focus just on generating problems, not on how important they are or how to solve them.

Next, all the problems found by all the evaluators are compiled into a single list, and the evaluators rate the severity of each problem. We'll see one possible severity scale in the next slide. Evaluators can assign severity ratings either independently or in a meeting together. Since studies have found that severity ratings from independent evaluators tend to have a large variance, it's best to collect severity ratings from several evaluators and take the mean to get a better estimate.

Finally, the design team and the evaluators meet again to discuss the results. This meeting offers a forum for brainstorming possible solutions, focusing on the most severe (highest priority) usability problems.

When you do heuristic evaluations in this class, I suggest you follow this ordering as well: first focus on generating as many usability problems as you can, then rank their severity, and then think about solutions.

## Severity Ratings

---

- Contributing factors
  - Frequency: how common?
  - Impact: how hard to overcome?
  - Persistence: how often to overcome?
- Severity scale
  1. Cosmetic: need not be fixed
  2. Minor: needs fixing but low priority
  3. Major: needs fixing and high priority
  4. Catastrophic: imperative to fix

Fall 2005

6.831 UI Design and Implementation

10

Here's one scale you can use to judge the severity of usability problems found by heuristic evaluation. It helps to think about the factors that contribute to the severity of a problem: its **frequency** of occurrence (common or rare); its **impact** on users (easy or hard to overcome), and its **persistence** (does it need to be overcome once or repeatedly). A problem that scores highly on several contributing factors should be rated more severe than another problem that isn't so common, hard to overcome, or persistent.

## Evaluating Prototypes

- Heuristic evaluation works on:
  - Sketches
  - Paper prototypes
  - Buggy implementations
- “Missing-element” problems are harder to find on sketches
  - Because you’re not actually using the interface, you aren’t blocked by feature’s absence
  - Look harder for them

Fall 2005

6.831 UI Design and Implementation

11

A final advantage of heuristic evaluation that’s worth noting: heuristic evaluation can be applied to interfaces in varying states of readiness, including unstable implementations, paper prototypes, and even just sketches. When you’re evaluating an incomplete interface, however, you should be aware of one pitfall. When you’re just inspecting a sketch, you’re less likely to notice missing elements, like buttons or features essential to proceeding in a task. If you were actually *interacting* with an active prototype, essential missing pieces rear up as obstacles that prevent you from proceeding. With sketches, nothing prevents you from going on: you just turn the page. So you have to look harder for missing elements when you’re heuristically evaluating static sketches or screenshots.

## Writing Good Heuristic Evaluations

- Heuristic evaluations must communicate well to developers and managers
- Include positive comments as well as criticisms
  - "Good: Toolbar icons are simple, with good contrast and few colors (minimalist design)"
- Be tactful
  - Not: "the menu organization is a complete mess"
  - Better: "menus are not organized by function"
- Be specific
  - Not: "text is unreadable"
  - Better: "text is too small, and has poor contrast (black text on dark green background)"

Fall 2005

6.831 UI Design and Implementation

12

Here are some tips on writing good heuristic evaluations. First, remember your audience: you're trying to communicate to developers. Don't expect them to be experts on usability, and keep in mind that they have some ego investment in the user interface. Don't be unnecessarily harsh.

Although the primary purpose of heuristic evaluation is to identify problems, positive comments can be valuable too. If some part of the design is *good* for usability reasons, you want to make sure that aspect doesn't disappear in future iterations.

## Suggested Report Format

---

- What to include:
  - Problem
  - Heuristic
  - Description
  - Severity
  - Recommendation (if any)
  - Screenshot (if helpful)

12. Severe: **User may close window without saving data** (error prevention)

If the user has made changes without saving, and then closes the window using the Close button, rather than File >> Exit, no confirmation dialog appears.

Recommendation: show a confirmation dialog or save automatically



## Cognitive Walkthrough: Another Inspection Technique

- Cognitive walkthrough = expert inspection focused on learnability
- Inputs:
  - prototype
  - task
  - sequence of actions to do the task in the prototype
  - user analysis
- For each action, evaluator asks:
  - will user know what subgoal they want to achieve?
  - will user find the action in the interface?
  - will user recognize that it accomplishes the subgoal?
  - will user understand the feedback of the action?

Fall 2005

6.831 UI Design and Implementation

14

**Cognitive walkthrough** is another kind of usability inspection technique. Unlike heuristic evaluation, a cognitive walkthrough is particularly focused on evaluating learnability – determining whether an interface supports learning how to do a task by exploration.

In addition to the inputs given to a heuristic evaluation (a prototype, typical tasks, and user profile), a cognitive walkthrough also needs an explicit sequence of actions that would perform each task. This establishes the *path* that the walkthrough process follows. The overall goal of the process is to determine whether this is an easy path for users to discover on their own.

Where heuristic evaluation is focusing on individual elements in the interface, a cognitive walkthrough focuses on individual actions in the sequence, asking a number of questions about the learnability of each action.

•Will user try to achieve the right subgoal? For example, suppose the interface is an e-commerce web site, and the overall goal of the task is to create a wish list. The first action is actually to sign up for an account with the site. Will users realize that? (They might if they're familiar with the way wish lists work on other site; or if the site displays a message telling them to do so; or if they try to invoke the Create Wish List action and the system directs them to register first.)

•Will the user find the action in the interface? This question deals with visibility, navigation, and labeling of actions.

•Will the user recognize that the action accomplishes their subgoal? This question addresses whether action labels and descriptions match the user's mental model and vocabulary.

•If the correct action was done, will the user understand its feedback? This question concerns visibility of system state – how does the user recognize that the desired subgoal was actually achieved.

Cognitive walkthrough is a more specialized inspection technique than heuristic evaluation, but if learnability is very important in your application, then a cognitive walkthrough can produce very detailed, useful feedback, very cheaply.