Steps toward a Discipline of Cognitive Engineering

A three-day exposition of the *didi* methodology, in which interaction design applies engineering rigor to become human optimization

Synopsis

This class is intended for practicing business analysts, new product development people, service delivery managers, and front-end programmers, but presents a "human technology," not coding, so business-side people responsible for defining and delivering new tools are a key part of the intended audience. It is a three day time-compressed version of the COMS 6998-005 seminar class in the Department of Computer Science, Columbia University; the result of a 2003 invitation to W. Bradford Paley to teach his interaction design methodology to graduate students, acquainting them with an alternative to traditional user interface design. It builds upon an understanding the mechanics of composing standard widgets into windows, but its main focus is on meaning: we try to ensure that application structure is directly driven by domain experts. It provides design processes that ensure widgets—both traditional UI buttons and grids and the innovative task-specific ones this methodology can produce—more closely map the way experts actually think while doing their jobs. The results are tools that lets people integrate more ideas, more naturally—lessening "information overload," improving efficiency, reducing risk—and letting people get back to doing the jobs they care about, rather than coping with their systems.

It details easily-applied strategies that let attendees work with business-side domain experts to extract and understand the behavior, decision-making steps, and the shared mental models that make them experts. Different strategies address different steps in the process:

- an *understanding of the problem domain and tasks* is captured in structured outline form,
- the *mapping from an expert's mental representation to widgets* is facilitated by rapid sketching with experts,
- *information layering* is driven by a careful application of visual perception findings, and
- *early but rigorous testing* is done with "paper prototypes."

Results from using the methodology in the financial domain over the last decade include interfaces that are immediately recognized by experts as illustrations of their workflow; resulting in measured 20:1 speedups in some critical tasks, 10:1 or greater information density improvements in data-rich displays, greater ease in comprehending and manipulating complex relationships among decision factors, and improvements in the ease—even fun—which people experience while executing their tasks. Though the most tangible results of the process are screen-based designs it is not strictly a graphic design or branding exercise. It is an optimization process: optimizing the link between human and computer, thereby optimizing the human's effectiveness and comfort. Nonetheless, the final screens telegraph the character of the tasks clearly, revealing real business understanding, clarity, and structure—they show that the designer really knows the business and cares about the comfort of the audience, bolstering visual branding with real meaning.

The class is structured to allow some normal work to be accomplished by attendees; it typically meets 10:00-12:30 and 1:30-4:00, in six sessions across three days. It has a few chapters of assigned background reading, and two light required exercises. It is presented as a seminar, and works best with seven to ten attendees, though more can be accommodated. Time is factored in to allow questions and reflections on current projects; this both locks the ideas into attendees' minds and paves the way toward adding immediate value to existing projects.

Topics

Session 1: A Lightning Run-Through

The entire contents of the class are discussed in reverse order: backwards with respect to how they are applied in practice. We start with easily understood visual layering effects and build to more abstract ideas concerning experts' mental models. The presentation is intentionally quick; it serves to lay out and integrate the dozen or more disciplines that provide the underpinnings of the methodology—from Typography to Vision's object recognition, from Attention and Memory to Psycholinguistics' jargon development. The remaining sessions are then presented in the order they will be applied on real-world projects.

Session 2: Extracting the Expert's Mental Model

Understanding the actual workflow—directly from the expert—provides the absolutely essential foundation for designing tools that properly support people executing their tasks. This session steps through social aspects of engaging with domain experts; from doing "homework" before meeting experts, to interviews, to on-site observation, to an active apprenticeship. Structured outlining techniques are taught; they capture the goals, tasks, steps, and cognitive states/transitions that drive the creation of schematic "candidate windows" in the next session.

Session 3: Paper and Pencil: the Ultimate Rapid Prototyping Tools

It is no easier to properly use paper and pencil than it is to stand before an audience and give a talk, or write a book; but fortunately the medium is familiar—and once learned extremely expressive. We teach how to use standard office tools (paper, pencil, eraser, tape, pens, highlighters) in a structured way to facilitate a conversation with an expert in which the *expert* is the one driving the creation of the interface. A quickly-acquired graphical ability supports the development of candidate window designs as "schematics": images that contain everything necessary to accomplish a task, sketched as layouts of "information objects." Schematics are often innovative, even visionary, as they capture the way information is organized for a given task in the expert's mind (rarely a spreadsheet-like grid). The shapes and contents of information objects are also approximated in these schematics in an attempt to make them maximally *recognizable*, *distinguishable* from one another, and *expressive*, in terms of the information needed to finish a window's target task.

Session 4: Information Layering: Managing Attention

Schematics hold all the functional parts of the interface, but final renderings as bitmaps develop the information layering that makes sure the functions are easy to find, recognize, and act upon. Perceptual distinctions drive our attention in many different ways: drawing the eye to new information; hiding expected, rarely needed labels and affordances in plain sight; grouping objects; hierarchically segmenting windows into simple, possibly re-usable modules; and clearly associating labels with what they label without distracting. Concepts such as pop-out search versus sequential search, flagging functional distinctions with distinct "visual vocabularies," and rapid categorization of objects based on feature recognition are discussed and tied to design. The visual cues that let experts feel at home and act effectively in the new tools we build are drawn from cues that let do so in the real world; for guidance we consider artists as "expert observers and visual translators" and learn to draw representation ideas from their works.

Session 5: Paper Prototyping

One of the most cost-saving parts of the methodology is the ability to test in more realistic situations much sooner than most development processes can. "Version 1.0s" are rarely fully usable, and coping mechanisms for the failure of early testing—such as a slowly staged roll-out and agile design—have become institutionalized. PowerPoint decks which get business-side sign-offs don't actually protect the business or developers from releases that don't contain everything needed to accomplish a task because they don't engage the actual intellectual mechanisms that are in play when a task is being accomplished. Paper Prototypes combine renderings of every significant state of a system with the intelligent paper-shuffling of the system designer in front of an expert. It is possible to not only engage the actual task-completion mechanisms before coding (even before specification writing), but immediately identify what's missing or wrong *and fix it* in a single session—saving huge amounts of development effort, costs, and calendar time, as well as improving business/developer relations. Bug detection and correction can be literally hundreds of times faster than we see in most current development processes, assuming a 10:1 time difference between designing and specification times another 10:1 time difference between spec-writing and coding.

Session 6: Expanding the Range of the Methodology and Concrete Review

The methodology is typically aimed at the well-understood, almost rote processes an expert in a specific domain of practice applies in their day-to-day business (like a NYSE floor broker doing a "refill"), and it reaps great benefits there. But many of the techniques can support more exploratory, open-ended work (like that of an analyst or portfolio manager), or work that is done by a wider pool of people (spanning several roles), or even new work processes that have never been offered before. In this wrap-up session we consider concrete examples suggested by the attendees to see which parts of the methodology can be applied or extended to unusual application domains.

Professor Bio

W. Bradford Paley wrote his first computer graphics and computer game code in 1973, pioneered exotic computer input devices and display techniques in the early 90s, and applied what came to be known as Information Visualization techniques to financial data as early as 1988. A self-described "lapsed engineer," he works to apply engineering rigor to optimize human communication mediated by computers; codifying techniques taken from areas as diverse as mind science, fine arts, evolutionary biology and typography.

He has a Phi Beta Kappa degree in Economics from the University of California at Berkeley, and teaches his Cognitive Engineering methodology as a graduate seminar at the Columbia University Department of Computer Science as an Adjunct Professor. He has given keynote or invited talks at most of the major conferences in Information Visualization and User Interface Design, and has won numerous awards including ID Magazine's Design Distinction award and Grand Prize [non-interactive] at the Japan Arts Media Festival. His work has been displayed in MoMA, he is a NYFA fellow and major NYSCA grantee, and his designs are in constant use in several major banks and two exchanges including the NYSE.

Presentations and Deployed Applications of the Methodology

Parts of the methodology have been presented in keynote talks for InfoVis 2003, CHI 2005's "alt.chi" subconference, the Ninth International Conference on Information Visualization (2005), Institutional Investor Financial Technology Forum (Winter Workshop, 2006), the Association for Computing Machinery's Interactive Tabletops and Surfaces (21010) and at numerous invited talks and panel presentations.

The methodology has been successfully applied for deployed projects at Goldman Sachs (GS's own NYSE Broker Handheld in 2000, written up as a feature article in the June 2000 ID Magazine), Morgan Stanley (Foreign Exchange system, both sales and trading; and an Internal organization chart/HR/permissioning/access control system), Merrill Lynch (Equities Sales Trading System), New York Stock Exchange (Specialist Workstation for Hybrid Market, partially deployed), New York Stock Exchange (Broker Handheld, design and prototype, partially deployed), Lehman Brothers (Prime Brokerage Web site, client Report Access system), Rosenblatt Securities (Custom Advanced Charting Package, Imbalance Tracker Tool), Earl Industries (Intelligence Community Analyst's Workstation, deployed as prototype), BIDS Trading (Trade Blotter: trade input, tracking, alerting), JP Morgan Asset Management (Trading blotter, execution optimizer, and related tools, deployed as prototype) and others.

Contact

W. Bradford Paley
Digital Image Design Incorporated (*didi*)
brad@didi.com
(917) 686-3274