

The Demonstrational Interfaces Project at CMU

Brad A. Myers,
Francesmary Modugno, Rich McDaniel, David Kosbie, Andrew Werth,
Rob Miller, John Pane, James Landay, Jade Goldstein, and Matthew A. Goldberg

Human Computer Interaction Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
<http://www.cs.cmu.edu/~bydemo>
bam@cs.cmu.edu

October 27, 1995

Abstract

The Demonstrational Interfaces Project at CMU has been investigating various aspects of demonstrational interfaces for the last eight years. During this time, we have created six interactive software building tools that use demonstrational techniques, as well as an architecture to support demonstrational programming in general. In addition, we have created a demonstrational Visual Shell (iconic interface to a file system, like the Macintosh Finder), a demonstrational text formatter, and a demonstrational charting tool. There are three fundamental research questions we explore through these tools: how to give the user appropriate feedback and control over inferencing, appropriate algorithms for inferencing, and which domains are appropriate for demonstrational techniques. This paper summarizes our activities, approach and lessons learned.

1. Introduction

In the User Interface Software Project and the Demonstrational Interface Project in the Human-Computer Interaction Institute at Carnegie Mellon University, we have created a variety of demonstrational, interactive programs. Like most demonstrational projects, the main goal is to allow non-programmers to achieve the effect of programming without needing to learn a textual programming language. We also wrote a number of general papers about Demonstrational Interfaces [16, 11]. All of our demonstrational systems have used heuristics to try to automatically generalize from the user's examples, unlike systems such as SmallStar [2] and KidSim [20]. Some of our systems, including Pursuit and Marquise, take the conventional approach of recording a script of the user's actions, but most of our systems generalize from the *results* of a demonstration and so are not sensitive to the particular sequence of actions to construct the examples.

This paper gives a quick summary of our various demonstrational systems, and then discusses our lessons learned.

2. Summary of Systems

As part of the large-scale User Interface Software Project, which developed Garnet [12] and now Amulet [18], we created a number of interactive demonstrational tools. Another set of systems were created as part of the Demonstrational Interfaces Project. In chronological order, the systems are Lapidary, Jade, C32, Gilt, Tourmaline, Marquise, Pursuit, Katie, Gold, Silk and Gamut.

Lapidary allows new widgets and new application-specific graphical objects to be created by drawing pictures and then specifying the behaviors using dialog boxes. The demonstrational part is that the objects are automatically generalized into prototypes which can be instantiated at run-time. Feedback in Lapidary includes a novel iconic representation for showing constraints, and dialog boxes for confirming inferences.

Jade automatically generates dialog boxes from a list of their contents [21]. The primary demonstrational component is that some of the layout and design rules can be specified by example.

C32 is a spreadsheet interface that allows complex constraints to be specified [13]. The main demonstrational aspect of C32 is that when constraint formulas are copied from one place to another, C32 generalizes and transforms the constraint appropriately to its new context.

Gilt is an interface builder that provides widget layout. Two demonstrational additions were made to Gilt. First, Gilt will infer *graphical styles* from an example dialog box, including the placement of objects and the properties to use [3]. The placement was based on *graphical tab stops* which provide a direct way to manipulate the alignment. The second demonstrational component was a technique for eliminating “call back procedures” which are used to connect widgets to applications [14]. It turns out that many call-backs perform fairly mundane operations, such as making other widgets become active or inactive, and Gilt allows these to be replaced by demonstrating the desired behavior.

Tourmaline is a text formatter that allows “macrostyles” to be specified by example [15, 22]. Unlike the styles in conventional editors like Microsoft Word, Tourmaline allows a macrostyles to contain different formatting for different parts of a header, such as the title, author and author’s affiliation. Tourmaline uses heuristics to distinguish the role and formatting of the parts.

Marquise is used to create graphical editors by demonstration. It supports demonstrating how the low-level “rubber-band feedback” looks, how selection handles are drawn and how they behave, how objects are created, and how palettes control what objects are created and how they look. As the user demonstrates events that cause an action (the “stimulus”), Marquise drops icons on the window showing where the events occurred. The user can then refer to these icons directly when demonstrating the response. This technique was subsequently used in Grizzly Bear [1], and will be substantially generalized

in our new Gamut system. Marquise also invented a new form of feedback for inferences where phrases in textual sentences serve as buttons that pop up alternative options.

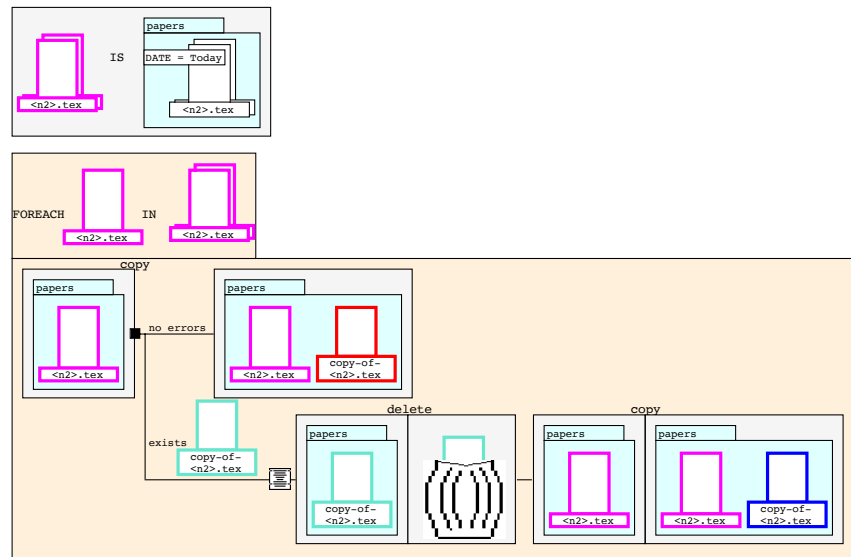


Figure 1: A complex program in the comic-strip language of Pursuit containing an explicit loop and a branch. The program copies each .tex file in the papers folder. The initial copy operation either succeeds (the upper branch) or fails (the lower branch). This conditional is depicted graphically by the branch (i.e., the little black square) and predicates after the first panel for the copy operation. Pursuit generated this program as the user demonstrated the actions on two actual file objects: one in which the copy executed successfully, and one in which the output file already existed in the papers folder.

Pursuit was the PhD thesis of Francesmary Modugno, and is a visual shell (an iconic interface to a file system, like the Macintosh finder) [10, 7, 8]. As the user demonstrates a program by executing example commands, Pursuit builds a visual language representation of operations and inferences. The visual language is based on the “comic strip metaphor:” the panels show the relevant data, and changes from one panel to the next represent the operation (see Figure 1). This language is innovative because, unlike most other textual and visual languages, it represents the data and leaves the operations implicit. The language provides a single medium for verifying and correcting inferences, for reviewing completed programs later, and for editing programs. Formal human factors experiments showed that the language and system were successful and usable.

Katie is the PhD thesis of David Kosbie, and is an architecture for supporting script-based demonstrational programs [9, 4]. An important innovation in Katie is the support for “aggregate events” which means that user’s actions are recorded at multiple levels: the low level event stream as well as the higher-level actions resulting from these events. Katie investigates how users might decide at what level

scripts containing these aggregate events might be played back. Katie inspired the novel hierarchical command architecture now used in Amulet [19].

Gold allows custom business charts to be specified by demonstration [17]. Gold supports column charts, stacked column charts, line charts, pie charts, and many forms of scatter charts. The user draws a few examples of the charting elements, such as rectangles or circles, and the system generalizes to make elements for all of the data in a spreadsheet. Graphical “link boxes” show the inferred association of the graphics to the spreadsheet.

Silk is the PhD research of James Landay, and it allows graphic designers to sketch an interface with a pen on a computer tablet [5, 6]. Silk recognizes the widgets as they are drawn and allows the interface to be exercised for testing. Storyboards can be sketched to show the temporal behavior of the interface. The inferred type of objects are shown by highlighting the name in a button panel, and the user can cycle to the next guess or explicitly select the correct choice.

Gamut is the PhD research of Rich McDaniel, and is our newest demonstrational system. Gamut will allow non-programmers to construct complete games and educational software by demonstration. The main focus of Gamut is on new interaction techniques and metaphors to allow the system to infer the complex rules that control real games.

3. Lessons Learned

In creating all of these systems, we are researching three major questions: what are useful interaction techniques for specifying the demonstration and for providing feedback about what was inferred, what are the appropriate algorithms for inferencing, and which domains are appropriate for demonstrational techniques.

3.1 Feedback and Control

The most important question is how the users will control, understand, and correct the inferences made by the system. Any system that generalizes from examples will occasionally guess wrong, and it is important that users know what the system is doing so they will feel comfortable and in control. This is related to a number of issues in Artificial Intelligence systems, including confirmation dialogs in speech systems and explaining the reasoning in knowledge based systems.

In the early Peridot system, I used a conventional question-and-answer dialog to confirm inferences. Users found this disruptive and tended to answer “yes” without thinking. In the Lapidary interactive user interface builder and in the Tourmaline text formatting system, we used dialog boxes to confirm inferences, which seemed more successful. In Marquise, we invented a novel form of natural language

feedback, where the user could click on phrases representing different parts of the inference to get a pop-up menu of alternatives. In the Gold custom charting system, text-input boxes appear near charting elements to show the inferred relationships.

Pursuit incorporates our most successful and novel mechanism for feedback. Despite skepticism from other researchers that its graphical programming language would be usable, formal user studies showed that nonprogrammers could create fairly complex programs, and that the visual language was more effective than an equivalent textual language.

The general lesson seems to be that graphical presentations work better than textual ones, and that users prefer a passive presentation of the inferences over one that interrupts with questions. For example, the graphical icons used for the constraints in Lapidary, the position of the mouse events in Peridot and Marquise, and the graphical language in Pursuit were more successful than the various question-and-answer techniques. In future systems, including Gamut, we will be exploring other graphical techniques for feedback and control.

3.2 Inferencing Algorithms

The second fundamental research question we address is what are appropriate representations and inferencing algorithms. For a system to be acceptable, it must guess right most of the time. Most of our systems, including all of our early ones, used straightforward rule-based techniques and pattern matching that were empirically tuned to give acceptable performance. These resulted in highly predictable user interfaces, and this technique has been adopted by most other demonstrational systems. In Gamut, we are exploring a more elaborate algorithm to provide powerful inferences that can take into account the user's hints.

All of these algorithms use models of the domain that underlie applications. For example, in Peridot and Lapidary, we used constraint and behavior models, in the Tourmaline text formatter, we used a model of the most common forms of section headers, in Gold, we have a model of the most common business charts, in Pursuit, we used a model of the common operations in a visual shell, and in Gamut we will use a model of the common properties of board games. These models constrain the inferences and significantly increase their accuracy.

3.3 Domains

The third research question is to which domains can demonstrational interfaces be successfully applied. We have identified useful aspects of user interface construction, text formatting, business charting, file manipulation in a Visual Shell, and educational game construction that are appropriate for being demonstrational. They share the properties that the natural way a person would describe the problem to

another person is by drawing examples, and that domain knowledge can be used to narrow the range of possibilities for generalizing from the examples. Each domain also illuminates new issues for feedback and representations.

In the future, we will be working to further develop demonstrational interfaces in new areas. In the Gamut tool mentioned above, we will significantly expand the range of what can be created by demonstration, by inventing new interaction techniques, metaphors and algorithms. We will also expand the work on creating custom business charts and data visualizations by demonstration. Other areas we will look into include computer-aided manufacturing and the world-wide web.

4. Conclusions

The Demonstrational Interfaces Project has created many interesting systems in different domains. We believe that there is a great potential for demonstrational interfaces to be a “step beyond direct manipulation” [16] and we look forward to collaborating with AI researchers to incorporate more elaborate and accurate inferencing algorithms.

Acknowledgements

The Demonstrational Interfaces Project has been primarily funded by NSF under grants IRI-9319969 and IRI-9020089, and by the Hertz Foundation. The User Interface Software Project has been primarily funded by NCCOSC under Contract No. N66001-94-C-6037, ARPA Order No. B326, and by the Avionics Lab, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U. S. Air Force, Wright-Patterson AFB, OH 45433-6543 under Contract F33615-90-C-1465, Arpa Order No. 7597. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

References

1. Martin R. Frank. *Model-Based User Interface Design by Demonstration and by Interview*. Ph.D. Th., College of Computing, Georgia Institute of Technology, 1995. In progress.
2. Daniel C. Halbert. *Programming by Example*. Ph.D. Th., Computer Science Division, Dept. of EE&CS, University of California, Berkeley, CA, 1984. Also: Xerox Office Systems Division, Systems Development Department, TR OSD-T8402, December, 1984.
3. Osamu Hashimoto and Brad A. Myers. Graphical Styles for Building User Interfaces By Demonstration. ACM SIGGRAPH Symposium on User Interface Software and Technology, Proceedings UIST'92, Monterey, CA, Nov., 1992, pp. 117-124.
4. David Kosbie. *Hierarchical Event Histories in Graphical User Interfaces*. Ph.D. Th., Computer Science Department, Carnegie Mellon University, 1996. In progress.

5. James Landay and Brad A. Myers. Interactive Sketching for the Early Stages of User Interface Design. Human Factors in Computing Systems, Proceedings SIGCHI'95, Denver, CO, May, 1995, pp. 43-50.
6. James Landay. *Interactive Sketching for the Early Stages of User Interface Design*. Ph.D. Th., Computer Science Department, Carnegie Mellon University, 1996. In progress.
7. Francesmary Modugno and Brad A. Myers. A State-Based Visual Language for a Demonstrational Visual Shell. 1994 IEEE Workshop on Visual Languages, IEEE Computer Society, St. Louis, MO, Oct., 1994, pp. 304-311.
8. Francesmary Modugno, T.R.G. Green and Brad A. Myers. Visual Programming in a Visual Domain: A Case Study of Cognitive Dimension. Proceedings of Human-Computer Interaction '94, People and Computers IX, Glasgow, Scotland, Aug., 1994, pp. 91-108.
9. David S. Kosbie and Brad A. Myers. Extending Programming By Demonstration With Hierarchical Event Histories. In Brad Blumenthal, Juri Gornostaev and Claus Unger, Ed., *Human-Computer Interaction: 4th International Conference EWHCI'94, Lecture Notes in Computer Science, Vol. 876*, Springer-Verlag, Berlin, 1994, pp. 128-139.
10. Francesmary Modugno. *Extending End-User Programming in a Visual Shell with Programming by Demonstration and Graphical Language Techniques*. Ph.D. Th., Computer Science Department, Carnegie Mellon University, 1995. Computer Science Technical Report CMU-CS-95-130.
11. Brad A. Myers. Invisible Programming. 1990 IEEE Workshop on Visual Languages, IEEE Computer Society, Chicago, Ill, Oct., 1990, pp. 203-208.
12. Brad A. Myers, Dario A. Giuse, Roger B. Dannenberg, Brad Vander Zanden, David S. Kosbie, Edward Pervin, Andrew Mickish, and Philippe Marchal. "Garnet: Comprehensive Support for Graphical, Highly-Interactive User Interfaces". *IEEE Computer* 23, 11 (Nov. 1990), 71-85.
13. Brad A. Myers. Graphical Techniques in a Spreadsheet for Specifying User Interfaces. Human Factors in Computing Systems, Proceedings SIGCHI'91, New Orleans, LA, April, 1991, pp. 243-249.
14. Brad A. Myers. Separating Application Code from Toolkits: Eliminating the Spaghetti of Call-Backs. ACM SIGGRAPH Symposium on User Interface Software and Technology, Proceedings UIST'91, Hilton Head, SC, Nov., 1991, pp. 211-220.
15. Brad A. Myers. Text Formatting by Demonstration. Human Factors in Computing Systems, Proceedings SIGCHI'91, N.O., LA, Apr, 1991, pp. 251-256.
16. Brad A. Myers. "Demonstrational Interfaces: A Step Beyond Direct Manipulation". *IEEE Computer* 25, 8 (August 1992), 61-73.
17. Brad A. Myers, Jade Goldstein, and Matthew A. Goldberg. Creating Charts by Demonstration. Human Factors in Computing Systems, Proceedings SIGCHI'94, Boston, MA, April, 1994, pp. 106-111.
18. Brad A. Myers, Rich McDaniel, Alan Ferency, Andy Mickish, Alex Klimovitski, and Amy McGovern. The Amulet Reference Manuals. Tech. Rept. CMU-CS-95-166, Carnegie Mellon University Computer Science Department, June, 1995. also Human Computer Interaction Institute CMU-HCII-95-102. WWW = <http://www.cs.cmu.edu/~amulet>.
19. Brad A. Myers and David Kosbie. Reusable Hierarchical Command Objects. Submitted for Publication.
20. David Canfield Smith, Allen Cypher and Jim Spohrer. "KidSim: Programming Agents Without a Programming Language". *Comm. ACM* 37, 7 (July 1994), 54-67.

- 21.** Brad Vander Zanden and Brad A. Myers. Automatic, Look-and-Feel Independent Dialog Creation for Graphical User Interfaces. Human Factors in Computing Systems, Proceedings SIGCHI'90, Seattle, WA, April, 1990, pp. 27-34.
- 22.** Andrew J. Werth. Tourmaline: Formatting Document Headings by Example. Master Th., Information Networking Institute, Carnegie Mellon University, Oct. 1992.