MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

Panel: Compelling Intelligent User Interfaces – How Much AI?

Larry Birnbaum, Eric Horvitz, David Kurlander, Henry Lieberman, Joe Marks, Steve Roth

TR96-28 December 1996

Abstract

Efforts to incorporate intelligence into the user interface have been underway for decades, but the commercial impact of this work has not lived up to early expectations, and is not immediately apparent. This situation appears to be changing. However, so far the most interesting intelligent user interfaces (IUIs) have tended to use minimal or simplistic AI. In this panel we consider whether more or less AI is the key to the development of compelling IUIs. The panelists will present examples of compelling IUIs that use a selection of AI techniques, mostly simple, but some complex. Each panelist will then comment on the merits of different kinds and quantities of AI in the development of pragmatic interface technology.

Proceedings of the 1997 International Conference on Intelligent Interfaces

MERL – A MITSUBISHI ELECTRIC RESEARCH LABORATORY
http://www.merl.com

Panel: Compelling Intelligent User Interfaces – How Much AI?

Larry Birnbaum
ILS/Northwestern

Eric Horvitz
Microsoft Research

David Kurlander
Microsoft Research

Henry Lieberman
MIT Media Lab

Joe Marks (chair)
MERL

Steve Roth
CMU

TR-96-28 October 1996

Abstract

Efforts to incorporate intelligence into the user interface have been underway for decades, but the commercial impact of this work has not lived up to early expectations, and is not immediately apparent. This situation appears to be changing. However, so far the most interesting intelligent user interfaces (IUIs) have tended to use minimal or simplistic AI. In this panel we consider whether more or less AI is the key to the development of compelling IUIs.

The panelists will present examples of compelling IUIs that use a selection of AI techniques, mostly simple, but some complex. Each panelist will then comment on the merits of different kinds and quantities of AI in the development of pragmatic interface technology.

1. First printing, TR96-28, October 1996. To appear in the Proceedings of the 1997 International Conference on Intelligent Interfaces.

Position Statements

Larry Birnbaum
It's the Content, Silly!

Intelligent systems only perform as well as their representations of the task they are trying to perform and of the world they are trying to perform it in. If these representations are reasonable, then even extremely simple algorithms can result in useful performance. If they aren't, then no algorithm, no matter how sophisticated, can yield good performance.

This observation suggests that in building intelligent interfaces, our main focus should be on content – the task being performed jointly by user and system, and the information being manipulated in this performance – and not algorithms – e.g., constraint propagation, or deduction, or belief networks, or tree search, or neural networks, or whatever other magic bullet is currently obsessing people. In my view, the extent to which these techniques work in a given application is mostly a reflection of the extent to which a good job has been done in modeling the task and domain. However, by failing to focus on these modeling issues explicitly, an opportunity to build powerful semantic technologies is being missed.

This presentation will describe our efforts to develop such semantic technologies for building structured hypermedia systems, educational software, interface construction tools, and performance support systems. It will focus on ASK Systems, a family of hypermedia case libraries developed at the Institute for the Learning Sciences to capture the "war stories" of human experts in video form, and make them available in a conversational mode of interaction implemented simply in graphical point and click interfaces. Over the past few years, dozens of these systems have been built, and many fielded, in domains as varied as military transportation planning, economic history, business-process reengineering, drug therapy for ALS, fluid-coupling design, and air-campaign planning. They are currently in commercial production by the Learning Sciences Corporation.

My talk will reflect joint work with Ray Bareiss, Gregg Collins, Roger Schank, and several other colleagues.

Eric Horvitz
AI and the User Interface:
On the Value of Inference

Researchers attempting to enhance the computer-user interface have often pursued opportunities with the use of complex inferential machinery. Unfortunately, the sophisticated machinery frequently does not deliver great value, and can be mimicked by simpler techniques. At times, inelegant application of automated reasoning may even interfere with human-computer interaction.

In the context of the current dominant interaction schemes, there is great opportunity to make user interfaces more compelling by focusing effort on better design of layout, controls, and functionality of user interfaces, and, where necessary, weaving into the designs relatively straightforward automation such as basic event monitoring coupled with simple rules and control heuristics, and straightforward pattern matching, similarity metrics, and

search techniques. Overall, we can make great strides by focusing on better UI design, taking into consideration the possibility of integrating simple automation procedures into the functionality of the user interface. We must remain alert to attempts to use sophisticated inference simply to get around poor design, or in lieu of better design combined with simple automation techniques.

That said, I would also like to make a strong case for continuing, in parallel, research and development on more sophisticated user models and reasoning machinery for building compelling intelligent user interfaces. More sophisticated reasoning machinery such as logical and probabilistic inference can provide us with the means for detecting complex sequences of events, for grappling with uncertainty, learning usage patterns, and for controlling tradeoffs in the functioning of user interfaces.

A computer system is often faced with uncertainty about a user's goals and needs. Although good designs can minimize the uncertainty in a user's intentions, methods for reasoning and decision making under uncertainty about user goals—in response to such evidence as a history of user actions—can be important in the operation of compelling interfaces. For a variety of tasks, it may be important to employ intelligent-reasoning machinery for disambiguating the goals of the user, and determining the best action to take (i.e., selecting the action(s) that will maximize the expected utility of the user) in the face of inescapable uncertainty about a user's goals.

Consider the problem of guiding the generation of compelling dialogue with a user. Effective dialogue will likely hinge on natural language processing, in conjunction with methods that can make decisions given uncertainty about a user's goals, the sense of words used by the user, the context at hand, and the long-term and short-term sequence of interface events. Another example where sophisticated inference may often be crucial comes from the domain of monitoring applications. It can be be critical to employ theoretically sound methods for managing the complexity of information displayed to a user in monitoring time-critical, high-stakes systems. Such techniques must balance the costs and benefits of hiding information via pruning, abstraction, or use of other form of summarization.

I will highlight the value of inference by describing some details of Lumiere, an experimental system developed by the Decision Theory and Adaptive Systems Group at Microsoft Research. Lumiere continues to reason about the goals and needs of users as they work with software by considering multiple user actions over time with Bayesian models. Lumiere also employs a probabilistic analysis of words in a user's query and integrates the results from analyses of actions and words to identify the informational needs of users.

There is much to be done in the realm of design and creative innovation with simple automation methods. However, as demonstrated by Lumiere, there is also rich opportunity to extend our understanding of the role of sophisticated automated reasoning for enhancing the user interface, and, more generally, for enhancing the overall human–computer interaction experience. We need to continue vigorous work in both of these realms.

David Kurlander
Intelligence in the Interface:
Often More Harm Than Good

Although the application of AI techniques to user interfaces still shows a great deal of promise, researchers in intelligent UI need to take a step back and gain perspective on the design tradeoffs that must be balanced in building real interfaces.

In designing intelligent interfaces, it is critical to weigh the advantages of using AI techniques versus the benefits of doing things traditionally. First, intelligent interfaces often make mistakes, and the cost of verifying and correcting inferences can be prohibitive. Second, AI techniques are often slow, and can make what otherwise might be an interactive interface seem unresponsive. Third, users need a clear mental model of how the computer will respond to their input, and some uses of intelligence in the interface cloud this model. Fourth, users often want an explanation of why the system did something they way it did, and some forms of AI decision processes are difficult to convey to the users.

All of these disadvantages of using AI in the interface can be overcome. In many cases, it is a question of where to apply intelligence, so that the benefits outweigh the problems. In other cases, conventional interface techniques can be applied to support the use of intelligence in the interface, and make it actually pragmatic. Ironically, people often apply AI to interfaces to make them more natural, but without crafting the interface to support the intelligence, the AI component makes the interface far less usable.

It is impossible to tell the true worth of an interactive technique or a new piece of interface technology, without getting it in the hands of real people. The previous IUI conference was attended by a mix of people from both the AAAI and CHI communities. UI researchers already know (or should know!) that interface research is worthless without seeing how people truly respond to the innovations. A large component of UI research is empirical. Although I am not suggesting that all interface research needs formal user studies, it is still critical to make research systems available for people to try out, so they can see for themselves what works. It will be interesting to learn how many of the research systems being reported at this conference are actually available for people to try!

Comic Chat is a graphical Internet chat program that I developed at Microsoft Research. It exploits some AI technologies in its interface, and is being used by thousands of people. I will briefly discuss some of what we learned from exposing our research to so many individuals, and the effect of various design decisions on the user experience.

Henry Lieberman
AI in User Interfaces:
Making Every Little Bit Count

Incorporating artificial intelligence capabilities such as reasoning or learning into user interfaces can be viewed as helpful or viewed as annoying. Some have suggested that the best approach is just to be very conservative in putting AI capabilities in user interfaces, so as not to risk the wrath of the user. A key is to realize that many user interface situations are

very underconstrained – the user can be presented with a wide variety of choices, and there may be no a priori reason for preferring one over another.

In these situations, some AI software, often perceived by the user as an "agent" can help by intelligently making suggestions that help the user choose. I introduce the concept of "agent defaults" to describe this notion. Using AI to help the user deal with underconstrained user interface situations casts the software in the role of a helpful assistant rather than an omniscient problem solver.

I'll illustrate this with *Letizia*, an agent that helps the user browse the Web. It acts as an "advance scout" recommending pages. Some key interface principles for successful integration of AI are:

- Don't disturb the user's interaction. It should always be possible for the user to ignore the agent. Suggest rather than act.
- Operate in real time. Much of the benefit of the agent comes from acting while the user is "busy"
- Watch what the user is doing. Take advantage of "free" information implicit in user actions.

Steve Roth
Iterating Between AI and UI Design: Letting the Product be the Guide

It seems obvious to say that user interfaces must be judged by the ease and effectiveness with which they are used by people to perform tasks. This means they usually need to be developed in concert with many other facets of a product - including how they coordinate with other applications, workspaces and interfaces. UI design must be driven by a clear picture of the product purpose.

Much of the research on intelligent interfaces, like other AI work, has been driven by different criteria and strategies, for example demonstrating the completeness with which a problem has been represented computationally or the correctness of an inference mechanism. Our experience in the SAGE project has been that when we shifted attention towards product development and focused on usability issues, the kinds of problems that needed to be addressed changed dramatically. Some changes required better direct manipulation techniques (sometimes eliminating the need to apply AI techniques). Other problems created new interesting IUI research. The question is not how much AI or what kind. It's how it can be selectively applied as needed as part of more complete product. I'll illustrate these points with experiences in the design of automatic presentation systems.