

Looking for the Bright Side of User Interface Agents
Ben Shneiderman *User interface agents offer exciting new opportunities, but progress would be greater if goals and terms were clarified.*

Promoters of agent-oriented interfaces, as I understand them, wish to allow users to carry out complex actions at later times on multiple, and possibly, remote machines. As with most user interface designs, the challenge is to create a mechanism that is powerful, yet comprehensible, predictable and controllable. A well-designed agent interface would enable users to specify tasks rapidly, be confident that they will get what they want, and have a sense of accomplishment when the job is done.

If users are unsure about what the agent will do, if they cannot repeat previously successful actions, or if they cannot understand the instructions, dialog boxes, and error messages, they will not tolerate the interface. There is growing danger that agents will be a deception and an empty promise. I fear that many promoters of 'intelligent agents' have

inadequately considered their designs, have not learned the lessons of history, and will omit adequate user testing to refine their vague notions.

What is an agent?

The first question I would ask a proponent of agents is - what is an agent (and what is not an agent)? Is a compiler an agent? How about an optimizing compiler? Is a database query an agent? Is the print monitor an agent? Is email delivered by an agent? Is a VCR scheduler an agent? Most agent promoters that I ask say no to these questions because they have some vague dream that an agent is more 'intelligent', probably anthropomorphic (human form), and somehow more adaptive than these applications.

Further questioning reveals a disturbingly vague set of responses. The usual example of having the computer monitor previous database queries so that it can automatically find relevant items in the future is often offered under the term 'knowbots.' However, the detailed mechanisms for such an agent are not well enough specified to allow an implementer to go to work. While many scenarios, including Apple Computer's famous 1987 videotape on the Knowledge Navigator, include an anthropomorphic representation of an agent, there are some designers who recognize the unwanted and unnecessary deception.

An appealing example emerged in recent stories on General Magic's Magic Cap. Your purchase of an airline ticket generates a planned action that checks the airline schedule two hours before departure and sends you email if there is a delay. This is comprehensible, predictable, and specific, and can be invoked or modified with a simple template or dialog box. Maybe the designers and marketeers like calling this an agent, but I'll bet that the public won't.

Looking at the components

My autopsy on agents reveals these components:

- anthropomorphic presentation
- adaptive behavior
- accepts vague goal specification
- gives you just what you need
- works while you don't
- works where you aren't

The first three seem appealing at first, but have proven to be counterproductive. The latter three are good ideas but can be achieved more effectively with other interface mechanisms. Our results with dynamic queries, user controlled information visualization, query templates, direct manipulation, triggers, control panels, scheduled procedures, planned actions, and other concepts have been demonstrated to be effective, whereas the agent scenarios have a quarter century history of failure. Many of the same people who have made exaggerated promises for artificial intelligence, natural language processing, voice and hand writing recognition, and robots are now pushing agents.

Some designers continue to be seduced by anthropomorphic scenarios, even though these have repeatedly been rejected by consumers. The effective paradigm for now and the future is comprehensible, predictable, and controllable interfaces that give users the sense of power, mastery, control, and accomplishment.

At the end of the day, most users want to feel that they have done a good job, not that some machine has done their work for them.

The control panel on your computer like the cruise control on your car (have you noticed that it is not called the intelligent driving assistant nor the auto agent?) emphasizes the ways you, the users, are in control. The users have goals and make choices, and the computer responds promptly to carry out the users instructions. Users want to be in control.

Similarly bank terminals have matured through the early phase in which anthropomorphic designs (Tillie the Teller, Harvey Wallbanker, BOB the Bank of Baltimore) dominated. The anthropomorphic styles are cute the first time, silly the second time, and an annoying distraction the third time. The designs have also passed through the second phase, focusing on technology, in which terms like electro, auto, and compu dominate. The third, and more durable, phase emphasizes names that convey the service provided to the user: CashFlow, MoneyMover, Money Exchange, 24-hour Bank, etc. Even the generic term Automated Teller Machines (ATMs) has become Advanced Transaction Machines.

Anthropomorphic terms and concepts have continually been rejected by consumers, yet some

Ben Shneiderman
Human-Computer
Interaction Laboratory
and Institute for
Systems Research
University of Maryland
College Park, MD 20742
e-mail: ben@cs.umd.edu

designers fail to learn the lesson. Talking cash registers and cars, SmartPhone, SmartHome, Postal Buddy, Intelligent Dishwasher, and variations have all come and gone.

Even the recent variation of Personal Digital Assistant had to give way to the more service oriented name now used in the Apple Newton ads: MessagePad. We'll leave it to the psychoanalysts to fathom why some designers persist in applying human attributes to their creations.

Adaptivity under the hood of your car is positive (the computer adjusts the carburetor based on air and engine temperature, etc.), but it is more tricky if it interferes with your choices. If your car decided that since your windshield wipers were on, you should be driving 20% slower you would be quite unhappy. Similarly, you probably would not like cars to monitor your driving habits and if you had too many abrupt stops, it would slow you down or suggest that you take a rest. On the other hand simple, predictable linkages such as turning on the interior lights when you open the door are an acceptable automation. The lesson is to be cautious in choosing how automatic or adaptive systems should be. A few mistaken choices and the users are angry — high accuracy is necessary. There is a Gresham's Law of Interaction: Bad experiences drive out good users.

Vague goal specification is dangerous, just ask the Sorcerer's Apprentice. On the other hand, high level goals that translate into many smaller predictable actions are great, for example, dragging a directory to the trash can conveniently deletes all the contents with a single action (reversible where possible). Who would use a system that allowed a vague command like 'Delete all old files' or even vaguer 'Delete all useless files'. I believe that specific predictable actions are desired. Allow users to display files in order by age or by frequency of use and then allow them to invoke specific predictable actions.

Back to a scientific approach

I close by invoking the basic principle of user interface research: move past first intuitions to scientific evaluations.

If you think you have an idea for a successful agent, build one and test it against viable alternatives, such as direct manipulation on action templates. Measure learning time, speed of performance, error rates, retention over time, and subjective satisfaction. Then we can talk some more.

But, possibly, just possibly, all this heated debate is excessive and agents will merely become the Pet Rock of the 1990s — everyone knows they're just for fun. 