Introduction to the Special Issue on "Usable AI"

Anthony Jameson, Aaron Spaulding, and Neil Yorke-Smith

Abstract

When creating algorithms or systems that are supposed to be used by people, we should be able to adopt a "binocular" view of users' interaction with intelligent systems: a view that regards the design of interaction and the design of intelligent algorithms as interrelated parts of a single design problem. This special issue offers a coherent set of articles on two levels of generality that illustrate the binocular view and help readers to adopt it.

Who Needs a Special Issue on "Usable AI"?

The term "usable" isn't heard all that much in discussions among AI people. You're more likely to hear it when listening to folks who are interested in the human side of computer use—such as people in the field of human-computer interaction (HCI).

But how much distance is there between these two fields? Maybe not that much. After all, the algorithms developed in AI research are often intended to be deployed in systems that involve some sort of interaction with users. The AI may contribute to the basic functionality of the system, such as recommendation or the support of task execution; or it may enhance the interfaces of a system, as with systems that enable human-like forms of communication between the user and the system. We will refer to interactive systems that incorporate some sort of AI technology (or technology that at one time was viewed as belonging to AI) as *interactive intelligent systems*.

Systems that are supposed to be used by people ought to be *usable*, taking into account human needs, capabilities, and the contexts of use. The field of HCI has accumulated a large repertoire of methods and principles for designing systems that fulfill this criterion.

So do people contributing AI components to interactive systems need to concern themselves with HCI? The answer can be "no", if one of the two following strategies is applied:

Strategy 1: Work on the technical optimization of algorithms of a type that has already been successfully deployed in usable interactive systems.

In many areas, it is known (from research and/or experience) that a system component that achieves particular technical goals can be put to good use in interactive systems (e.g., accurate methods for information retrieval, recommendation, or machine translation). AI researchers can therefore concentrate on improving their algorithms in terms of accepted metrics, without thinking constantly about users and usability. This general strategy has proved immensely useful—and in many cases probably inevitable—for the improvement of AI technology for interactive systems.

But there are limitations to what AI can contribute to interaction in this way. This approach manages to factor users out of the picture by making some assumptions about the forms that user-system interaction takes and the criteria for its success. When we want to deploy AI in new scenarios, with different success criteria for the AI components, we need to think explicitly about the impact that the AI will have on users. A second strategy often comes into play here:

Strategy 2: Develop AI algorithms that can help to realize an apparently beneficial new form of interaction; leave it to HCI people to design and test usable interfaces.

AI researchers often believe that some technology that they have created can lead to a new and improved functionality or interaction styles that can benefit users. They may then produce compelling demonstrators that seem to require only the intervention of skilled interaction designers (if even that) before they can be deployed successfully with users.

This strategy has the benefit of giving an AI-technology push to the advancement of interactive systems, exploiting what AI people know about what is now technically possible with AI. But it also has serious limitations:

When someone does in fact try to deploy the algorithms in question in a system that is really used by people, they are likely to discover that some changes to the technology are required before the system becomes truly usable and useful: For example, if an intelligent algorithm for the scheduling of personal activities is involved, it may turn out that users of personal scheduling systems have requirements that cannot be met using the algorithm in question. The algorithm may be based on unrealistic assumptions about how users schedule events in their personal lives or about the extent to which users want to provide explicit input to the system and to be able to understand and to second-guess the system. The algorithm may miss the opportunity to provide the support that users would appreciate most, providing instead functionality they consider relatively unimportant. Any of these reasons may be sufficient to leave the algorithm languishing as an infertile research prototype on the researchers' demonstration computers.

A Binocular View of Interactive Intelligent Systems

Both of the strategies just mentioned share the property of focusing almost entirely on the technical aspects of the intelligent system, making more or less explicit and specific assumptions about how users would interact with a system employing the technology. This focus can be called a *monocular view* of interaction with intelligent systems. As Figure 1 illustrates, people within an HCI background may similarly be inclined to take a different monocular view that focuses on user interaction while making assumptions about the underlying intelligent technology.

[Figure on last page.]

Figure 1. Illustration of monocular and binocular views of interactive intelligent systems.

A general theme of this special issue is that in many cases a *binocular view* is more effective: The questions of how a system's intelligence should be realized and how users should be able to interact with that intelligence are addressed simultaneously. In the binocular view, then, the search space is the cross-product of the technical design space and the interaction design space. As a result, new *combinations* of interaction design and intelligent technology may be discovered that yield desirable forms of

user-system interaction even though each part of the combination, seen in isolation, might seem ill-motivated.

The articles in this special issue illustrate many of the diverse forms that research and design within the binocular view can take. More specifically, each of them addresses one or more of the following questions about usability and AI:

1. How can the incorporation of AI enhance the usability of interactive systems?

   It would be naive to expect that making an interactive system more intelligent automatically enhances its usability; but there are many theoretically founded and empirically documented ways in which the incorporation of AI can help a system to fulfill usability criteria. AI people can benefit from an awareness of these opportunities.

2. In what ways can the incorporation of AI unintentionally diminish a system's usability, and how can these challenges be met successfully?

   It is well known that AI in an interactive system can have negative usability side effects such as diminished predictability and controllability. Strategies for preventing or mitigating such side effects are available, but many of them require adjustments to the AI technology itself, not just to the system's user interfaces.

3. How should the methods that are employed to ensure an interactive system's usability be selected, adapted, and applied to take into account the special demands introduced by the use of AI?

   Although a lot of the generally accepted methodology for designing and testing usable systems is applicable without much change to systems that involve AI, there are some differences that it is worthwhile for AI people to know about.

What's in This Special Issue?

The special issue addresses the three questions just listed on two levels of generality:

It begins on the more general level, with a historical reflection on the relationships between AI and HCI, followed by three *theme articles*, each of which addresses

one of the three questions just listed, introducing concepts and summarizing general lessons learned.

On the more specific level, there are seven *case studies*, each of which reports on experience with a particular intelligent interactive system, or a group of related systems, in such a way as to illustrate the general themes introduced in the theme articles.

Readers who would like to start with concrete examples are advised to read the case studies first, in any order. These articles vary in length from bite-sized summaries of lessons learned from previously published research to articles with the length and degree of detail of full conference papers. Most of the case studies include marginal notes that explain how the case study is related to the theme articles.

Readers who are already familiar with examples of the issues raised by the use of AI in interactive systems may want to start with the theme articles, which they may find to offer a novel perspective on this area. The theme articles include numerous references to the case studies.

How Did This Special Issue Come About?

This special issue originated with a workshop on "Usable AI" at the 2008 conference on Computer-Human Interaction (CHI 2008) in Florence. (Several AAAI Spring and Fall Symposia in recent years had addressed related topics.) Participants were a representative sample of researchers who had addressed in their own work some of the issues introduced above. Most of the workshop was devoted to attempts to synthesize their experience with different classes of systems. The discussion continued for about a year after the workshop, via a wiki and phone conferences, the goal being to produce a publication with the tightly knit structure described above. We are grateful to *AI Magazine* for its willingness to host this result of that work, and we hope that the readers of the Magazine will enjoy reading it.
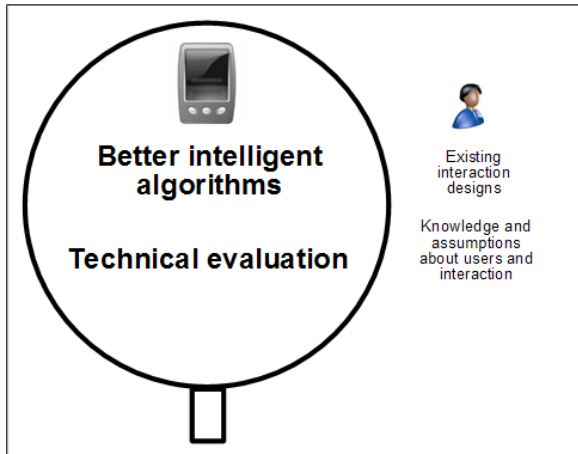
Short Bios

Anthony Jameson is a principal researcher at DFKI, the German Research Center for Artificial Intelligence. He has been conducting research at the intersection of AI and HCI for 30 years, being among the first researchers to address topics like conversational recommender systems, Bayesian approaches to user modeling, adaptation to situational resource limitations, recommendation via decision-theoretic planning, and the temporal evolution of users'

choices regarding interactive intelligent systems. He holds a PhD in Psychology from the University of Amsterdam.
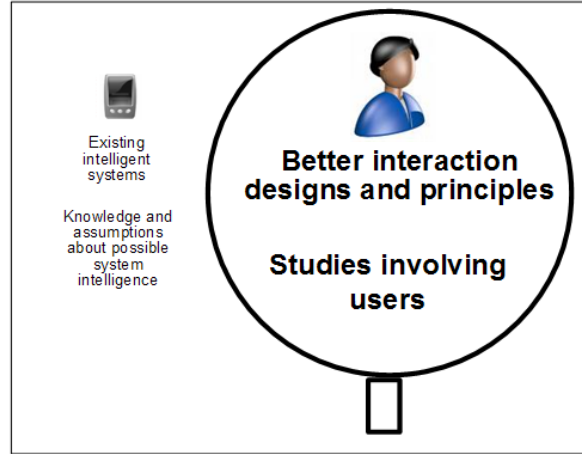
Aaron Spaulding is a computer scientist and interaction designer at SRI's Artificial Intelligence Center. His work centers on developing usable interfaces for AI systems that meet real user needs. He was an organizer of the CHI 2008 Workshop on "Usable Artificial Intelligence" and the CHI 2009 Special Interest Group on "Usable Intelligent Interactive Systems".

Neil Yorke-Smith is an Assistant Professor of Business Information and Decision Systems at the American University of Beirut, Lebanon, and a Research Scientist at SRI International's Artificial Intelligence Center. His research aims to help people make decisions in complex situations, with interests that include planning and scheduling, preferences, advisable agents, intelligent interfaces, and constraint programming, and their real-world applications. He holds a Ph.D. from Imperial College London.
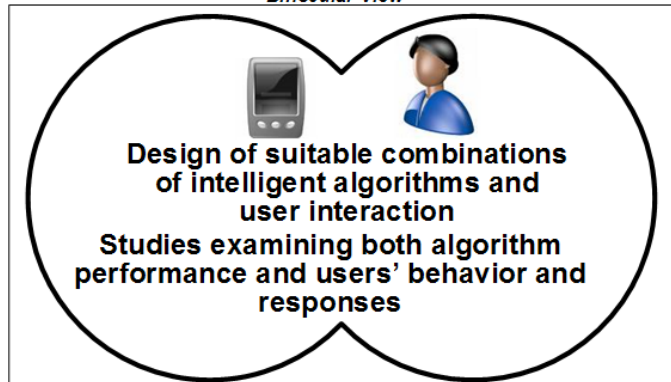
Figure 1. Illustration of monocular and binocular views of interactive intelligent systems.