

Introduction to the Special Issue on “Usable AI”

*Anthony Jameson, Aaron Spaulding,
and Neil Yorke-Smith*

■ *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.*

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 the provision of recommendations or the support of task execution; or it may enhance the interfaces of a system, as with systems that enable humanlike 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 or experience) that a system component that achieves particular technical goals can be put to good use in interactive systems (for example, 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 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, he or she is 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 suffi-

cient 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 with 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.

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.

First, 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.

Second, 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.

Third, how should the methods that are employed to ensure an interactive system's usability be selected, adapted, and applied to take into

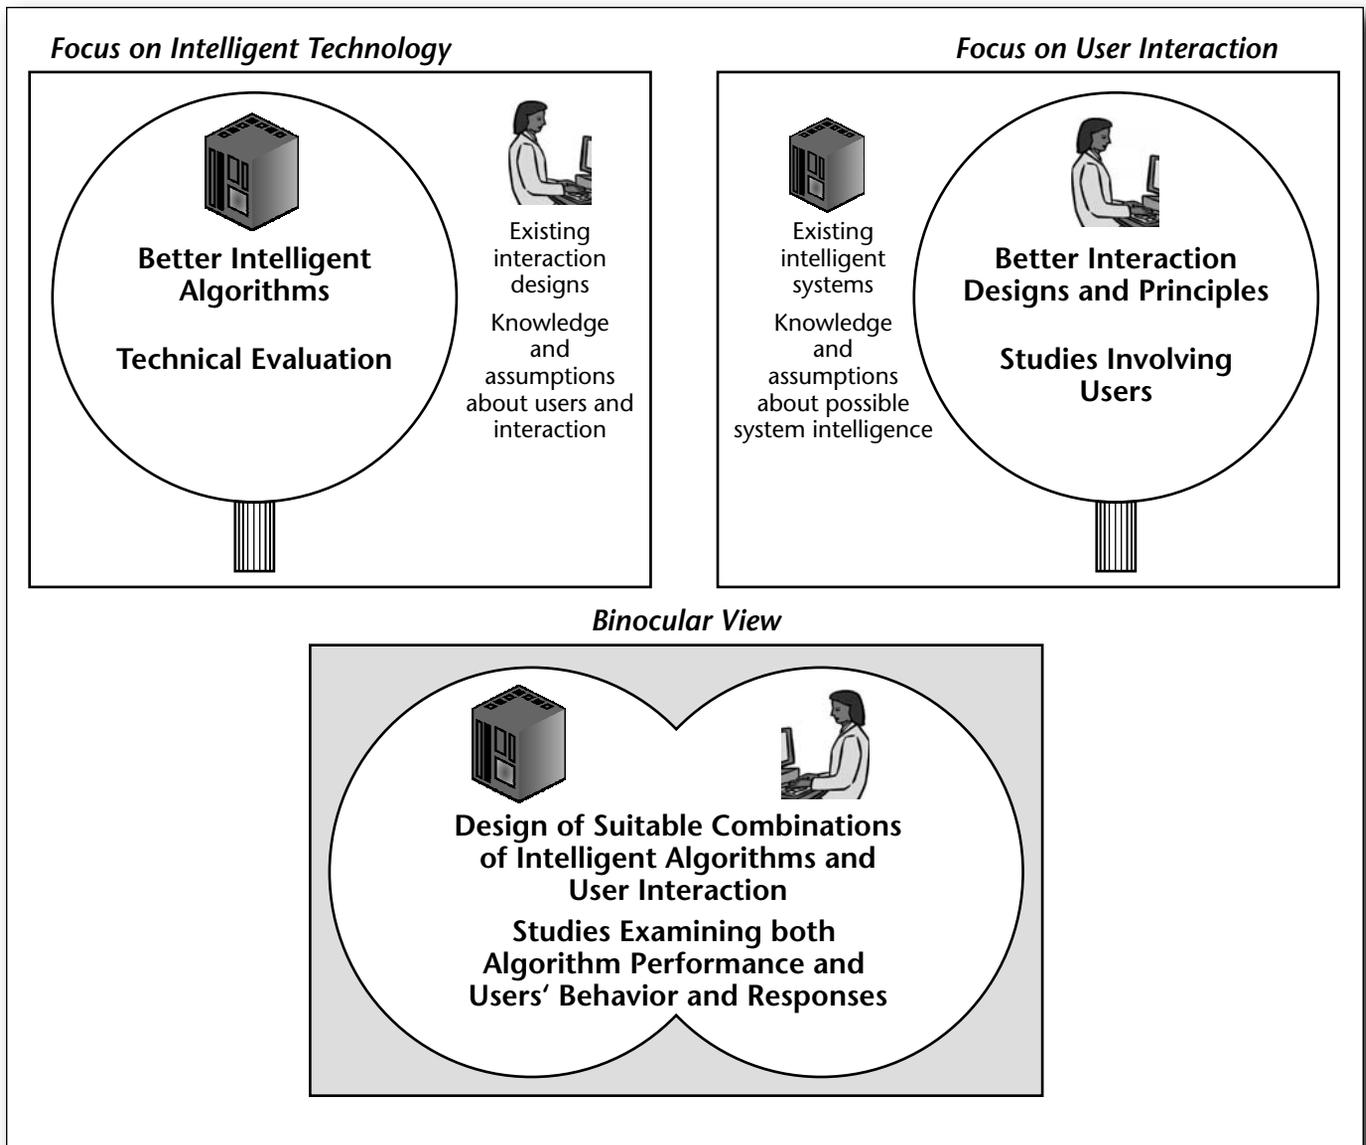


Figure 1. Illustration of Monocular and Binocular Views of Interactive Intelligent Systems.

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 straightforwardly applicable 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?

This 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 interactive intelligent 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 pre-



10th International Symposium on Smart Graphics

June 24-26, 2010 in Banff, Canada

The Tenth International Symposium on Smart Graphics will bring together researchers from computer graphics, visualization, interactive art, graphic design, cognitive psychology and artificial intelligence, all working on different aspects of computer generated graphics. This year's meeting will be held in the vibrant and creative mountain setting of the Banff Centre in Alberta, Canada.

Advances and breakthroughs in computer graphics have made visual media the basis of the modern user interface, and it is clear that graphics will play a dominant role in the way people communicate and interact with computers in the future.

Recent advances in this field have allowed AI researchers to integrate graphics in their systems, and on the other hand, many AI techniques have matured to the point of being easily used by non-specialists. These very techniques are likely to be the vehicle by which both principles from graphics design and the results of research into cognitive aspects of visual representations will be integrated in next generation graphical interfaces.

Important Dates:

Feb 14 Submission deadline
March 8 Notification of review results
March 20 Camera-ready copy due
June 24-26 Smart Graphics

Organizing Committee:

Robyn Taylor (University of Alberta)
Pierre Boulanger (University of Alberta)
Andreas Butz (University of Munich)
Marc Christie (IRISA/INRIA Rennes)
Brian Fisher (Simon Fraser University)
Antonio Krueger (DFKI)
Patrick Olivier (Newcastle University)

Proceedings:

Published as Springer LNCS

Hosted by:

University of Alberta

Venue:

Banff Centre
Banff, Alberta, Canada

In cooperation with:

AAAI, ACM SIGART, Eurographics

Symposium website:

www.smartgraphics.org

viously published research to articles with the length and degree of detail of full conference papers. Most of the case studies include 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, through a wiki and telephone 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.

Acknowledgements

The work of Anthony Jameson on this special issue was supported by funding from the targeted research unit Prevolution (code PsychMM), funded by the Autonomous Province of Trento, while the author was visiting FBK-irst. Anthony Jameson and Aaron Spaulding also benefitted from support from the project Halo 2, funded by Vulcan, Inc. The work of Aaron Spaulding and Neil Yorke-Smith on this special issue was supported in part by the Defense Advanced Research Projects Agency (DARPA) under Contract No. FA8750-07-D-0185/0004. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA, or the Air Force Research Laboratory.

Anthony Jameson, Ph.D. 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 through decision-theoretic planning, and the temporal evolution of users' choices regarding interactive intelligent systems. He is the founding coeditor (with John Riedl) of *ACM Transactions on Interactive Intelligent Systems*, whose purpose is to encourage and disseminate research that takes the binocular view described in this introduction.

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.

 The MIT Press

The Art of Agent-Oriented Modeling

Leon S. Sterling and Kuldar Taveter

"In *The Art of Agent-Oriented Modeling* readers will find an answer: a thorough description of all the ideas behind agent-oriented software engineering and a new approach to modeling that can fit many different methodologies. Far from being a painful set of definitions and procedures, it will be a pleasure to read." — Maurizio Martelli, Università di Genova

Intelligent Robotics and Autonomous Agents series
408 pp., 141 illus., \$38 cloth

Now in Paper

Online Stochastic Combinatorial Optimization

Pascal Van Hentenryck and Russell Bent

A framework for online decision making under uncertainty and time constraints, with online stochastic algorithms for implementing the framework, performance guarantees, and demonstrations of a variety of applications.

248 pp., 75 illus., \$18 paper

Probabilistic Graphical Models

PRINCIPLES AND TECHNIQUES

Daphne Koller and Nir Friedman

"This landmark book provides a very extensive coverage of the field, ranging from basic representational issues to the latest techniques for approximate inference and learning. As such, it is likely to become a definitive reference for all those who work in this area." — Kevin Murphy, University of British Columbia

Adaptive Computation and Machine Learning series
1208 pp., 399 illus., \$60 cloth

Now in Paper

Constraint-Based Local Search

Pascal Van Hentenryck and Laurent Michel

"This book presents a powerful new programming language paradigm for combinatorial optimization, uniting the power of local search with the declarativeness of constraint programming. This book will become an important reference for students and practitioners of combinatorial optimization." — Andrew J. Davenport, IBM T. J. Watson Research Center

448 pp., 102 illus., \$22 paper

Chess Metaphors

ARTIFICIAL INTELLIGENCE AND THE HUMAN MIND

Diego Rasskin-Gutman

translated by Deborah Klosky

"This book, in an accessible but profound way, approaches difficult but essential questions about the function of two types of intelligence destined to coexist as parent and child: human intelligence and artificial intelligence."

— Miguel Illescas Córdoba, International Chess Grand Master, Director, Chess Education and Technology, Spain

232 pp., 58 illus., \$24.95 cloth

To order call 800-405-1619 • <http://mitpress.mit.edu> • Visit our e-books store: <http://mitpress-ebooks.mit.edu>

2009 Proceedings from AAAI Press

International Conference on Automated Planning and Scheduling 2009

www.aaai.org/Press/Proceedings/icaps09.php

Twenty-Second International FLAIRS Conference

www.aaai.org/Press/Proceedings/flairs09.php

Third International Conference on Weblogs and Social Media

www.aaai.org/Press/Proceedings/icwsm09.php

International Conference on Computational Cultural Dynamics

www.aaai.org/Press/Proceedings/icccd09.php

Fifth Artificial Intelligence and Interactive Digital Entertainment International Conference

www.aaai.org/Press/Proceedings/aiide09.php

The Eighth Symposium on Abstraction, Reformulation, and Approximation

www.aaai.org/Press/Proceedings/sara09.php

www.aaai.org/Press