Integrative Conceptions of Interaction With Intelligent and Autonomous Systems

Anthony Jameson
DFKI, German Research Center for Artificial Intelligence
Saarbrücken, Germany
jameson@dfki.de

The ongoing research summarized here is being supported by the 7th Framework EU Integrating Project GLOCAL, Event-based Retrieval of Networked Media, under grant agreement 248984.

Abstract

This position statement in parallel (a) summarizes the author's experience with issues of end-user interaction with intelligent and autonomous systems and (b) sketches his integrative conceptions of these issues and the benefits of such conceptions.

Keywords

Intelligent systems, autonomous systems, interaction, usability

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User interfaces – Evaluation / methodology

[1] http://www.dfki.de/web/welcome?set_language=en&cl=en
[2] http://dfki.de/~jameson

Introduction

This position paper aims to do two things in parallel: (a) summarize my experience in the area of the workshop; and (b) introduce and motivate the perspective that I would like to represent at the workshop.

From 1977 through 2000, I worked on the topics of this workshop (among others) in academia as a student, instructor, researcher, and project director at several universities. Since 2001, I have worked in a mixed academic/industrial environment at the German Research Center for Artificial Intelligence (DFKI[1]), a research institute whose primary mission is to transfer innovative research on intelligent technologies to practical contexts in industry and public life.

My work in the area of intelligent and autonomous systems has addressed the following types of systems, in roughly chronological order: natural language dialog systems and conversational recommender systems; numerical uncertainty management for user and student modeling; on-line modeling of the consequences of working memory load and time pressure; decision-theoretic planning for interactive systems; language technology systems; recommendation to groups; recommendation based on life

logs; and systems for the exchange and discussion of media and documents that are organized in terms of the events that they concern. (Publications will be found on my web homepage.[2])

Limitations of the Predominant View of the Goals of Research in This Area

Whereas up through about the year 2000 my focus was on making interactive systems more intelligent, by that time I had come to believe that the major bottleneck concerns not the intelligence of the systems themselves but rather ways of ensuring that this intelligence is put to good use in the interaction with users—a goal that requires not only knowledge from computer science but also interdisciplinary integration of knowledge from other fields such as psychology and human-computer interaction. Initially, I adopted the point of view which predominates up to today: that of addressing particular goals such as understandability, controllability, and privacy protection whose fulfilment is necessary and largely sufficient if good interaction with intelligent and autonomous systems is to be ensured.

My first empirical study along these lines ([5]) was originally intended to provide empirical support for a typical assumption: that, for users, having control over adaptation (e.g., in an adaptive conference scheduling system) is preferable to having less control. I found that the results were more complex: Increased control was seen by users as being associated with a surprising variety of benefits but also with a surprising variety of drawbacks; and these were weighted differently by different users. From this point on, I saw a need to adopt a more comprehensive view of goals such as controllability, which (a) considers why and when they are worth pursuing and (b) how they are interrelated.

One reason why a broader view is required is the problem of trade-offs: As in most other contexts, meeting a given goal often requires concessions with regard to other goals; and finding a suitable trade-off requires an understanding of the importance and larger meaning of each goal.

Second, each of these goals can be pursued in different ways. For example, there are many alternative ways of trying to make a system's behavior comprehensible to users, which can lead to different types of comprehension. Deciding what type of comprehensibility to strive for depends on understanding the purposes of the comprehensibility.

Conception of Threats to Interaction with User-Adaptive Systems

My first integrative conception concerned a (large) subset of intelligent and autonomous systems: those which adapt to their users in some way. The section on "Usability Challenges" in my chapter on such systems in the 2003 *Human-Computer Interaction Handbook* ([1]), which was followed by updates in two later versions (e.g., [4]) is centered around five frequently discussed threats to the quality of interaction with user-adaptive systems: *diminished predictability and comprehensibility*, *diminished controllability*, *obtrusiveness*, *infringement of privacy*, and *diminished breadth of experience*. It embeds these threats in a larger framework by explaining for each one, (a) what *typical properties* of an intelligent system can give rise to it; (b) what *preventive measures* can be taken to prevent the threat from arising in the first place, and (c) what *remedial measures* can be taken to mitigate any negative side effects if prevention is unsuccessful. This framework supports the analysis of trade-offs by showing how measures intended to prevent or remedy one threat (e.g., diminished controllability) can tend to aggravate a different threat (e.g., obtrusiveness).

This framework, which can be summarized compactly on a single page, has proven useful over the years in consulting

with industrial partners, because it makes it possible to summarize many of the lessons learned from research and practice in this area within a comprehensive but comprehensible framework.

A More Inclusive Conception With a Temporal Dimension

I introduced a different sort of integrative framework as part of a special issue of *AI Magazine* that I coedited in 2009 titled *Usable AI*. Since my own discussion of "usability side effects of intelligent processing" ([2]) was intended to cover all forms of interactive system that make use of intelligent technology, it addressed a broader range of usability issues, including consequences of partly inaccurate processing, as well as practically important issues that are generally considered theoretically too uninteresting to merit discussion, such as the fact that novel intelligent technologies often run in computing environments which are incompatible with the familiar ones that are in everyday use.

Second, this overview devoted systematic attention to the often overlooked temporal dimension: the fact that the usability issues in question often take systematically different forms in early and later use of a given system by a given user. The article included a set of hypotheses about how the causes and consequences of particular usability threats change with increasing experience in operating the system, and it included some discussion of how such changes can be taken into account in the design of such systems.

A Choice and Decision Making Perspective

The two frameworks just summarized did not include carefully articulated accounts of the psychology of interaction with intelligent systems. My current work in this area is part of my more general effort to introduce into the HCI field a synthesis of knowledge about how people make (small) choices and (larger) decisions while using information technology. This perspective is relevant to intelligent and autonomous systems in two ways:

1. In many if not most cases, the function of such a system is to help people make good choices of one sort or another (e.g., which documents to read; how to configure an interface), by performing actions on behalf of the user entirely autonomously, by recommending choices to the user, or by supporting the choice process in some other way.
2. Even where the primary function is not to help users make choices in a given domain, intelligent and autonomous systems almost invariably require users to make nontrivial choices about how to interact with the system (e.g., to what extent to rely on the intelligent functionality). If these choices are made poorly by users, the quality of the interaction with the intelligent system can be impaired just as much as by poor system design or poor performance of the intelligent technology.

Despite this ubiquity and importance of user choices, research in this area has made little use of the vast amount of scientific knowledge that exists about how people make choices and decisions. This neglect is understandable given that the relevant research is scattered over hundreds of works that have appeared in numerous areas of psychology which often show little connection with each other and which hardly ever pay attention to the issues that arise with interaction with technology. In order to make the most relevant concepts and results of this research accessible to HCI researchers and practitioners, I have contributed a synthesis of it in a chapter in the 3rd edition of the *Human-Computer Interaction Handbook* ([3]), in a course offered at CHI 2011 and CHI 2012, and in a much longer multiauthored article that is currently in preparation for the journal *Foundations & Trends in HCI*.

Initial applications of this most recent conceptual framework to intelligent systems—specifically, recommender systems—occurred in workshop keynote addresses at relevant workshops at UMAP 2011[3] and RecSys 2011[4] as well as in a plenary panel discussion at the latter conference.[5]

Use of Frameworks for Application-Oriented Analysis, Design, and Evaluation

Since the late 1990s, I have aimed to transfer results of research on intelligent and autonomous systems to industry and other practical contexts in two complementary ways: (a) by consulting or participating in projects for companies like Daimler Benz, Deutsche Telekom, Vulcan, Inc., and many smaller companies; and (b) by presenting numerous tutorials at international conferences such as CHI, IJCAI, and IUI. The integrative frameworks have proven to be helpful in these contexts: People who are working on a particular system or problem do not in general want to know specifically, for example, how to improve the controllability of a given intelligent system. They want first to grasp the big picture of how people (might) interact with the system in question, including any issues and tradeoffs that may arise. An integrative framework is a good way to achieve this type of understanding, which can then be followed by analyses of specific problems whose relationships to the whole picture is understood.

Intended Contribution to the Workshop

At the workshop, I would like (a) briefly to explain the importance of integrative conceptions such as those sketched here and (b) use a suitable conception of this sort to help to organize and place in context the various contributions to the workshop.

References

[1] A. Jameson. Adaptive interfaces and agents. In J. A. Jacko and A. Sears, editors, *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*, pages 305–330. Erlbaum, Mahwah, NJ, 2003.

[2] A. Jameson. Understanding and dealing with usability side effects of intelligent processing. *AI Magazine*, 30(4):23–40, 2009.

[3] A. Jameson. Choices and decisions of computer users. In J. A. Jacko, editor, *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*. CRC Press, Boca Raton, FL, 3rd edition, 2012.

[4] A. Jameson and K. Z. Gajos. Systems that adapt to their users. In J. A. Jacko, editor, *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*. CRC Press, Boca Raton, FL, 3rd edition, 2012.

[5] A. Jameson and E. Schwarzkopf. Pros and cons of controllability: An empirical study. In P. De Bra, P. Brusilovsky, and R. Conejo, editors, *Adaptive Hypermedia and Adaptive Web-Based Systems: Proceedings of AH 2002*, pages 193–202. Springer, Berlin, 2002.

[3] http://www.di.uniba.it/~swap/DM/programme.html
[4] http://recex.ist.tugraz.at/RecSysWorkshop/keynote.jsp
[5] http://acmrecsys.wordpress.com/2011/10/25/panel-on-the-filter-bubble/