

## Editorial

Research and development in intelligent user interfaces is a fledgling discipline of computer science. It certainly touches on central technical and scientific challenges of our time, but at the same time its interdisciplinary nature has caused a slow, yet constant evolution of the field. By definition, intelligent interfaces encompass issues of artificial intelligence applied in the context of human–computer interaction. Both of these fields have unique characteristics that impact the growth of intelligent–interfaces research. On one hand, artificial intelligence enjoys a solid theoretical foundation although it has seen its progress impeded by years of boom or bust. On the other hand, human–computer interaction is a very young discipline gathering ever increasing interest from scientists, but lacking at times the sound theoretical underpinnings of more formal areas of informatics. Clearly, intelligent interfaces combine the good and bad sides of its underlying technologies. Thus, it is not surprising that we have witnessed a timid development in the intelligent interfaces area until recently.

The situation, however, has now changed for the better and a substantial body of work is starting to emerge. This work is attacking not only the central theoretical challenges of intelligent interfaces but it is also producing promising state-of-the-art applications of the corresponding technology. Naturally, this newly created scientific output demands an appropriate forum for presentation and discussion. Unfortunately, the main conferences in artificial intelligence and human–computer interaction are not geared towards interdisciplinary work. As a result, a group effort was started in 1995 to develop a conference series for intelligent user interfaces research and applications. Before then, two workshops had been held that demonstrated the increasing need for a full conference. Thus, in January 1997, the first International Conference in Intelligent User Interfaces (IUI97) was held in Orlando, Florida (USA). I had the pleasure, along with Professor Ernest Edmonds, of chairing the event.

IUI97 had a number of key themes, all of a highly interesting nature. Model-based interface development, agents, knowledge-based generation of presentations, evaluation, planning-based approaches for interaction, and intelligent, web-based applications were among the most relevant topics represented at the conference. If there was one overall thesis, however, that was prevalent throughout the conference it was this: The application, even if modest, of knowledge-based and knowledge engineering principles to the

process of human–computer interaction can greatly enhance the capabilities and effectiveness of user interfaces. I would emphasize the word modest here as many of the published papers showed that the simple application of basic fundamental principles of artificial intelligence were enough to open new interaction paradigms.

In the same spirit of breadth and depth of IUI97, we have gathered six of the outstanding conference papers and revised them as journal papers for this special issue of the Knowledge-Based Systems journal. These papers will provide the reader with a good understanding of some of the central challenges in intelligent interfaces as well as a view into some of the interesting applications that have emerged.

The paper by Robert St. Amant and Paul Cohen (Interaction with A Mixed-Initiative System for Exploratory Data Analysis) is an excellent example of the use of planning approaches to improve interaction. Their system, AIDE, assists the user in the exploration of statistical data by exploiting a library of exploration plans. Notice the word assists in the previous sentence. Increasingly, many of the applications of intelligent interfaces are shying away from brute-force, purely automated mechanisms and moving into a paradigm of supporting the activities and decision-making process of users. AIDE provides empirical evidence of the potential success of the latter approach.

Michelle Zhou and Steve Feiner (Top-Down Hierarchical Planning of Coherent Visual Discourse) explore the concept of declarative visual discourses. Borrowing from ideas of model-based interface development technology, this paper defines mappings between task decompositions, domain objects, and visual interaction elements. The result is a computational foundation for the generation of visual discourses that effectively improves over previous methodologies.

Hirohiko Sagawa, Masaru Takeuchi, and Masaru Ohki (Description and Recognition Methods for Sign Language Based on Gesture Components) make significant inroads into the problem of recognizing sign language gestures. Considerable progress is achieved here by the simple application of basic knowledge representation principles. The authors define the structure and organization of sign gestures and exploit it via knowledge engineering techniques. The evaluation of the resulting interface is very encouraging.

The issue of interface agents is presented through two papers. The first one, by Douglas Moran, Adam Cheyer,

Luc Julia, and David Martin (Multimodal User Interfaces in the Open Agent Architecture), highlights a number of applications based on sets of small, flexible, and reusable interface agents. Using a knowledge-based architecture as a working framework, the authors show how the concept of openness can be exploited in the realm of intelligent agents. The second paper, by Milind Pandit and Sameer Kalbag (The Selection Recognition Agent: Instant Access to Relevant Information and Operations), applies the idea of interface agents who observe the actions of users and can detect interesting patterns and offer relevant help. In this case, the agent recognizes words and phrases in the text that users are typing or browsing. The agent is at all times under user control, again moving away from the problem of automation of processes to that of assistance to users.

The paper from Kristina Höök (Evaluating the Utility and

Usability of an Adaptive Hypermedia System) touches on a seldom-discussed area in intelligent interfaces, namely their usability. The author presents a user study on an adaptive hypermedia system that produces very positive results. There are a number of good lessons to be learned here regarding the evaluation of this type of interfaces.

Overall, this collation of papers represents not only the latest advances of the field but also gives the pulse of the philosophy with which artificial intelligence techniques are being applied to human-computer interaction. Together, they signal the beginning of maturity for the field of intelligent interfaces and augur the coming of additional important future work.

**Angel R. Puerta**  
**Guest Editor**  
**Stanford University**