PREFACE

Virtually every application of artificial intelligence requires the use of context. At the very least, the application domain itself provides an important context for reasoning, and AI applications can be more effective and efficient if they exploit this contextual knowledge. Often, the application must work in several different contexts. For example, an interface will have to work with different users, a medical diagnosis system will have to deal with different patients and different disease presentations, and an autonomous underwater vehicle will need to work effectively in different geographical locations and terrain types.

For applications to take their context into account and to profit from contextual knowledge, context and contextual knowledge must be presented explicitly. Many existing approaches in AI, both theoretical and applied, take context into consideration implicitly. That is, contextual knowledge may be present in the approach, but it is not explicitly identified as distinct from other kinds of knowledge, nor is context identified as an entity about which inferences can be made. This, like any other sort of implicit representation of knowledge, denies the reasoner access to its own knowledge about context (e.g., for learning or reasoning about its adequacy), leads to redundant representation, and makes it difficult for humans to maintain the knowledge base.

Recently, a number of researchers have been working on explicitly modeling and using context. Other researchers have likely found themselves including context in their applications out of necessity, which may lead to important insights to help understand context in a general way.

The goal of this workshop is to gather together researchers concerned with the explicit representation and use of context in AI applications. This narrow focus will allow the exploration of how theoretical results can be used for specific applications and how the use of context in a particular application might be generalized to help others. The workshop provides the participants with a forum for presenting and discussing the issues pertinent to their work with others working on very similar problems.

The participants were required to address questions such as:

- What are appropriate representations of context for AI applications?
- How can contextual information be used in AI programs?
- How does the application and domain affect the representation of context?
- How can existing applications be extended to explicitly represent context?
- How can contextual information be learned by AI programs, and how can existing contextual information be updated from experience?
- What general principles can be developed based on the use of context and contextual knowledge to solve particular problems in an application?
- Can we arrive at a common, useful definition of "context" and "contextual knowledge" for AI applications?

Nineteen works have been selected for the workshop. Each of them directly concerns the modeling and use of context in a real-world application in different domains. An original feature of this workshop is the existence of an electronic discussion among the authors during the writing of the papers. This sometimes vigorous discussion facilitated a harmonization of the different views presented during the workshop.

Persons interested in more information about the workshop or about research on context in general may visit the Context in AI Web page, http://cdps.umcs.maine.edu/Context.