

Preface and Acknowledgments

Though agents are a rapidly growing area of research in the artificial intelligence and Java communities, most past definitions of agenthood are behavioral—agents are programs that exhibit certain hard to specify behaviors. Principled software techniques to build agents and converting existing programs (which may have been developed in PASCAL, C, LISP, etc.) and which may have used diverse data structures did not exist. In this book, we provide a theoretically clean, yet practically realizable way to build and deploy agents on top of legacy and/or specialized data structures and code bases. We believe (and hope) this book will be of interest to anyone interested in building software agents in a principled manner and deploying them on a network such as the Internet. The chapters in the book consist of three basic parts.

1. All readers should read chapters 1–4 and 6. These chapters deal with basic ideas and requirements underlying agent design, together with basic information on how agents may be programmed, how agents may declare the services they offer, and what functionalities must be supported by an agent infrastructure supporting interoperation between different agents.
2. Readers with specific interest in implementing agents and agent systems should read chapters 5, 12, 13. These chapters specify how agents and the underlying agent interoperability infrastructure interact. They also specify details about development environments and applications of agent technology.
3. Readers with interests in advanced agent reasoning should read chapters 7–11. These chapters describe extensions of the basic agent framework so that agents can reason about beliefs, about time, about uncertainty, and about security requirements. Complexity results are also described.

The work described in this book is an international effort for which funds were provided, in part, by the Army Research Laboratory under contract DAAL01-97-K0135, the National Science Foundation under grants IRI-93-57756, IIS-9820657 and IIS-9907482, the Austrian Science Fund Project N Z29-INF and a grant from Germany's DAAD.

We note that the following chapters have been derived, in part, from our published or to be published papers referenced below:

1. Chapter 4 and 6 are derived from reference *T. Eiter and V. S. Subrahmanian and G. Pick. Heterogeneous Active Agents, I: Semantics, Artificial Intelligence 108(1-2):179–255, 1999.*
2. Chapter 7 is derived in part from *J. Dix and V. S. Subrahmanian and G. Pick. Meta Agent Programs, accepted for publication in Journal of Logic Programming.*
3. Chapter 11 is derived from *T. Eiter and V. S. Subrahmanian. Heterogeneous Active Agents, II: Algorithms and Complexity, Artificial Intelligence 108(1-2):257–307, 1999.*

4. Chapter 12 is derived from *T. Eiter and V. S. Subrahmanian and T. J. Rogers. Heterogeneous Active Agents, III: Polynomially Implementable Agents, Artificial Intelligence 117(1):107–167, 2000.*

We are grateful to T. J. Rogers, Carolyn Gasarch, John Benton and Jason Ernst for extensive work on the implementation of different parts of the *IMPACT* system. In particular, T. J. was the guiding force behind the *IMPACT* system implementation.

We received extensive technical comments from researchers worldwide including Jim Hendler, Dana Nau, Luis Pereira, Bob Kowalski, Keith Clark, Terry Swift, Alex Dekhtyar, Hector Munoz, Jeff Horty, Daniel Veit, Leon Sterling, and Mirco Nanni.

We have benefited from numerous discussions on agents with folks from the U.S. Army including Jagdish Chandra, Phil Emmerman, Larry Tokarcik, Dave Hislop, Joseph Schafer, Jack Marin, Steve Choy, Swati Allen, Tim Gregory, Uma Movva and Paul Walczak. Mike Gaughan of Lockheed Sanders was a valued collaborator.

In addition, we would like to thank Bob Prior, Michael Sims, and Judy Feldmann of the MIT Press for extensive assistance.