Preface

Pei WANG, Ben GOERTZEL, Stan FRANKLIN

The field of Artificial Intelligence (AI) was initially directly aimed at the construction of "thinking machines" – that is, computer systems with human-like general intelligence. But this task proved more difficult than expected. As the years passed, AI researchers gradually shifted focus to producing AI systems that intelligently approached specific tasks in relatively narrow domains.

In recent years, however, more and more AI researchers have recognized the necessity – and the feasibility – of returning to the original goal of the field. Increasingly, there is a call to focus less on highly specialized "narrow AI" problem solving systems, and more on confronting the difficult issues involved in creating "human-level intelligence," and ultimately general intelligence that goes beyond the human level in various ways. "Artificial General Intelligence (AGI)", as this renewed focus has come to be called, attempts to study and reproduce intelligence as a whole in a domain-independent way.

Encouraged by the recent success of several smaller-scale AGI-related meetings and special tracks at conferences, we took the initiative to organize the very first international conference on AGI. Our goal in doing so was to give researchers in the field an opportunity to present relevant research results and to exchange ideas on topics of common interest.

The response to AGI-08 was even stronger than we expected. We received many interesting papers addressing a wide range of AGI-relevant topics and exploring various kinds of theoretical and technical ideas. Given the complexity and difficulty of the problems the field is facing, we believe it is crucial to encourage the exchange of ideas and opinions with various degrees of maturity, ranging from position papers to descriptions of mature mathematical and cognitive AGI theories, and practical work with implemented AGI architectures.

In this collection, the AGI-08 papers are organized into three groups. The conference papers are divided into full-length papers (12 pages, with a few exceptions) and short position statements (5 pages). Also included are the papers presented in the post-conference workshop on the sociocultural, ethical and futurological implications of AGI.

We believe meetings like AGI-08 are important, not only because of their presentations and discussions, but also because of the potential they have to help catalyze the self-organization of a vital AGI research community. Together, we will continue to directly challenge one of the most difficult and essential problems in human history, the creation of human-level artificial general intelligence.

We thank the following referees for devoting their valuable time to reviewing the submitted papers: Sam Adams, James Anderson, Mike Anderson, Eric Baum, Mark Bickhard, Henry Brighton, Nick Cassimatis, Hernan Castro, L. Andrew Coward, Hugo de Garis, Wlodzislaw Duch, Richard Duro, David Friedlander, John Hall, David Hart, Marcus Hutter, Cliff Joslyn, Randal Koene, Moshe Looks, Bruce MacLennan, Don

Perlis, Matthias Scheutz, Juergen Schmidhuber, Lokendra Shastri, Boris Velichkovsky, Karin Verspoor, Paul Vogt, and Mark Waser.

We would also like to thank the other members of the AGI-08 Organizing Committee, Sidney D'Mello, Bruce Klein, and Lee McCauley, for the thought and effort they expended in preparing the conference; and Bruce Klein and Natasha Vita-More for their help with the post-conference workshop which resulted in a number of the papers contributed to this volume.