

Preface

Ben GOERTZEL

The topic of this book – the creation of software programs displaying broad, deep, human-style general intelligence -- is a grand and ambitious one. And yet it is far from a frivolous one: what the papers here illustrate is that it is a fit and proper subject for serious science and engineering exploration. No one has yet created a software program with human-style or (even roughly) human-level general intelligence – but we now have a sufficiently rich intellectual toolkit that it is possible to think about such a possibility in detail, and make serious attempts at design, analysis and engineering. This is the situation that led to the organization of the 2006 AGIRI (Artificial General Intelligence Research Institute) workshop; and to the decision to pull together a book from contributions by the speakers at the conference.

The themes of the book and the contents of the chapters are discussed in the Introduction by myself and Pei Wang; so in this Preface I will restrict myself to a few brief and general comments.

As it happens, this is the second edited volume concerned with Artificial General Intelligence (AGI) that I have co-edited. The first was entitled simply *Artificial General Intelligence*; it appeared in 2006 under the Springer imprimatur, but in fact most of the material in it was written in 2002 and 2003. It is interesting to compare the material contained in the present volume, which was written in 2006, with the material from the previous volume. What is striking in performing this comparison is the significant movement toward practical realization that has occurred in the intervening few years.

The previous volume contained some very nice mathematical theory (e.g. by Marcus Hutter and Juergen Schmidhuber) pertaining to AGI under assumptions of near-infinite computational resources, some theory about the nature of intelligence as pertaining to AGI, and some descriptions of practical AGI projects at fairly early stages of development (including the NARS and Novamente systems developed by Pei Wang and myself respectively). The current volume, on the other hand, seems to represent significant progress. To take just a few examples: In the current volume, there is theoretical work (Eric Baum's and Moshe Looks' papers) that takes up Hutter's and Schmidhuber's emphasis on algorithmic information, and ties it in with practical suggestions regarding near-term AI design. My own Novamente system, which was described in fairly abstract terms in the earlier volume, is here represented by several papers by various authors reporting specific mathematical and experimental results, concretizing some (though by no means all, yet!) of the speculations made in the paper on Novamente in the previous volume. And, here we have a sufficient number of AGI design proposals, depicted in sufficient detail, that we have considered it worthwhile to include a chapter specifically comparing and contrasting four of the designs presented herein (Novamente, NARS, and the designs proposed by Stan Franklin and Alexei Samsonovich in their chapters).

In sum, what seems evident in comparing the prior volume with this one is that, while the end goal of the AGI research programme has not yet been achieved (and the

proximity of achievement remains difficult to objectively predict), the field is gradually broadening its scope beyond mathematical and conceptual ideas, and becoming more of a practical pursuit.

And I am confident that if there is another edited volume in another 2 or 3 years time, the field will appear yet further dramatically advanced. The “AGI Winter” is thawing, and the AI field is now finally making sensible progress toward its original goal of creating truly thinking machines. The material presented here only scratches the surface of the AGI-related R&D work that is occurring around the world at this moment. But I am pleased to have had the chance to be involved in organizing and presenting at least a small percentage of the contemporary progress.

Finally, thanks must be extended to those who helped this volume, and the workshop that inspired it, to come into being. Bruce Klein deserves the lion’s share of thanks, as the 2006 AGIRI Workshop would not have come into being without his extraordinary vision and dedication. Everyone who attended the workshop also deserves a piece of gratitude, and especially those who spoke or participated in panel discussions. Anya Kolonin did a fine job of reformatting the manuscript for publication. And finally I must extend heartfelt thanks to my co-editor Pei Wang for his part in helping to pull together this book, and the scientific program of the workshop that inspired it. In my work with him over the years I have found Pei to display a combination of good sense, insight and reliability that is distressingly rare in this world (populated as it is by mere humans ... for now...).