

Artificial Knowledge

An Evolutionary Approach

Finbarr Vincent McMullin

A thesis submitted for the degree of Ph.D.

**OLLSCOIL na hÉIREANN
THE NATIONAL UNIVERSITY OF IRELAND**

University College Dublin
Department of Computer Science

October 1992

Head of Department: Professor F. Anderson
Supervisor: Dr. John Kelly

DEDICATION

For Colette: partner, wife, lover, mother of our children—but, more important than any of these, a true and enduring friend.

Contents

Abstract	1
Acknowledgements	2
1 Setting Out	4
1.1 Introduction	4
1.2 On Criticism	4
1.3 Popper's Problem	5
1.4 Making a Mind	6
1.5 Knowledge and Its Growth	6
1.6 On Darwinism	10
1.7 The Genesis of Artificial Life?	13
1.8 Conclusion	15
2 Artificial Mentality	16
2.1 Introduction	16
2.2 Three Hypotheses	17
2.3 A Personal Bias	18
2.3.1 <i>Why</i> Physicalism is (Still) Repugnant	19
2.3.2 Some Contrary Views	21
2.3.3 But: does it <i>really</i> matter?	27
2.4 Refuting Computationalism?	29
2.4.1 Searle's Chinese Room	30
2.4.2 Dualist Interactionism	34
2.4.2.1 Criticism by Dennett	35
2.4.2.2 Eccles Neurophysiological Perspective	36
2.4.2.3 Popper on AI	37
2.4.2.4 The Open Universe	38
2.4.2.5 Arguing Against Physicalism	39
2.4.2.6 Arguing For Dualism	42
2.5 Conclusion	49
3 Artificial Knowledge	51
3.1 Introduction	51
3.2 The Turing Test	53
3.2.1 Definition	53
3.2.2 Sufficiency?	53
3.2.3 Necessity?	54
3.2.4 An Informal Test	58
3.3 The Problem Situation in AI	61

3.4	On Cognitive Architecture	63
3.5	On Computational Semantics	70
3.6	On the “Engineering” of Knowledge	92
3.7	Building a Baby	96
3.8	The Growth of Knowledge	98
3.8.1	Evolutionary Epistemology	98
3.8.2	On “Random” Variation	101
3.8.3	UVSR and AI	105
3.9	Conclusion	115
4	Artificial Darwinism	116
4.1	Introduction	116
4.2	Von Neumann’s <i>Theory of Automata</i>	118
4.2.1	Background	118
4.2.2	Von Neumann’s Problem (P_v)	124
4.2.3	Alan Turing: the A_T -system	128
4.2.4	On “Universal” Construction	130
4.2.4.1	Universal the First	132
4.2.4.2	Universal the Second	132
4.2.4.3	Universal the Third	137
4.2.4.4	And So?	139
4.2.5	von Neumann’s Solution	140
4.2.5.1	The Kinematic Model	140
4.2.5.2	Some Notation	141
4.2.5.3	The Core Argument	143
4.2.5.4	A Minor Blemish(?)	147
4.2.5.5	Loose Ends(?)	153
4.2.6	Critique	161
4.2.7	The Von Neumann Myth	171
4.3	A New Problem Situation	182
4.3.1	P_a : The Problem of <i>Autonomy</i>	182
4.3.2	The Genetic Algorithm	189
4.3.2.1	Holland’s Problem (P_h)	190
4.3.2.2	P_v Again.	191
4.3.2.3	What is the Genetic Algorithm?	194
4.3.2.4	What good is the Genetic Algorithm?	194
4.3.3	Constraining the Interactions	197
4.3.4	Autopoiesis: The Organisation of the Living?	202
4.4	Conclusion	205
5	Artificial Genesis	207
5.1	Introduction	207
5.2	The Universe α_0	213
5.2.1	Outline	213
5.2.2	A Little Formality	216
5.2.2.1	The Elements	217
5.2.2.2	The Bond States	218
5.2.3	The Primitive Operators	219
5.2.3.1	Bond Modification (BM)	219

5.2.3.2	Exchange (EX)	220
5.2.4	The Emergent Operators	224
5.2.4.1	The <i>Codon String</i> Function $\pi()$	225
5.2.4.2	The <i>Binding</i> Function $\alpha()$	226
5.2.4.3	The <i>Decoding</i> Function $\gamma^{-1}()$	227
5.2.4.4	Searching for Raw Materials	229
5.2.4.5	Outline E-OP Syntax	230
5.2.4.6	The CP E-OPs	231
5.2.4.6.1	CP Operand Class 0	233
5.2.4.6.2	CP Operand Class 1	234
5.2.4.6.3	CP Operand Class 2	235
5.2.4.6.4	CP Operand Class 3	236
5.2.4.6.5	CP Operand Class 4	237
5.2.4.7	The DC E-OPs	238
5.2.4.7.1	DC Operand Class 3	240
5.2.4.7.2	DC Operand Class 4	241
5.3	“Life” in α_0 ?	241
5.4	AV0: A Realisation of α_0	247
5.4.1	The Programs	248
5.4.2	The Disk Images	248
5.4.3	The State String	249
5.4.4	Pseudo-random Number Generator	250
5.4.5	Primitive Operators	250
5.4.6	Emergent Operators	255
5.4.7	Tracking Complexes	256
5.5	Playing God	257
5.5.1	The Predictions	257
5.5.2	Parameter Values	257
5.5.3	Experiment 1: The Seed Complex	259
5.5.4	Experiment 2: The <i>Modified Seed</i> Complex	260
5.5.5	Experiment 3: The PartSR Complex	261
5.5.6	Experiment 4: The FullSR Complex	262
5.5.7	What’s going wrong?	263
5.5.8	Can We Fix it?	267
5.6	Conclusion	271

6 Rainbow’s End? 273

Bibliography 281

Abstract

I present a new analysis of the *problem situation* in Artificial Intelligence (AI), grounded in a Popperian epistemology.

I first review arguments purporting to establish that no purely “computational” system can realise *genuine* mentality. I conclude that the question is still open; but that the more pressing question is whether such a system can even exhibit intelligent *behaviour*. Attention is thus directed at the computational embodiment of *knowledge*, and its *growth*. I suggest that much of the work in this area incorporates a flawed, naïve empiricist, epistemology. I adopt Popper’s view that the growth of knowledge is possible *only* through a process of *unjustified variation and selective retention*. In particular, the innate knowledge of biological organisms has arisen by such a process, in the form of Darwinian evolution.

I review previous work on the realisation of Darwinian processes in computational systems. In particular, I present a critical reinterpretation of von Neumann’s pioneering work in this area. I conclude that no system to date has exhibited substantive growth of artificial knowledge via a process of Darwinian evolution. More importantly, I argue that this problem is deeper than is generally recognised, requiring the effective integration of *autopoiesis* with *evolvability*. To achieve this it may ultimately be necessary to realise something analogous to the *genesis of life*. I review one proposal for such a phenomenon: Holland’s so-called α -Universes. I present an *implementation* of a specific α -Universe and review the (largely negative) results of empirical tests carried out on it.

I conclude with the claim that the problem of realising the spontaneous genesis of “artificial life” is of great difficulty, but that its solution may yet prove to be an essential prerequisite for the realisation of anything deserving to be called “artificial intelligence”.

Acknowledgements

Whatever there may be of value here, it was made possible only through the most generous help from many different individuals and organisations. I am grateful to all who have helped me over the years.

I must thank my parents firstly: they gave me, and my brothers Paul, Paddy, Michael and Brian, the great gifts of freedom and opportunity, while somehow managing to ensure that we did not abuse them. I should do so well for my own children!

In my education and professional life I have been helped by very many people; but I must mention, in particular, Jim Lacy who let me become an Engineer in my own way; Michael Healy and Pat O'Donoghue who actually paid me to do it; Mark Dargan, Brendan Curtin and Brendan McMahon who helped me have fun doing it; Stephen Brennan who told me about HAHN while it was a fun place to be; Sean O'Riordan who let me sign up, and still let me go on doing things my own way; all the gang who built the original *autonomous systems* (complete with the famous *berserk mode*)—Michael Brogan, Jon Bruty, John Clancy, Nada Couzi, Arthur Cremin, Frank Devitt, Joe Fitzpatrick, James Kelly, Paul McHugh, Joe McManus, Ray Moran, Ciaran Mulloy, Eileen Noonan, Robert O'Dea, Niall O'Donnell, Bridget Walsh—to name only a very few (if we ever have another reunion, it will have to be in the Holiday Inn, Irvine); and Charles McCorkell, of DCU, for giving me yet another chance to go my own way.

In my five years in Dublin City University, I have benefited greatly from discussions with very many colleagues both in DCU and elsewhere. I am indebted to all the staff of the School of Electronic Engineering in DCU for continuing encouragement, even as their eyes glazed over; Noel Murphy and Paul Whelan have been especially tolerant, far above and beyond the call of duty. Paul McKeivitt made useful comments on earlier drafts of the Thesis, while a Research Fellow in DCU. A special thanks is due to my uncle, Ernan McMullin, for sending me the *Daedalus*, Winter 1988, special issue on AI: it was more help than he can have imagined; thanks also for resisting the temptation to tell me how far wrong I was probably going. Gabriel McDermott, then in UCD, helped greatly when I was trying to find a starting point. John Holland of the University of Michigan, Chris Langton of Los Alamos National Laboratory, Steen Rasmussen of the Technical University of Denmark, Tom Ray of the University of Delaware, and Peter Cariani of the State University of Binghamton were all most generous in responding to my correspondence. Francisco Varela of CREA, Paul Bourguine of CEMAGREF, and John Stewart of CNRS, provided very helpful advice also.

The following are people I've never met or communicated with, but, in the most literal possible sense, this Thesis could never have appeared on paper without them. So, thanks to Don Knuth for T_EX; Leslie Lamport for L^AT_EX; Oren Patashnik for B_IB_TE_X; Eberhard Mattes for emT_EX; Tomas Rokicki for dvips; David Rhead for the author-date style file; Stephen Page for double-spacing etc.; and Niel Kempson for MS-DOS B_IB_TE_X.

DCU has been generous and enlightened in its material support, in particular paying the NUI examination fee, and all fees associated with my enrollment in University College Dublin. And speaking of UCD, John Kelly has been the very model of a Ph.D. supervisor, giving me great enormous lengths of rope, and still, incredibly, managing to stop me hanging myself.

Thanks also to all my family-in-law, who long since adopted me as one of their own.

Emma Jane McMullin was born just as I was setting out on the road that would eventually lead to this tome; David John joined in a little later. They have enriched my life immeasurably. They never understood why Daddy spent so much time locked up in his "study", but they took it with extraordinarily good humour, and never ever interrupted.

Well, almost never.

I have dedicated this work with all my love to my wife, Colette McNamee: absolutely above and beyond everything and everyone else, it simply could not have happened without her.