Supermachines and Superminds

ERIC STEINHART
Department of Philosophy, William Paterson University, Wayne NJ 07470, USA; E-mail: esteinhart1@nyc.rr.com; steinharte@wpunj.edu

Abstract. If the computational theory of mind is right, then minds are realized by machines. There is an ordered complexity hierarchy of machines. Some finite machines realize finitely complex minds; some Turing machines realize potentially infinitely complex minds. There are many logically possible machines whose powers exceed the Church–Turing limit (e.g. accelerating Turing machines). Some of these supermachines realize superminds. Superminds perform cognitive supertasks. Their thoughts are formed in infinitary languages. They perceive and manipulate the infinite detail of fractal objects. They have infinitely complex bodies. Transfinite games anchor their social relations.

Key words: complexity, divine mind, supertask, infinite mind, infinite computer

1. Introduction

1.1. Complexity hierarchies

Once upon a time every living thing had some degree of cognitive power. The classical story of cognition – the story told from Plato up to but not including Descartes – said that plants, animals, humans, and stars all have souls of different kinds and different degrees of power. The classical story of cognition painted a picture of *an ordered complexity hierarchy*. The "great chain of being" (Lovejoy, 1936) rises link by link from the lifeless and mindless through rank upon rank of living thinking things – a towering city of minds. Plants had nutritive souls. Animals add perceptive and locomotive souls. Humans add rational souls with various intellectual powers. Stars were thought to engage in more perfect forms of cognition. Above them there were hierarchies of angels engaged in increasingly powerful kinds of thought. The classical story tells us that above all created minds there is an uncreated Absolute Mind, the mind of God.[1] God is omniscient; better, God's mind has that cognitive power than which no greater is possible.

I doubt the reality of the classical complexity hierarchy. I don't doubt the reality of the modern complexity hierarchy. It's a hierarchy of ever more complex logically possible objects. Some of these objects are computing machines with cognitive powers. The finite levels of this hierarchy are familiar. All known artifacts and organisms – even human animals – are likely to be *finite state machines* (digital computers with finite memory). But finite state machines are less powerful than *Turing machines* (TMs). TMs are digital computers with infinite memory. Some writers say that human animals are TMs (though that does not seem biologically plausible). The physical actuality of TMs is controversial. They are nevertheless

Minds and Machines **13:** 155–186, 2003.

physically possible. They are at least concrete objects in other physically possible universes. And while the cognitive powers of TMs are greater than those of finite machines, they are not the most powerful machines.

There are many discussions of machines far more powerful than TMs. These are the *supermachines*. Supermachines can perform computational supertasks – actually infinite calculations. It is doubtful that these are actual physical things. If they exist, they inhabit physically possible universes far more complicated than ours. I want to talk about the cognitive powers of supermachines. Most of the supermachines I discuss are infinitely complex physical things. If they are too complex to be physical in any ordinary sense, I think of them in terms of analogical extensions of physical categories into the infinite. If they are too complex to be physical in any sense, then they are purely mathematical objects. If infinitely complex machines exist, and if minds really are computational, then some of these machines realize infinitely powerful minds. These infinitely powerful minds are the *superminds*. They are minds that perform cognitive supertasks – actually infinite mental operations. I aim in what follows to sketch a theory of superminds.

1.2. The Plan of the Argument

Any attempt to talk both about minds and about the infinite is bound to be difficult. My discussion proceeds through 13 sections. Section 2 describes the modern complexity hierarchy. Sections 3 through 6 deal with the finite. Section 3 discusses finite state machines (FSMs) and defines finitary universes (finite networks of FSMs) and finitary games. Section 4 argues that cells are FSMs so that organisms are also FSMs. Section 5 deals with cognitive and moral psychology in game-theoretic terms. Section 6 discusses the hierarchy of finite minds from bacteria through humans. Sections 7 through 12 deal with infinite complexity. Section 7 describes infinite state machines (ISMs). These are the supermachines. It defines infinitary universes (infinite networks of ISMs) and infinitary games. Section 8 talks about infinitary organisms. Section 9 starts the discussion of superminds. It deals with cognitive supertasks and it analyzes super-intentionality in terms of infinitary logical languages. Section 10 looks at the psychology of superminds in terms of super-perception, super-thought, and super-will. Section 11 looks at the farther reaches of the hierarchy of superminds. Section 12 speculates on the existence and nature of Absolute Minds – minds than which no more cognitively powerful are possible. I relate these to various conceptions of God. Section 13 concludes with analogies between the classical story of cognition and the modern computational story. I argue that both stories posit transfinitely endless hierarchies of machines, organisms, and minds.

2. The Ontological Background

2.1. THE MATHEMATICAL COMPLEXITY HIERARCHY

I want a unified logical account of machines and their minds. Such an account requires the definition of an ordered complexity hierarchy in which all machines are located and in which they all have ranks. One way to obtain such a hierarchy is to posit an endless series of increasingly complex physically possible universes. However: (1) it is difficult to rank such universes relative to one another; (2) some machines are so complex they do not seem to be physical in any sense. So I need a more uniform and extensive hierarchy.

Modern mathematics – like classical philosophy – depicts reality as an ordered complexity hierarchy. The bottom level V_0 of the mathematical hierarchy contains one simple concrete object Ø. Nature abounds with cases in which a single object forms a more complex object. One proton P composes the hydrogen nucleus {P}; one flourine atom F composes the flourine molecule {F}; one amoeba cell A composes the organism {A}. So the object Ø composes the more complex object {Ø}. Since {Ø} is more complex than Ø, we put {Ø} on the next higher level V_1. Any comprehensive study of complexity has to consider all logically possible combinations of simpler objects into complexes. For every level V_n, there exists a higher level V_{n+1} that contains all the logically possible combinations of all the objects on the lower levels. So V_2 contains {{Ø}} and {Ø, {Ø}}. Complexity accumulates as the levels rise higher.

For every endless series of levels of the hierarchy, there exists a limit level above them all that contains every object on any level of the series. Since V_0, V_1, V_2, and so on form an endless series of levels, there is a limit level above them that contains all finitely complex objects (for "finitely complex", I'll just say "finitary"). If we let ω be a limit number that is greater than every finite number,[2] then the first limit level is V_ω. The level above it is $V_{\omega+1}$. $V_{\omega+1}$ is the first level that contains infinitely complex objects (for "infinitely complex", I'll just say "infinitary"). The levels of the hierarchy rise endlessly. The name for the whole hierarchy is V. V contains all the objects on all levels.[3] V is an absolutely infinite World that contains all logically possible structures. The objects in V are known as pure sets (aka pure classes). V is the iterative hierarchy of sets.

2.2. PHYSICAL THINGS IN THE MATHEMATICAL HIERARCHY

I like to think of the sets in V as the ultimate theoretical entities of deep science. One very good reading of physical science says that physical things just *are* sets (Quine, 1969, pp. 147–152, 1976, 1978, 1981, pp. 15–18). You can find the standard mathematical constructions of physical objects in the physics books. Physics books say our space-time is R^4 where R is the real numbers. Since I'm a scientific realist, and since I read scientific language literally (following Tarski), I accept the equation of space-time with R^4. The points in R^4 are sets. Some sets in V stand in spatio-temporal relations to one another. So long as space-time is geometric, and so long as geometry can be done algebraically, and so long as algebra can be done set-theoretically, space-time is ultimately just a set.

A space-time more generally is a pair (P, d) where P is a set of points and d is a distance relation. The distance relation is a set of ordered pairs of points. The ordered pairs are sets. One good way to define points is to identify them with their coordinate tuples. For a four-dimensional space-time, a point is a tuple (x, y, z, t) whose components are the positions of the point on the for coordinate axes. Given a space-time you can define functions that link its points to scalars (e.g. mass-densities), vectors (e.g., electromagnetic forces), matrices (e.g., local gravitational curvatures), or operators (e.g., that say how the values of fields change along the time axis). Unified field theory is clearly consistent with set-theoretic ontology. Smart (1972, p. 510) says: "if a unified field theory is successful, our ontology may consist simply of point-instants, classes of them, classes of classes of them, and so on, and physical objects will be definable in terms of all of these". An event is just a vector consisting of a point's coordinates plus all its field values. One set of events causes another set of events. A pool table is a space-time R^4 with mass-density field that assigns masses of 0 or 1 to points. A ball is a set of points whose mass-densities are all 1, whose space-shape is round, and whose time-shape is linear. The balls X and Y are 4D space-time worms. Their paths converge, collide, diverge. Their convergent paths are two 4D worms early-X and early-Y. They form the set {early-X, early-Y}. Their divergent paths are two 4D worms late-X and late-Y. They form the set {late-X, late-Y}. If you think of causality mechanically, then the set {early-X, late-X} causes the set {late-X, late-Y}. So some sets stand in causal relations to one another.

As the complexity of sets in V increases, all possible physical objects and structures emerge. All discrete spatio-temporal-causal systems inhabit the finite levels of V. There we find cellular automata whose space-times are finite and whose points are finite state machines. All possible finitary physical universes (with finitary organisms and minds) occupy the finite levels of V. As we pass beyond the limit level V_ω into the transfinite, infinitary physical universes (with finitary organisms and minds) appear. All dense and continuous spatio-temporal-causal systems appear in the lower infinitary levels of V. There we find cellular automata whose space-times are infinitely divided and whose points are infinite state machines. So our actual universe (whether finitary or infinitary) occurs low down in V. All possible physical universes occur in V (Rucker, 1995, pp. 200–202). The order of V stratifies these into a *hierarchy of all possible physical universes*. As the complexity of the sets in V increases, we ascend to levels whose objects cease to be physical in any ordinary sense. We must analogically extend the notions of space, time, and causality into these high levels. As physicality fades away we rise to the extremely complex levels of V. Many writers say these levels are inhabited by machines (and organisms and

Table I. Operations of a soda-dispensing FSM

Current state	Input	Coin toss	Next state	Output
S_0	Nickel	Heads	S_5	None
S_0	Nickel	Tails	S_5	None
S_0	Dime	Heads	S_0	Output soda; say "thanks"
S_0	Dime	Tails	S_0	Output soda; say "enjoy"
S_5	Nickel	Heads	S_0	Output soda; say "thanks"
S_5	Nickel	Tails	S_0	Output soda; say "enjoy"
S_5	Dime	Heads	S_0	Output soda and nickel; say "thanks"
S_5	Dime	Tails	S_0	Output soda and nickel; say "enjoy"

minds) of extraordinary power. At the upper limits of the World (of the hierarchy of sets) we find the Absolute Machines and Absolute Minds.

3. Finite Physical Complexity

3.1. Finite state machines

A *finite state machine* (FSM) has a finite set I of possible inputs, a finite set S of possible states, a finite set O of possible outputs, and a finite set R of possible values for a random variable. A *feature* of an FSM is any whole whose ultimate parts are in its states, inputs, or outputs. An FSM has a transition function F that maps its current (input, state, random) configuration onto its next state and a transition function G that maps its current (input, state, random) configuration onto its output. So the list (I, S, R, O, F, G) specifies an FSM type of which concrete FSMs are tokens. The system of possible histories of an FSM is a tree whose branches are series of configurations. An FSM is *deterministic* iff F and G are the same for all values of the random variable. An FSM is *non-deterministic* otherwise. Table 1 displays the operation of an FSM that dispenses sodas (Doyle, 1991, p. 51). The coin is a random variable.

3.2. Finite networks of finite state machines

A network of FSMs is *closed* iff every FSM in that network gets input from and gives output to an FSM in that network. A network of FSMs is *open* otherwise. A *finitary universe* is a closed finite network of FSMs. Finitary universes inhabit the finitary levels of V. If U is some finitary universe, then the FSMs in U are the points of U; the states of these FSMs carry field values (e.g. values of mass-density or force fields); their input–output relations realize the geometry of space-time; their transition functions realize causal laws. Examples of finite networks of FSMs include cellular automata like the game of life (Poundstone, 1985) and lattice gas automata (Wolf-Gladrow, 2000).

Let U be some finitely complex universe. Split U into two open subnetworks. Call one the *agent* (A) and the other the *environment* (E). The interaction between A and E is a *finitary game*. They interact like this: (1) E produces an output event which it sends to A; (2) A receives the event as input, changes its state, and produces an output event which it sends to E; (3) E receives the event as input, changes its state, and produces an output event which it sends to A. The cycle repeats until either A or E enter a halting state (in which either one fails to produce any further changes or outputs). If A and E are non-deterministic, many changes are possible at any time. So A and E are embedded in a *game tree* of branching histories. Each linear path in the tree is a possible game over U. Let W be the set of all possible games over U. A *proposition* over U is any set of possible games of U. So the collection P of propositions over U is the power set of W.

For example: A and E are both chess playing machines. Each takes a chess move as input and produces a chess move as output. States are possible legal chess boards (a board is an arrangement of chess pieces on an 8 by 8 grid). The values of the random variable encode the preferences of the player whenever the board permits the player to make many legal moves. Function F maps each (move, board, random) configuration onto a new board. Function G maps each (move, board, random) pair onto a next legal move. G is a strategy function. Each player works like this: (1) it gets an input move; (2) it applies its F to the move and its current board to produce a new board; (3) it applies its strategy function G to the move and its current board to produce an output move. A chess game is any series of interactions that (1) starts from the initial board; (2) alternates A and E moves; (3) ends with checkmate or stalemate. The game tree for these chess players is a network whose nodes are boards and whose links are legal moves. Each linear path from the initial board to an ending board is chess game. Any set of chess games is a chess proposition.

4. Finitely Complex Organisms

An organism is a living physical system.[4] All known organisms are networks of cells. The best scientific evidence implies that cells are only finitely complex (Lodish et al., 1995).[5] Every cell is an FSM. Since a finite network of FSMs is an FSM, and since an organism is a finite network of cells, every organism is an FSM.

At any instant, any cell takes in only finitely many physical quanta (photons, atoms, molecules) as input and produces only finitely many physical quanta as output. I focus on the molecular inputs and outputs to cells (though I don't mean to exclude energetic inputs or outputs). There are only finitely many distinct possible finite sets of molecular inputs or outputs to any cell at any time. So, the set of all possible inputs to some cell is I; the set of all possible outputs is O; both I and O are finite sets of molecules.

The internal states of cells are connect-the-dots networks whose dots are atoms and whose lines are chemical bonds or interactions. Every cell contains only finitely many atoms. Its atoms form molecules that fall into finitely many discrete chemical (e.g. conformal) states. They have definite thresholds of binding (they are activated or not). Molecules interact as locks and keys which either fit or do not fit. Although a key may take on infinitely many positions within a lock, when it is turned it either opens or fails to open the lock. The lock-key nature of biochemical interactions makes them discrete (Bhalla and Iyengar, 1999). So at any time any cell is in one of finitely many possible finitary cytoplasmic states. The finite set of all cytoplasmic states of some cell is S.

When a cell in some cytoplasmic state gets some input molecules (they cross its membrane), its internal *molecular regulatory network* (its MRN) determines both its next cytoplasmic state and the molecules it produces. The MRN is made of genes, RNAs, and proteins. At any time, the cell's MRN associates each (input, state) pair with some next state and with some next output. The set of possible transformations of the MRN determines the two functions F and G. The MRN is the cell's computer (Bray, 1995; Scott and Pawson, 2000; Steinhart, 2001). It is a (non-deterministic) digital machine. The MRN of any cell is only finitely complex. The functions F and G of any cell are finitary.

Since every cell has finitely many finite inputs, outputs, and states, and since it can perform only finitely many molecular transformations of (input, state) pairs into next states and next outputs, every cell instantiates some FSM type. Since thermal noise and fluid turbulence may operate as finitary random variables in cells, cells are likely to be non-deterministic FSMs. Theories of cells as FSMs are increasingly standard in biology (Yockey, 1992; McAdams and Shapiro, 1995; Cuthbertson et al., 1996; Somogy and Sniegoski, 1996; Yuh et al., 1998). Any network of FSMs is also an FSM. Organisms are networks of cells. If this biological reasoning is right, then every organism is a non-deterministic FSM. Organisms (natural or artificial) are living machines.

5. Finitely Complex Minds

5.1. Finite degrees of intelligence

A finitary universe divides into an agent A and environment E. Let A^* be all the features of A. Let E^* be all the features of E. The features in E^* are environmental contexts. Let P be all propositions over the universe composed of A and E. Let R be a set of intentional relations. Such relations include: perceives, believes, wills.

An intention function is a table with four columns. The first column holds features from A^*. The second column holds features from E^*. The third holds intentional relations from R. The fourth holds propositions from P. So any row in such a table is a list (agent feature, environmental context, intentional relation, proposition). Example: q is a retinal-neural feature of an agent; G is a normal context and b is an abnormal context; p is a proposition; so: (q, G) is a perception that p while (q, b) is a hallucination that p. If a mind sees that P we need not add that it perceives that P. Such redundancies are removable. Any (agent feature, environmental context) pair is associated with at most one (intentional relation, proposition) pair. Precisely: an *intention function* is any map f from $A^* \times E^*$ onto $R \times P$. The science of such functions is *psychosemantics*.[6] A psychosemantically sound intention function satisfies the requirements of some true psychosemantic theory. Say that A is *intelligent* in E with respect to the intentional relations R iff there is exactly one psychosemantically sound intention function f from $A^* \times E^*$ onto $R \times P$. Analysis of the intention function for some agent leads to the replacement of its features with sentences in some language of thought (Fodor, 1975). The signs in these sentences are mapped onto parts of the World by (intensional) model-theoretic semantics.

An FSM is intelligent iff it is in some environment in which it is intelligent. Any intelligent FSM has some finite degree of intelligence and is a *finite mind*. There are many ways to rank the degrees of intelligence of finite minds. One way is to rank the intelligence of finite minds in terms of their intention functions: mind x is at least as intelligent as mind y iff the intention function of x includes that of y. Another way (Dennett, 1996, ch. 4) is to divide finite minds into four genera of increasing cognitive power: (1) Darwinian minds; (2) Skinnerian minds; (3) Popperian minds; and (4) Gregorian minds. More generally: x is at least as intelligent as y iff x is able to make every inference that y can make. A third way is to rank the intelligences of minds in terms of the *expressive powers* of their languages of thought. Suppose the expressive powers of logical languages are ranked like this: (1) any propositional calculus; (2) any propositional calculus plus temporal operators; (3) any predicate calculus (with times); (4) any predicate calculus with modal operators; (5) any predicate calculus with modal and attitude operators. Minds whose languages of thought are more powerful are more intelligent. All languages of thought of finite minds are *finitary languages*. Technically: they are $L(\omega, \omega)$ languages.[7] The more powerful languages of thought of superminds are *infinitary languages*.

5.2. Value theories in terms of games

A game tree is a set of possible histories of the interactions between players A and E. Each player has a *utility table* with two columns. The first column is a list of the fixed-points (the final or limit stages of any play) in the histories of the game tree. The second is the utility of that fixed-point for the player. Example: for a chess game the fixed-points are the checkmates and stalemates. The utility table for White associates fixed-point S with +1 if White mates (wins); it associates S with -1 if White is mated (loses); it associates S with 0 if S is stalemate. The utility table for Black is the inverse of that for White.

Positive utilities are pleasurable (good). Negative utilities are painful (bad). The moral, social, and emotional modalities of minds can be defined in terms of the utilities. The agent A values (loves, desires) positive utilities and disvalues (hates, avoids) negative utilities. Players A and E cooperate insofar as they have shared utilities and compete insofar as they have opposed utilities. Players fear outcomes with negative utilities and hope for outcomes with positive utilities (these outcomes are goals). The utility tables facilitate belief-desire psychologies: the white player wants to move his pawn to the back row because he believes it will defeat the black player. The computational theory of emotions developed by Ortony et al. (1990) can be used to extend this analysis. Axelrod (1984) and Danielson (1992) develop moral theories in the context of games played by machines. Theories of virtues and vices may be developed in terms of strategies, preferences, and utilities.

A game typically involves many players. It is social. A universe can be divided into agents and environments in many ways: me and my environment; you and your environment Games enable us to think about the social experiences of players along with their those moral qualities involving other agents. Game theory is an important part of economics and politics. We use moral-political verbs in games: the rook threatens the queen; the bishop defends the king; the White player sacrifices his knight; the Black player moves her queen very aggressively. I will not presently discuss the emotional, moral, social, or political aspects of minds or superminds. Since transfinite value theory seems in its infancy (Sorensen, 1994), I leave these topics for future work.

6. The Hierarchy of Finite Minds

6.1. Finite minds from bacteria to humans

There are increasingly complex FSM networks in nature.[8] The level of complexity at which intelligence first emerges in natural FSM networks is hardly clear. Doyle (1991, p. 69) argues that "Possible minds may be as simple as a soda machine, or as complex as Lev Tolstoy." But I doubt that rocks or soda machines are intelligent. On my reading of the scientific literature, it seems most reasonable to say that intelligence emerges within the MRNs of living or almost-living things (Steinhart, 2001). Viruses (like the lambda phage) have MRNs that guide their activity after they infect cells (Ptashne, 1996). Bacteria have MRNs that exhibit minimal intelligence (Koshland, 1980). An intelligent FSM network is at least as complex as a living or almost-living thing. More generally: it is as complex as an *adaptive autonomous agent* (Maes, 1995; Russell and Norvig, 1995).

If single cells have non-zero intelligence, then increasingly complex networks of cells have increasing intelligence. Intelligence adds up. I believe the best available scientific evidence implies that natural human intelligence is only finite. No natural part or process of any human animal is infinitely complex. Every human animal contains only finitely many cells. If cells are FSMs (as I have argued), then human animals are FSMs. So human animals are finite minds. We are not as powerful as any machine with potentially or actually infinite complexity. I find no dishonor in earthly finitude – but others seem to find it offensive. On the basis of present scientific evidence, we are at most FSMs.[9] The intelligence of human animals is finite.[10] If this evidence changes, I'm happy to change my view.[11] It's better to be promoted than demoted. My discussion of the complexity hierarchy of minds does not depend on where human animals are located in that hierarchy. As the complexity of artificial FSMs (whether alive or not) reaches that of living systems, we may expect to see a parallel emergence of intelligence.

6.2. Superhuman minds and their bodies

It is surely possible for there to be organisms or artifacts far more intelligent than human animals. Whether or not human minds lie within the class of finitely complex minds, it would be absurd to argue that the cognitive power of the human animal exceeds every degree of complexity. There is some degree of cognitive and computational complexity that serves as an upper bound for minds realized by human bodies. Doyle (1991, p. 41–44) argues that "Human beings and Turing-equivalent machines need not exhaust the range of entities in which to realize psychologies". Moravec (1988, 2000) argues for the possible existence of machines far more intelligent than human animals. He speculates (1988, p. 74) that machines may evolve intellects 10^{30} times more powerful than human minds. While Moravec's claims may be overly dramatic, the cognitive power of human flesh is hardly the upper bound for computational or cognitive complexity.

It is likely that a superhuman organism would have a superhuman body. One sort of superhuman body extrapolates the growth pattern of the human flesh. This growth pattern is recursive. It is algorithmic. The human body consists of a central stick S_0; on this central stick are mounted four other swiveling sticks at level S_1 (arms and legs); on each of the sticks in S_1 there are five swiveling sticks at level S_2 (fingers and toes). It is apparent that we could have fingers on our fingers. Our limbs could be further articulated. Moravec (1988, p. 102–108, 2000, p. 150–154) describes *robot bushes* that have many levels of branching limbs. I prefer to think of them as *animal bushes* (super-intelligent versions of basket starfish bodies). Their ultimate fingers (at level S_n for some large n) are sufficiently small and precisely coordinated to perform highly detailed physical tasks. They are equipped with tiny detectors so that each finger is a sense organ.

7. Transfinite Physical Complexity

7.1. Countably infinite state machines

An *infinite state machine* (ISM) has a set I of possible inputs, an set S of possible states, a set of random values R, and a set O of possible outputs. Its input, random, and output sets may be finite or infinite. Its set of possible states is infinite.[12] An ISM has a function F that maps its current (input, state, random) triple onto its next state and a function G that maps its current (input, state, random) triple onto its next output. So the list (I, S, R, O, F, G) specifies an ISM type of which concrete ISMs are tokens.

A good introduction to ISMs is an *infinitary game* in which two *infinite state machines* compete by forming an infinitely long series of digits (Hamilton, 1982, p. 189). The two players are the agent (A) and environment (E). Their inputs and ouputs are the digits 0 to 9. Their states are series of digits of any finite length. Since there are ω many finite series of digits, each player has ω many states. The next state function F maps each (input digit d, series s) pair onto the series made by appending d to the end of s. F is realized by means of a memory register M of length ω. The output function G maps each (digit, series) pair onto an output digit. G encodes the machine's strategy. Each player operates as follows: (1) it gets an input d; (2) it looks up (M, d) in its strategy table to produce an output digit; (3) it appends d to the end of its memory M to make the series Md; (4) it sends its output to the other player. Game trees, possible histories, and propositions are defined for these players by analogy with the finite case.

Each appending or look-up operation involves only finitely many operations. However: there is no finite upper bound to the length of these operations. For any n, each ISM has to look up a series of length n in its strategy table. If we want each machine to complete its operation in one unit of time (one clock tick), then we need to make each machine *accelerate*. Acceleration involves the *Zeno compression* of ω acts into continuous but finitely extended time. Let [0,1] be some continuous unit period of time (e.g., 1 s). The first act is done in 1/2 s. The next act is done in 1/4 s so that two acts have been done in 3/4 s. Generally: the n-th act is done in $1/2^n$ s so that n acts have been completed in $(2^n - 1)/2^n$ s. Example: if agent A finds digit d_n in $1/2^n$ s, then within 1 s A can find the entry for any series in its strategy table. Although these ISMs accelerate, they never take limits. They are only *countably* complex. Each player applies its functions F and G in one time unit (one clock tick). It makes its move in one clock tick. Each move generates a digit. If the n-th clock tick is an even number, then A produces digit d_n; if it is odd, then E produces digit d_n. The result in ω clock ticks is an infinite series of digits (it is a real number $0.d_1d_2d_3d_4...$).[13] If you want the game to be over in finite time, then just make the clock accelerate.

7.2. UNCOUNTABLY INFINITE STATE MACHINES

A *supertask* is an actually infinite series of operations done in some finite region of space-time. Many supertasks have consistent recursive definitions and converge to well-defined objects at transfinite limits (Earman and Norton, 1996; Koetsier and Allis, 1996). An uncountably complex ISM is one that performs computational supertasks. It generates limits (Steinhart, 2002). Although there are many different kinds of uncountably complex ISMs,[14] I focus here on *multi-tape accelerating Turing machines* (MATMs).

An MATM is a classical multi-tape Turing machine that can do supertasks (Copeland, 1998a; Davies, 2000; Hamkins and Lewis, 2000). Since it can perform computational supertasks, an MATM is more powerful than a classical multi-tape Turing machine (Copeland, 1998b). It can produce an infinite series of configurations and take the limit of that series. An MATM has: (1) some input registers; (2) some internal state registers; and (3) several output registers. Although I won't mention the random variables, I'm including non-deterministic MATMs. Each register is like the tape of a classical Turing machine. It has ω many one bit variables (values 0 or 1). A register configuration is the list of the values of its variables plus the position of its read/write head. The action of an MATM is controlled by an internal FSM just like a classical Turing machine. From one moment to the next, the FSM reads from and writes to the various registers and changes its own internal state. An MATM's configuration is the list of its register configurations plus the state of its FSM. So an MATM has uncountably infinitely many configurations.

An MATM accelerates to limits. It starts at time t=0 in its initial configuration S_0. It produces its first configuration S_1 at time 1/2. It makes S_2 at time 3/4. Generally: it makes S_n at time $(2^n - 1)/2^n$. At the limit time t=1, an MATM is in its limit configuration S_ω. The limit configuration is the list of the limit configurations of its registers. The limit of any register is some function of the entire previous infinite series of configurations of that register. There are many ways to define limits for registers. I describe three. Let R_n be the configuration of register R at time $(2^n - 1)/2^n$. We may let R_ω be the infinite disjunction of all the R_n. Formally: $R_\omega = \bigvee_{n=1}^{\omega} R_n$. We may let R_ω be the infinite conjunction of all the R_n. Formally: $R_\omega = \bigwedge_{n=1}^{\omega} R_n$. If we think of R_n as a real number between 0 and 1, then we may let R_ω be the calculus $\varepsilon - \delta$ limit of all the R_n. Formally: $R_\omega = \lim_{n \to \omega} R_n$.

An input configuration for an MATM is the list of its input register configurations. The input set I for an MATM is the set of its possible input configurations. Analogously, the state set S is the set of the MATM's possible internal state configurations and its output set O is the set of its possible output configurations. An MATM's accelerating operations realize functions F and G. If i is an input in I and s is a state in S, then F maps each (i, s) pair onto a next state in S, and G maps each (i, s) onto an output o in O. Each application of F and G is a supertask. The operations of an MATM apply F and G simultaneously. For any countable ordinal

κ, an MATM can perform κ operations in one time unit (one clock tick).[15] So a group of MATMs can be synchronized by common clock.

7.3. Infinite networks of infinite state machines

An infinite network of infinite state machines (ISMs) is the basis for an infinitary physical universe. The network is a space–time whose point-instants are the configurations of its ISMs. The input-output connections among the ISMs determine the geometry of space-time. Their internal state registers carry field values. Their transition functions realize the causal laws. Example: if the output registers of MATMs are linked to the input registers of other MATMs, the result is a network of interacting infinite state machines. Any network of ISMs is a set in the hierarchy V. Infinitary universes inhabit the infinitary levels of V. They may contain (copies of) finitary universes as subnetworks.

An infinitary universe U is a closed infinite network of ISMs. Let U be some infinitary universe. Split U into two open subnetworks. Call one the *agent* (A) and the other the *environment* (E). The interaction between A and E is an *infinitary game*. They interact infinitely many times. The agent (A) and the environment (E) interact by alternating moves. Each move happens in one clock tick. Each move is a supertask in which the player accepts an input, changes its state, and produces an output. The output of each player is the next input to the other player. The players generate an infinite series of moves. For infinitary games, any infinite series of moves has a limit. For example: the limit of an infinite series of digit choices is a real number. The configuration of each player at each limit ordinal is the limit of that player's previous configurations. Game trees, possible histories, and propositions are defined by analogy with the finite case.

An infinitely complex universe (an infinite network of ISMs) has an infinitely divided and infinitely extended space and time. It also has infinitely divided causality or materiality (i.e., fields whose intensity values are rational or real numbers). If points (ISMs) carry the instantaneous values of fields, then infinitely divided space-time suffices for the definition of things with infinitely precise detail. Such things are familiar as *fractals* (Mandelbrot, 1978). Many supertasks produce fractal things as their end products. I am particularly interested in transfinite machines that can take the limits of *Lindenmayer systems* (L-systems; Prusinkiewicz and Lindenmayer, 1990). L-systems are recursive rules for making drawings or shapes or structures. Example: (1) start with a Y; (2) replace each branch of the Y with a smaller Y. Figure 1 shows several iterations of this L-system. One standard way to make L-systems is by means of *turtle geometry* (Abelson and diSessa, 1984). A turtle is a mathematical creature that lives in Euclidean space (e.g. on the flat Euclidean plane). Although classical TM heads move on 1D tapes, you can think of turtles as accelerating TM heads that move with infinite precision on 2D planes or in 3D volumes. A TM turtle that moves in many dimensions and that accelerates can make infinitely detailed black and white drawings or carve infinitely detailed

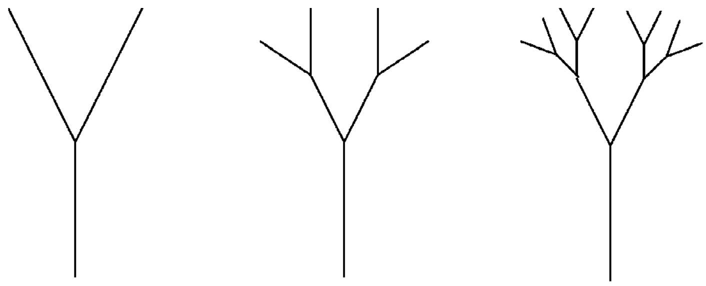

Figure 1. Three iterations of a simple L-system.

sculptures. The limits of L-systems are typically fractals (Koch curves and islands; dragon curves, etc.). Since they are infinitely detailed space-time structures, the limits of L-systems are appropriate objects of perception, thought, and action for superminds.

7.4. AN INFINITARY GAME THAT MAKES A FRACTAL OBJECT

Let U be an infinite closed network of ISMs split into open subnets A and E. The topology of E is a continuous flat square whose sides have length 1. Formally this square is denoted $[0,1]^2$. You can think of $[0,1]^2$ as a square piece of infinitely divisible graph paper. Each point in E has coordinates (x, y) where x and y vary between 0 and 1. Each ISM in E supports a 1-bit color value so that E supports a black-and-white color field. If we identify each ISM in E with its (x, y) location in E, then any color field over E is some function from $[0,1]^2$ to $\{0,1\}$. You can think of each function from $[0,1]^2$ to $\{0,1\}$ as a black-and-white drawing. For any such function f, the point (x, y) in $[0,1]^2$ is black if f(x,y) is 1 and white otherwise. The agent A has its eye on one side of E and its hand on the other side. On the eye side, the input registers of A are arranged in a flat square like an infinite retina. Each retinal cell of A sees one color bit on E. On the hand side, the output registers of A are arranged in a flat square like an infinite grid of fingertips. Each fingertip of A sends 1 (mark with black ink) or 0 (erase with white ink) to its corresponding point in E. The remaining ISMs in A link A's retina to its fingertips.

Let S be the set of all functions from $[0,1]^2$ to $\{0,1\}$. So S is the set of all black and white drawings on the paper $[0,1]^2$. Let S_0 be the drawing made by dividing a blank piece of paper in half horizontally and vertically and writing | in the top row of the paper. If S_n is any drawing, let S_{n+1} be the drawing produced by copying S_n and then dividing the bottom row of S_n in half and the right column of S_n in half and writing N| in the top half of the blank part of the paper. Any two drawings are merged by taking their point-by-point disjunction.[16] Any series of drawings has a disjunction. Let S_ω be the disjunction of S_n for all finite n. Suppose that S_0 is made in 1/2 unit of time and that if S_n was made in any fractional unit of time, then S_{n+1} is made twice as fast. This is a consistent supertask that converges in 1 unit of time

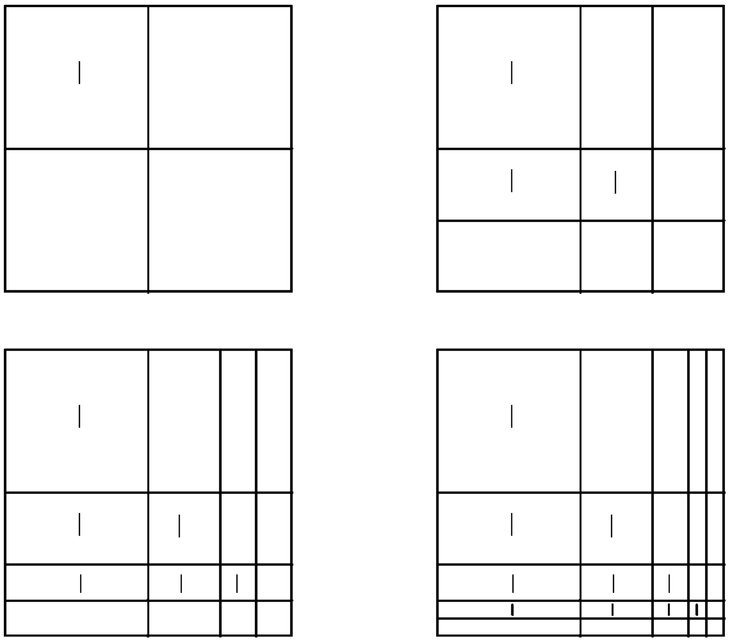

Figure 2. A few iterations towards the Hilbert Paper.

to a limit drawing S_ω such that | is on S_ω and if N is on S_ω then N| is on S_ω. S_ω is the *Hilbert Paper*. A simple infinitary game between A and E makes S_ω. A's n-th move consists of sending outputs to E that make drawing S_n. In this game, E cooperates. E's *n*-th move consists of changing the states of its machines so that S_n is made. The limit of the infinite series of moves is the drawing S_ω that contains every finite stroke-series. Figure 2 shows a few iterations of this game.

8. Transfinitely Complex Organisms

Since the forms of finitary organisms resemble the finite iterations of L-systems, I suppose that the forms of infinitary organisms resemble the limits of L-systems. For example: just as the forms of finite plants are the finite iterations of L-systems, so the forms of infinite plants are the limits of L-systems. Since infinitely complex universes contain the limits of all sorts of L-systems, I suppose they contain super-organisms.

Transfinite organisms in super-universes are infinitely complex networks of cells. These cells may be finitely or infinitely complex. It's clearly more interesting if the cells are infinite state machines (ISMs). The growth pattern of a super-organism generates an infinitely complex fractal system of cells (e.g., the limit of an L-system). A super-organism grows through transfinitely many stages. It starts

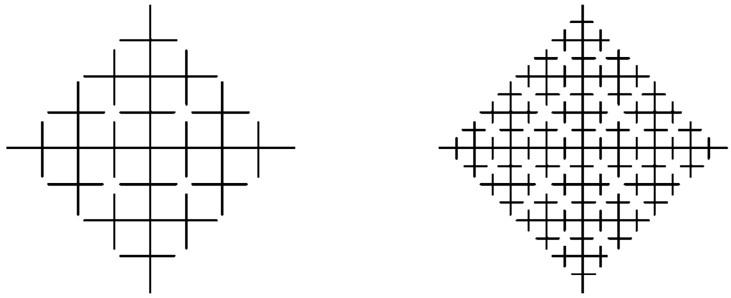

Cross S_2 with 2 nestings. Cross S_3 with 3 nestings.

Figure 3. The growth pattern of a fractal super-organism.

with some initial stage S_0; it develops through successor stages S_n for every finite n< ω; it ends with a limit stage S_ω.

For example: consider a super-organism whose growth pattern is an L-system whose limit is an infinitely self-nested cross. Figure 3 shows a few iterations of this growth pattern. The initial stage S_0 of the super-organism is defined like this: the super-organism starts as a red cell in the center of $[0,1]^2$; in 1/2 s, it grows four arms of length 1/2 minus the endpoints; each arm is a line of black cells except for a blue cell at the midpoint of the arm. For each finite n $>$ 0, the successor stage S_n of the super-organism is defined like this: in $1/2^{n+1}$ seconds, each blue cell in a line of length $1/2^n$ changes to a black cell and grows 2 arms of length $1/2^{n+1}$ perpendicular to the line that contains it; each arm is a line of black cells except for a blue cell at the midpoint of the arm. For the limit ordinal ω, the limit stage S_ω of the super-organism is the union of all the cells in all S_n for n finite. The super-organism is a fractal – an infinitely deeply self-nested cross.

It is possible to generalize Moravec's bush robots or bush animals to infinitely complex superbushes. A *superbush* is an animal bush with infinitely branching limbs. A superbush is described by a transfinite program: (1) it has an initial limb S_0; (2) for every finite n, each endpoint of each limb at level S_n branches into two shorter and thinner limbs at level S_{n+1}; (3) the whole superbush S_ω is the union of S_n for all finite n. Each limb is half as long and half as thick as the limb on which it is mounted. A superbush agent has as many fingertips as there are real numbers. It can move its limbs and fingertips to manipulate and produce infinitely detailed concrete objects in its environment. It can perceive, think about, and willfully make the fractal limits of L-systems.

Superbushes are specific instances of supergraph animals. A *supergraph animal* (just a "supergraph" for short) is a living connect-the-dots structure with infinitely many dots and connections. The dots are joints and the connections are limbs. Finite graph animals are depicted in the artificial life program called "Framsticks" (Mandik, 2002). Supergraph animals are generalizations of Framsticks to the transfinite.

9. Transfinitely Complex Minds

9.1. Cognitive supertasks

For every machine, there is a supermachine able to everything done by the machine. More: there is a supermachine able to do everything done by the whole class of machines (e.g. a super-Turing machine that can do everything any Turing machine can do). So if minds are intelligent machines, and if supermachines exist, then superminds exist.

A supermind has infinitely great cognitive power. It is a concrete object that can perform *cognitive supertasks*. As an example of a cognitive supertask, consider reading a book with infinitely many pages (Borges, 1964, p. 58). The book is organized like this: its front and back covers have some thickness T; if p is any page after the front cover, then the thickness of the page after p is half the thickness of p; if p is any page before the back cover, then the thickness of the page before p is half the thickness of p. If the front cover is at point 0 and the back cover is a point 1, then the book is a pair of series of pages ({0, 1/4, 3/8, 7/16,...}, {...9/16, 5/8, 3/4, 1}). Each page has finitely many words. Zeus reads the book like this: at time t=0, he reads the first page; at time $t = 1/2$, he reads the last page; at $t = 3/4$ he reads the page after the first page; at $t = 7/8$ he reads the page before the last page; he goes on like this for 1 min. At time $t = 1$, he has read the whole book. Zeus has read infinitely many words in 1 min – a cognitive supertask. If Zeus is to remember what he has read, he needs an infinite memory.

A supermind can play infinite epistemic games. Some propositions require infinitely much data to verify them (e.g., “Beyond every star there is another”). Science is often thought of as making a series of increasingly accurate theories that progress towards some perfectly accurate limit theory. The limit theory may result from an epistemic supertask. Example: a super-scientist A plays an infinitary game with an environment E that obeys fixed laws (Juhl, 1995): (1) E gives some data to A; (2) A tries to give E some theory of E that explains the data; (3) E tries to present A with some data that the theory does not explain; (4) A tries to provide a better theory. The cycle of data and theory goes to the limit. If there is some limit theory that explains all the data, A wins; if not, E wins. Brams (1983) investigates whether we are actually playing such games with infinitary minds.

Superminds have infinitely many perceptions, beliefs, volitions, and so on. More importantly: their cognitive states have infinitely rich content. They have (1) *super-perception;* (2) *super-thought;* and (3) *super-will*. They perceive and represent their environments to infinite depths of detail. They perform infinitely many inferences on infinitely detailed representations. They have infinitely deep and subtle plans. A supermind changes its environment to some infinite depths of detail. For example: a supermind can super-perceive all the infinitely many colors of superlight; it can super-compute the structure of an infinitely detailed arrangement of supercolors; it can super-manipulate the field-values at point-instants with infinite precision. So can paint and see an infinitely detailed colored picture whose beauty infinitely exceeds every finite degree of beauty. Superminds have infinitely rich psychological (intellectual, emotional, moral) lives. They have infinitely complex characters with virtues and vices. Communities of supergraphs instantiate all sorts of infinitely complex social and political theories.

9.2. Transfinite intentionality

A supermind is an intelligent ISM. Psychosemantics for superminds parallel those for finitary minds. Let U be an infinitary universe divided into agent A and environment E. So A^* is the infinite set of all features of A and E^* is the infinite set of all features of E. Let R include at least all finitary intentional relations (perceives, believes, wills, etc.). The set P of propositions over U is infinite. An intention function f for the agent A is any psychosemantically sound map f from $A^* \times E^*$ onto $R \times P$. Intelligence is defined as in the finite case. The salient difference is that f is infinitary. Hence any supermind is more intelligent than every finite mind. There are *transfinite degrees of intelligence* that increase along with the complexities of supermachines. Any supermind that lives in an infinitary universe (such as a continuously divided spatio-temporal-causal network of ISMs) is easily thought of in physical terms. It is a concrete physical whole in the set-theoretic hierarchy V. As the complexity of the objects in V increases beyond the continuous, these physical interpretations become less visualizable.

An analysis of the intention function for any supermind leads to the systematic replacement of its features with sentences in some language of thought (Fodor, 1975). The signs in these sentences are mapped onto parts of the World by (intensional) model-theoretic semantics. While the mental sentences for finite minds are formulae in finitary languages of thought, the mental sentences for superminds are formulae in *infinitary languages of thought*. They are *super-sentences*. Sentences in infinitary languages may have infinite series of quantifiers, infinite conjunctions or disjunctions, and relations that have infinitely many places. If the intentionality of a mind is measured by the complexity of its language of thought, then any supermind has super-intentionality.

Super-sentences are stored in super-memories able to hold infinitely detailed grammatical networks of signs. They are manipulated (e.g., arranged into infinitely long arguments) by super-computations. Example: let G(x) mean that number x is the sum of two primes. At time 1/2, the supermind Achilles forms the sentence S_1 = G(2) and knows that it is true. At time 3/4, Achilles forms S_2 = (G(2) & G(4)) and knows that it is true. At the *n*-th moment in a Zeno compression, Achilles forms the sentence S_n = (G(2) &...G(2n)) and knows that it is true. At the limit time 1, Achilles forms the limit of all these sentences. This is the super-sentence S_ω. It is an infinite conjunction (G(2) & G(4) & G(8) & G(10) & ...). At the limit time 1, Achilles knows whether this super-sentence is true or false. Since S_ω is equivalent to (∀n)(if n is even, then G(n)), it isn't a very rich super-sentence. Richer super-sentences come with infinite sequences of quantifiers, infinitely detailed internal grammatical structure, and relations with infinitely many places. Example: $(\exists x_1, x_2, \ldots x_n, \ldots)((x_1 < x_2)\&(x_2 < x_3)\&\ldots(x_n < x_{n+1})\&\ldots)$.

There are endlessly many infinite numbers greater than ω. Example: ω_1 is an uncountably infinite number that towers over ω. For any numbers κ and λ greater than ω, there are infinitary $L(\kappa, \lambda)$ languages. Although we need not go into technical details here, the expressive power of any language like $L(\omega_1, \omega)$ and $L(\omega_1, \omega_1)$ is far greater than that of $L(\omega,\omega)$. Karp (1964) describes objects that cannot be fully described by finitary sentences but that can be fully described by super-sentences.[17] Examples include: the natural numbers under the successor operation; the class of well-ordered systems.

If the complexity of the content of a mental representation corresponds to that of some mentalese sentence, then a supermind can truly and completely represent any finitary object in one synchronic thought. You and I can see the entire part-whole structure of small finitary patterns without counting or performing a series of analytic acts. A supermind can see all the detail of any finitary structure in one synchronic vision. If a "gods-eye" view is an infinitary vision, then we can rank such views using infinitary languages.

10. Some Cognitive Powers of Superminds

10.1. SUPER-PERCEPTION

A finite mind has only finitely many input states. The retina of a finite mind is finitely divided and finitely extended grid of receptors for atomic sense data. A supermind has infinitely many input states. The retina of a supermind is an infinitely divided grid of receptors. Between any two receptors there is always another. So a supermind can see (or hear, or feel) infinite detail. Since a supermind's receptors are ISMs, these receptors may register light intensity or color with infinite precision. As we see only finitely many (three) dimensions, so a supermind may see infinite dimensional spaces.

A mind with super-perception is able to perceive infinite detail in finite time. It does this by performing (completing) perceptual supertasks. A supermind is able to perceive the limits of L-systems (it can see, in one limit act, all the infinite detail of any Koch curve or fractal plant). I illustrate super-perception with Royce's perfectly accurate map of England within England (Royce, 1927, pp. 506–507). A part of the map perfectly accurately depicts the part of England that contains the map. It's an infinitely nested series of pictures. Royce's self-nested map of England is an infinitely rich physical fractal object. Royce's definition is recursive: the structure of England is endlessly repeated within England. Royce's map is easily formalized as a connect-the-dots diagram (a graph). The map at level 0 is denoted M_0. Suppose M_0 is a square containing a cross (England's not what it used to be). For any finite n, the next map M_{n+1} is M_n plus a copy of M_0 in its smallest lower right-hand

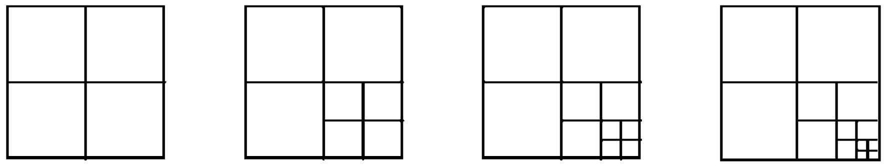

Figure 4. The first four iterations of a Roycean self-nested map.

square. The limit map M_ω is infinitely rich. It is the union of all M_n for n finite. Royce's perfect map of England is a map of the form M_ω. Figure 4 shows the first four iterations of a trivially simple version of Royce's self-nested map.

A mind with powers of super-perception successively perceives (sees) the members of an infinite series of increasingly rich perceptual objects and converges in the limit to the limit object of that series. Suppose Athena is a super-perceiver who plans to completely see the Royce map. At time 1/2, Athena sees M_0. At time 3/4, Athena sees M_1. As time goes from 0 to 1, Athena brings each map M_n into focus. You may think of her eyes as microscopes able to focus in on ever-greater detail at ever-higher levels of resolution. At each successive time 1/2, 3/4, and so on, she "zooms-in" on the next level of detail of the map while keeping the size of her visual field the same. Athena forms an internal mental representation of map M_n at the *n*-th visual act. For all $n < \omega$, she sees map M_n at time $(2^n - 1)/2^n$. As time passes, Athena's accelerating series of visual acts converges to the limit vision. A time 1, Athena sees the limit map M_ω. So at time 1, she completely and truly and in that instant sees the infinitely complex map. Athena has an internal infinitely complex mental representation of map M_ω at time 1. If Athena is an MATM, then she is able perform the whole infinite operation in any non-zero interval. So Athena is able to perceive countably many countably complex objects in any finite time.

10.2. SUPER-THOUGHT

A supermind needs infinitely many internal states in order to store infinitely detailed mental representations of infinitary objects. It needs a *super-memory* to record its infinitely detailed perceptions and ideas (e.g., its reading of the Borges book or vision of the Royce map). A minimal super-memory is an infinitary register with ω many bits. Such a register can be squeezed into a finite continuous space by Zeno compression. The memory of an infinite von Neumann machine is an array of ω many infinitary registers. A supermind can store a complete true description of itself in its super-memory. A supermind (and only a supermind) can have perfect self-awareness and self-knowledge.[18]

A supermind can perform infinitely many computations on the mental representations in its super-memory. Some of these transfinite computations are *super-inferences*. Super-inference involves the solution of problems not solvable by finite means. For example: a super-mind running on an MATM is able to compute the halting function for all classical TMs; it is able to fill out a table whose *i*-th row and *j*-th column contains 0 if the *i*-th classical TM halts on input *j*. It can calculate an infinite look-up table that lists the recursive total functions from the positive integers to {0,1} and it can do anti-diagonalization on that table to obtain a non-recursive function (Shagrir, 1997, pp. 328–329). It can perform infinitely long deductions. It can find the truth-values of arithmetical propositions by sheer enumeration. For example: it can decide the truth-value of Goldbach's conjecture just by exhaustive search. For any non-constructive proof, a sufficiently powerful supermind ought to be able to find examples of the objects or processes asserted by the proof. For example: the Banach-Tarski theorem asserts that "a closed three-dimensional solid ball may be split into finitely many pieces which can be rearranged without distortion to form two solid balls of the same size as the original" (Hamilton, 1982, p. 186). A sufficiently powerful supermind ought to be able to figure out a way to cut a solid ball that satisfies the theorem, so that it can take the ball apart and put its pieces back together into two balls.

A supermind can do anything that can be done by any less powerful machine. It can simulate (imagine) all possible finite computations. So it can exactly simulate any finitely complex universe. Its simulation is not an approximation; it is an exact reproduction. So a supermind can contain a copy of any finitary universe (any closed network of FSMs). If there are persons in some finitary universe, the supermind can exactly replicate their lives (Moravec, 1988, pp. 178–179). If exactness is needed for being "real" (rather than being fake or inauthentic), then the copy of a universe in a supermind is as real as the original. A supermind that simulates within itself every possible finitary universe will generate within itself a data structure like Leibniz's *Palace of the Fates*. Leibniz says that all possible worlds are present in the mind of God as ideas. He describes the system of ideas in a conversation between the mythical Athena and Theodorus:

> [Athena said to Theodorus]: The pyramid you see here is the Palace of the Fates ... It contains representations not only of that which actually happens but also of all that which is possible. At the beginning of time, Zeus surveyed all these possibilities and sorted them into alternative possible worlds; ... I have only to speak, and we shall see a whole world that my father Zeus might have produced ... These worlds are all here as ideas in the divine mind of Zeus. ... After saying this, the goddess led Theodorus into one of the halls of the Palace of the Fates; when he was inside, it was no longer a hall, it was a world presented to his mind as if he were watching a dramatic play happening on a stage. ... Looking around in the Palace, Theodorus saw a collection of documents bound up into a book. ... The goddess told him: this is the history of the world which we are now visiting; you are looking at the book of its Fate. ... put your finger on any line in this book, and you will see represented actually in all its detail every fact which is described by that line of divine writing. ... Athena took Theodorus into other halls in the Palace, hall after hall, world after world, in which they saw

endlessly many ways things might have been, or might yet be. (Leibniz, 1996, Sections 414–417)

10.3. SUPER-WILL

One analysis of willing says that agent A in environment E wills that p iff A works in E to make p true. If willing is methodical rather than accidental, then an agent A wills all the outcomes of its plans or strategies. Example: if a chess player plays with a fixed strategy, then she wills the conjunction of all the outcomes of that strategy. A player who plays according to a fixed strategy wills all the wins and all the losses that it entails. An agent has a more perfect will that p iff that agent's plan has a greater chance of making p be true. A plan is an algorithm. An algorithm has a set of final states (halting states or fixed-point states to which it converges in the limit). Each execution of an algorithm by A in E to a final state determines a history of the universe in which A and E reside. So the execution of an algorithm S determines a proposition (the set of all the possible histories in which A runs S to some final state). So A in E wills that p iff there is some algorithm S such that A runs S and p is the set of all histories containing final states of S. If T is a theory, then agent A wills that T iff it runs an algorithm whose finals states are all models of T.

A finite mind has a finite will. It can only will the halting states of finite algorithms. It is only able to will the fixed-points to which finite algorithms converge. Example: a finite mind can make finite theories true by building their models. A supermind has *super-will*. It can at least will the halting states of Turing machines. An intelligent MATM or other supermachine that takes limits can will the final states of supertasks. It can will the states to which those supertasks converge. A supermind that can take limits can will the existence of (i.e., it can make or construct) all kinds of infinitary fractal objects. It can draw Royce's perfect map of England or construct the limits of L-systems in finite time.

An infinitary organism (e.g., a superbush) wills the existence of an infinitary object by completely running an algorithm that converges to that object in the limit. For example: a superbush wills the existence of the *Sierpinski carpet* (Mandelbrot, 1978, pp. 166–167) by completely running the following algorithm: (1) the initial situation S_0 consists of a flat square continuous material surface on which the superbush is resting its finger tips (exactly one finger tip rests on each point on the square); (2) for any n, the superbush transforms S_n into the successor situation S_{n+1} by dividing every square in S_n into a 3 by 3 grid like a tic-tac-toe board and then punching out the middle square to make a hole (punched-out points are not in S_{n+1}); (3) the limit situation S_ω is the intersection of all S_n for n finite. The limit situation is the Sierpinski carpet. Any project guided by a transfinite algorithm is a *super-project*. Making the Sierpinski carpet is a super-project. A more physicalistic super-project (it involves digging) is making *the Lakes of Wada*. It takes place in a universe in which space, time, and matter are all continuous. The Lakes of Wada is a simple supertask. Koetsier and Allis (1997) describe it like this:

> Imagine an island in the ocean on which there are two lakes. We carry out a project on the island during which the three different kinds of water are kept separated. On the first day we construct dead end canals starting from the ocean and from the two lakes in such a way that each point of dry land is at a distance of less than 1 kilometer from the sea and from the water of both lakes. The three kinds of water remain separated. On the second half day we extend the canals in such a way that each dry point is at most 1/2 kilometer from the three kinds of water. On the following quarter day we continue until each dry point is less than 1/4 kilometer from the three kinds of water. After two days of work the island will have been turned into a curve that has the amazing property that it separates the three kinds of water – they do not mix, while at the same time each point of the curve can be approached arbitrarily close from each of the three kinds of water. (p. 293)

A supermind can will theories whose models are infinitary objects. A sufficiently powerful supermind in a sufficiently rich environment can will the truth of Zermelo-Fraenkel-Choice (ZFC) set theory by building a concrete model of ZFC. An agent able to do this would be able to empirically verify the continuum hypothesis. The environment for such an agent would have to be at least as complex as the least stage of the hierarchy V that contains a model of ZFC. It would have to be at least as complex as some level of that hierarchy whose index is an inaccessible cardinal (Hamilton, 1982, pp. 230–233). Such agents and their environments are physical in analogically extended senses (e.g., the space-time in which the agent works is even more divisible than the continuum). A sufficiently powerful supermind can make sentences in infinitary languages be true.

11. The Hierarchy of Transfinite Minds

11.1. Super-physical universes

A finitary universe is any closed finitary network of FSMs. A finitary universe splits into an agent A and environment E. The interactions between A and E form a finitary game. If A is intelligent, then A is a finitary mind. Finitary minds play finitary games. Finitary games generally have ordinary physical interpretations. Example: a chess board is a familiar 2D space and chess moves form a familiar 1D time. The causal laws of chess are the rules for changes of positions (motions) of the pieces (things).

An infinitary universes is any closed infinitary network of ISMs. An infinitary universe splits into an agent A and an environment E. The interactions between A and E form an infinitary game. If A is intelligent, then A is a supermind. Superminds play infinitary games. Some games played by superminds have ordinary physical interpretations. They are played in continuous space-times with infinitely detailed material structures. Causal laws are like those for finitary games. Example: if A and E play a game by choosing an infinite series of digits, we may visualize them as writing down digits one after another faster and faster and smaller and smaller on a paper tape.

It is possible for superminds to play games in universes that are too complex to be physical in any ordinary sense. These universes are extremely complex objects in the set-theoretic hierarchy V. They are super-physical systems. A super-physical system has generalized space, generalized time, and generalized causality. An ordinary space is a set of points with a metric (a distance function). More generally: a space is any set with a metric. A *generalized space* is any collection that contains every object needed for playing any possible history of the game. An ordinary time is a series of instants. More generally: a time is some number line (e.g., the natural, rational, or real number lines). A *generalized time* is any linearly ordered series. Generalized times include the whole ordinal line and the surreal number line (Conway, 2001). Kitcher (1984, pp. 146–147) talks about *supertime* as a highly superdenumerable "medium analogous to time, but far richer than time". Generalized motions are curves in generalized space-times. Generalized causal laws are just the rules of any transfinite game. These analogical extensions of ordinary physical concepts are not likely to extend to arbitrarily complex wholes in V. Superminds may play games in environments so complex they are purely mathematical. This is not a shift from the "concrete" to the "abstract". It is a shift from the simpler to the more complex.

11.2. Some superminds in super-physical universes

Most infinite games are mathematical devices for the construction of infinite series of situations. Gale and Stewart (1953) is the classical source. Freiling (1984) gives many examples. The universes for these games are sets. They are infinitely detailed lifeworlds in which superminds interact. Since these games are presented in the rather dry terminology of set-theory, they may seem psychologically shallow. This shallowness may be a false impression. Since we are finite, we have little epistemic or emotional access to infinitely complex systems. Consequently: we are not likely to find much drama in the careful selection of sets of real numbers. But an infinite agent, one able to experience all the properties of these sets as we experience our own phenomenal world, might live a psychologically rich life in a game world that seems shallow to us. For such an agent, sets of real numbers might be extremely beautiful or ugly; the selection of a beautiful set might be an extremely pleasurable act; the selection of an ugly set very painful. The selection of a certain set according to a certain strategy is an exercise of skill. It is possible to think of these shallow numerical or set-theoretical games as the skeletal versions of infinitely complex games in which the players live psychologically rich lives.

Jech (1984) defines a game he calls a "cut and choose" game. The game is interesting because – unlike games mentioned before – the players perform different operations. The White player cuts a set into two disjoint subsets (these are two portions of the original set); the Black player chooses one portion; the White player then cuts the set again; the Black player chooses another portion. The game continues to infinity with alternating cutting and choosing. White wins if and only if the infinite product of Black's choices is 0; otherwise Black wins. The strategies of infinite games are functions that define the moves of players. It takes infinitely extended memory plus infinite computational power to apply any strategy that maps arbitrary sequences (of previous moves) onto the player's next move. A player who lacks infinitely extended memory has to use strategies that depend only on finitely many previous moves. The skill of a player (his or her strategy) depends on his or her memory and computational power. Ciesielski and Laver (1990) and Scheepers (1993) describe infinite games in which infinite memory makes a difference.

12. The Absolute Minds

I've discussed a variety of endless series of ever greater (ever more complex) objects. Say an endless series *is bounded above by* x iff x is greater than every object in the series. Example: 1/2, 3/4, 7/8 is bounded above by 1. Say machine x *is greater than* machine y iff x can do everything that y can do but y cannot do everything that x can do. So there is an endless series of finite state machines (FSMs). The endless series of FSMs is bounded above by the universal Turing machines (UTMs). Any UTM can do what any FSM can do; but any UTM can do things that no FSM can do. I say there's an endless series of ever greater infinitary machines. I don't know whether or not the endless series of supermachines is bounded above or not. If no, then for any machine there is a greater (more powerful) machine above and beyond it. If so, then there are some machines greater than every machine in the endless series of infinitary machines. Say an object of some kind is *absolute* iff it is an instance of the kind than which no greater is possible. Any attempt to define a greater instance of that kind will be logically inconsistent. I follow Cantor's distinction of (1) the finite; (2) the infinite; and (3) the absolute. So a machine greater than every finitary and infinitary machine is an *Absolute Machine*.

I tend to think there are Absolute Machines and Absolute Minds. For any ordinals κ and λ, there is an infinitary language $L(\kappa, \lambda)$. There are games on arbitrarily large ordinals. As ordinals go, so go languages and games; so also go machines and minds. The series of machines and minds goes on like the series of ordinals. It is endless in a very strong sense. I'm arguing by analogy for Absolute Machines and Minds: just as the series of ordinals is bounded above (by the proper class Ω), and just as the series of sets is bounded above (by the proper class V), so also the series of machines is bounded above by an Absolute Machine and the series of minds is bounded above by an Absolute Mind. Just as a UTM can do whatever any other Turing machine can do (including any other UTM), so an Absolute Machine can do whatever any logically possible machine can do (including any other Absolute Machine). Absolute Machines are computationally universal in a very strong sense. An Absolute Machine contains every finitary and infinitary mind. It can do whatever any of those minds can do. Absolute Machines are Absolute Minds. The power of an Absolute Mind is so extreme that it does everything that every other mind does. If this is right, then Absolute Minds are cognitively or psychologically indiscernible. They are exact mental copies of one another. If there are many distinct Absolute Minds, each *is the same mind as* every other and each *is the same person as* every other.

Absolute Minds are those minds than which there are none more powerful or intelligent. Every Absolute Mind possesses all cognitive perfections. Taliaferro (1985, p. 139) defines omniscience as "supreme epistemic excellence". He says "X is *omniscient* if and only if it is impossible for there to be a being with greater cognitive power and this power is fully exercised" (Taliaferro, 1994, p. 287). Royce (1987) gives a psychologically rich analysis of such omniscience. A notion of omniscience as supreme cognitive power seems superior to the traditional definition of omniscience as merely knowing that p iff p is true. It is more dynamic. An Absolute Mind has the power to produce all truths (by making models of all consistent theories). If there are Absolute Minds, then all the cognitive operations of all human and superhuman minds are parts of these Minds.

An Absolute Mind realizes every other mind. It has parts that are indiscernible from those other minds. Somewhat more precisely, we may (and must) characterize Absolute Minds by *reflection principles:* if M is omniscient, and if p is any psychological property of M, then there exists some x such that x is a proper part of M and x instantiates p. Just as reflection principles are central to the study of extremely complex sets, so computational and cognitive reflection principles are likely to be important tools for studying very complex machines and minds.

The classical Neoplatonists (Plotinus, Proclus) argued for a trinitarian conception of God according to which God's three main hypostases are (1) Unity; (2) the Divine Mind; and (3) the World Soul. For these Neoplatonists, cognition stops with the Divine Mind; Unity is above the Divine Mind and so is beyond intelligence. My view of the hierarchy of minds accords well with this Neoplatonism. The community of Absolute Minds is a system of indiscernible intellects at the most complex level of mathematical existence. If Unity is construed as the ground of being (as Being rather than some being), then indeed Unity has no intelligence. It surpasses all sets, all machines, all minds. Traditional theists are not likely to approve of this picture. They are likely to argue that the Mind of God is superior to all minds. A trinitarian Christian might try to argue that there are exactly three Absolute Minds. I leave the difficulties of the Trinity and the Mind of God to others. One need not believe Anselm's "Proof" to accept the Anselmian definition of God as that than which no greater is possible. Since the Anselmian definition of God is the only one that makes any sense to me, I infer that God surpasses even the

Absolute Minds. For both theist and Neoplatonist, I think the Absolute Minds must be less than God.

13. Conclusion

According to classical metaphysics, reality is an *ordered complexity hierarchy*. Aristotle pictured a hierarchy of living thinking substances: plants, animals, humans, celestial intelligences. Above them all is an Absolute Mind: the Prime Mover. This Aristotelian divinity is an intellect that thinks about thinking. The Neoplatonists talked about endless hierarchies of intelligences (Proclus, 1992). Above them all is an infinitary Absolute Mind: the Divine Mind (Nous). The Medieval Christians posited choirs of angels. Above them all is the Absolute Mind of God. The story of the Mind of God is developed in modern thought by Berkeley, Spinoza, and Royce, among others. Although science has falsified many details of the classical story, its logical structure lives on.

Many features of classical metaphysics survive in modern mathematics. Mathematics pictures the World as an ordered complexity hierarchy. The World is the iterative hierarchy V of pure sets. The World consists of rank upon rank of increasingly complex objects. It starts with one simple particular object at its bottom level. The objects at the next level are all the logically possible combinations of all the objects at all the lower levels. These sets are all concrete particulars. Simpler sets are members of more complex sets. As sets grow ever more complex they form points, regions, vectors, fields, space-times, particles, systems of particles in motion according to causal laws. All possible spatio-temporal-causal systems (including our actual universe) exist within V. Sets are the ultimate theoretical entities of deep science – as real as quarks and stars.

I do not oppose mathematics to physics. I advocate a pythagorean ontology according to which the existence of V is the best explanation for the order and structure of our physical universe. If V exists, then a *hierarchy of all logically possible machines* exists. If all logically possible machines exist, then there are machines of arbitrarily high complexity that satisfy the biological requirements of life. If V exists, then a *hierarchy of all logically possible organisms* exists. If all logically possible machines exist, then there are machines of arbitrarily high complexity that satisfy the psychological requirements of intelligence. If V exists, then a *hierarchy of all logically possible minds* exists. If my reasoning is right, then V contains a community of Absolute Minds that live and think supremely far above all other minds. The modern mathematical World V looks very much like the classical cosmos – a glorious city of intelligent living machines. If my reasoning is right, then we live in a very richly and beautifully populated World.

Acknowledgements

I especially thank Jack Copeland and Jim Moor for their help. Many thanks as well to Pete Mandik, Jim Fetzer, and Pat Grim. The support of the philosophy departments at William Paterson University and Dartmouth College is greatly appreciated.

Notes

[1]Leslie (2001) continues the classical story in modern scientific terms. I learned of his book too late for this article. I hope to incorporate his ideas in future work.

[2]The finite ordinal numbers are the whole numbers 0, 1, 2, 3, and so on. Each ordinal n is the set of all ordinals less than n. So 0 = { }, 1 = {0}, 2 = {0, 1}, 3 = {0, 1, 2} and so on. The least infinite ordinal is the set of all finite ordinals. So the set ω = {0, 1, 2, 3,...} is the least infinite ordinal. The next infinite ordinal is $\omega + 1$ = {0, 1, 2, 3, ...ω}. There are endlessly many greater infinite ordinals beyond ω.

[3]I work within the universe of *sets* and *proper classes* defined by Von Neumann–Bernays–Godel *class* theory (VBG) and the von Neumann theory of *ordinal numbers* developed in VBG (Hamilton, 1982: chs. 4 and 6).

[4]Farmer and Belin (1991, p. 818) list these features of living things: existence not as a thing but as a pattern in space-time; self-reproduction; information storage of a self-representation; metabolism; functional interactions with the environment; interdependence of parts; stability under perturbations; the ability to evolve; growth. These features are easily extended to the transfinite (e.g., infinitely many genes; infinite genetic algorithms). Any thing, finitary or infinitary, ought to have features like those listed by Farmer and Belin to be living.

[5]I am not aware of any biologically sound argument for infinite complexity in actual living systems. Siegelmann (1996), Boucher (1997), and Penrose (1991) are not biologically plausible. Quantum mechanics (so far from revealing any natural infinitude) seems to prove that human animals are merely digital machines. Deutsch (1985) argues that quantum computation is just parallel Turing computation. The quantum mechanical theory of information (Bekenstein and Schiffer, 1990) implies that cells are only finitary (hence not even Turing machines) Moravec (2000, p. 166) uses Bekenstein's quantum mechanical theory of information to calculate the maximum amount of information in any human animal. He says we each contain about 10^{45} bits – a small finite number.

[6]A psychosemantic theory defines the constraints that any psychologically sound intention function must satisfy. One kind of psychosemantics says that state S of system x means that P iff S optimally causally covaries with P (Stalnaker, 1984, p. 24; Tye, 1995, p. 101). Other approaches include teleosemantic theories (Millikan) and asymmetric dependency theory (Fodor, 1987, 1990). Glasgow and Papadias (1992) show how psychosemantic theories might handle imagination. Ortony et al. (1990) show how to handle emotions. I prefer to analyze the attitudes using a physicalized version of Hintikka's modal analysis (Hintikka, 1969, 1975). For example: x sees that p = there is some retinal-neural state q such that x is in state q, and in all possible universes in which x is in state q, it is the case that p. The game tree for an agent and environment serves as the system of possible universes that supports the Hintikkan truth-conditions.

[7]If n and M are ordinal numbers, then an L(n, M) language has logical operator sequences of length less than n and quantifier sequences of length less than M. Since every number less than ω is finite, all sentences in an L(ω, ω) language have finitely long sequences of logical operators and finitely long sequences of quantifiers. The first-order predicate calculus is the prime example of an L(ω, ω) language. The lengths of sentences in L(ω, ω) languages can be of any finite length, they are

potentially infinite. Finite minds (like human animals) use some L(N, M) with both N and M finite. The language of thought of any supermind is *infinitary language* L(κ, λ) with κ and λ both greater than ω.

[8]I take a resolutely physicalistic attitude to the ontology of finite intelligence. This attitude is deeply anti-Cartesian. Although I deny the existence of Cartesian minds (*res cogitans*), I do not thereby deny the reality of the soul. I take an Aristotelian position: the soul is the form of the body at a level of abstraction sufficiently high to avoid any reference to particular materiality and sufficiently detailed to distinguish the body from all other possible organisms. The soul is the *algorithmic* form of the body. If someone were to raise religious worries about the immortality of the soul, then I would reply (1) since the soul is a property, it is eternal; (2) the resurrection of the body is a kind of personal immortality far more perfect than any disembodied cogitation (especially if that resurrection body were an infinitely complex organism – a physical supermachine with a supermind).

[9]Finitary minds understand the infinite by means of finitary descriptions (such as recursive definitions). Poincare (1952, p. 11) says: "reasoning by recurrence ... is the only instrument which enables us to pass from the finite to the infinite".

[10]Although the intelligence of human animals is finite, some will argue that human *animals* are merely proper parts of human *persons*. They will say that human animals are the material parts of human persons but that these persons have immaterial parts (Cartesian minds) that are infinitely intelligent. Since I reject Cartesian dualism, I can see only one way that human animals can be proper parts of human persons. It is that human animals are initial finite parts of human persons while human super-animals are later transfinite parts of human persons. One way to work this out is to use something like Hick's (1976) resurrection theory. Accordingly: earthly human animals are only finitary machines; earthly human animals are resurrected into infinitary resurrection bodies (into supermachines). Human persons are infinitary things whose earthly animal parts are all only fintary. You may be infinitely complex in the future; you are not infinitely complex now.

[11]Copeland (2000, p. 9) says: "A *wide mechanist* ... holds that the mind is a machine but countenances the possibility of information-processing machines that cannot be mimicked by a universal TM, and allows in particular that the mind may be such a machine". I'm a wide mechanist. I'm open to the possibility that earthly human bodies are infinitary.

[12]A machine is countably infinitely complex iff the least upper bound of the cardinality of its set of possible states is ω. The following machines are countably infinitely complex: push-down automata with actually infinitely deep stacks; linear bounded automata (classical Turing machines with a finitely long tapes determined by the unbounded sizes of their inputs); classical Turing machines whose tapes are finite but always extendible; classical Turing machines; classical Turing machines that make non-recursive moves (Giunti, 1997) or that operate on non-recursively defined data (Shagrir, 1997).

[13]Since this is a game, we need some utility function that assigns outcomes to the players. Let S be any set of real numbers. If the number $0.d_1d_2d_3d_4$ ... is in S, then A wins; if not, then E wins. Utility functions ground the preferences, goals, emotions, attitudes, and the moral and social values of the players of infinite games.

[14]Other transfinite machines are analog shift maps (Siegelmann, 1995); real random access machines (Preparata and Shamos, 1985, p. 1.4); flowchart machines (Blum et. al, 1998).

[15]Zeno compressions can be nested. You can put ω acts between times 0 and 1/2. Put the first act at 1/4, the second at 3/8, the third at 7/16 and so on. You can put ω more acts between 1/2 and 3/4, and so on. See Rucker (1995, pp. 66–69).

[16]For any drawings A and B, and for any point (x, y), let the color of the point (x,y) in the merged drawing (A * B) be the disjunction of the color of (x,y) in A and B. So point (x, y) is 1 in (A * B) iff (x,y) is 1 in A or (x,y) is 1 in B.

[17]We may divide objects and tasks into those that have finitary descriptions and those that have only infinitary descriptions. Finitary minds like ours have at most finitary representations of infinitary things (e.g., recursive definitions in L(ω,ω)). We cannot describe any object whose complexity ex-

ceeds finitary definition in L(ω, ω). At this boundary the superminds themselves take off into the pythagorean heavens, leaving us in the finitary dust. While we can argue *that* superminds are, we have little understanding of *what* they are. The ability of superminds to perceive, think, and will in languages greater than L(ω,ω) places them far beyond our comprehension.

[18]Any infinitary object contains a proper part that is isomorphic to the whole. An infinitary object is *self-reflecting*. For instance: in the case of the infinitely self-nested cross, one of the corner crosses completely reflects the whole. Since no finitary object contains any part with equal complexity, no finitary object is able to contain a complete and true description of itself. *Self-consciousness* is the complete mental self-reflection (self-representation) of the whole mind by a part of the mind. No finitary mind is completely self-reflecting; hence no finite mind is completely self-conscious; all and only infinitary minds are self-reflecting; hence all and only infinitary minds are truly self-conscious.

References

Abelson, H. and diSessa, A. (1984), *Turtle Geometry: The Computer as a Medium for Exploring Mathematics*, Cambridge, MA: MIT Press.

Axelrod, R. (1984), *The Evolution of Cooperation,* New York: Basic Books.

Bekenstein, J. and Schiffer, M. (1990), ‘Quantum Limitations on the Storage and Transmission of Information’, *International Journal of Modern Physics*, 1 pp. 355–422.

Bhalla, U. and Iyengar, R. (1999), ’Emergent Properties of Networks of Biological Signaling Pathways’, *Science* 283, pp. 381–387.

Blum, L., Cucker, F. Shub, M. and Smale, S. (1998), *Complexity and Real Computation*, New York: Springer-Verlag.

Boolos, G. and Jeffrey, R. (1980), *Computability and Logic* (2nd edition), New York: Cambridge University Press.

Borges, J.L. (1964), *Labyrinths*, New York: New Directions.

Boucher, A. (1997), ’Parallel Machines’, *Minds and Machines* 7, pp. 543–551.

Brams, J. (1983), *Superior Beings,* New York: Springer.

Bray, D. (1995), ‘Protein Molecules as Computational Elements in Living Cells’, *Nature* 376 (6538), pp. 307–312.

Ciesielski, K. and Laver, R. (1990), ‘A Game of D. Gale in which One of the Players has Limited Memory’, *Periodica mathematica Hungarica* 22, pp. 153–158.

Conway, J. (2001), *On Numbers and Games* (2nd edition), Natick, MA: A. K. Peters.

Copeland, B.J. (1998a), ‘Super-Turing Machines’, *Complexity* 4(1), pp. 30–32.

Copeland, B.J. (1998b), ‘Even Turing Machines can Compute Uncomputable Functions’, in C. Calude, J. Casti, and M. Dinneen, eds., *Unconventional models of Computation,* New York: Springer.

Copeland, B.J. (2000), ‘Narrow Versus Wide Mechanism’, *Journal of Philosophy* 95(1), pp. 5–32.

Cuthbertson, R., Paton, R. and Holcombe, M. (1996), *Computation in Cellular and Molecular Biological Systems*, River Edge, NJ: World Scientific.

Danielson, P. (1992), *Artificial Morality: Virtuous Robots for Virtual Games*, New York: Routledge.

Davies, E.B. (2000), ‘Building Infinite Machines’, *British Journal for the Philosophy of Science* 52(4), pp. 671–682.

Dennett, D. (1996), *Kinds of Minds,* New York: Basic Books.

Deutsch, D. (1985), ‘Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer’, *Proceedings of the Royal Society,* Series A, 400, pp. 97–117.

Doyle, J. (1991), ‘The Foundations of Psychology’, in R. Cummins and J. Pollock, eds., *Philosophy and AI: Essays at the Interface,* Cambridge, MA: MIT Press, pp. 39–78.

Earman, J. and Norton, J. (1996), ‘Infinite Pains: The Trouble with Supertasks’, in A. Morton and S. Stich, eds., *Benacerraf and His Critics,* Cambridge, MA: Blackwell, pp. 231–262.

Farmer, J. and Belin, A. d'A. (1991), 'Artificial Life: The Coming Evolution', in C. Langton, C. Taylor, J. Farmer, and S. Rasmussen, *Artificial Life II*, Reading, MA: Addison-Wesley, pp. 815–840.

Fodor, J. (1975), *The Language of Thought*, Cambridge, MA: Harvard University Press.

Fodor, J. (1987), *Psychosemantics*, Cambridge, MA: MIT Press.

Fodor, J. (1990), *A Theory of Content and Other Essays*, Cambridge, MA: MIT Press.

Freiling, C. (1984), 'Banach Games', *Journal of Symbolic Logic* 49(2), pp. 343–375.

Gale, D. and Stewart, F.M. (1953), 'Infinite Games With Perfect Information', in H. Kuhn and A. Tucker, eds., *Contributions to the Theory of Games Vol. 2,* Annals of Mathematical Studies 28, Princeton, NJ: Princeton University Press, pp. 245–266.

Giunti, M. (1997), *Computation, Dynamics, and Cognition*, New York: Oxford University Press.

Glasgow, J. and Papadias, D. (1992), 'Computational Imagery', *Cognitive Science* 16, pp. 355-394.

Hamilton, A. (1982), *Numbers, Sets, and Axioms: The Apparatus of Mathematics*, New York: Cambridge University Press.

Hamkins, J.D. and Lewis, A. (2000), 'Infinite Time Turing Machines', *Journal of Symbolic Logic* 65(2), pp. 567–604.

Hick, J. (1976), *Death and Eternal Life*, New York: Harper and Row.

Hintikka, J. (1969), *Models for Modalities,* Dordrecht: Reidel.

Hintikka, J. (1975), *The Intentions of Intentionality,* Dordrecht: Reidel.

Jech, T. (1984), 'More Game-Theoretic Properties of Boolean Algebras', *Annals of Pure and Applied Logic* 26, pp. 11–29.

Juhl, C. (1995), 'Is Gold-Putnam Diagonalization Complete?', *Journal of Philosophical Logic* 24, pp. 117–138.

Karp, C. (1964), *Languages with Expressions of Infinite Length*, Amsterdam: North-Holland.

Kitcher, P. (1984), *The Nature of Mathematical Knowledge*, New York: Oxford University Press.

Koetsier, T. and Allis, V. (1997), ''Assaying Supertasks', *Logique et Analyse* 159, pp. 291–313.

Koshland, D.E., Jr. (1980), 'Bacterial Chemotaxis in Relation to Neurobiology', *Annual Review of Neurosciences* 3, pp. 43–75.

Leibniz, G.W. (1710/1996) *Theodicy*, trans. E.M. Huggard with modifications by E. Steinhart, Peru, IL: Open Court.

Leslie, J. (2001), *Infinite Minds: A Philosophical Cosmology*, New York: Oxford.

Lewis, D. (1986), *On the Plurality of Worlds*, Cambridge, MA: Blackwell.

Lodish, H., Baltimore, H., Berk, A., Zipusky, S., Matsudaira, P. and Darnell, J. (1995), *Molecular Cell Biology*, 3rd edition, New York: W.H. Freeman.

Lovejoy, A. (1936), *The Great Chain of Being*, Cambridge, MA: Harvard University Press.

Maes, P. (1995), 'Modeling Adaptive Autonomous Agents', in C. Langton, ed., *Artificial Life: An Overview*, Cambridge, MA: MIT Press, pp. 135–162.

Mandik, P. (2002), 'Synthetic Neuroethology', in T. Bynum and J. Moor, eds., Metaphilosophy Special Issue: *CyberPhilosophy: The Intersection of Philosophy and Computing, Metaphilosophy* 33(1/2), pp. 11–29.

Mandlebrot, B. (1978), *Fractals: Form, Chance, and Dimension*, San Francisco: CA: W.H. Freeman.

McAdams, H. and Shapiro, L. (1995), 'Circuit Simulation of Genetic Networks', *Science* 269, pp. 650–656.

Moravec, H. (1988), *Mind Children: The Future of Robot and Human Intelligence*, Cambridge MA: Harvard University Press.

Moravec, H. (2000), *Robot: Mere Machine to Transcendent Mind*, New York: Oxford University Press.

Ortony, A., Clore, G., and Collins, A. (1990), *The Cognitive Structure of Emotions*, New York: Cambridge University Press.

Penrose, R. (1991), *The Emporer's New Mind*, New York: Penguin Books.

Poincare, H. (1952), *Science and Hypothesis*, New York: Dover Publications.

Preparata, F. and Shamos, M. (1985), *Computational Geometry: An Introduction*, New York: Springer.
Proclus (1992), *The Elements of Theology*, trans. E.R. Dodds, New York: Oxford University Press.
Prusinkiewicz, P. and Lindenmayer, A. (1990), *The Algorithmic Beauty of Plants*, New York: Springer.
Ptashne, M. (1996), *A Genetic Switch: Phage (Lambda) and Higher Organisms,* 2nd edition, New York: Blackwell Science.
Quine, W.V. (1969), *Ontological Relativity and Other Essays,* New York: Columbia University Press.
Quine, W.V. (1976), 'Wither Physical Objects?', *Boston Studies in the Philosophy of Science* 39, pp. 497–504.
Quine, W.V. (1978), 'Facts of the Matter', *Southwestern Journal of Philosophy* 9(2), pp. 155–169.
Quine, W.V. (1981), *Theories and Things*, Cambridge, MA: Harvard University Press.
Royce, J. (1927), *The World and the Individual, First Series, Supplementary Essay*, New York: The Macmillan Company.
Royce, J. (1987), 'The Conception of God', in J. Royce, J. Le Conte, G. Howison and S. Mezes, *The Conception of God,* New York: The Macmillan Company.
Rucker, R. (1995), *Infinity and the Mind*, Princeton, NJ: Princeton University Press.
Russell, S. and Norvig, P. (1995), *Artificial Intelligence: A Modern Approach*, Englewood Clifs, NJ: Prentice Hall.
Scheepers, M. (1993), 'Variations on a Game of Gale (I): Coding Strategies', *Journal of Symbolic Logic* 58(3), pp. 1035–1043.
Scott, J. and Pawson, T. (2000), 'Cell Communication: The Inside Story', *Scientific American* 282(6), pp. 72–79.
Shagrir, O. (1997), 'Two Dogmas of Computationalism', *Minds and Machines* 7(3), pp. 321–344.
Siegelmann, H. (1995), 'Computation beyond the Turing Limit', *Science* 268, pp. 545–548.
Siegelmann, H. (1996), 'The Simple Dynamics of Super Turing Theories', *Theoretical Computer Science* 168, pp. 461–472.
Smart, J.J.C. (1972), 'Space', in P. Edwards, ed., *The Encyclopedia of Philosophy*, New York: Macmillan Publishing, Vol. 7, pp. 506–511.
Smolgyi, R. and Sniegoski, C. (1996), 'Modeling the Complexity of Genetic Networks', *Complexity* 1(6), pp. 45–63.
Sorensen, R. (1994), 'Infinite Decision Theory', in J. Jordan, ed., *Gambling on God*. London: Rowman and Littlefield, pp. 139–159.
Stalnaker, R. (1984), *Inquiry*, Cambridge, MA: MIT Press.
Steinhart, E. (2001), 'Persons vs. Brains: Biological Intelligence in the Human Organism', *Biology and Philosophy* 16(1), pp. 3–27.
Steinhart, E. (2002), 'Logically Possible Machines', *Minds and Machines* 12(2), pp. 259–280.
Taliaferro, C. (1985), 'Divine Cognitive Power', *International Journal for Philosophy of Religion* 18, pp. 133–140.
Taliaferro, C. (1994), *Consciousness and the Mind of God*, New York: Cambridge University Press.
Tye, M. (1995), *Ten Problems of Consciousness*, Cambridge, MA, MIT Press.
Yockey, H. (1992), *Information Theory and Molecular Biology*, New York: Cambridge University Press.
Yuh, C.-H. et al. (1998), 'Genomic cis-Regulatory Logic', *Science* 279, pp. 1896–1902.