



{...}

*Laugh-a while you can, monkey boy (9/11/00)*¹

... As for our other misadventures:

I did find the Joy piece in the *Wired* archive, but tired after half an hour or so of slogging through extraneous downloads and never got to the end of it. (Surfing has now become roughly as exciting as standing in line at the Post Office.) Meaning, I guess, that I got through the part where the machines are taking over, but never quite got to the

¹ To [RS]. — This was originally provoked by an article in *Wired* written by Bill Joy on the subject of the threat posed by artificial intelligence and the necessity of regulating its development. The difficulties in obtaining the complete text notwithstanding, it was obvious from a cursory inspection of his argument, such as it was, that he thought he and his tech-bro buddies could call a halt to something that (it was clear) none of them understood. I found this fantasy of command and control enormously irritating; and though far more serious thinkers, viz. Stephen Hawking, have since expressed similar sentiments, I still do.

part where we all learn kung fu like Keanu and beat them up in virtual reality while looking very cool in black leather and mirrorshades. — It is profoundly depressing to compare this freshman-essay composition with the philosophical pieces we were accustomed to receiving from scientists a generation ago: Dyson, Monod, even Bethe or Feynman; let alone Bohr, Einstein, Born, Heisenberg, Schrödinger. Or (more directly to the point) Turing or Von Neumann. — Mr. Joy doesn't seem to be able to reason on his own, preferring, apparently, a technique of random quotation introduced by namedropping anecdote; the result, I presume, of a career spent thinking in buzzwords. I don't care how much money the guy has made; he's an illiterate cretin.²

On the other hand I've read one or two occasional pieces of Kurzweil's, and he's very smart. He wrote an article in the *American Scientist* a few years ago, for instance, that indicated an appreciation of the fact that the historical development of artificial intelligence has recapitulated the development of analytical ("linguistic") philosophy; even quoting Wittgenstein, as I recall. (Of course this won my heart at once.) Whatever he has said on the subject is probably much more interesting. But is this yet in print?

(The significance of Wittgenstein in this connection is that he selfconsciously took an engineer's attitude toward the philosophy of logic; thus the invention of truth tables — the original "mechanical procedure" — which, carried to its logical conclusion, led to Turing's treatment of abstract machines. Insofar as there's a philosophy of mind in Wittgenstein's *Tractatus*, it's a theory of automata.)

² Doubtless an exaggeration, but it cannot be sufficiently emphasized that computer engineers, as a class, are persons with extremely narrow educations who have a hypertrophied ability to keep track of minutiae and a very poor ability for handling abstractions.

I flipped through the responses to Joy's article in a subsequent issue while standing in line at the grocery store, but couldn't find much substance; indeed, it looked as though they'd simply printed the comments of the people who sounded most impressive; still more namedropping. (Again, the corrupting influence of the principle that all journalism is celebrity journalism.) — I have somewhere the Xerox of an entire issue of one of the AI journals devoted to ritual denunciations of Roger Penrose, published shortly after the appearance of *The Emperor's New Mind*; no less than forty separate polemics, and every one of them sucked. — I wouldn't expect much from any public forum, in other words; even if the remarks were not constrained to sound-bite length.

As for the end of the world and the destruction of humanity, all this is true, sure enough. But since nothing can be done to stop it, and since the exact course events will take is impossible to foresee (for the usual reason, i.e., even if the initial and final states and even the change in free energy driving the transition were known, the chain of intermediate catalytic reactions would not be), there's not much point in worrying about it.

(This is a slightly stronger version of the usual observation on unpredictability: not simply instability of the evolution under perturbation of the initial conditions, but instability of the path connecting initial and final states under slight perturbations of either. Or slight perturbations of the form of the Lagrangian, for that matter.)

There was a period when (never mind exactly why) I went around polling people on their opinions regarding the identities of the Four Horsemen of the Apocalypse. This turned out to be one of those questions for which the officially received answer was by no means the best one. (E.g. Lankton: "Poverty, Flatulence, Horniness, and Atrophy.") — Anyway I never did get it straight, and found it easiest to stick with the first thing that had come into my own head —

predictably, Groucho, Chico, Harpo, and Zeppo.

I doubt we need a romantic lead, making Zeppo once again dispensable. This leaves us with three obvious threats: artificial intelligence; nanotechnology (specifically artificial life, the mean between the two extremes); and reprogramming the genome.

The problem you have in discussing these several Horsemen, then (something which certainly sails over the head of Joy, but when I need a whipping boy in these polemics I generally flog Minsky) (we stuck a fork in Skinner long ago), is just the problem that their writers used to have with the Marx Brothers: although it's gratifying to pretend for the benefit of journalists and babes at cocktail parties that you've scripted the performers' routines, in reality you never have any idea before the curtain goes up precisely what they're going to do. Obviously you remember the story about George S. Kaufman at a performance of *Animal Crackers* ("Wait a minute — I thought I heard one of the original lines"); guys like Joy and Minsky, alas, are never so honest.

This problem is intrinsic, you can't discuss it without confronting the fallacious idea that you can bottle intelligence, which is (psychologically, anyway) a corollary of the hacker obsession with fantasies of control, and, accordingly, most of the salient points fall out in the first third of the discussion; which, therefore, is the longest. (And anyway it figures that Groucho would do most of the talking.) — Actually the same points keep falling out over and over again in slightly different form, but since even now no one seems to understand them, redundancy in exposition probably serves a useful purpose. (I read only the other day that Minsky is still on tour claiming that once the Master Algorithm has been unearthed the human brain can be replaced by a hundred-megahertz CPU: amazing stupidity.)

At any rate:

With regard to artificial intelligence, you are doubtless aware that the progress of the research program, at least as classically conceived, has been vastly exaggerated.

It is amusing to contrast the brilliant advertising campaign with the relative triviality of the problems no one seems to be able to solve: e.g., finding the date fields in old computer programs; anyone with a way to automate this task would have made more money than Gates overnight, in 1999. — Or, more generally, the unsolved problem of the intelligent compiler. Of all the purportedly mechanizable intellectual activities with which I am familiar, hacking is the one that should be easiest to automate. And if you could mechanize the optimization of the translation of higher-level languages into assembly language, you could buy Microsoft with petty cash.

This doesn't mean that I put much stock in the traditional arguments against the possibility of machine intelligence (e.g. Searle's room that doesn't know Chinese, or the folklore arguments Penrose attempted to summarize from the Gödel incompleteness theorem.) These have never made sense to me, and I presume that they're wrong (in fact not even wrong in an interesting fashion.)

It seems to me instead that, though machine intelligence is obviously possible, there has been an irrational insistence on trying to develop it in exactly the wrong way. (Or: with the wrong definition of "machine".) And that this insistence has been motivated by illusions about rationality, precision, predictability, understanding, and the ability to manipulate and control.

You observe, for instance, that the apparent triumphs of the program, e.g. the computer's victory over Kasparov, on closer examination only reveal its futility. For here you have on one side a machine evaluating billions of positions before it makes a move; and on the other a guy

who doesn't seem to be examining more than a couple of dozen (and who then loses only because he chokes). Obviously whatever he is doing — whatever the workings of his natural/all-too-natural intelligence — bears no resemblance to the brute-force machine algorithm, and (at least on the face of it) is infinitely more efficient.

Moreover (and maybe most important) what Kasparov does seems to scale better, in the sense that if you increased the complexity of the problem and asked the computer to evaluate ten times as many positions, its performance would deteriorate much more dramatically than Kasparov's.

And, really, the classical AI idea isn't much more than this: a brute-force tree-search, with some (essentially arbitrary) evaluation function at the leaves. The implicit thesis is that the function of intelligence can be reduced to deterministic algorithms (ideally, to just one master algorithm — like resolution in Prolog); the implication is always that even if this isn't the way thinking is done in a state of nature, this is the way to do it with a computer; and anyway that it is more “exact”, and therefore better.

A sympathetic reading of this approach might compare it with Carnap's interpretation of the logic of induction as a method of justification, not of discovery: you itemize a set of rules, and brush away protests that these do not represent the real methods of the human scientist with the rejoinder that you aren't concerned with that; and anyway they *ought* to.

This is the distinction between, e.g., how you know that 25 is the next number in the sequence 1, 4, 9, 16, ..., and how you would try to prove that this statement is correct; or between the act of recognition of a picture as a picture of Richard Nixon, and the assignment of a probability or degree of truth to the statement that the picture in question is a picture of Richard Nixon. Somehow though the former is

easy, and anyone can do it, the latter is difficult, and only a logician can even try; moreover (since the act of recognition is a real process in a real brain, and the “degree of truth” is a mathematical fantasy) somehow in the name of logical reconstruction you replace science with science fiction. But when I was reading Carnap it made a twisted sort of sense.

Another of the “classical” approaches, rulebased systems, involved enumerating an enormous number of facts and principles for a class of problems and trying to deduce the solution for any individual instance by brute force. — The analogy with civil law is interesting. — The difficulty that the set of “facts” and “axioms” is never internally consistent was always an embarrassment. — This worked fairly well with things like medical diagnosis (EMYCIN) and very badly with things like trying to teach robots to drive. No one ever attempted to address the obvious question, namely, if we only have a hundred thousand genes, how do we code all these rules? and how did the rules evolve?

But (anyway) I find it hard to give classical AI the sympathetic reading. Everything about it suggests an appeal to the fantasies of control so dear to the hacker psyche.

About which, so far as I can recall, the original reference is Weizenbaum’s book *Computer Power And Human Reason*: I haven’t looked at this in years, but seem to remember that he stressed the idea that most AI applications, e.g. his own ELIZA, were based on very simple tricks, and emphasized the importance of revealing the geek behind the curtain, as it were; moreover, and surprisingly for an MIT guy, he expressed grave reservations about hacker culture before anybody outside Cambridge even knew what it was.³

³ Having looked at it since, I can report that this summary is entirely accurate.

The peculiar inflexibility of the doctrine is the giveaway: it suggests some unstated (“unconscious”, if you like) set of motivations or reasons which are the actual supports of the structure. (Some sort of compulsion.)

That is, there is an assumption, as it were unquestionable, a priori, that the possible moves of an intelligence can be enumerated, catalogued, before the fact, and searched, surveyed, examined. — After all, this is how computer programs are written. You anticipate everything in advance. For each situation, you craft a response. — If you’re writing a computer algebra program, you have a list of procedures you can employ to transform an expression, and, basically, you try all of them; if you’re trying to integrate an expression involving elementary functions, you list all the procedures you can employ to transform that; if you’re writing air-traffic-control software, you list all the paths the planes can follow and try to keep them from running into one another (yeah, right); etc., etc.

Obviously it ought to be an embarrassment that the strategy is always the same but the lists of procedures are always contrived ad hoc. Still, if you have a problem like chess for which you can handtune the search, the results can be impressive.

But in consequence, in the design of computer programs employed by others, the designer tends to develop the attitude that the user (or “luser”, as the MIT guys always loved to say) is just a rat in a maze, whose every move has been anticipated in advance by the (godlike) programmer. And presently the designer begins to think that he can anticipate anybody’s moves, anytime, anywhere; and that the fact that his computer programs never acquire the true autonomy that is characteristic of real intelligence is an asset, not a liability — the program need do no more than bottle the intelligence of the programmer, as it were, and this should be good enough. For after all the programmer can code anything. Thus inevitably he does not

address the question of what problem-solving is so much as he attempts to solve all the problems in a given class in advance and then bottle the results; though presumably the classes are getting bigger as the project moves along, and he expects someday to arrive, as it were, at root.

This always reminds me of the story about the late scion of the degenerate Habsburg line who saw his first train pulled by an engine and assumed there had to be a horse inside. But in fact in traditional AI programs (in computer programs generally) there *is* no engine; and there always *is* a programmer inside.

Again, this is not very different from what Carnap thought: that human reasoning could be codified, regularized, and then (of course) improved upon. — I always thought it was amusing that Carnap kept backing away from the implication that one would be able with his inductive logic to evaluate the probabilities of truth of competing theories and thus, as it were by calculation, determine which was most correct — resolving all scientific controversies “scientifically” — obviously this should have been the whole point of the exercise; the idea goes back at least to Leibniz. — Though to be fair only recently has it become possible to mechanically verify even mathematical proofs of any complexity.

For the hacker, at any rate, megalomania is not simply an occupational disorder, but a kind of methodological imperative.

I’m never really sure whether I should take all this seriously. I’m always reminded of the way that Rotwang the inventor, with wild hair and mad glowing eyes, describes his robot in Fritz Lang’s *Metropolis*: “I have created a machine in the image of man, that never tires or makes

a mistake!”⁴ Then I remember that the first thing the (emphatically female) robot actually does is to mesmerize a club full of bankers with an erotic striptease. — Isn’t this the story of the Internet, after all.

Anyway. All this provides an illustration of what I call the Napoleonic fallacy: the idea that you can sit on a white horse on top of a high hill overlooking the battle, looking down at the noise, the smoke, and the confusion, and with a swift commanding glance of your godlike eagle eye compose the whole into a unifying vision; and then, presumably, summon an orderly, and commit your reserves to the charge.

But Napoleon is like Maxwell’s Demon: when you examine what he has to do in order to function as you imagine that he would, you find that he can’t. — The “swift commanding glance” presupposes gathering the (noisecorrupted) data and then piping it through some (bandwidth-limited) conduit back to Napoleon — exactly the kind of situation for which Backus invented the phrase “von Neumann bottleneck”;⁵ and “composing the whole into a unifying vision” usually involves a combinatorially impossible problem in pattern recognition — which, as it were by definition, you’re pretending will be handled by a single processor.

Implicitly you presuppose a pyramid — a hierarchical structure, a chain of command — with Napoleon at the top, evaluating everything that’s piped up to him. — But whatever Napoleon’s speed might be, it is finite. Given that, and the simple fact that the size of the pyramid he sits upon must grow exponentially, its maximum height must be proportional to the logarithm of his speed.

⁴ This title card does not appear in the most recent (and complete) print I have viewed, though of course that means nothing; the sentiment was clearly there.

⁵ Cf. John Backus, “Can Programming Be Liberated From the Von Neumann Style?” *Communications of the ACM*, Volume 21, Number 8 (August 1978). (Turing Award Lecture 1977.)

So you know you can't do this as a serial computation. But once you start to try to analyze this as a parallel computation, you start deconstructing Napoleon! — and he factors into a network of less-omniscient generals, on grayer horses on lower hills, who are arguing with one another via dispatches that only get through once in a while because the orderlies transmitting them keep getting shot ... presently nobody is taking orders at all, and you have a sort of procedural chaos from which order emerges only from the bottom up, not by imposition from the top down. Or percolates outward from the middle of the network, which meanwhile you have realized isn't really a hierarchical tree at all, but a very unhierarchical general graph — in which, you perceive, the characteristic size of the neighborhood over which any individual node can maintain a coherent picture of the global state is, again, fixed by that node's speed.

What *works* in nature, in other words, is something much closer to anarchy, and involves local communication and local control. It is important to think like a physicist here, and not like an engineer: to realize that this state of affairs is dictated not by the frivolous accidents of history, but by mathematical necessity. Nature never bothered with the “rational” approach because, actually, it doesn't work.

On the usual argument there'd be no point in implementing biological evolution, for instance; adopting the role of Napoleon, you'd simply “look at” the possible genomes and then select the best one. (Actually, and even harder: “best” — maybe in the sense of “largest mutually compatible” — subset of them. Compare Leibniz 'idea that the best of all possible worlds was the largest mutually compatible set of possibilities, and translate this into biology.) Part of the fallacy, obviously, lies in supposing you could examine and rank, say, $2^{6000000000}$ possible genomes; another, curiously enough, lies in the

failure to recognize that the process that constitutes your “internal examination” of the set of genomes is essentially isomorphic to the external, experimental process of evolution — that it is exactly this calculation, in other words, that Nature is performing. (It’s a related observation that with problems of this order of complexity you find yourself thinking, literally, that God himself could not tell you beforehand what the answer is, and that the world represents some kind of simulation He’s running to find out. Reminding you of the old observation about Augustine’s idea that the world might consist only of thoughts in the mind of God, and that this wouldn’t make any difference to anything.)

The point is a trifle subtler, of course: it would be absurd to suppose that you could examine a zillion genomes one by one, but it might not be absurd to assume the existence of an algorithm that permitted you, as it were, to prune the tree of choices fast enough to allow a deterministic calculation in a reasonable amount of time. But the first thing you learn from complexity theory is that almost any interesting problem (e.g., Boolean satisfiability, the existence of a Hamiltonian circuit, three-colorability, the travelling salesman) that appears to grow in difficulty exponentially in the size of the input actually does; that it cannot be pruned effectively, that the exponential fanout of the tree is irreducible.

The futility of parallelism (which is apparently less obvious) then follows. The number of processors you can stuff into a machine grows (since we don’t live in Hilbert space) only as the cube (or in arbitrary fixed finite dimension as a polynomial) of its size; the number of possibilities that have to be examined in a general tree-search goes up exponentially. Accordingly, though for instance like everyone else I stand in awe of Adleman’s⁶ ingenuity in inventing a technique for

⁶ Leonard M. Adleman, “Molecular Computation of Solutions to Combinatorial Problems”, *Science* (New Series), Volume 266, Issue 5187 (Nov. 11, 1994), 1021-1024.

solving the canonical combinatorial optimization problems by coding their representations into DNA strands and, in effect, using every molecule in a sample of macroscopic size as an independent processor to test a trial solution, with a mole⁷ of DNA you'd be able to solve the travelling salesman problem for (let's wave our hands) 24 cities. If you filled all the oceans with DNA you'd be able to solve it for something like 37 or 38 cities. If you filled the physical universe, you might be able to solve it for 60 cities. (That's 83209871127413901442763411832233643807541726063612459524492776964096000000000000 routes you have to examine, incidentally.)(Thanks. I'll be here all week.)

This assumption (of the irreducibility of fanout) is essentially the assumption that, as per conjecture, P doesn't equal NP (so far as I'm concerned this, not Gödel's, is the theorem that says intelligence isn't mechanizable in the traditional sense); an open question, strictly speaking, but equivalent to the assertion that, in any nontrivial formal system, finding a proof of a theorem is inherently harder than verifying that proof — which is, to say the least, intuitively sound. — Moreover when you think about it you realize that it is just some such a point that must vitiate the mechanization program: remembering, e.g., may be a matter of mechanical procedure, but discovery or invention is something dual to it, like remembering backwards in time; just the distinction between verifying a proof and finding one.

(Also, as I'm always observing, it's much harder to write than it is to read.)

There's something right about that idea: that reversing the direction of time is, mainly, computationally prohibitive. Compare Maxwell's Demon, which seemed to have a computational version. — Maybe

⁷ I.e., Avogadro's Number of molecules; about 6.02 times 10²³.

this: given the macrostate in which the gas is divided evenly between the halves of the box, the computational difficulty of identifying the microstate (out of an enormous number N) that puts all the molecules back on one side is huge; the Demon has something like a sorting problem. The number of steps then probably just goes as $\log(N)$, which is to say, the entropy is the length of the sort. — More succinctly: consider the computational difficulty, as, say, a picture-puzzle, of putting Humpty-Dumpty back together again.

One obvious (though never publicly admitted) corollary is the impossibility of writing long computer programs without mistakes. Writing a program is equivalent to finding a proof, in a suitable formal system. (Naively, it's a proof that a given function is recursive, though I think there's better theory on this point, cf. the literature on typed lambda-calculi, particularly on the Curry-Howard correspondence; and it's a bit of a mystery how intuition provides the alternative description of the function. — Part of the problem in real life, actually, is that intuition simply doesn't; the process of writing the program is largely a process of inventing and refining the specification.) The difficulty should go up exponentially with length — and, if you examine the empirical evidence honestly, this is obviously the case. — The difficulties of matching a series of program segments don't add, as by some mental/optical illusion you think they ought to: they multiply. — The limits of human capability have long since been reached and exceeded. (And thus no wonder nothing ever really works.) There is no hill high enough, no eagle-eye sharp enough, no horse white enough to finish debugging Windows 95. — Or the operating system for the IBM 360, for that matter; the *locus classicus* (cf. *The Mythical Man-Month*) which seemed most apropos to the problem when I started thinking about this in connection with the problem of validating the SDI code.

On a less cosmic scale, this (the impossibility of Napoleon) is exactly the problem with a command economy. First, as a practical matter the

people pretending to run it can't acquire and assimilate all the information they need fast enough to make decisions (I remember hearing a story about a midlevel Soviet apparatchik who was supposed to try to adjust sixty thousand prices a month); second, if you ask yourself exactly how you're supposed to compute, e.g., the greatest good of the greatest number, you realize immediately that (a) this is an impossibly difficult combinatorial optimization problem and (b) this is (modulo the arbitrariness of the utility function) what the market is doing for you, anyway — in a distributed-computational model much more efficient than anything you can design from the top down. (However, contra the usual free-market mythology, the invisible hand is only guaranteed to find a local, not a global minimum; and it's obvious that large corporations and governmental bureaucracies are inherently inefficient in exactly the same way.)

Ah, these are great times to be an anarchist. You find yourself winning arguments you hadn't even thought to contest.

Anyway this suggests the approach that anyone familiar with mathematical physics would have thought of in the first place — namely, when you have a seemingly intractable problem, you take the hint from the way that Nature solves it.

This is exactly the fascination of protein folding, since the problem is *prima facie* combinatorially impossible (Levinthal's paradox): you might have 10^{200} ways of folding a polypeptide of 200 amino acids; the natural cycle time of the system (at 300 degrees Kelvin) is about $1.6 \cdot 10^{-13}$ seconds, suggesting that it can't sample much more than a few trillion configurations in the time it observably takes to fold; how does it find the ground state? — Literally the question is: how does the protein make this computation?

Historically this motivated the idea of trying a theory of automata more closely modelled on the workings of the brain; this is, of course, the theory of neural networks, which began in the Forties with McCullough and Pitts, languished in the Horse Latitudes for several decades, and was then revived in the Eighties with the work on associative memory models advanced, e.g., by Hopfield.

To make a long story short, this is the first real advance in the philosophy of mind since Hume, and it's obviously the correct approach; though the details are formidable, only details remain. So it's clear that a real thinking machine will not be deterministic, that its output on given input will not be predictable, that it will solve problems by guessing, that (just as evolution does) it will depend on error to function, that it will not in the ordinary sense be programmable ("parallel programming" is almost a contradiction in terms); that it will, in short, be selforganizing, near-chaotic, undesigned, autonomous. — Also (and this is a critical realization) that it won't ever solve the problems it's presented with "exactly", not because of the imperfection of its design, but because exact solutions for those problems don't really exist.

So, anyway, the idea of dissecting the human mind on a laboratory bench, itemizing its parts and their functions, enumerating its possible outputs for each of its possible inputs, and then building a well-oiled and wholly deterministic machine that will reproduce not its actual behavior but the behavior it is "supposed" to exhibit is delusional. (But note that this fantasy is also shared to some extent by true believers in behaviorism and psychoanalysis.)

Still this doesn't mean you can't design and build (or at least grow) something like a human brain. In fact you should be able to model it efficiently in (nondeterministic) software; and then improve upon it. It's just that this doesn't imply that you'll know how it works.

There is already at least one guy I've heard of who builds neural network chips that evolve their own programming to solve the specific problems he assigns them using genetic algorithms, and then, like any other biologist, spends most of his time trying to figure out how the hell they do it. — In a nutshell, this is the future of computer programming.

In fact it's not even obvious that you'd know when you'd done it. Gibson's original prediction was that artificial intelligences would emerge more or less by accident on the global network, and that it would take, as it were, the best part of three novels to figure out exactly what had happened. This still seems plausible.

Having made this fairly overwhelming case against determinism, I should admit that it's possible that strict algorithms might be able to run in polynomial time on quantum computers. — That is, it's possible that quantum computation is inherently more powerful than classical computation; and that, perhaps, in this model P equals NP , Church's thesis is modified, combinatorial optimization problems can be solved exactly within reasonable time limits, and Napoleon might be able to climb as high a hill as he likes. As evidence in favor of this conjecture you can wave your hands generally at the native parallelism of quantum-mechanical dynamical evolution (the particle doesn't traverse the classical path of least action but all paths simultaneously with an amplitude whose phase is proportional to that action) and specifically at Shor's quantal algorithm for factoring integers.⁸ As evidence against it you can wave your hands at the apparent necessity of preserving reversibility (aka unitarity) in quantum computations,

⁸ Which, however, isn't exponentially faster than the "classical" algorithms whose slowness ensure the security of RSA encryption.

which seems to indicate that the previous handwaving argument linking the polynomial/exponential distinction to the irreversibility of time wasn't complete bullshit; and guess that the size of a reversible quantum computer might have to grow exponentially with the size of the problem.

Penrose, one should note, did half-seriously propose that the brain is inherently quantum-mechanical in its operation (as did Eddington⁹ long before him, actually, though no one seems to recall it); however his arguments weren't very convincing, and anyway the point seems to be that quantum mechanics doesn't banish determinism but rather reinstates it. But I expect the correct analysis of quantum computation will show the way to solve the P/NP problem; which is to say, that none of this will really be understood until the theorem has been proven. Meanwhile I'll stick with my conclusions.

Anyway. The assumption — of Minsky, of Frankenstein, of Rotwang — and of Joy, were he smart enough to be that interesting — is that in the process of replacing all the messy wetware of the human brain with the rule-based systems that it imperfectly attempts to implement, the riddle of the nature of intelligence will be laid bare to a select few, the dudes on the white horses on the high hill, the philosopher-kings, the hacker Napoleons; and that these guys will have it within their power, as it were by construction (as a mathematician would put it) to control the behavior of their creations. (And the behavior of all those messy imperfect humans as well; but let's not go there. — Yet.) But the artificial brains won't work that way, the riddle is though transparent nonetheless impenetrable, and control as always is an illusion. — You can fantasize yourself the master puppeteer; but no one, not even Cusack,¹⁰ can pluck a trillion strings.

⁹ In *The Nature of the Physical World* [Gifford Lectures, 1927].

¹⁰ Cf. *Being John Malkovich* [Spike Jonze, 1999].

Thus though machine intelligence has no apparent bounds, hacker intelligence has very obvious limits. And though it will not be that long before it's possible to fabricate a machine smarter than Von Neumann, it won't be a Von Neumann machine; will not be programmed since not programmable; and its behavior, being more complex than that of its creators, will be even more impossible to predict in advance.

(Von Neumann himself was, alas, rather easily impressed by bozos in uniform; robots, being really alien, will be harder not easier to tame.)

So the fantasies of control are delusional. What emerges will be something inherently uncontrollable: insofar as they know what it's doing, it won't work, and insofar as they don't, it will. It won't even be possible to direct the research in such a way as to avoid unintended consequences. And nobody gets to play the short dead dude.

In fact that's probably the best summary: nobody gets to play Napoleon; as usual, everybody gets to play Bill and Ted.

As for the second menace: I don't know much of the literature. In principle it's possible to manipulate things at the atomic level (not that we don't do this already, cf. chemistry), and this may have applications to the fabrication of materials; fine. What might be construed as controversial or alarming is the idea of (physically-realized) artificial life. As best I can determine from flipping through the papers in a couple of volumes of the Santa Fe proceedings, the people proposing such experiments, Drexler, for example, are not entirely stupid; but nowhere near as smart as, say, Oppenheimer and Bethe, whose minor oversights nearly precipitated the end of civilization.

For instance Drexler seems to see the necessity of drawing a clear conceptual distinction between nanomachinery designed for some specific (and necessarily very narrow) purpose and, say, uncontrollable synthetic viruses that will breed until they devour the world; but doesn't really succeed in doing it.¹¹ This is not encouraging.

The problem with infinitesimal machines, obviously, is that (contra *Fantastic Voyage*) they aren't good for much if manufactured one at a time and operated singly by human telepresence; they have to be autonomous, and probably they have to be self-reproducing or at least self-modifying. (They'd have to be hand-tooled for any given site or application, and you'd need trillions for every job. What kind of assembly line could turn them out?) — Skipping a few steps, what this means, mathematically, is that you have to imagine some kind of generative grammar (in the sense of Lindenmayer's variation on Chomsky) which produces the little suckers as the endproduct of a recursive development process using a fixed set of rules from an initial quasigenomic string of specifications (no accident this sounds like morphogenesis); if the end result is supposed to be nontrivial, these have the complexity of computer programs, and the outputs of computer programs are (see above) wholly unpredictable. Protests that the programs can be debugged before the critters are set loose are (I say yet again) simply fatuous. (E.g., the problem of predicting what language is generated by a general phrase-structure grammar — literally: whether or not the language is nonvoid — is essentially the same as the halting problem for Turing machines; not simply NP-hard, in other words, but recursively unsolvable.) Moreover if the point is to

¹¹ This is the howler: "Because nanoreplicators will differ fundamentally from biological systems, there is reason to believe that novel and remarkably dangerous systems could be constructed — but are they likely to appear by accident? Several facts make such accidents easy to avoid and difficult to cause. The most obvious and least fundamental of these is that, since these systems will be designed, their parts and structures will be known; moreover...the relationships among their parts will be designed and fixed." — Apparently Drexler never programmed a computer.

manufacture these machines for applications which involve their interaction with real living things, e.g. scrubbing sclerotic arteries or killing tumor cells, their programming will have to be extremely flexible: the most obvious prototypes of such gadgets are the antibodies of the immune system, and these adapt to attack foreign intruders by mutating at a prodigious rate until they develop binding sites specific to the alien objects they want to recognize and negate.

I suppose there's some fantasy entertained about directing this army of little boogers by remote control; but this is just Napoleon again, obviously — and the smaller the critters are, the more literally Napoleon and Maxwell's Demon look alike. So forget that.

But to restate the most obvious objection yet again: you well remember your first happy adventures in computer programming — that harmless-looking little procedure that crashed the operating system, the input-output routine that wrote zeroes throughout core and copied itself in fragments all over the disk drive, those entertaining embarrassments that recalled to you Mickey's exploits as the Sorcerer's Apprentice in *Fantasia*. — So who's backed up the biosphere? because some imbecile will certainly find a way to erase it, if this line of research is pursued.

In conclusion, though some kind of miscegenation between the organic and the not-yet-organic is certain to occur, I doubt it will take the form of reinventing bacteria and letting them eat us. Who could be that stupid?

(Don't answer that.)

Let's pause to state the computational challenge problem which, if we translate it back into existing biology, this line of thought suggests: given the genome for an unknown organism, to generate a (complete)

simulation of it. — And, so long as we're imagining impossibilities, the inverse problem: from a (necessarily incomplete) description of an organism, to produce the code that will generate it. (Suppose, e.g., that you wanted to make dinosaurs, but you couldn't find Crichton's mosquitoes trapped in amber.) — The forward problem (suitably constrained) can't be impossible, even though it properly contains relative trivialities like protein folding; you can always grow the organism (i.e., Nature can perform the computation.) But the inverse might be. (Again, this is the difference between following a proof and finding one; in this case, the "proof" that a given organism can be obtained by morphogenesis.)

(I'm haunted by an idea I first found in an old story by Keith Laumer, *Worlds of the Imperium*, which described a magic television that allowed you to channelsurf (as it were) sideways in time; he illustrated this with a vivid description of a scene of a farmer plowing a field behind a couple of oxen which morphed repeatedly as the viewers tuned their way from the original settings through a series of variations in which the animals changed into alien forms, the farmer's skin turned purple, he grew more fingers and extruded antennae from his forehead, the landscape rippled into hills, the sun and the sky changed color, etc., etc. I remembered this not long ago when I was trying to figure out some way to steer the changes in a generic picture of a face: if you could put the right controls on the software, you could, maybe, produce some kind of generalized police-artist that would allow you to draw people from memory — or, maybe more interesting, make them up. (Actually I think the cops do have something like this now, but I don't know how it works, or how well.) — Similarly you might imagine a tuner that would steer you through possible variations on a genome and generate simulations of the corresponding creatures for you in real time. — The problem, in both cases, is trying to figure out some happy mean between your naive vision of a single tuning knob and the nasty reality, which is that you're trying to navigate your way

around a Hilbert space, and need an infinite number of them. Or at least a few billion.)¹²

As for the third menace, the reprogramming of the genome, this will certainly happen. Since it will probably start happening in the very near future, people are already talking (in the pages of *Business Week*!) about legal restrictions, etc. — as if there were some distinction you could clearly enunciate in a courtroom between eliminating genetic defects and introducing genetic assets. Even if there were legal restrictions, they would not apply universally (I'm happy that we do for the moment run the world, but there is an alarming hubris in this unspoken assumption that American law is the only law) — offshore, in Asia, if need be on the Moon — and the competitive advantages of ignoring these constraints will be so enormous that a way will be found to circumvent them.

For instance at the first moment that someone finds it advantageous to breed supermen for specialized purposes, or simply to clone the most obvious candidates, it will happen. — In the Gibsonian scenarios (and Gibson is our surest guide here) the ruling classes end up owning them, and there's certainly ample historical precedent: Leibniz and Bach and Euler were all kept as pets.

On the other hand the lead time for delivery of a crop of Ed Wittens is at least twenty years, and a lot can happen in that time.

At the other end of the food chain, it's hard enough already to figure out what separates men from apes; very modest augmentations would suffice to turn chimpanzees into a new servant class. Try to stop that.

¹² Note added later: the underlying skeletal model in *Shrek* is supposed to have five hundred forty joints; so in some sense you can see from the verisimilitude of the animation that a few hundred knobs — the first few hundred dimensions of the Hilbert space — would suffice, insofar as reproducing human motion is the object. I'm still not sure about faces.

(A career at the butt end of the service industries has taught me the great truth that makes the world go round: people are cheaper than machines. — And if people were only cheaper, it would all go round that much faster. Just ask the editors of the *Wall Street Journal*.)

In general it's so easy to think of mechanisms that would accentuate the internal differentiation of the species that you can't imagine that it won't happen; more, you have to suspect there's some kind of principle that dictates this out of natural necessity.

But the motives that drive the initial applications will likely be more prosaic. I remember having being struck by Gibson's observation (now a couple of decades old) that whole gangs of disaffected youth could decide to look like James Dean; and thinking that this had the stink of the truth about it. — Indeed a sort of improved cosmetic surgery is a natural first step. But where does this stop?

None of this is at all new; it's just that it's finally within reach. Most of the obvious points were made by J. D. Bernal in the Twenties in a speculative essay entitled *The World, The Flesh, and The Devil*; a work which I looked up recently and read again, just to confirm the suspicion that all this had been foretold long ago.

Bernal was no mean stylist. The first sentence is memorable: "There are two futures, the future of desire and the future of fate, and man's reason has never learnt to separate them."¹³ — So much, I want to say, for science fiction. — But, skipping over a number of interesting speculations about materials science, the first descriptions of photon sails, space stations, solar panels, etc., etc., and cutting to the

¹³ Of this Freeman Dyson said "I do not know of any finer opening sentence of a work of literature in English," and arguing counterexamples — "Call me Ishmael." — "A screaming comes across the sky." — "The sky was the color of television, tuned to a dead channel." — etc., etc., — makes for great cocktail party sport.

biological chase, you find the piquant summary: “It is quite conceivable that the mechanism of evolution, as we know it up to the present, may well be superseded ... after all it is only nature’s way of achieving a shifting equilibrium with an environment; and if we can find a more direct way by the use of intelligence, that way is bound to supersede the unconscious mechanism of growth and reproduction.” — Pointing out that this began, in effect, with the invention of tools, he continues: “Normal man is an evolutionary dead end; mechanical man, apparently a break in organic evolution, is actually more in the true tradition of a further evolution.” — Though Bernal’s explicit vision of this is a trifle old-fashioned, a sort of brain-in-a-barrel idea (reading this over again I realized this was the origin of the old classic Thirties scifi story *Professor Jameson’s Satellite*, about a scholar who dies, has his body put into orbit, and then wakes up a couple of million years later when a party of exploring robot-dudes find the orbiting casket and transplant his brain to one of their bodies — presumably the origin of the fantasy which apparently governs the decision by the Alcor-bracelet dudes to have their brains frozen when they die and stored in Scottsdale), his idea is, as it were, mechanism-independent, and it’s obviously correct.

Bernal must obviously have had a direct literary influence on Huxley; presumably it extended farther. This sounds, actually, like another urtext of cyberpunk, which is in large part the elaboration and development of the theme of the interpenetration of the animate and the inanimate; regarded variously with paranoiac alarm (in Pynchon), fascination (in Gibson), or with a sort of grisly playfulness (in Cronenberg).

Bernal concludes that a class division will appear in humanity depending on whether they do or do not embrace these changes; and suggests that the ones who do will be the more intelligent, adventurous, etc., and that they’ll probably end up living somewhere off the planet. — I too expect that this will happen, but expect instead

that the division will fall along pre-existing class lines: i.e., the ones who will have the money then to pay for cosmetic alterations, to have their children augmented, and to remove themselves from the reach of terrestrial law will be essentially the same ones who have the money now for cosmetic surgery, to place their children in private schools, and to buy their way out of murder raps; but as more power is concentrated in fewer hands, and the already considerable competitive advantages of the wealthy over the disadvantaged become qualitative biological differences, the gap will become a phase boundary. — It's a kind of Scott Fitzgerald joke now, when people refer to the very rich as a different species; but presently this will literally become the case.

And then they'll all go to war and smoke all the rest of us trying to get at one another. (Like 1914, only worse.) Not exactly a cheerful prospect.

But probably the machines get smart and take over first.

Let's make a parenthetical note that, if it should prove economical (in the broadest sense) for humans (and not intelligent robots designed specifically for the purpose) to do things like colonize Mars and mine the asteroids, they will certainly be drastically modified and thoroughly re-engineered humans. (But, pace Dyson, it seems as if a robot would be a better idea all around: something that could live in vacuum directly off sunlight.) And of course it's been obvious for a long time that interstellar flight is completely impractical for organic lifeforms; a million-year lifespan would be the minimum requirement. I think you need something stabler than DNA to pull that off.

So, inverting the argument, if we were to be visited by extraterrestrials, it's unlikely we'd recognize them as organic lifeforms. The little green men would be wholly superfluous; the flying saucers themselves would be the aliens. — Or worse, something like Hoyle's

Black Cloud, some nanotechnological virus descending in a swarm to devour the Earth. — It never ceases to amaze me what a fucking Pollyanna Carl Sagan was.)=

With regard to the other parlor stunts, e.g. cloning, it seems inevitable that some Gatesian megalomaniac will try it; the temptation to hand over the empire of Microsoft, for instance, not simply to a designated heir, but to one's self, would certainly be irresistible. — Think of it this way: this is what the evil mastermind in a James Bond movie would do; and there are now any number of overnight billionaires who would like to think they're the kind of guys for whom Doctor No is the only appropriate role model. — But in view of what's now possible, the idea already sounds retro.

In summary:

Wittgenstein said once that a philosophical work could be written entirely in the form of a series of jokes; and it must have occurred to you, in your meditations on the nature of comedy, that not only are there jokes which express something really deep, but that you can't imagine any other way of stating what the joke expresses that doesn't destroy its meaning in the translation; even if you knew what it was. (Donald Richie on Kurosawa: "While quite ready to talk about lenses, or acting, or the best kind of camera-dolly, he is unwilling to discuss meaning or aesthetics. Once I asked what a certain scene was really about. He smiled and said: 'Well, if I could answer that, it wouldn't have been necessary for me to have filmed the scene, would it?'") — So it's the point, finally, that Douglas Adams' gag about the purpose of life on earth — that it's all some kind of enormous computation, by some kind of organic massively-parallel computer, all run to find the answer to the question (or the question to the answer) of life, the universe, and everything — is exactly right; it is exactly that. (And the

idea that the output might be trivial could be the best part.) — Obviously we don't understand what the computation is, or what it's for, or whether calling it a computation is really the best way to look at it, or whether we are even allowed to figure any of this out (but, dig we must) — and that's why the most direct statement you can make about it takes the form of a joke. But it is something like this. And that's why it's a great joke.

But then it seems obvious that (as they say) the software is independent of the hardware it's running on (this seems to be the engineer's version of Platonism: platform-independence), indeed that it's continuously redesigning the hardware it's running on; and that if there's some radically different direction it can take to continue the optimization of whatever metaphysical function it's trying to maximize, it will certainly take it.

In fact arguably it's changed platforms before: there's some entertaining science fiction about templates stored in clays preceding organic life, and there was almost certainly an RNA world which preceded the DNA/RNA/protein era governed by the central dogma.

Nor does it matter a great deal whether we (as individuals, or even collectively) can predict what direction it will take, or figure out in more than the vaguest terms what it is doing; we can understand this well enough to see that it is by definition something that we cannot design, harness, or control. Not all the King's horses nor all the King's men could reassemble Humpty Dumpty; and something here is being put together that no one can take apart. (There's a missing principle, something that's been disguised by the way we usually look at statistical mechanics: not simply the entropic but also the organizing principles are irreversible and inexorable.)

When I write this out and look at it, it doesn't seem terribly original:

“evolutionary” philosophies are old and fairly lame; none of this is a whole lot different from what Bergson said, or Whitehead (when that mood was on him), or even a hack like Herbert Spencer; and all of that derives (cf. Arthur Lovejoy) from temporalizing the idea of the Chain of Being, which dates from the philosophical stone age. The difference is that those guys weren’t in the position of swimming in the shallow water, watching the amphibians march away onto the land.

As usual, Nietzsche understood it better than anybody else: it’s less the idea of some sort of life force or *elan vital* than an abstract will-to-power: the dual to entropy; a force with the properties of an ineluctable necessity or an immutable Fate. — The story (as it were) is not about us, but about that; it is silly to suppose that we can rewrite it.

But, look on the bright side: the *Übermenschen* would have killed us off as a minor corollary of their competition; at worst our machines will end up keeping us as pets.

.....

A couple of afterthoughts:

Regarding the automatic generation of theory, etcetera: the ancestry of the idea actually extends farther into the past than Leibniz: I may still have a little volume by Martin Gardner on *Logic Machines* which traces the notion of the grand *Ars Combinatoria* back to medieval times;¹⁴ I seem to recall the name of Ramon Lull, for instance. — It keeps getting reinvented. E.g., writing on the semiotics of the cinema,

¹⁴ Martin Gardner, *Logic Machine and Diagrams*. New York: McGraw-Hill, 1958.

Christian Metz refers to someone's proposal of "a 'permutational art' in which poetry, discarding the chaste mystery of inspiration, will openly reveal the portion of manipulation it has always contained, and will finally address itself to computers...The 'poet' would program the machine, giving it a certain number of elements and setting limitations; the machine would then explore all the possible combinations, and the author would, at the end of the process, make his selection." — And compare Swift, of course, when Gulliver visits Laputa; etc., etc.

There's another monograph here. — Of course it never seems to occur to anyone just how many "all the possible combinations" are. (Except Umberto Eco, I now recall; who traces this idea back to some tradition about the Torah of improbable antiquity.)¹⁵

(I once had an elaborate scheme for musical composition by automata that sounded like this "computational poetry". Fortunately that was before I had computers.)

In re "evolutionary philosophies": I happened across a copy of Dawkins' *The Selfish Gene* in a used bookstore a while ago and (for the first time) read it; what he means by a "gene" is difficult to figure out, but he probably intends to promote some kind of Platonism that reifies chunks of biological programming. If you ignore the propagandizing, this is more or less correct. — Some decent ideas in this, but not enough to explain why he got to marry the blonde Romana.¹⁶

And, incidentally: though it is presumed that there is no (classical) polynomial-time algorithm for finding the factors of an integer, from a plausible hypothesis you can show the validity of a (nonconstructive)

¹⁵ Cf. *Foucault's Pendulum*. Eco himself had obviously played with programming an early personal computer.

¹⁶ The inimitable Lalla Ward.

test for compositiveness that runs in something like quintic time. — The “plausible hypothesis” is a generalized Riemann hypothesis, which gives you some idea as to the difficulty of this kind of question.¹⁷ — In general the formidability of the mathematical machinery that has to be brought to bear to prove even the simplest propositions about running times is daunting. I sat through a three hour talk by Smale not all that long ago, into which he inserted progressively more and more impressive apparatus until finally invoking results about the cohomology of the braid group (I burst out laughing in the middle of the lecture — I mean, how very) to establish some relatively trivial bounds on the difficulty of locating the zeroes of polynomials. — The mathematical theory of algorithms is largely terra incognita.

Later.

¹⁷ A polynomial-time test for primality was later found that did not require GRH.

{...}

Addenda

Jonathan Swift, from *A Voyage to Laputa*:

The first Professor I saw was in a very large Room, with forty Pupils about him. After Salutation, observing me to look earnestly upon a Frame, which took up the greatest part of both the Length and Breadth of the Room, he said perhaps I might wonder to see him employed in a Project for improving speculative Knowledge by practical and mechanical Operations. But the World would soon be sensible of its Usefulness, and he flattered himself that a more noble exalted Thought never sprung in any other Man's Head. Every one knew how laborious the usual Method is of attaining to Arts and Sciences; whereas by his Contrivance, the most ignorant Person at a reasonable Charge, and with a little bodily Labour, may write Books in Philosophy, Poetry, Politicks, Law, Mathematicks and Theology, without the least Assistance from Genius or Study. He then led me to the Frame, about the Sides whereof all his Pupils stood in Ranks. It was twenty Foot Square, placed in the middle of the Room. The Superficies was composed of several bits of Wood, about the bigness of a Dye, but some larger than others. They were all linked together by slender Wires. These bits of Wood were covered on every Square with Paper pasted on them, and on these Papers were written all the Words of their Language, in their several Moods, Tenses, and Declensions, but without any Order. The Professor then desired me to observe, for he was going to set his Engine at Work. The Pupils at his Command took each of them hold of an Iron Handle, whereof there were forty fixed round the Edges of the Frame, and giving them a sudden turn, the whole Disposition of the Words was entirely changed. He then commanded six and thirty of the Lads to read the several Lines softly as they appeared upon the Frame; and

where they found three or four Words together that might make part of a Sentence, they dictated to the four remaining Boys who were Scribes. This Work was repeated three or four Times, and at every turn the Engine was so contrived that the Words shifted into new Places, as the Square bits of Wood moved upside down.

(A plate is inserted here illustrating “The Literary Engine”.)

Six Hours a-day the young Students were employed in this Labour, and the Professor shewed me several Volumes in large Folio already collected, of broken Sentences, which he intended to piece together, and out of those rich Materials to give the World a compleat Body of all Arts and Sciences; which however might be still improved, and much expedited, if the Publick would raise a Fund for making and employing five hundred such Frames in Lagado, and oblige the Managers to contribute in common their several Collections.

He assured me, that this Invention had employed all his Thoughts from his Youth, that he had emptied the whole Vocabulary into his Frame, and made the strictest Computation of the general Proportion there is in Books between the Numbers of Particles, Nouns, and Verbs, and other Parts of Speech.

Compare John Stuart Mill’s account of his depression:

After the tide had turned, and I was in process of recovery, I had been helped forward by music I at this time first became acquainted with Weber’s *Oberon*, and the extreme pleasure which I drew from its delicious melodies did me good by showing me a source of pleasure to which I was as susceptible as ever. The good, however, was much impaired by the thought that the pleasure of music (as is quite true of such pleasure as this was,

that of mere tune) fades with familiarity, and requires either to be revived by intermittence, or fed by continual novelty. And it is very characteristic both of my then state, and of the general tone of my mind at this period of my life, that I was seriously tormented by the thought of the exhaustibility of musical combinations. The octave consists only of five tones and two semi-tones, which can be put together in only a limited number of ways, of which but a small proportion are beautiful: most of these, it seemed to me, must have been already discovered, and there could not be room for a long succession of Mozarts and Webers, to strike out, as these had done, entirely new and surpassingly rich veins of musical beauty.

{...}

Gibbon:

The philosophic spirit consists in being able to go back to basic ideas; to perceive and to bring together first principles. The view of the philosopher is exact, but at the same time extensive. Placed on a height, he comprehends a great expanse of territory, of which he forms a clear and unique image, while other minds just as exact, but more confined, discern only a part of the expanse.¹⁸

¹⁸ Translated from the French, actually; see the introduction, page xxi.

{...}

The Napoleonic fallacy (continued)

By way of eccentric illustration: even in the primitive media environment of the turn of the (last) century the success of Wells' *The War of the Worlds* spelled franchise, and an American newspaper commissioned the astronomer Garrett P. Serviss to write a sequel for serialization. The result was *Edison's Conquest of Mars* [1898], a remarkable fantasy in which the legendary inventor, assisted by assorted famous scientists of the day, reverse-engineers the wreckage left by the invasion, designs a fleet of electric spaceships armed with really cool-sounding ray guns, and sets off across interplanetary space to even the score. After the element of surprise is lost, the war isn't going all that well, but at a critical juncture Edison and his posse are advised of the existence of a central power house which controls the system of locks and gates for the famous Martian canals,¹⁹ and, staging a stealth assault by night, they enter it, intending to flood the planet. Overcoming the guards, they find the control room:

The chamber was an immense square room at least a hundred feet in height and 400 feet on a side, and almost filling the wall opposite to us was an intricate display of machinery, wheels, levers, rods and polished plates. This we had no doubt was one end of the engine which opened and shut the great gates that could dam an ocean. ... There were scores of knobs and handles, all glistening in the electric light, any one of which, so far as the uninstructed could tell, might have been the master key that controlled the whole complex apparatus.

¹⁹ The single greatest moment of romantic disillusionment in living memory was that which attended the reception of the first photographs of the Martian surface by Mariner IV in July, 1965; rather than possessing a landscape covered by vegetation, irrigated by a network of waterways natural or artificial, Mars was revealed to be a dead world, cratered like the Moon. In an instant several generations of speculative fiction went the way of Prospero's cloud-capped towers; leaving, alas, not a rack behind.

“Quick,” said Mr. Edison, “where is it?”

But their native guide is clueless. What to do?

It was at this critical moment that the wonderful depth and reach of Mr. Edison’s mechanical genius displayed itself. He stepped back, ran his eyes quickly over the whole immense mass of wheels, handles, bolts, bars and levers, paused for an instant, as if making up his mind, then said decidedly, “There it is,” and stepping quickly forward, selected a small wheel amid a dozen others, all furnished at the circumference with handles like those of a pilot’s wheel, and giving it a quick wrench, turned it half-way around.

— with which the planet is flooded and, as the title of the ensuing chapter puts it, “Vengeance is ours” — but — there is no better description of the swift commanding glance than this.