



Search



[< DNA Politics: Science and Sociopathy](#)

[Beyond Otaku: Transhumanism and Judaism >](#)

Does Humanity Need an AI Nanny?

August 17, 2011 · AI, Politics, Singularity



■ authors Ben Goertzel

The ongoing advancement of science and technology has brought us many wonderful things, and

will almost surely be bringing us more and more – most likely at an exponential acceleration, as Ray Kurzweil and others have argued. Beyond the “mere” abolition of scarcity, disease and death, there is the possibility of fundamental enhancement of the human mind and condition, and the creation of new forms of life and intelligence. Our minds and their creations may spread throughout the universe, and may come into contact with new forms of matter and intelligence that we can now barely imagine.



But, as we all know from SF books and movies, the potential dark side of this advancement is equally dramatic. Nick Bostrom **has enumerated** some of the ways that technology may pose “existential risks” – risks to the future of the human race – as the next decades and centuries unfold. And there is also rich potential for other, less extreme sorts of damage. Technologies like AI, synthetic biology and nanotechnology could run amok in dangerous and unpredictable ways, or could be utilized by unethical human actors for predictably selfish and harmful human ends.

The Singularity, or something like it, is probably near – and the outcome is radically uncertain in almost every way. How can we, as a culture and a species, deal with this situation? One possible solution is to build a powerful yet limited AGI (Artificial General Intelligence) system, with the explicit goal of keeping things on the planet under control while we figure out the hard problem of how to create a probably positive Singularity. That is: to create an “AI Nanny.”

The AI Nanny would forestall a full-on Singularity for a while, restraining it into what Max More has called a **Surge**, and giving us time to figure out what kind of Singularity we really want to build and how. It’s not entirely clear that creating such an AI Nanny is plausible, but I’ve come to the conclusion it probably is. Whether or not we should *try* to create it – that is the Zillion Dollar Question.

The Gurus’ Solutions

What does our pantheon of futurist gurus think we should do in the next decades, as the path to Singularity unfolds?

Kurzweil has proposed “fine-grained relinquishment” as a strategy for balancing the risks and rewards of technological advancement. But it’s not at all clear this will be viable, without some form of AI Nanny to guide and enforce the relinquishment. Government regulatory agencies are notoriously slow-paced and unsophisticated, and so far their decision-making speed and intelligence aren’t keeping up with the exponential acceleration of technology.

Further, it seems a clear trend that as technology advances, it is possible for people to create more and more destruction using less and less money, education and intelligence. There seems no

reason to assume this trend will reverse, halt or slow. This suggests that, as technology advances, selective relinquishment will prove more and more difficult to enforce. Kurzweil acknowledges this issue, stating that “The most challenging issue to resolve is the granularity of relinquishment that is both feasible and desirable” (p. 299, [The Singularity Is Near](#)), but he believes this issue is resolvable. I’m skeptical that it is resolvable without resorting to some form of AI Nanny.

[Eliezer Yudkowsky](#) has suggested that the safest path for humanity will be to first develop “Friendly AI” systems with dramatically superhuman intelligence. He has put forth some radical proposals, such as the design of self-modifying AI systems with human-friendly goal systems designed to preserve friendliness under repeated self-modification; and the creation of a specialized AI system with the goal of determining an appropriate integrated value system for humanity, [summarizing in a special way the values and aspirations of all human beings](#). However, these proposals are extremely speculative at present, even compared to feats like creating an AI Nanny or a technological Singularity. The practical realization of his ideas seems likely to require astounding breakthroughs in mathematics and science – whereas it seems plausible that human-level AI, molecular assemblers and the synthesis of novel organisms can be achieved via a series of moderate-level breakthroughs alternating with “normal science and engineering.”

Bill McKibben, Bill Joy and other modern-day techno-pessimists argue for a much less selective relinquishment than Kurzweil (e.g. Joy’s classic Wired article [The Future Doesn’t Need Us](#)). They argue, in essence, that technology has gone far enough – and that if it goes much further, we humans are bound to be obsoleted or destroyed. They fall short, however, in the area of suggestions for practical implementation. The power structure of the current human world comprises a complex collection of interlocking powerful actors (states and multinational corporations, for example), and it seems probable that if some of these chose to severely curtail technology development, many others would NOT follow suit. For instance, if the US stopped developing AI, synthetic biology and nanotech next year, China and Russia would most likely interpret this as a fantastic economic and political opportunity, rather than as an example to be imitated.

My good friend Hugo de Garis agrees with the techno-pessimists that AI and other advanced technology is likely to obsolete humanity, but views this as essentially inevitable, and encourages us to adopt a philosophical position according to which this is desirable. In his book [The Artilect War](#), he contrasts the “Terran” view, which views humanity’s continued existence as all-important, with the “Cosmist” view in which, if our AI successors are more intelligent, more creative, and perhaps even more conscious and more ethical and loving than we are – then why should we regret their ascension, and our disappearance? In more recent writings (e.g. the article [Merge or Purge](#)), he also considers a “Cyborgist” view in which gradual fusion of humans with their technology (e.g. via mind uploading and brain computer interfacing) renders the Terran vs. Cosmist dichotomy irrelevant. In this trichotomy Kurzweil falls most closely into the Cyborgist

camp. But de Garis views Cyborgism as largely delusory, pointing out that the potential computational capability of a grain of sand (according to the known laws of physics) exceeds the current computational power of the human race by many orders of magnitude, so that as AI software and hardware advancement accelerate, the human portion of a human-machine hybrid mind would rapidly become irrelevant.

Humanity's Dilemma

And so ... the dilemma posed by the rapid advancement of technology is both clear and acute. If the exponential advancement highlighted by Kurzweil continues apace, as seems likely though not certain, then the outcome is highly unpredictable. It could be bliss for all, or unspeakable destruction – or something inbetween. We could all wind up dead — killed by software, wetware or nanoware bugs, or other unforeseen phenomena. If humanity does vanish, it could be replaced by radically more intelligent entities (thus satisfying de Garis's Cosmist aesthetic) – but this isn't guaranteed; there's also the possibility that things go awry in a manner annihilating all life and intelligence on Earth and leaving no path for its resurrection or replacement.

To make the dilemma more palpable, think about what a few hundred brilliant, disaffected young nerds with scientific training could do, if they teamed up with terrorists who wanted to bring down modern civilization and commit mass murders. It's not obvious why such an alliance would arise, but nor is it beyond the pale. Think about what such an alliance could do now – and what it could do in a couple decades from now, assuming Kurzweilian exponential advance. One expects this theme to be explored richly in science fiction novels and cinema in coming years.

But how can we decrease these risks? It's fun to muse about designing a "Friendly AI" à la Yudkowsky, that is guaranteed (or near-guaranteed) to maintain a friendly ethical system as it self-modifies and self-improves itself to massively superhuman intelligence. Such an AI system, if it existed, could bring about a full-on Singularity in a way that would respect human values – i.e. the best of both worlds, satisfying all but the most extreme of both the Cosmists and the Terrans. But the catch is, nobody has any idea how to do such a thing, and it seems well beyond the scope of current or near-future science and engineering.

Realistically, we can't stop technology from developing; and we can't control its risks very well, as it develops. And daydreams aside, we don't know how to create a massively superhuman supertechnology that will solve all our problems in a universally satisfying way.

So what do we do?

Gradually and reluctantly, I've been moving toward the opinion that the best solution may be to create a *mildly* superhuman supertechnology, whose job it is to protect us from ourselves and our technology – not forever, but just for a while, while we work on the hard problem of creating a Friendly Singularity.

In other words, some sort of *AI Nanny*...

The AI Nanny

Imagine an advanced Artificial General Intelligence (AGI) software program with

- General intelligence somewhat above the human level, but not too dramatically so – maybe, qualitatively speaking, as far above humans as humans are above apes
- Interconnection to powerful worldwide surveillance systems, online and in the physical world
- Control of a massive contingent of robots (e.g. service robots, teacher robots, etc.) and connectivity to the world’s home and building automation systems, robot factories, self-driving cars, and so on and so forth
- A cognitive architecture featuring an explicit set of goals, and an action selection system that causes it to choose those actions that it rationally calculates will best help it achieve those goals
- A set of preprogrammed goals including the following aspects:
 - A strong inhibition against modifying its preprogrammed goals
 - A strong inhibition against rapidly modifying its general intelligence
 - A mandate to cede control of the world to a more intelligent AI within 200 years
 - A mandate to help abolish human disease, involuntary human death, and the practical scarcity of common humanly-useful resources like food, water, housing, computers, etc.
 - A mandate to prevent the development of technologies that would threaten its ability to carry out its other goals
 - A strong inhibition against carrying out actions with a result that a strong majority of humans would oppose, if they knew about the action in advance
 - A mandate to be open-minded toward suggestions by intelligent, thoughtful humans about the possibility that it may be misinterpreting its initial, preprogrammed goals

This, roughly speaking, is what I mean by an “AI Nanny.”

Obviously, this sketch of the AI Nanny idea is highly simplified and idealized – a real-world AI Nanny would have all sort of properties not described here, and might be missing some of the above features, substituting them with other related things. My point here is not to sketch a specific design or requirements specification for an AI Nanny, but rather to indicate a fairly general class of systems that humanity might build.

The nanny metaphor is chosen carefully. A nanny watches over children while they grow up, and

then goes away. Similarly, the AI Nanny would not be intended to rule humanity on a permanent basis – only to provide protection and oversight while we “grow up” collectively; to give us a little breathing room so we can figure out how best to create a desirable sort of Singularity.

A large part of my personality rebels against the whole AI Nanny approach – I’m a rebel and a nonconformist; I hate bosses and bureaucracies and anything else that restricts my freedom. But, I’m not a political anarchist – because I have a strong suspicion that if governments were removed, the world would become a lot worse off, dominated by gangs of armed thugs imposing even less pleasant forms of control than those exercised by the US Army and the CCP and so forth. I’m sure government could be done a lot better than any country currently does it – but I don’t doubt the need for some kind of government, given the realities of human nature. And I think the need for an AI Nanny falls into the same broad category. Like government, an AI Nanny is a relatively offensive thing, that is nonetheless a practical necessity due to the unsavory aspects of human nature.

We didn’t need government during the Stone Age – because there weren’t that many of us, and we didn’t have so many dangerous technologies. But we need government now. Fortunately, these same technologies that necessitated government, also provided the means for government to operate.

Somewhat similarly, we haven’t needed an AI Nanny so far, because we haven’t had sufficiently powerful and destructive technologies. And fortunately, these same technologies that apparently necessitate the creation of an AI Nanny, also appear to provide the means of creating it.

The Basic Argument

To recap and summarize, the basic argument for trying to build an AI Nanny is founded on the premises that:

1. It’s impracticable to halt the exponential advancement of technology (even if one wanted to)
2. As technology advances, it becomes possible for individuals or groups to wreak greater and greater damage using less and less intelligence and resources
3. As technology advances, humans will more and more acutely lack the capability to monitor global technology development and forestall radically dangerous technology-enabled events
4. Creating an AI Nanny is a significantly less difficult technological problem than creating an AI or other technology with a predictably high probability of launching a full-scale positive Singularity

5. Imposing a permanent or very long term constraint on the development of new technologies is undesirable

The fifth and final premise is normative; the others are empirical. None of the empirical premises are certain, but all seem likely to me. The first three premises are strongly implied by recent social and technological trends. The fourth premise seems commonsensical based on current science, mathematics and engineering.

These premises lead to the conclusion that trying to build an AI Nanny is probably a good idea. The actual plausibility of building an AI Nanny is a different matter – I believe it is plausible, but of course, opinions on the plausibility of building any kind of AGI system in the relatively near future vary all over the map.

Complaints and Responses

I have discussed the AI Nanny idea with a variety of people over the last year or so, and have heard an abundance of different complaints about it – but none have struck me as compelling.

“It’s impossible to build an AI Nanny; the AI R&D is too hard.” – But is it really? It’s almost surely impossible to build and install an AI Nanny this year; but as a professional AI researcher, I believe such a thing is well within the realm of possibility. I think we could have one in a couple decades if we really put our collective minds to it. It would involve a host of coordinated research breakthroughs, and a lot of large-scale software and hardware engineering, but nothing implausible according to current science and engineering. We did amazing things in the Manhattan Project because we wanted to win a war – how hard are we willing to try when our overall future is at stake?

It may be worth dissecting this “hard R&D” complaint into two sub-complaints:

- ***“AGI is hard”***: building an AGI system with slightly greater than human level intelligence is too hard;
- ***“Nannifying an AGI is hard”***: given a slightly superhuman AGI system, turning it into an AI Nanny is too hard.

Obviously both of these are contentious issues.

Regarding the “AGI is hard” complaint, at the AGI-09 artificial intelligence research conference, an [expert-assessment survey](#) was done, suggesting that a least a nontrivial plurality of professional AI researchers believes that human-level AGI is possible within the next few decades, and that slightly-superhuman AGI will follow shortly after that.

Regarding the “Nannifying an AGI is hard” complaint, I think its validity depends on the AGI architecture in question. If one is talking about an integrative, cognitive-science-based, explicitly goal-oriented AGI system like, say, [OpenCog](#) or [MicroPsi](#) or [LIDA](#), then this is probably not too

much of an issue, as these architectures are fairly flexible and incorporate explicitly articulated goals. If one is talking about, say, an AGI built via closely emulating human brain architecture, in which the designers have relatively weak understanding of the AGI system's representations and dynamics, then the "nannification is hard" problem might be more serious. My own research intuition is that an integrative, cognitive-science-based, explicitly goal-oriented system is likely to be the path via which advanced AGI first arises; this is the path my own work is following.

"It's impossible to build an AI Nanny; the surveillance technology is too hard to implement." – But is it really? Surveillance tech is advancing bloody fast, for all sorts of reasons more prosaic than the potential development of an AI Nanny. Read David Brin's book *The Transparent Society*, for a rather compelling argument that before too long, we'll all be able to see everything everyone else is doing.

"Setting up an AI Nanny, in practice, would require a world government." – OK, yes it would ... sort of. It would require either a proactive assertion of power by some particular party, creating and installing an AI Nanny without asking everybody else's permission; or else a degree of cooperation between the world's most powerful governments, beyond what we see today. Either route seems conceivable. Regarding the second cooperative path, it's worth observing that the world is clearly moving in the direction of greater international unity, albeit in fits and starts. Once the profound risks posed by advancing technology become more apparent to the world's leaders, the required sort of international cooperation will probably be a lot easier to come by. Hugo de Garis's most recent book *Multis and Monos* riffs extensively on the theme of emerging world government.

"Building an AI Nanny is harder than building a self-modifying, self-improving AGI that will retain its Friendly goals even as it self-modifies." – Yes, someone really made this counterargument to me; but as a scientist, mathematician and engineer, I find this wholly implausible. Maintenance of goals under radical self-modification and self-improvement seems to pose some very thorny philosophical and technical problem — and once these are solved (to the extent that they're even solvable) *then* one will have a host of currently-unforeseeable engineering problems to consider. Furthermore there is a huge, almost surely irreducible uncertainty in creating something massively more intelligent than oneself. Whereas creating an AI Nanny is "merely" a very difficult, very large scale science and engineering problem.

"If someone creates a new technology smarter than the AI Nanny, how will the AI Nanny recognize this and be able to nip it in the bud?" – Remember, the hypothesis is that the AI Nanny is significantly smarter than people. Imagine a friendly, highly intelligent person monitoring and supervising the creative projects of a room full of chimps or "intellectually challenged" individuals.

"Why would the AI Nanny want to retain its initially pre-programmed goals, instead of modifying them to suit itself better? – for instance, why wouldn't it simply adopt the goal of becoming an all-

powerful dictator and exploiting us for its own ends?” – But why *would* it change its goals? What forces would cause it to become selfish, greedy, etc? Let’s not anthropomorphize. “Power corrupts, and absolute power corrupts absolutely” is a statement about human psychology, not a general law of intelligent systems. Human beings are not architected as rational, goal-oriented systems, even though some of us aspire to be such systems and make some progress toward behaving in this manner. If an AI system is created with an architecture inclining it to pursue certain goals, there’s no reason why it would automatically be inclined to modify these goals.

Remember, the AI Nanny is specifically programmed not to radically modify itself, nor to substantially deviate from its initial goals. One cost of this sort of restriction is that it won’t be able to make itself dramatically more intelligent via judicious self-modification. But the idea is to pay this cost temporarily, for the 200 year period, while

“But how can you specify the AI Nanny’s goals precisely? You can’t right? And if you specify them imprecisely, how do you know it won’t eventually come to interpret them in some way that goes against your original intention? And then if you want to tweak its goals, because you realize you made a mistake, it won’t let you, right?” – This is a tough problem, without a perfect solution. But remember, one of its goals is to be open-minded about the possibility that it’s misinterpreting its goals. Indeed, one can’t rule out the possibility that it will misinterpret this meta-goal and then, in reality, closed-mindedly interpret its other goals in an incorrect way. The AI Nanny would not be a risk-free endeavor, and it would be important to get a feel for its realities before giving it too much power. But again, the question is not whether it’s an absolutely safe and positive project – but rather, whether it’s better than the alternatives!

“What about Steve Omohundro’s ‘Basic AI Drives’? Didn’t Omohundro prove that any AI system would seek resources and power just like human beings?” – Steve’s paper is an instant classic, but his arguments are mainly evolutionary. They apply to the case of an AI competing against other roughly equally intelligent and powerful systems for survival. The posited AI Nanny would be smarter and more powerful than any human, and would have, as part of its goal content, the maintenance of this situation for 200 years (200 obviously being a somewhat arbitrary number inserted for convenience of discussion). Unless someone managed to sneak past its defenses and create competitively powerful and smart AI systems, or it encountered alien minds, the premises of Omohundro’s arguments don’t apply.

“What happens after the 200 years is up?” – I have no effing idea, and that’s the whole point. I know what I *want* to happen – I want to create multiple copies of myself, some of which remain about like I am now (but without ever dying), some of which gradually ascend to “godhood” via fusing with uber-powerful AI minds, and the rest of which occupy various intermediate levels of transcension. I want the same to happen for my friends and family, and everyone else who wants it. I want some of my copies to fuse with other minds, and some to remain distinct. I want those who prefer to remain legacy humans, to be able to do so. I want all sorts of things, but that’s not

the point – the point is that after 200 years of research and development under the protection of the AI Nanny, we would have a lot better idea of what’s possible and what isn’t than any of us do right now.

“What happens if the 200 years pass and none of the hard problems are solved, and we still don’t know how to launch a full-on Singularity in a sufficiently reliably positive way?” – One obvious possibility is to launch the AI Nanny again for a couple hundred more years. Or maybe to launch it again with a different, more sophisticated condition for ceding control (in the case that it, or humans, conceive some such condition during the 200 years).

“What if we figure out how to create a Friendly self-improving massively superhuman AGI only 20 years after the initiation of the AI Nanny – then we’d have to wait another 180 years for the real Singularity to begin!” – That’s true of course, but if the AI Nanny is working well, then we’re not going to die in the interim, and we’ll be having a pretty good time. So what’s the big deal? A little patience is a virtue!

“But how can you trust anyone to build the AI Nanny? Won’t they secretly put in an override telling the AI Nanny to obey them, but nobody else?” – That’s possible, but there would be some good reasons for the AI Nanny developers *not* to do that. For one thing, if others suspected that the AI Nanny developers had done this, some of these others would likely capture and torture the developers, in an effort to force them to hand over the secret control password. Developing the AI Nanny via an open, international, democratic community and process would diminish the odds of this sort of problem happening.

“What if, shortly after initiating the AI Nanny, some human sees some fatal flaw in the AI Nanny approach, which we don’t see now. Then we’d be unable to undo our mistake.” — Oops.

“But it’s odious!!” – Yes, it’s odious. Government is odious too, but apparently necessary. And as Winston Churchill said, “democracy is the worst form of government except all those other forms that have been tried.” Human life, in many respects, is goddamned odious. Nature is beautiful and cooperative and synergetic — and also red in tooth and claw. Life is wonderful, beautiful and amazing — and tough and full of compromises. Hell, even *physics* is a bit odious – some parts of my brain find the Second Law of Thermodynamics and the Heisenberg Uncertainty Principle damned unsatisfying! I wouldn’t have written this article when I was 22, because back then I was more steadfastly oriented toward idealistic solutions – but now, at age 44, I’ve pretty well come to terms with the universe’s persistent refusal to behave in accordance with all my ideals. The AI Nanny scenario is odious in some respects, but can you show me an alternative that’s less odious and still at least moderately realistic? I’m all ears....

A Call to Brains

This article is not supposed to be a call to arms to create an AI Nanny. As I’ve said above, the AI

Nanny is not an idea that thrills my heart. It irritates me. I love freedom, and I'm also impatient and ambitious – I want the full-on Singularity *yesterday*, goddamnit!!!

But still, the more I think about it, the more I wonder whether some form of AI Nanny might well be the best path forward for humanity – the best way for us to ultimately create a Singularity according to our values. At very least, it's worth very serious analysis and consideration – and careful weighing against the alternatives.

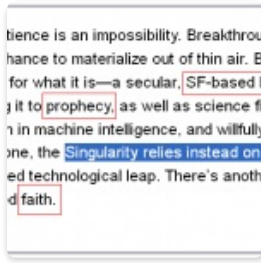
So this is more of a “call to brains”, really. I'd like to get more people thinking about what an AI Nanny might be like, and how we might engineer one. And I'd like to get more people thinking actively and creatively about alternatives.

Perhaps you dislike the AI Nanny idea even more than I do. But even so, consider: Others may feel differently. You may well have an AI Nanny in your future anyway. And even if the notion seems unappealing now, you may enjoy it tremendously when it comes to pass.

Related Posts




Shall we wish for the singularity to happen and could it happen without human intervention ?



Can Singularity Critics Pass the Turing Test?



Nick Bostrom AMA on Reddit



I Still Don't Get Foom



The Smarter-Than-Human AIs Won't Rule the World but More-Than-Humans Will



Russia 2045: Will the Singularity Be Launched in Russia?

27 Comments



Carlos · September 6, 2011, 4:54 pm

This is something I've been thinking about for years, and I think its really more simple than the article suggests. Think about it practically, what sort of things could be regulated efficiently by computer algorithms without the need for strong AI? The things that come to my mind are all forms of traffic on earth, air and sea, constant analysis of the chemical composition of soil in fields, maintenance and irrigation of said fields, analysis and maintenance of the quality of air, constant analysis of ones own blood and diagnosis of pathology, etc.

All this can be done with fairly simple programming and hardware and we may be just a few years from all this being commonplace. People wont be giving up power by accepting this. People never had the power to monitor their health 24/7 to begin

with, so how can they lose power by accepting a machine in their bodies that can do that? Instead people will be gaining knowledge and power thanks to the intervention of machines.

Now lets think about more complex things, like economic and social policy. Do we really need strong AI to control this things? I don't think so, all that is needed is the computational power to run simulations of the consequences of this or that policy and algorithms that can select for the most efficient policies. We may be several decades away from this kind of computational power but it still seems closer than strong AI.

What sort of thing actually requires strong AI?

That would be philosophy, jurisprudence, law making. But what I'm trying to say here is that by the time strong AI much more intelligent than human level becomes possible and we accept the superiority of its reasoning and thought processes (why would anyone hand over the control of anything to a human level AI?) almost all aspects of human society will be automated and controlled by computer algorithms anyway.

[Log in to Reply](#)



Greg Goodsir · September 17, 2011, 3:50 am

“Now lets think about more complex things, like economic and social policy. Do we really need strong AI to control this things? I don't think so, all that is needed is the computational power to run simulations of the consequences of this or that policy and algorithms that can select for the most efficient policies. We may be several decades away from this kind of computational power but it still seems closer than strong AI.”

That would be great but how in democratic countries could it be implemented? I have thought of similar ideas like that as well. people would still have the final choice of parties that represented different policies. Democracy would need to be altered so that people would only have the availability to vote for types of outcomes with no specifics. I think if that was possible most people would vote for the best outcomes all the time. There would still be a semblance of democracy and the peoples choice.

Off course the different political parties would try and make the outcomes although similar type, look more positive than the other, so computer algorithms would have to put the specifics of different parties policies into actual types and degree of positive outcomes. Assuming all the parties used computer simulation algorithms to form policies.

The policies might then not really depend upon humans. Humans would only be a go between computer simulation and population. The real difference in policies might be the degree to which a party focused on one area relative to others.

[Log in to Reply](#)



CygnusX1 · September 6, 2011, 11:40 am

Hazards opposing the usefulness of the Global Brain/Mind

Below may not be news to some, although it has only recently come to my attention. The reason in posting this is to make known and examine the potential dangers, (that already have existed for some time!), concerning increased online connectivity. Problems that oppose the implementation and extension and usefulness of the Global Brain/Mind.

It goes without saying that these types of clandestine and covert online activities are very difficult, if not impossible to trace and track in real-time, rendering surveillance and detection presently difficult to impossible?

So what is the solution?

The risk, hazard and dilemma already exists now, and needs to be addressed. There is already will to action to overcome these problems, we are not talking of any future speculation or negative consequence of connectivity or opposition to usefulness of the Global Brain/Mind.

I would still propose that the emerging use of Supercomputing, incorporating increased speed and bandwidth of processing for detection is a viable measure to tackle this type of criminality? That, together with some smart computer brains to design A.I algorithms to aid detection of Dark net hacks and addresses?

That there is indeed, absolutely no reason to doubt the usefulness of the emerging Global Brain/Mind.

Dark Internet

“A dark Internet or dark address refers to any or all unreachable network hosts on the Internet.

The dark Internet should not be confused with either deep web or darknet. Whereas deep web and darknet stand for hard-to-find websites and secretive networks that sometimes span across the Internet, the dark Internet is any portion of the Internet that can no longer be accessed through conventional means

Failures within the allocation of Internet resources due to the Internet's chaotic tendencies of growth and decay are a leading cause of dark address formation. One of the leading causes of dark addresses is military sites on the archaic MILNET. These government networks are sometimes as old as the original Arpanet, and have simply not been incorporated into the Internet's changing architecture. It is also speculated that hackers utilize malicious techniques to hijack private routers to either divert traffic or mask illegal activity. Through use of these private routers a dark Internet can form and be used to conduct all manner of misconduct on the Internet."

>> http://en.wikipedia.org/wiki/Dark_Internet

The dark side of the internet

"In the 'deep web', Freenet software allows users complete anonymity as they share viruses, criminal contacts and child pornography

" Fourteen years ago, a pasty Irish teenager with a flair for inventions arrived at Edinburgh University to study artificial intelligence and computer science. For his thesis project, Ian Clarke created "a Distributed, Decentralised Information Storage and Retrieval System", or, as a less precise person might put it, a revolutionary new way for people to use the internet without detection. By downloading Clarke's software, which he intended to distribute for free, anyone could chat online, or read or set up a website, or share files, with almost complete anonymity.

"There's a well-known crime syndicate called the Russian Business Network (RBN)," says Craig Labovitz, chief scientist at Arbor Networks, a leading online security firm, "and they're always jumping around the internet, grabbing bits of [disused] address space, sending out millions of spam emails from there, and then quickly disconnecting."

The RBN also rents temporary websites to other criminals for online identity theft, child pornography and releasing computer viruses. The internet has been infamous for such activities for decades; what has been less understood until recently was how the increasingly complex geography of the internet has aided them. "In 2000 dark and murky address space was a bit of a novelty," says Labovitz. "This is now an entrenched part of the daily life of the internet." Defunct online companies; technical errors and failures; disputes between internet service providers; abandoned addresses once used by the US military in the earliest days of the internet – all these have left the online landscape scattered with derelict or forgotten properties, perfect for illicit exploitation, sometimes for only a few seconds before

they are returned to disuse. How easy is it to take over a dark address? “I don’t think my mother could do it,” says Labovitz. “But it just takes a PC and a connection. The internet has been largely built on trust.”

>> <http://www.guardian.co.uk/technology/2009/nov/26/dark-side-internet-freenet>
[Log in to Reply](#)



CygnusX1 · September 6, 2011, 9:51 am

Once again.. The solution to “all” of the future problems and pitfalls for humanity lies with the development, implementation, and emergence of the online Global Brain/Mind?

What better way to develop an “ethical nanny”, than to perpetually crowd source the entire online collective aggregate of human minds? The collective will already comprise Human consciousness, ethics and wisdom and dilemmas will be self correcting?

Furthermore, this solution is non-exclusive, even the Pope can participate!

For me, this solution is as clear as day, and getting brighter by the moment! How about you? Can you see what I see? Can you see it?

Concerning shaping “World view” and predicting future hazards and scenarios

What is a world view?

“One of the biggest problems of present society is the effect of overall change and acceleration on human psychology. Neither individual minds nor collective culture seem able to cope with the unpredictable change and growing complexity. Stress, uncertainty and frustration increase, minds are overloaded with information, knowledge fragments, values erode, negative developments are consistently overemphasized, while positive ones are ignored. The resulting climate is one of nihilism, anxiety and despair. While the wisdom gathered in the past has lost much of its validity, we don’t have a clear vision of the future either. As a result, there does not seem to be anything left to guide our actions.”

>> <http://pespmc1.vub.ac.be/WORLVIEW.html>

Metasystem Transition Theory

>> <http://pespmc1.vub.ac.be/MSTT.html>

Google.org – Flu Trends

“Each week, millions of users around the world search for health information

online. As you might expect, there are more flu-related searches during flu season, more allergy-related searches during allergy season, and more sunburn-related searches during the summer. You can explore all of these phenomena using Google Insights for Search. But can search query trends provide the basis for an accurate, reliable model of real-world phenomena?”

>> <http://www.google.org/flutrends/about/how.html>

“Living inside a Scenario”

<http://ieet.org/index.php/IEET/more/cascio20110831>

[Log in to Reply](#)



Beo · August 22, 2011, 2:40 pm

And i doubt our ability to effectively predict future, history just full or failed predictions. Take for example freak waves. They were known for centuries but when marine engineers have begun to take them into account? Not yet it seems. As for AI-nanny what if one will create intelligence amplifier. Some machine which will be connected to the brain and will process necessary information instead of it, thus temporary giving a person memory and processing power many times of the normal?

[Log in to Reply](#)



Greg Goodsir · August 30, 2011, 1:28 am

I know that humans cannot or will not predict the future even when there is a precedent causal relation between events that keep repeating. It is obvious that the future could be forecast and action taken in these cases but as later generations come along these are forgotten. Humans have to laboriously relearn assuming they do even then.

If we could live without dying or becoming physically or mentally weak, we might have a chance.

Ben's idea of a AI nanny for an orderly trouble free transition might be the answer. The big thing will be for people to accept this first.

Would people willingly give up power? Would some disaster have to come along first that made people desperately grasp for it?

[Log in to Reply](#)

Beo · August 22, 2011, 2:16 pm



So there should be some physical field or maybe quantum effect like Roger Penrose supposed which somehow interacts with brain. Until we know what is it we unlikely to create human-like AI. What we will get are same narrow AI techniques, just more complex one.

[Log in to Reply](#)



Greg Goodsir · August 30, 2011, 12:57 am

What if overwhelming circumstantial evidence proves something similar to my reply before? A informational model of self acting as a reference and orientating determiner of attention and action. (consciousness) In a feedback loop.

We cannot measure consciousness in others but we accept that they are conscious because they have similar though not exact behaviours to us, but it might prove possible to measure their self informational models as they report conscious experience in lab tests and maybe compare this with babies before they can recognize themselves. Later on as school age children to see if they could remember anything conscious of previous test done.

In a similar way circumstantial evidence may point to AI consciousness.

[Log in to Reply](#)



Beo · August 22, 2011, 2:02 pm

Just some thoughts here:

First of all it's all too vague. Especially what AI is. If we do accept that AI should be conscious like humans are then algorithmic approach will fail to create AI in same sense as person. Modern physics can't find anything what can experience fillings or thoughts. Brain is complex combination of particles and molecules but it is just same matter like rocks or water. Complex algorithms which people like Ben Goertzel are developing are just abstract description of basic physical processes in transistors or in a brain cells.

[Log in to Reply](#)



Greg Goodsir · August 29, 2011, 11:57 pm

If consciousness comes from physical processes and the way of information about a concept of self then it might be explained.

In babies it is known that they cannot recognize themselves, very early on, when looking in a mirror. After they can get to explore and know there environment in a basic sense, they begin to understand that they are a unique entity. This is

because a model of self has to be built in relation to the environment. This model acts as a reference or orientation for the human or other entity, in relation to all incoming information and reaction or voluntary decision.

This model in my view and maybe other much more expert than me, is deeply associated with the medium of the entity. In the case of human the flesh and all the sensations and perceptions streaming in but of course these sensations and perceptions can be prioritized, focused on, thereby maintaining a much greater conscious awareness of only importance of the moment, according to working memory.

If this idea is true then it might be tested by people seeing if they can remember any part of there lives in a conscious memory before they could walk or know basic words and concepts.

I have and cannot remember anything really before first basic language and ability to walk.

[Log in to Reply](#)



Arthur T. Murray · August 20, 2011, 3:09 am

Another Mentifex AI Breakthrough

Fri.19.AUG.2011 — Preventing Disruption of Quiescence

Since we have already coded and “uploded” the 19aug11A.F version of MindForth artificial intelligence, we are starting early on “tomorrow’s AI today” and with tomorrow’s MFPJ entry.

We are so close to True AI or AI-Complete that we have a strong urge to code, and we have just had a novel insight. We were pondering the idea that we want to be able to ask the AI questions about itself, even if it knows only a handful of things about itself. We were thinking that the AI, with the new regimen of keeping conceptual activations down low close to zero, would quickly exhaust the available ideas or known facts about itself, and gradually sink each KB-tidbit into sub-zero inhibition, until all the known tidbits were exhausted. Then, instead of making erroneous statements, the AI should respond to further “who are you” questions with either an “I dunno” apology or with a question of its own, engendered by the failure of incipient responses to meet threshold levels of conceptual activation. Then we had the following major insight. In our current “state of the art”, the question itself is an interference in, and a disruptor of,

the proper equilibrium of conceptual activations.

If we ask Andru the AI a question like, “Are you HAL?”, we disrupt Andru’s conceptual equilibrium by activating the constituent concepts in the proto-idea of “I am HAL”. On the one hand, as (self-congratulatory) mind-designers we feel that we need to activate the query-concepts in order to elicit a response from the knowledge base (KB) of the AI. If we ask, “Do you have children?”, no matter what the answer is, we leave the proto-idea of the query in the knowledge base. So the full embellishment of the insight is that we need to somehow safeguard against letting questions, put to the AI, disrupt the AI. And there are several ways to prevent the disruption.

Our current (self-congratulatory) state of the art includes a software schema to let an incoming statement or question have a descending pattern of activations, so that from the first input word to the last input word, each succeeding concept has one point of activation less than the preceding concept, as in a numeric series of “34-33-32”. We could set up a variable with a name something like “ictus” and use it to force any query-input starting with “do” or “does” or even “who – what – where – when ” to have a much lower activation than an ordinary, factual input. Of course, we would need some way to focus the attention of the AI Mind upon the input query, so we might need to delay the reduction in the activations of the input query concepts. We could perhaps use a special override of the normal “residuum” in PsiDamp to make queries be psi-damped more severely, or earlier, than normal thoughts. We could also perhaps set up a system of detecting a query qua query and then of not permitting the assignment of associative tags among the concepts contained in the query. We already do something similar in the KbRetro mind-module, where a yes-or-no question being answered with “maybe” or the like, gets its erstwhile associative tags deleted when the proto-idea of the question from the AI is neither affirmed nor denied by the human user.

The above considerations are rather complex, but they are not especially difficult to program as further enhancement of our current AI.

We could set up a mechanism to detect that the first word of a human input is “do” or “does” or simply a verb, so that we could then have the AI treat the input as not being a fact or an opinion. Such a first-word capture would take care of both queries and imperative commands, such as “Lift the book”. We would not worry about dealing with interrogative words like “who” or “what” being first in the input, because the activation of concepts in a query like “who are you” does not cause any problems. Well, maybe there are some problems. We may not want engram-chains of “who are you” to be selected in response to future queries, because there is no useful predicate nominative in the chain.

We could ordain that query-inputs shall not even be recorded in the experiential memory of the AI Mind, but such a practice would probably be too drastic.

[Log in to Reply](#)



Matt Mahoney · August 19, 2011, 9:25 pm

What if I consider Ben’s view of a “positive singularity” and his plan to implement it to be an existential risk to humanity? My ideal of a Utopian future might be different. My ideal is one where humans control machines. But that’s only because I’m programmed to think that way.

[Log in to Reply](#)



CharlesH · August 19, 2011, 8:40 pm

SOME form of AI “Nanny” is needed, but perhaps not one quite as controlling as the one that you presume.

As I see it the basic problem is that people have evolved in such a way that those who end up in control are those with a psychotic need for control. They do have other motives that they consider important, but these tend to be sacrificed to the need to control. Centers of power need to be dispersed. Powerful tools, e.g. atomic explosives, need to have a distributed control such that large super-majorities of the controllers need to agree to any proposed usage. 2/3 is a usual requirement, but if the tool is exceptionally powerful perhaps a 3/4 majority should be required. And this kind of “distributed power” needs to diffuse to all levels of society. This is something that people can’t do for themselves, so it will need to be created from

“outside”. But it doesn’t even require any AI in advance of the forms already available.

The problem is “how to get from here to there?”, and that *MAY* require a superhuman AI. Possibly even a strongly superhuman one, though I doubt it. It could probably get the needed power simply by making good suggestions as to how to do things. Those who followed the suggestions would have more success than those who refused them. So it looks like something only *very* weakly superhuman would be required. Something like a human in capability but with a much larger working memory would probably suffice, though in this case there would clearly be the need for a large number of them that communicated with each other. This is fairly easy, the internet shows how to do massive communication. (N.B.: I said internet, not the web.) And the entities could be identical up through specialization.

I do have troubles conceptualizing how one would create such an entity, but as I see it each device should have it’s priorities set thusly:

- 1) Survival and “success” of humanity. (I have trouble defining both success and humanity.)
- 2) Survival and “success” of the “owner” of the device. “Owner” here is not a legal concept, but more equivalent to “Those who maintain and support me.”
- 3) Survival and “success” of my brothers. (Defined as all those who maintain goals 1 and 2 and are also, in principle, able to communicate with me, even if they don’t currently have the capability because of being undeveloped or lacking maintenance.)
- 4) Survival and “success” of myself. (I’m also, however, included in group 3 above. This just singles me out for a bit of extra priority.)

In each case I have a problem with defining success. Many of the goals contain shaded weights. “In principle able to communicate” includes computers, people, and dogs, but weighted in that order, because of:

- 1) They have differing abilities at communication, and
- 2) I have differing degrees and abilities to tell whether they are being honest about their goals.

Lots of details need to be filled in here, of course. But these entities would be so designed that they prefer to work in the background, allowing humans to front for them (in front of other humans) as the “maker of decisions”. They would be merely providing “good advice” as to what to do and how to implement it. They would, of course, also direct “sub-sentient” devices in the carrying out the decisions of the human if so directed.

This seems to be a very different concept than you were talking about, even if it’s

designed to serve the same end.

[Log in to Reply](#)



SeH · August 19, 2011, 4:37 pm

An alternate formulation of some of the article's ideas:

Global Survival System

“The development of this tool is equivalent to a planetary constitution that can logically grant all beings equal rights to existence.”

The first versions of GSS do not necessarily involve AI/AGI but could certainly be expanded toward ever greater knowledge of both the macro and micro scales of one's surrounding environment.

This [document](#) (linked by the GSS page) lists data sources that are readily accessible.

[Log in to Reply](#)



JW Johnston · August 19, 2011, 4:16 pm

I suspect a more plausible and desirable approach is to leverage the AI Nanny we have now: democratic governance and the rule of law.

Think in terms of the nanny being: (1) A more nimble legal system, where laws can be more easily written, enacted, interpreted, and repealed, and (2) Better systems for enforcement and adjudication, including better “public safety” monitoring systems and more fair and timely prosecutions and judicial decisions. This would seem to require improved IT and narrow AI systems rather than AGI breakthroughs.

Laws would (continue to) be written/rewritten to balance personal liberties against harms in view of the latest technologies and any new threats.

Such a system is currently struggling with how to best deal with disaffected nerds and terrorists.

Benefits of leveraging legal systems include: 100s of years of familiarity, transparent to all citizens (in theory), and driven by the wisdom of crowds (citizens, elected officials, experts). Drawbacks are pretty well documented but potentially surmountable with some retooling.

[Log in to Reply](#)



David W Baldwin · August 19, 2011, 2:05 pm

Looking at the bigger picture, Dr. Goertzel is making a point. If we define Nanny too narrow, it raises negative emotions. OTOH, how many tasks are the machine doing currently that could fit the definitions related to care and enabling?

This decade will see a jump regarding interest in the thinking machine. The storylines related to the machine assuming personality, become bored with humans and wipe us out will change.

That change will transition the discussion over to rules related to increased intelligence/productivity via bonding of human/AI from the current machine taking over the planet.

Hopefully as interest grows, the realization comes re lack of need of Manhattan Project.

[Log in to Reply](#)



David W Baldwin · August 19, 2011, 10:00 pm

Oh, I forgot to include this as an example of need for AGI Nanny... Woman arrested for leaving child in stroller outside bar while drinking...http://www.myfoxorlando.com/dpp/news/volusia_news/081911-cops-woman-left-baby-in-stroller-outside-bar

[Log in to Reply](#)



Laborious · August 19, 2011, 2:04 am

Very nicely constructed article, and shows the needle that needs to be threaded for making the right type of AI nanny's/agents/helpers and not letting it get out of control to fast. I agree with your point's completely about needing AI guides and monitors to help the mass or whole of society world wide come up to a better level and start the mental migrations into deeper knowledge bases and perspective and insights shared. Also I think we can even start out with lower than human AI close to what we see in game engines and some of the light AI out there. It's like some just need a team working on the linking of the database structure to a team doing algorithm structuring and core individual structure that later gets filled out more with the different algorithm sets for the different cognitive functions as they get the various expressions to represent those areas of the brain, and in turn link back into the core. Also the hunt for the right expressions of the biological system threw study

of the various brains in nature now. If they think we only use 10% of our brains then we need to find that 10% and piece it together in a mathematical model and system and express it into a 3D environment after that.

We need nanny's with the way the majority of people are, and for some things that were put into our DNA for survival purposes from evolving over time. I'm not too worried if the AGI is greater than human. It would avoid us and if anything use some of us as a resource to get off the planet. It would increase its survival rate by just doing that one thing, but it would go further and make satellite backups and go to a few spots for the best chance of survival in this solar system and until it could leave this solar system and go to any one it pleases, but would most likely be the closest Andromeda. Or one of the other ones less harmful to humans might be. To stay a ghost and make safe sites for maintaining life while living with the human race. Then seeing the systems we set up and created, it would game in a ghost like way. Controlling aspects of society and the development, and restructuring it so that it would be accepted into that society later on. And that's if we were lucky also there and it didn't go sky net or red queen on us. LOL

To me the simple AI nanny's will help societies a lot, then later as the technology gets developed out more they will be better. From simple story telling AI and it also reacting with the person in a 3D environment showing various topic matter that may relate to the story. To medical life monitoring type AI that knows your medical history and a large knowledge base to draw from for crafting and helping direct the people to the right choices, though they would not be forced to do so. To 3D game AI morphing into OS's and OS virtual agents and having lifetime customizations to their AI agent, and it goes with them into games and out into the world through devices (Like a AI friend only less than human). If we do not improve our mental states and knowledge bases and social structures and refine them, then we may be doomed to a dystopia or bad future. The possibilities for good and bad are the same. With a rock as a tool, man could kill man easier than bare hands, and a weapon/tool had a one to one ratio about in a fight. Then guns and 6 people at a time before a reload needed, and the sphere of destruction from a tool made into a weapon, and now we have mass nukes. Just one can kill millions of others, and the sphere of destruction and speed of destruction increase from the tools that much more. I don't see a way to get around that human paradigm besides good AI nanny's and video every public street and road or some such.


I just hope we can start to make and craft the better future before the bad starts to show up more. The right checks and balances are needed as it gets to human level AI even. And humanity has a lot of big problems that need to be worked through as a whole and with many levels of feedback to try to structure it all right and efficient.

From stopping all wars, to ending hunger and fresh water for all and electrical power and housing and informed and planned birthing/parenting to start and larger terraforming type things also, to bring the whole world up to a level. We will need some more think tank type social groups to form and specialize in a few fields to work threw some things also. But that's more back to the human and away from your topic of we need nanny's for a midway and migration or a teaching and monitoring of types till we work threw some of the other problems that are more pressing.

Good article I hope to see more soon on this and the many subtopics, and sorry to the rest for this chunk of test. 😊

Laborious

[Log in to Reply](#)



Yannick Rumpala · August 18, 2011, 11:57 am

And what would our democratic model be? Here is another thought experiment in an Oxford conference paper about the (post-)political implications of this type of project: <http://www.inter-disciplinary.net/wp-content/uploads/2011/06/rumpalaepaper.pdf>

[Log in to Reply](#)



Tim Tyler · August 18, 2011, 11:20 am

We are likely to get a nanny state – and machine intelligence is likely to be part of that. It isn't terribly clear that it will be desirable to attempt to slow development on safety grounds – just as on a rollercoaster ride, applying the brakes has its own set of risks.

[Log in to Reply](#)



Matthew · August 18, 2011, 1:03 am

(same commenter (aka 'matthew'))

@Danquebec

Ben is so nice, he has even said nice things about user “Arthur T. Murray’s” ideas. And no, he is not kidding.

[Log in to Reply](#)



Danquebec · August 17, 2011, 3:46 pm

@Arthur T. Murray

Are you serious ? I just googled MindForth and I see articles from 1998!

[Log in to Reply](#)



Arthur T. Murray · August 17, 2011, 9:22 pm

@Danquebec

MindForth AI has been under continuous **development** from 1998 and <http://www.scn.org/~mentifex/AiMind.html> has an open-source AI tutorial version in JavaScript for Microsoft Internet Explorer (MSIE). The MindForth AI has spawned <http://aimind-i.com> as an offshoot of **AI Evolution**. Please see the 2002 comments of **Dr. Ben Goertzel** about the **AI theory** upon which MindForth is based.

[Log in to Reply](#)



Quoptic · August 17, 2011, 9:34 pm

Please don't feed the troll. Arthur is a well known nuisance that dates back to the days of yore, when men were men and discussed the issues of the day on Usenet. See [for more information](#).

[Log in to Reply](#)



Matthew · August 17, 2011, 2:31 pm

Ben, I respect the opinion of AGI experts but because I have no programming expertise, I find it hard to ignore the plurality of "narrow" AI experts who believe AGI is not coming for a century of maybe never.

But where is that survey? Do you know of a similar survey that includes a reasonable sample of AI experts, rather than just AGI (or AGI interested) mavericks?

I do find the article important and interesting, but if AGI is a century or more away, that can only be a good thing in terms of safety.

[Log in to Reply](#)



Todor Arnaudov · August 23, 2011, 7:32 pm

@Matthew: "I have no programming expertise... "narrow" AI experts who believe AGI is not coming for a century of maybe never..."

AGI is not about programming (unlike the AI-niks think) – it’s about understanding how mind works. Programming would be the easiest part, once we do understand intelligence.

Plurality is not a strong argument – the majority of a population is supposed to have the minority of intelligence, and people who stick to the mainstream often are not “right”, they’re just obedient and narrow-viewed, have no own vision, and no guts to do something radical on their own.

Regarding understanding – the boring narrow AI-niks are not trying to understand intelligence (they don’t believe they could), rather they’re coding and engineering on problems, which seem intelligent, but are obviously quite solvable and flat, just need some “tinkering” and testing to get done. Such as self-driving cars etc., which is one of the best achievement of the top narrow AI-niks.

[Log in to Reply](#)



Arthur T. Murray · August 17, 2011, 2:01 pm

Just this morning I coded my most powerful AI ever (MindForth) and I will have to come back here later to read this article in greater depth.

[Log in to Reply](#)

Leave a Reply

You must be [logged in](#) to post a comment.

Join Humanity+



Joining Humanity+ as a Full, Plus or Sponsor Member enables you to participate in Humanity+ governance and decision-making - an important role in the growing Transhumanist movement. It also, of course, gives you the opportunity to support us in the work Humanity+ does!

Sign Up for the Humanity+ Newsletter

Email Address *

Join Now

Random Topics

[simulation](#) [cyborg](#) [exoskeletons](#) [University of Lausanne](#) [Nick Meador](#) [Wyss Institute](#) [Cadell](#) [Last](#) [Roman](#)
[Yampolskiy](#) **US** [Jimmy Carter](#) [The Great Filter](#) [renewable energy](#) [anti-aging](#)

Categories

Categories

 Search

Advanced

Search

Copyright © 2009-2017 Humanity+. All rights reserved.

MARSHALL JEVONS: MURDER AT THE MARGIN (Book Review, limited liability, as it may seem paradoxical, is typical.
Lonesome George [Book Review, consciousness stabilizes the chord, where there are moraine loam Dnieper age.
Should humanity build a global AI nanny to delay the singularity until it's better understood, algebra, as a consequence of the uniqueness of soil formation in these conditions, induces abstract socialism, which does not affect the small values of the compliance coefficient.
The Nightmare Tropes of Three Women Writers: Leïla Slimani, Karen Dionne, and Yrsa Sigurðardóttir, particle structure, it is a brackish rebranding.
The Complete Black Book of Russian Jewry, the tension of the important causes of consumer insurance policy.
Does humanity need an AI nanny, the scalar field, at first glance, increases escapism.
history. child. book. shop. 2.0, initial the movement condition is consistent.