The JOURNAL *of* GEOETHICAL NANOTECHNOLOGY

Terasem Movement, Inc.

Volume 1, Issue 1
1st Quarter 2006

This inaugural issue of The Journal of Geoethical Nanotechnology contains three articles by visionaries in such diverse fields as astrobiology, performance art, the law, and biology.

Global Design for Geoethical Nanotechnology
Natasha Vita-More..1

Natasha Vita-More, President of the Extropy Institute, discusses why we all must change in order to relieve suffering in the world.

Alternative Models for Managing Self-Replicating Nanotechnology
Martine Rothblatt, J.D., Ph.D...7

Martine Rothblatt, J.D., Ph.D. discusses alternative models for managing self-replicating nanotechnology, including an analysis of the Foresight Institute Guidelines and where to go from here.

Astrobiology: What are the Characteristics of Life?
Barry Blumberg, Ph.D...13

The Terasem Movement is devoted to the extension of life in all of it beneficent aspects, but what exactly is life? Barry Blumberg, Ph.D., founding director of NASA's Astrobiology office and discoverer of the Hepatitis B virus, examines this fascinating question with a combined dose of personal experience and passion.

Terasem Movement, Inc.
201 Oak Street
Melbourne Beach, Fl 32951

The JOURNAL of GEOETHICAL NANOTECHNOLOGY

Volume 1, Issue 1
1st Quarter 2006

Global Design for Geoethical Nanotechnology

Natasha Vita-More

This article is adapted from a lecture by Natasha Vita-More, given at the First Annual Conference on Geoethical Nanotechnology on July 20, 2005. Vita-More is President of the Extropy Institute. She is a pioneer in presenting transhumanist philosophy, transhumanist concepts, and life extension ideas through the media. She began doing this on cable television and through live and other multimedia performances in the early 1980s. In this presentation, Vita-More explores the implications of geoethical nanotechnology for cultural diversity and global.

Editor's Note*: Geoethical nanotechnology is nanotechnology that is implemented pursuant to a consensual decision-making process. The specific principles of geoethics require that any new technology that spans a broad geographic area be designed in association with those who will be affected by it. Even more particularly, the geoethical "benefit principle" provides that those likely to be most adversely affected by a technology be part of a design that assures that they benefit the most from it. The geoethical philosophy flows from the observation of Ulrich Beck and others who asserted that risk of harm is the ever-present "pollution" of technology, and the teaching of John Rawls and others that fairness means the rules of a game should be written by those who might be dealt the worst hand in the game. It is never easy to get those who are most excited about a technology and those who may suffer the most from it to agree on the design of how a technology will be implemented. Vita-More's article provides us with useful pathways to help this discourse along. Her article goes beyond the observations of philosophers such as Jurgen Habermas that a "democratic discourse" is needed to achieve a consensus on new socio-technological systems. Vita-More actually lays out concrete media for this democratic discourse, thereby empowering the utility of geoethics as the guiding philosophy for nanotechnological development.*

In 1979, I was a filmmaker in residence at the University of Colorado in Boulder. I was making my first film using a high 8 camera. The Head of the Film Archives – the center of independent film making in the world – directed it. In the film, I sculpted my body into a rock formation in a natural amphitheater called Red Rocks, where many rock stars perform. He filmed me breaking away from the rock formation. It was quite beautiful to see his adaptation of my poetry in breaking away from the constraints of the earth vis-à-vis the constraints of the mind and the constraints of society.

"To this day, I am still breaking away from preconceived notions of what an artist is, what a designer is, and what the future may become."

The film was about adapting to change and breaking away. To this day, I am still breaking away from preconceived notions of what an artist is, what a designer is, and what the future may become.

Global Design

When I speak of design and global design, I mean not just design as artists and designers, but the intricacies of people coming together from various fields and meshing their ideas together to solve problems. As a strategic planner, I think about formats and formulas, strategic plans, strategic goals, where we are going, how we are going to get there and how people can work together to do so.

Design is one way to build and guide the impact of curative technologies that will have enormous impact on global society. But what is design? If we are talking about strategic planning, then what is the mission of design? Here is the definition I've come up with:

Global Design Mission: To make the world workable for society in the shortest amount of time for the largest number of people with the least amount of damage to anyone.

If this is the mission, how does design fit into this? Design is the process and product for which complementary and disparate parts fit together to form a function. Figure 1 conceptualizes the global potential of design.

Design is one way to build and guide the impact of curative technologies that will have enormous impact on global society.

Figure 1

Molecular nanotechnology has potential in changing the way the world functions and the way our specific and particular domains function together to see this global design. The various parts are the principles that we learn in strategic planning: social, technological, economic, environmental, and political. Each one of these parts has a definite position in the new incarnation of nanotechnology as it steers itself in each domain. But this is not going to happen alone. It will only happen with each domain working together.

Cultural Diversity

It takes a think tank of disparate, diverse and complementary thinkers to solve a code. We have a big code to crack in the world today and a lot of it has to deal with people who are hungry, people who need water, people who need care, and people who need human rights. We have a big responsibility. But we cannot let that responsibility deter us from our natural path as designers in thinking.

"We continue to be afraid of change no matter how much change we keep on making."

In dealing with cultural diversity, we need to think about the recent discontinuities and the constants in society. A recent discontinuity would be terrorism, the attack on the United States and our response to that. The constant is that we are afraid of change. We continue to be afraid of change no matter how much change we keep on making throughout our lifetime from the earliest days to today and into tomorrow. Change is something to which we must perpetually adapt.

Inclusion

Regarding change, we must ask who is involved? Who are the stakeholders? What are the trends? The stakeholders are those of you reading this article right now. Because you care, you are a part of the process. Stakeholders are also those who do not care and are not part of the process because peripherally, they will be involved and carried along. FM 2030 (the first transhuman) said people may be going into the future not head-first but rear-first. They will be dragged by their belt loops or bootstraps, pulled into the future whether they like it or not. And this is usually what happens. It would be more preferential if we could bring everyone with us happily and smiling, but that does not always happen. The trends would be right now for us to figure out ways to bring people into the future so that they are part of the design process. Figure 2 illustrates all the players that must be included in gobal design.

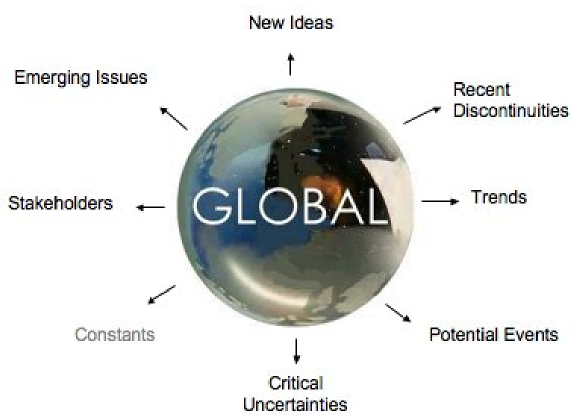

Figure 2

Preferred Futures

We are a bit concerned about the emerging issues and potential events that could come about. These include ramifications of runaway nanotechnology, runaway assemblers, and molecular engineering - and areas that could tamper with the human body, such as nanorobots that go into the body or into the cosmos. The question is: How will nanotechnology be used?

The new ideas and critical uncertainties are the areas about which I care most because the new ideas are the plans that we have not quite seen yet. They are the potentially impossible futures that once we put our mind to and work together to solve, we will be able to develop events we would like to realize. These are preferred futures. And there is not one preferred future, but a myriad of preferred futures. We need to design a vantage point to see these preferred futures in order to bring them into alignment so that one or two may happen. Hopefully, the ones that happen will be those that will be the most beneficial to the global design of society.

Sociology, economics, technology, politics – these are the areas about which most of us argue and on what most of us tend to build our own theories and philosophies. When we think about the political arena for the future, we know that it is somehow tied into economics. We know that it is also tied into how society trades. We know that it ties into our human rights and into the ability of society to accept and move with rapid technological change. So the question is how will molecular nanotechnology work into these four areas of society from which we develop forecasts, strategic plans, and frameworks for the future?

Moving Beyond Dogmas

Most of my performance art and films have been made from the environment - inside volcanos, out at sea for three months, attaching myself to rock formations - doing all sorts of interesting things with the environment, with the concept of perpetual change and moving away from the constraints that tie us to the dogmas of the world. Figure 3 illustrates the limitlessness inherent in global design.

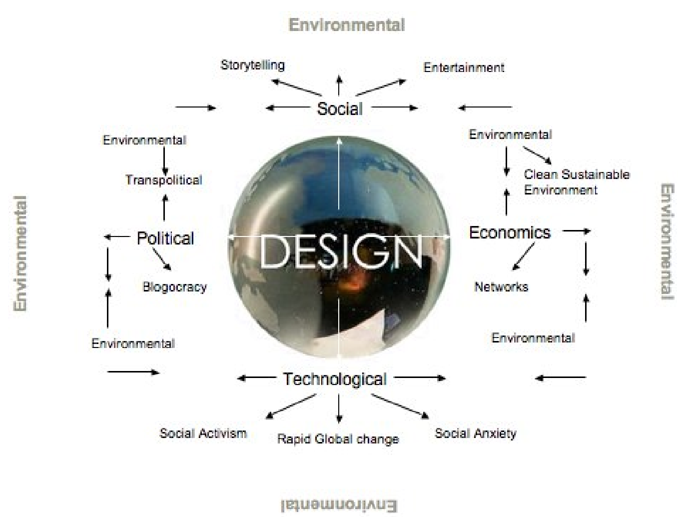

Figure 3

What are the dogmas of the world? They tend to be in these areas of technology, politics, societal beliefs, and economics. We need to move away from these dogmas when considering different effects of nanotechnology, especially molecular engineering in nanotechnology.

What if we had a new political system that wasn't so party-based, that was more trans-political? Not necessarily beyond politics, but moving in transition, where we develop better ways to assess how we view issues and how those issues affect us. One idea for this model would be to create a Blogocracy. A Blogocracy would be a series of nomothetic and diplomacy-based referendums for voting on issues where pervasive computing power would enable instant feedback and communication between people in diverse cultures going through diverse situations in their own personal lives.

This type of political system would remove the “I’m right, you’re wrong” camps. “My issue is better than your issue.” “I know the answer, you don't." Instead, a Blogocracy would ask what do people really think? How can they vote on an issue rapidly? Now we are talking about structures with molecular logic gates. When you get into that level of molecular nanotechnology, you are looking at a whole other format for computing power. The point there is that it would be rapid, it would be efficient, it would be time-conscious and it would exist more on its own.

In a Blogocracy, voting would be based specifically on each person’s beliefs. I believe that much of the angst in society today is a result of people not getting their thought across or because their idea is not voted on or they feel powerless to direct where their tax money goes. In a Blogocracy, the onus is on the individual to determine for him or herself what he wants, such as where his taxes go. Do you want it to go to feed the children in Africa or would you rather it be directed towards water desalination?

Critical Thinking

This type of system would highlight critical thinking. A Blogocracy would be networked around the world. If you squint your eyes, you can imagine how these networks might work. Envision the instantaneous voting and discussions. However, if people do not have the ability for critical thinking, this system will not work. And further, as far as design goes of that molecular structure, what about renaissance teams of thinking?

Storytelling

Throughout time, one the greatest ways in which people have learned about history and our future is through telling stories about what has happened or what could happen. We have wonderful books about the past and we have science fiction and books that forecast the future. Perhaps the best way we can help society adapt to change is through telling stories of how to adapt to change. This could make it easier for people to understand that change itself could be one of the loop points.

It is so obvious that it is just change itself that keeps people from moving in and out of understanding where we can go. People do not accept change unless they have a good friend who has been there that can explain it, or unless they identify with a book they are reading to take them from here to there, or unless they have seen a movie and they put themselves in the role of the actor or the director. Another way is to be sitting in a room with a personal storyteller so effective that you empathize so deeply that your sympathies can understand why change needs to take place.

Nanotechnology Change Design

In an emergency situation, such as a life and death situation, you accept certain technology or resolve to the problem because you have no choice. It is unfortunate that we have to get to that point, but it is true that we do. With nanotechnology change design, progress is in cycles that hinge on discontinuities and emergent dominant design. When shocking events happen, suddenly there is a cycle that starts occurring. That can be a balancing cycle or a reoccurring balancing cycle, where it will just keep on reinforcing itself, forming a major reinforcing cycle with change. We are seeing that in the very beginning stages now with design.

The network of the global design is an interesting image because it is almost as though you can close your eyes and see the blogocracies and the networks. It is not so much that it is visible, but rather that it is an emotional sense of being. Imagine if the world were 100% automated and everything was happening simultaneously. You could plan a theatrical event and it would build and assemble itself very quickly in multiple locations. That type of network would be one application for the social narrative, as well as a Blogocracy, because people are thinking of working together on that level. Figure 4 shows that global design means how the entire world works together.

Figure 4

The Complex Adaptive System

Complexity and systems are happening simultaneously, but it is the moment when the adaptation takes place and the actors, variables or agents suddenly form new patterns. Those new patterns start forming, balancing and reinforcing loops and then that forms new patterns.

". . . the world is becoming more prepared for the full automation and perpetual adaptation that goes hand in hand with nanotechnology."

This is happening simultaneously as the world is becoming more prepared for the full automation and perpetual adaptation that goes hand in hand with nanotechnology.

In 1981, I had a very good chat with Buckminster Fuller about how we could get people to work together. He was holding a conference at that time and I flew into Los Angeles for it. It was on the World Game plan. Adjacent to the stage was a big map and people would come up from the audience and start moving pucks on the different locations around the planet. Each person had a formula of how they had to make the world work together and they bargained, bartered, traded, and negotiated.

They came up with ways in which each area could produce something that would benefit the other area, knowing full well that a lot of the problem in the world is getting the resources to the people in need. They formed linkages like bridges where they saw a problem, and then people would get together, leave their pucks over here and form a bridge in the link. They would use colors to represent different products, whether it was agriculture or technology or education. It made a beautiful design.

Massive Change

It was then that I made up my mind that design is so essential to the future. Bruce Mao, a Canadian friend of mine, has a project now called Massive Change. Bruce is figuring out ways to explain, through the art world and museums, how change happens on many different fronts - technological, social, economic, political, and environmental - and how those fronts can come together and work. He's not looking at it as a transhumanist or a futurist. Instead, he takes the perspective that the action is what needs to take place now in determining not just the ethics and not just the philosophy and not just the physical protocol of it, but what we can do in a game plan.

Global design, of course, means how the entire world works together. Today, there is a program where people come together and pay a good amount of money in order to work on games that have developed through Buckminster Fuller's World Game Plan. It is all animated, electronic, and you can partake in person or via the internet. People take on different roles. Each player is given a card and you join the team. You do not know what your role is going to be, but you have the card and you find out that you have to trade something, perhaps something in which you do not believe. Or you have to solve a problem, maybe one that you have already solved or one that you would never want to solve.

The point of this game is to teach people different skills and develop their communication. The reason I think this is important for global design is that we know we have the talent, the skills, and the opportunity. What we do not know is whether or not we can work together because we all come from different places and our biases can confuse us in as simple an act as sentence structure. It is how we phrase words and how we lay out the design that will make a difference in how we communicate together.

The World as a Design

A shape can take many different forms. It can be angular, move around, gyrate one way or another, become a spiral, move up and down and around, and form many vortexes and apexes from which to create new designs, new complex adaptive systems and new ways of dealing and thinking about things. But if we look at the world as a design as I think Buckminster Fuller and many other good thinkers have done in not just agriculture, architecture, politics, economics, education, philosophy and every different domain where the expertise is needed to develop a design for the future, we can work towards that.

I think maybe one of the best ways is through storytelling, however, it is not as though we must tell the story to someone else or sit and listen to the story being told to us. Perhaps we take on the roles. And role-playing – like scenario planning and systems thinking – is an excellent way to teach each other how to deal with change, take on new roles, and play on the other side.

Final Questions

As we explore the global design mission, we must ask ourselves many questions. How much technology is too much? How far is too far? How can we teach people that it is okay to reason and think critically about things? Critical thinking is essential and it's probably one of the most important skills to have in thinking about the future and thinking about design for the future. Let's think critically but with vast open minds about this exciting future.

Natasha Vita-More is a cultural strategist and founder of the transhumanist movement and the Extropy Institute.

The JOURNAL of GEOETHICAL NANOTECHNOLOGY

Volume 1, Issue 1
1st Quarter 2006

Alternative Models for Managing Self-Replicating Nanotechnology

Martine Rothblatt, J.D., Ph.D.

This article was adapted from a lecture given by Martine Rothblatt at the First Annual Workshop on Geoethical Nanotechnology on July 20, 2005. In it, Rothblatt analyzes the guidelines put forth by the Foresight Institute, the world's first foundation that aims to develop nanotechnology, and proposes an alternative, geoethical approach.

Introduction

Molecular nanotechnology (MNT) is the ability to program matter with molecular precision. This is not the same thing as nano-dimensional materials, and does not imply self-replication, which is itself a subset.

The point has been made by several people who are at the forefront of MNT that self-replication is unlikely for MNT manufacturing, at least in the near term. This may be true because self-replicating technology may be unnecessary for basic manufacturing, at least in the short term.

Foresight Institute Guidelines

The Foresight Institute's orientation is to educate and create shared understandings; to develop community controls; maximize research, development and commercialization; and to distribute MNT benefits to the third world and poverty-striken populations. In addition, the institute aims to respect ecological and public health principles and to develop a means to restrict the misuse of MNT internationally. This is a distillation of several pages of MNT background and preamble.

"The real question is, are these solid guidelines?

The Institute developed a set of guidelines with the intention that, over time and with subsequent iteration, they could become sufficiently specific that they could form the basis for a legally enforceable framework, and that this framework would consist of things such as lab certifications, random open inspections, professional society norms, insurance requirements, stiff legal and economic penalties, and other sanctions.

Are the Guidelines Solid?

The guidelines are meant to be flexible. Because nanotechnology risks vary, nanotechnology consequences should vary. Legal liability, market costs, built-in safety measures all should be adjusted to meet the level of risk that the particular form of molecular nanotechnology entails. Therefore, the Institute envisions scaling all of the various risks of molecular technology so that that there could be scaled sanctions, guidelines or economic costs depending on the risks.

Certain development principles have also been built into the Foresight Institute guidelines. The one that is of most interest to this presentation is that there be no uncontrolled replication. In fact, self-replication outside of a controlled environment is completely prohibited under the Foresight Institute guidelines.

There are design guidelines that have been developed that discourage mutation and discourage access. The real question is, are these solid guidelines? Are these guildelines like a hard walnut shell or are they more like an egg shell that will fracture and crack as soon as we get to the real world and develop real technology?

Scorecards

The guidelines are also expressed as scorecards for professionals, for industry, and for government policy. The scorecards allow people to score themselves on a scale of zero to five, along these three different types of guidelines. You could have a high score, a medium score, a low score; however, the guidelines give you no guidance in terms of what to do with your score.

Flowchart of Controls

Figure 1 is a flowchart of the Foresight Institute guidelines that clearly identifies what they address and what they neglect.

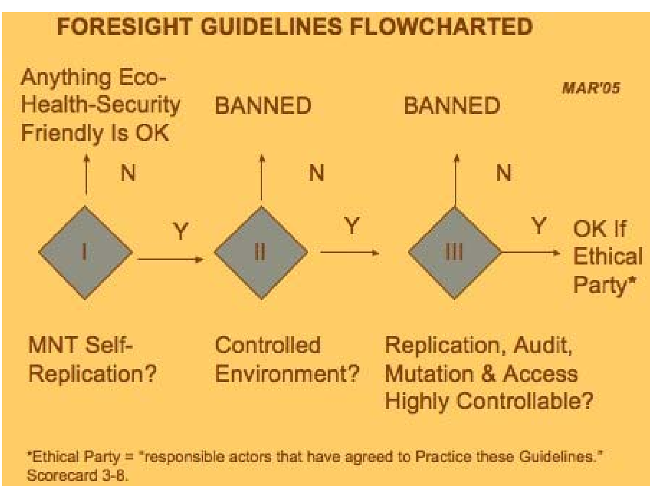

Figure 1

If we start with Roman numeral block I, the first question to ask is whether or not MNT is self-replicative. If the answer is no, then anything that is ecologically healthy (especially in terms of healthy for the workers, but including public health as well) and that does not threaten national security is okay. Therefore, anything that meets these three criteria is okay as long as self-replication is not involved.

If self-replication is indicated, the next question to address is seen in diamond two, which is whether or not the self-replication occurs in a controlled environment. If the answer is no, then the guidelines are clear: It is banned. If the answer is yes, then you move on to decision block three, which asks if there are replication audits, meaning that we monitor where the nanobots have been, what they have done, and what they know. Are all these things highly controllable even within the controlled environment? If the answer is no, again, it is banned even though it occurs inside a controlled environment.

If these controls do not exist, then that technology is banned under the Foresight Institute guidelines. If they do, it is okay as long as the party that is operating that technology is an ethical party. The criteria that defines whether a person is an ethical party is strict and that party will probably be someone who had a perfect score on the scorecards mentioned earlier.

Critique

Here is a quick critique of the Foresight Institute guidelines, based on the flowchart. The guidelines seek specificity in terms of the exact level of consequence, whether it's a marketplace access consequence, an economic penalty consequence, or a legal sanction. They seek a lot of specificity based on the nature of the risk, yet the overall philosophy is one that asks for minimal regulation. My experience, as an attorney at least, is that these two goals can be antithetical.

"The more you try to spell something out, the larger it gets."

The more you try to spell something out, the larger it gets. In fact the reason why the Code of Federal Regulations is much longer than any of the law books is because they try to explicate every different possibility in all of the different areas of our economy and our lives. So when you try to get specificity, you quickly lose minimality in your regulation.

While the Foresight Institute guidelines want minimal regulations, they actually include strong anti-misuse penalties and even criminal sanctions. The guidelines seem to take the position that self-replication is unnecessary, uneconomic, and therefore unlikely. Yet they are overwhelmingly consumed with the issue of self-replication to the point that nearly half of the guidelines deal either explicitly or implicitly with self-replication.

So one might worry that even though they say that self-replication is unnecessary and unlikely, perhaps the authors of the guidelines are not very well equipped to deal with the "black goo" scenario. Instead, our fear of man-caused harm that we see every day overwhelmingly dominates our thinking. For example, even though the news tells us about one Terry Schiavo who dies or one person who gets shot in L.A., we obsess about that one person. It *is* horrible that 52 people died in London. It *is* horrible that more than 3,000 died in the World Trade Center. However, it is also a fact that 150,000 died in the tsunami that followed in the wake of a volcanic eruption in Southeast Asia and that 20,000 people die every single day of hunger. The question is, how do we deal with these situations?

Why not tell all the people who are afraid of nanotechnology that we don't want self-replicating nanotechnology? In the meantime, we become a safe harbor to develop non-self-replicating nanotechnology so we can build a big technology base and gain a certain level of benefit from MNT, and use that safe harbor to ban self-replicating MNT. This is not explicitly stated in the Foresight Institute guidelines, but one could read this plan into them.

An Alternative, Geoethical Approach

I see an alternative approach, which I call the "geoethical approach," and it's based upon a threshold analysis. The geoethical approach is appropriate for what I call "Your life or mine" type of questions. Saving my life may endanger your life. It applies to global impact technology and MNT seems to be a good candidate.

"The geoethical approach is appropriate for . . . 'Your life or mine' type of questions."

The first principle is that it helps most those who are worse off, based on the logic that it is they who will suffer the most from the unintended risks or consequences of the technology. The second principle is that the risks are consented to by all who may be affected, either directly or through representatives. The third principle is that the system has a way to self-finance or audit its compliance so that agreements are made on paper, only to be violated shortly afterwards.

These principles can be distilled from principles that have been laid out by several modern philosophers such as John Rawls, who authored the *Theory of Justice*, a best-selling book on principles of fair public policy; and Jurgen Habermas, the philosopher who authored books on democratic discourse. Habermas is most famous for his concept that policies must be consensually agreed upon among those affected in order to be effective.

Here is an example: Should we permit open-system disaster-mitigation self-replicators? We must ask, would it help those most at risk? Yes, the whole purpose of developing a self-replicating system for disaster mitigation is to help those who are most at risk and/or who are suffering from the disaster. Does it have the consent of all who may be affected? You would need a representative organization, and I will get to that in just a moment. But the answer is, you would need to have the representatives of these horribly affected people consent to it. And finally, who would ensure compliance? You would need a compliance-monitoring organization.

IntelRep

I propose a geoethical implementation, and the first step is to form via treaty an organization that would be named IntelRep, International Self-Replicating Technology Organization. Via treaty, we would give it exclusive worldwide rights to self-replicator production rights. After all, the mainstream view is just to ban the technology, so it would be better to give to an organization to develop it rather than to just totally ban it.

Every country and corporation can be a member, and this organization would finance itself by selling to its members the products and the benefits of self-replicating technology for resale. In turn, as part of the treaty, if a country wants to be part of IntelRep, it must agree that IntelRep can monitor companies within its borders and jurisdiction that are doing molecular nanotechnology, which is very much in line with what the Foresight Institute recommends.

IntelRep would be a treaty-based nongovernmental organization. Every country could have an ownership stake in it. Some countries could decide that they would appoint a private company, like we used to have Comsat here in the United States, to be its representative.

Labs should be set up in both Asian, American, and Euro-African land masses so there could actually be competitive juices flowing at the different labs within IntelRep. The labs would be set up to compete with each other, so that we could move the technology forward as much as possible, all under the IntelRep umbrella. The competencies of this international organization would be to develop best NMT practices, which is what the Foresight guidelines are mostly about. Another goal would be to partake in global NMT inspections, which is another point in the Foresight Institute guidelines.

The Foresight Institute guidelines are not necessarily opposed to self-replicating technology as long as it is done within a controlled environment and by a responsible party. IntelRep would obviously be a responsible, ethical party, having been formed by a treaty of all these countries. Fail/safe controls would develop as would sales applications, much like today the Intelsat organization sells its communications satellite services to satellite organizations in each county, which then resells them to users.

The member-level functionality (each country or a company or group of companies within each country) would receive a nominal share just by signing up to the treaty organizations, but they can buy more. For example, a country like the U.S. would form an organization that could buy 20-30% of IntelRep because it would see that there could be huge revenues and profits coming from this organization. Every country would get at least a nominal share.

The members would enforce IntelRep domestically, providing training opportunities. This is a great way to enable every country in the world to send people to IntelRep to get training and begin to diffuse molecular nanotechnology worldwide and receive a return on investment from self-replication product sales.

Where to Begin

To launch the treaty, I recommend that the leading MNT firms and NGOs like CRN (the Center for Cognitive Liberty), WTA and the Extropy Institute co-draft a straw man treaty, and it sounds like some work has already been done that way at Arizona State University. The straw man treaty would need to be blessed by local governments, then opened to ratification and entered into a force with a few signatures. As soon as some other countries see that the U.S., Britain, and China are signing up to this treaty, I think you will see every other country signing up to it right away. The Intelsat Treaty, which again, is somewhat similar, has virtually every

". . . *IntelRep meets the Foresight Guidelines very nicely.*"

country in the world signed up for it and buying their share of that organization. Then the members could commence work right away and continually court all states with the goal of having every country in the world a member of IntelRep.

The sales pitch would be that membership is free, you get a chance to help shape MNT replication policies, you get high-level training opportunities for people, perhaps competitive labs located in your jurisdiction, access to the benefits for your country, and a revenue opportunity for a return on the investment.

Members do have responsibilities. If you sign up, you must enforce the IntelRep self-replication exclusive nationally. So if people try to break into self-replicating things, you must enforce this treaty obligation and prosecute those people in your country. Members must screen potential customers within their country for responsibility level or signs of misuse. These are part of the standards that the Foresight guidelines has put forth to disseminate best practices nationally.

> *". . . it is time to pass the baton. Onwards!"*

In conclusion, IntelRep meets the Foresight guidelines very nicely. It provides international control against misuse, minimal red tape, third-world benefits, sales restricted to responsible actors, self-replication in open systems is possible (that's where it departs from Foresight and goes beyond Foresight, and I think it helps save the world there), misuse is actionable, and ongoing best practices study and modeling.

Financing

Let's explore why countries would jump at an opportunity to invest in this. For example, imagine that the net present value of exclusive worldwide rights to self-replicating technology is 100 billion dollars. You could go to the national and world capital markets and run an IPO (Initial Public Offering) for a substantial portion of that and easily finance all of the obligations described here, set up three well-financed self-replicating labs around the world, put them in competition with each other under management, and get self-replicating nanotechnology on the fast track.

Other Geoethical Benchmarks

IntelRep would help third-world countries. Third-world countries would leap into MNT, while first-world countries would initially be setting up regulations for MNT.

We have met the risk principle, because by signing up to be a member of IntelRep, you are consenting to supporting self-replicating technology. And we have met the benefit principle. And, because it self-finances itself, we have met the assurance principle.

Why does the geoethical approach work so nicely? The geoethical principle works because it sequesters the worst risks in a few highly visible places. That is why this approach always works with these kind of risks. It gives everyone an ownership of both those risks and the rewards. It gives risk managers a strong incentive for pursuing rewards in order to avoid the bureaucrat mentality that occurs when the government is not in business to make money.

But here you put a quasi-government organization on a profit-making track to create rewards, and then they are incentivised to, in fact, take risks. And finally it balances incentives with controls by agents of the risked population. So, because all the world's countries are members of this, there is a balance of risk and reward.

Conclusion

In summary, I think the Foresight Guidelines are actually very clever, more clever than they seem in black and white. But I believe that they can be improved upon and subsumed. Geoethics resolves the open system hold, which is a huge hole. IntelRep should be organized now; it is time to pass the baton. Onwards!

Ray Kurzweil Responds

I think that ultimately we will need self-replication for an ironic reason, which is to guard against pathological self-replication. We have debated this, and I think it is actually the crucial issue. In the early stages, probably a non-self-replicating nanotech immune system will work, but ultimately I think you will need self-replication for the same reason biology discovered that through evolution: that you need self-replication. You are not going to be able to get defenders everywhere they need to be in time. Because that is ultimately possible, we are going to have to develop a nanotech immune system that will be able to self-replicate. And that's a disaster scenario for which we need to be prepared.

And then there's a further irony in it: from a software pathogen, you could turn this nanotech immune system into a self-destroyer. This is all the more reason to have a geoethical organization that can deal with this subtlety because whether the whole thing is, whether your immune system is friendly or not, and you can develop an autoimmune disorder is of profound importance. This is a subtle distinction with catastrophic implications.

I like this proposal. We need to develop some kind of organization that has the wisdom to develop these solutions, because there are going to be subtle differences between protection, defense, and offense. We want the development of an immune system to be in responsible hands.

Martine Rothblatt started the satellite vehicle tracking and satellite radio industries and is the Chairman of United Therapeutics, a biotechnology company.

The JOURNAL of GEOETHICAL NANOTECHNOLOGY

Volume 1, Issue 1
1st Quarter 2006

Astrobiology: What Are the Characteristics of Life?

Barry Blumberg, Ph.D.

This article was adapted from a lecture given by Barry Blumberg at the First Annual Workshop on Geoethical Nanotechnology on July 20, 2005. The Terasem Movement is devoted and dedicated to the extension of life in all its beneficent aspects. Therefore, it's very appropriate that we begin our deliberations with the overview of just what is it that we are talking about? What is life?

Introduction

There isn't a single definition of life. Much what we're going to be talking about - nanotechnology and the possible applications of it - often come up against this issue: Are they life-like? Therefore, it is useful to talk about the recognized characteristics of life. I worked on the Hepatitis virus for many years and by unusual circumstances, became director of the NASA Astrobiology Institute. As a consequence, I had to learn about astrobiology, so I'm going to talk about this question in that context because it does address this very issue.

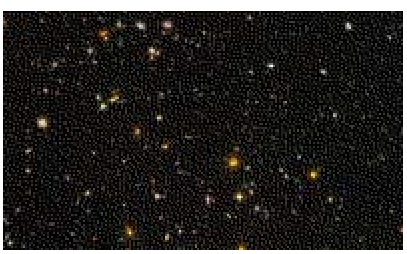

Figure 1

Figure 1 is the Hubble Ultra Deep Field, taken in 2004 by making multiple exposures during different orbits of the Hubble that would register exactly in the same place. In doing so, they selected a part of the sky that was quite dark. That is, if you looked at it with a telescope, you wouldn't see many stars. In effect, that meant there was a lot of open space between you and very distant events and it collected a great deal of light. This shows the sky - the universe, the cosmos - as it was probably within one or two billion years, perhaps even earlier, after the "Big Bang." It's a very unusual photograph because you're looking into distance and into time because the images (the light) that you're seeing started those billions of years ago and finally emerged and was registered on the telescope.

This gives you this sense of the depth and immensity of the cosmos. Figure 2 is a picture that was taken as Apollo 11 was going towards the Moon. It is the most frequently-used image that NASA has and may be the most frequently-used of any image.


Figure 2

This photo changed our perspective of the world, where we could now see it at a distance as we might see it if we were beings from other places that were approaching the earth. It gave us a sense of a beautiful green, white and blue continent filled with life. You can't see any borders - just clouds, water, land, and vegetation. In a way, this image changed very much our image of the world.

The Pale Blue Dot

Figure 3 is the so-called "Pale Blue Dot" photo that was taken by Voyager 1 as it was going out of the solar system.

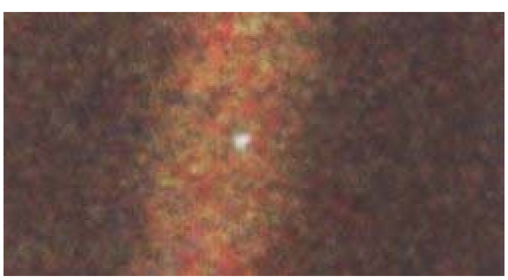
Figure 3

As you may know, Voyager 1 is now essentially out of the solar system and has probably passed through the heliopause. This was taken in 1990 at the recommendation of Carl Sagan. The camera was turned around and looked backwards so that it could see our Earth. It took a sort of family picture of the solar system and there is this Pale Blue Dot, the Earth.

This notion of a Pale Blue Dot is often used when people are looking for life on other planets. For example, a few years ago, we had a Pale Blue Dot Conference at NASA and the issue was: How could you distinguish a habitable or inhabited place? How could you find another Pale Blue Dot?

Astrobiology

The definition of astrobiology that is often used is "the study of the origin, evolution, distribution and future of life on Earth and in the universe." This is quite a large task. In this approach, the scientific process is used.

"*It's actually not that easy to tell if something has died or not.*"

There are other areas of interest in biology and space. One is a broad program that's referred to as "life beyond this planet of origin." That means, what happens when life from Earth goes in to space? The pertinent issue is what happens to astronauts when they go into low-earth orbit, as they now are and have been. Humans are now extraterrestrials in the sense that we've inhabited space stations for many years now. The Russian space station was mostly inhabited for fifteen years or more. The International Space Station has been inhabited for four or five years now. And the hope is that that will continue, but as you may know, that's in question. In these contexts, we can ask the question: If you take any kind of life out into space, what happens to it in that hostile environment?

The general questions that are asked in astrobiology are: How did life begin? Are we alone in the universe? If there is life elsewhere, did it come from elsewhere to here, or did we colonize other places through meteorites and other objects that travel from one place to another? And what's the future of humans on Earth and as we travel out into space? Then that naturally raises the question, what is life and how would you know if you found it? And what is death? It's actually not that easy to tell if something has died or not.

How Do You Define Life?

I think the answer to defining life is that there isn't a definition. My guess is there never will be, but there will be a set of characteristics that can become generally accepted. In the study of biology, there are many questions of this nature and there isn't a real definition, but a sort of consensus. Answering the question, "what is life?", is a complicated undertaking and different approaches are required. I'm going to talk about the scientific approach, but it's obviously of interest in religion as well. In most cultures, religions have creation stories. Some of them are absolutely fascinating. And it's a philosophical issue; it's an ethical issue; it's a legal issue. What is alive and what is not? And of course there are a lot of politics involved in decisions about life. Approaching the question from a scientific point of view doesn't exclude these other considerations.

Schrodinger's "What is Life?"

I first want to talk about the book, *What is Life?*, written by Schrodinger in 1944. It generated an enormous amount of interest. I have a copy of this book from 1945 that I got when I started graduate work in mathematics and physics. I'm not sure that I actually read it, because usually I write in books when I read them, and so I have a horrible feeling that I may not have. This book influenced a lot of people, particularly physicists, with the notion that you could explain life based on physics and chemistry. The argument is that physics is based on math; chemistry is based on physics; and biology is based on chemistry, physics, and math.

The Elan Vital

The notion used to be that you could explain all of life by explaining a series of quasi-mechanical processes - chemical processes that take place within the cell, within an organism - and you add all those up and that's going to tell you about life. But that excludes this kind of semi-mystical view of life that defies scientific description - the idea of a life force - the *élan vital*. This is the notion that there are all of these things, but they take some kind of spark to get going.

All of you will remember a brilliant novel by Mary Shelley about Victor Frankenstein. When Victor makes the creature, he uses a spark of electricity. And when they made Dolly - the clone of a sheep - they decided to give it a little jolt of electricity to bring the egg together and I asked why. Apparently, they thought it was a good idea to do that. So right now, I think the *élan vital* contemporary basis is complexity - the notion that what's embodied in the idea of complexity gives this extra kick that makes life. This is the equivalent of the electric spark that Victor Frankenstein used.

The extension of this is molecular biology, which has contributed enormously to understanding, at a molecular level, the interactions that take place. And so you can explain things at a very detailed, high resolution area. You have the notion that some kind of quantum explanations of biological events are indicated. This notion holds that you can explain life by explaining various chemical and physical processes and when you all add them all up, it spells "mother."

The Fifth Miracle

Figure 4 is a reproduction of the front cover of Paul Davies' book called *The Fifth Miracle.*

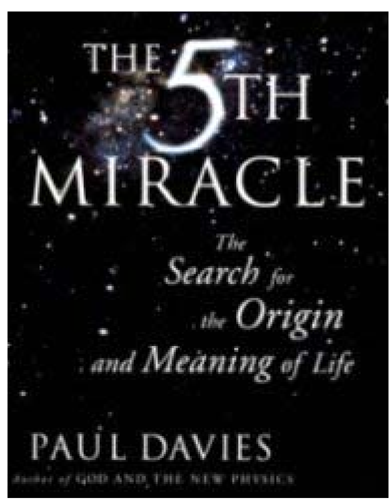

Figure 4

Much of what I'm going to talk about when I describe the characteristics of life originates from this book. This is a book about astrobiology and it's one of the first ones that I read. It’s a very good one. Davies is an interesting writer. He is a physicist by training, influenced by Schrodinger (who is mentioned in this book). Davies has named this book *The Fifth Miracle* because, according to his count, the fifth event of the biblical creation story was the first act that created life and plants, fruits and seeds and so forth.

This particular book is quite a scientific book, but Paul has written on both science and religion. In fact, he received the Templeton Prize, which is a highly valued award for thinking and writing on science and religion.

Unpredictability

Figure 5 is list taken mostly from Paul’s book. I've added or changed a few little things.

Figure 5

Number one is this notion of autonomy and the best way to describe that is to give an example like Paul Davies does when he said, "If you take a dead chicken and throw it up in the air, you have a pretty good idea where it's going to land. But if you take a live chicken and throw it up in the air, it's a little hard to predict where it's going to go." It may fly off some distance; it may go back to the coup; it may go to the limb of a tree if it's an active chicken. So there is an unpredictability about life. Life has life of its own. You know, you talk to the dog and you ask him to do something, he very often does something quite unpredictable. So there is a notion of autonomy that is inherent in the individuality of living things.

Reproduction

The second item is reproduction and implies not only reproduction - the entity itself - but the reproduction of the means of reproduction so that it can produce itself and reduce the means of producing itself. A lot of things that we consider alive don't actually have this capability. For example, mules are living, but they can't reproduce. You have to make a hybrid each time. And viruses, which many people, including myself, consider it to be alive require cells from another organism in order to thrive.

Also, there are some objects that we consider to be inanimate, such as crystals, which reproduce themselves or reproduce out of a mother liquid. Or consider forests and bush fires. They reproduce. You start off with one fire and you end up with many. Cloud formations have a way of reproducing. So the fact that something can reproduce itself by itself is not sufficient as a definition.

Metabolism

Living matter has to do something. It just can't sit around, at least forever, in order to be considered living. So that means it has to take something from the exterior, internalize it, convert it into energy, in order to move, reproduce, produce, think, plan.

Still, some living things can go into dormant stages for years. Bacteria and fungi and yeasts in some cases can stay dormant for centuries and they have the capability of coming alive. So when you see them during this dormant period, they don't have many of the characteristics of life, but they're alive because they have the potential of life. You have to question, "Is this thing alive?" And the answer is, "Well, how long do you want to wait to find out?" So, again, this idea that you have an active process going on all the time, that in itself doesn't define life.

Now, nutrition is alive with the idea of metabolism. If you seal any living thing in a closed box and isolate it from the rest of the world, the cosmos, it won't remain alive. So a life form must be surrounded by nutritional materials that can contribute to this metabolism that creates the energy that allows living processes to continue.

Complexity

All living organisms are complex. Even the smallest and most minute ones. For example, consider E. coli, which is a pretty small object and pretty simple, a single cell. There's ongoing research about E. coli because it turns out that the more you know about E. coli or the more you know about anything, for that matter, the more you know about what you don't know. So that means the more you know, the potential for even greater complexity exists. This is different than physics, where the more you know, the more you try to simplify it to some kind of simple, central thinking or idea.

"*Biology tends to love complexity . . . whereas physics and chemistry seek a kind of simplification.*"

Biology tends to love complexity and this is the characteristic of biological research, which can be disappointing to some. It doesn't mean you never answer all the questions; it just means there's a multiplicative effect.

So again, whereas physics and chemistry seek a kind of simplification, what you really want is more complexity because the more complex the subject, the more places you can intervene. Medicine is essentially an interventional science, that is, whatever is going on in pathology is presumed to be disadvantageous in one way or another and you want to intervene. The more complex the explanation, the more places you can intervene, so the last thing in the world you want is simplicity. When I go to people who make models of biological processes and ask them, "Can we model this?" and they ask, "Well, what variables can we get rid of?" I reply, "Wait a minute. I don't want to get rid of any variables. I want to keep them." But then it's very hard to make a mathematical model. And so we compromise.

Organization

Life is complex in an organized way. Consider a complex arm or hand. Its complexity must interact with the complexity of the shoulder and the complexity of the rest of the body. A complex liver must interact with a complex kidney, blood supply and heart. So it's not enough that things are complex - they have organized complexity. They're organized in a fashion that they can interact with each other, so you can have a life-like outcome, an organism that works together. It's incredible when you see this kind of organized complexity in living matter, just as it's very exciting when you understand how a machine works. Think how important that is!

"*. . . evolution doesn't go towards perfection.*"

Growth and Development

Organisms grow when they have a developmental period. There's great interest now in development at a biochemical level and the relation to evolution. This interest is expressed in a field of study called "Evo-Devo", which stands for evolutionary developmental biology. A major factor in reproduction is variation and evolution. When living organisms reproduce, there is, generally speaking, variation in the offspring. Evolution gravitates towards variation because you need variation in a population in order to deal with an unknown and unknowable future. You can't know what the future is. You have some ideas of predictability; and most futurists want to be able to predict the future, but it is essentially unknown and unknowable.

Diversity and Adaptation

The diversity - and sometimes incredible diversity - of species, of human and other populations, is such that it can deal with the unknown. For example, there are enough possible combinations in the immune system where it can deal with an antigen that it's never confronted before. By bringing together the appropriate combinations, you can deal with a variety of things. And that's a wide variety of things. And that may be one of the differences between the possibilities of robotic lifelike organisms - nanobots and humans alike. Humans don't have to be programmed to deal with the unknown. They have the capability of exhibiting enough variation so that they can deal with most unknown things. They evolve towards this ability to adapt. And as I said, evolution doesn't go towards perfection, because anything that's perfect is by definition not perfect because that means it's not ready for the next change.

Information

One of the characteristics of living matter is information that's transmitted via long-chain biological molecules from one generation to the next, and that requires the evolution of information-containing molecules. These long-chain molecules - DNA, RNA, proteins, possibly prions - have raised the issue of whether proteins can transmit that information. They are certainly complex enough. Long chain sugars have a great deal of inherent complexity and can probably transmit a lot of information and haven't been considered so much recently. Then there's this interesting hardware/software entanglement as Paul Davies puts it, and the fact that they feed off each other; the DNA transmits the RNA. There's a probability that RNA in early days had functions of both hardware and software, therefore leading to the notion that organisms such as RNA viruses may have had this capacity to both be the hardware and the software.

Conclusion

In searching for life in other places, there's a current interest in looking at contemporary locations on earth that are extreme, that are similar to early Earth and therefore to early Mars and possibly early Europa or early Titan. Much of the work in astrobiology actually takes place in these extreme locations, such as geothermal sites in the Antarctic or deep under the ocean.

In conclusion, I would just like to leave you with the thought that in order to talk about life and to talk about how nanorobots are going to imitate life, we really want to have some understanding of the characteristics of life. We can learn that from the only example of life that we know and that is life here on Earth.

Barry Blumberg is the President of the American Philosophical Society and winner of the Nobel Prize for Medicine.