

Minding the Planet

The Meaning and Future of the Semantic Web

Nova Spivack
CEO, Radar Networks
November 6, 2006

<http://www.mindingtheplanet.net>
nova@radarnetworks.com

Notes

- **Master Copy:**
http://novaspivack.typepad.com/nova_spivacks_weblog/2006/11/minding_the_pla.html
- **Last Update:** Tuesday, November 7, 2006, 10:17AM PST
- **License** -- This article is distributed under the [Creative Commons Deed](#). If you would like to distribute a version of this article, please link back to <http://www.mindingtheplanet.net> from your version, thanks.
- **Printable version** -- [Click here to download the printable PDF version of this article](#)
- **Illustrated Version** -- [See the version by the Lifeboat Foundation](#)

Prelude

Many years ago, in the late 1980s, while I was still a college student, I visited my late grandfather, Peter F. Drucker, at his home in Claremont, California. He lived near the campus of Claremont College where he was a professor emeritus. On that particular day, I handed him a manuscript of a book I was trying to write, entitled, "Minding the Planet" about how the Internet would enable the evolution of higher forms of collective intelligence.

My grandfather read my manuscript and later that afternoon we sat together on the outside back porch and he said to me, "One thing is certain: Someday, you will write this book." We both knew that the manuscript I had handed him was not that book, a fact that was later verified when I tried to get it published. I gave up for a while and focused on college, where I was studying philosophy with a focus on artificial intelligence. And I soon I started working in the fields of artificial intelligence and supercomputing at companies like Kurzweil, Thinking Machines, and Individual.

A few years later, I co-founded one of the early Web companies, EarthWeb, where among other things we built many of the first large commercial Websites and later helped to pioneer Java by creating several large knowledge-sharing communities for software developers. Along the way I continued to think about collective intelligence. EarthWeb and the first wave of the Web came and went. But this interest and vision continued to grow. In 2000 I started researching the necessary technologies to begin building a more intelligent Web. And eventually that led me to start my present company, [Radar Networks](#), where we are now focused on enabling the next-generation of collective intelligence on the Web, using the new technologies of the Semantic Web.

But ever since that day on the porch with my grandfather, I remembered what he said: "Someday, you will write this book." I've tried many times since then to write it. But it never came out the way I had hoped. So I tried again. Eventually I let go of the book form and created this weblog instead. And as many of my readers know, I've continued to

write here about my observations and evolving understanding of this idea over the years. This article is my latest installment, and I think it's the first one that meets my own standards for what I really wanted to communicate. And so I dedicate this article to my grandfather, who inspired me to keep writing this, and who gave me his prediction that I would one day complete it.

This is an article about a new generation of technology that is sometimes called the Semantic Web, and which could also be called the Intelligent Web, or the global mind. But what is the Semantic Web, and why does it matter, and how does it enable collective intelligence? And where is this all headed? And what is the long-term far future going to be like? Is the global mind just science-fiction? Will a world that has a global mind be good place to live in, or will it be some kind of technological nightmare?

I've often joked that it is ironic that a term that contains the word "semantic" has such an ambiguous meaning for most people. Most people just have no idea what this means, they have no context for it, it is not connected to their experience and knowledge. This is a problem that people who are deeply immersed in the trenches of the Semantic Web have not been able to solve adequately -- they have not found the words to communicate what they can clearly see, what they are working on, and why it matters for everyone. In this article I have tried, and hopefully succeeded, in providing a detailed introduction and context for the Semantic Web for non-technical people. But even technical people working in the field may find something of interest here as I piece together the fragments into a Big Picture and a vision for what might be called "Semantic Web 2.0."

I hope the reader will bear with me as I bounce around across different scales of technology and time, and from the extremes of core technology to wild speculation in order to tell this story. If you are looking for the cold hard science of it all, this article will provide an understanding but will not satisfy your need for seeing the actual code; there are other places where you can find that level of detail and rigor. But if you want to understand what it all really means and what the opportunity and future looks like -- this may be what you are looking for.

I should also note that all of this is my personal view of what I've been working on, and what it really means to me. It is not necessarily the official view of the mainstream academic Semantic Web community -- although there are certainly many places where we all agree. But I'm sure that some readers will certainly disagree or raise objections to some of my assertions, and certainly to my many far-flung speculations about the future. I welcome those different perspectives; we're all trying to make sense of this and the more of us who do that together, the more we can collectively start to really understand it. So please feel free to write your own vision or response, and please let me know so I can link to it!

So with this Prelude in mind, let's get started...

The Semantic Web Vision

The Semantic Web is a set of technologies which are designed to enable a particular vision for the future of the Web – a future in which all knowledge exists on the Web in a format that software applications can understand and reason about. By making knowledge more accessible to software, software will essentially become able to understand knowledge, think about knowledge, and create new knowledge. In other words, software will be able to be more intelligent – not as intelligent as humans perhaps, but more intelligent than say, your word processor is today.

The dream of making software more intelligent has been around almost as long as software itself. And although it is taking longer to materialize than past experts had predicted, progress towards this goal is being steadily made. At the same time, the shape of this dream is changing. It is becoming more realistic and pragmatic. The original dream of artificial intelligence was that we would all have personal robot assistants doing all the work we don't want to do for us. That is not the dream of the Semantic Web. Instead, today's Semantic Web is about facilitating what humans do – it is about helping humans do things more intelligently. It's not a vision in which humans do nothing and software does everything.

The Semantic Web vision is not just about helping software become smarter – it is about providing new technologies that enable people, groups, organizations and communities to be smarter.

For example, by providing individuals with tools that learn about what they know, and what they want, search can be much more accurate and productive.

Using software that is able to understand and automatically organize large collections of knowledge, groups, organizations and communities can reach higher levels of collective intelligence and they can cope with volumes of information that are just too great for individuals or even groups to comprehend on their own.

Another example: more efficient marketplaces can be enabled by software that learns about products, services, vendors, transactions and market trends and understands how to connect them together in optimal ways.

In short, the Semantic Web aims to make software smarter, not just for its own sake, but in order to help make people, and groups of people, smarter. In the original Semantic Web vision this fact was under-emphasized, leading to the impression that Semantic Web was only about automating the world. In fact, it is really about facilitating the world.

The Semantic Web Opportunity

The Semantic Web is one of the most significant things to happen since the Web itself. But it will not appear overnight. It will take decades. It will grow in a bottom-up, grassroots, emergent, community-driven manner just like the Web itself. Many things have to converge for this trend to really take off.

The core open standards already exist, but the necessary development tools have to mature, the ontologies that define human knowledge have to come into being and mature, and most importantly we need a few real “killer apps” to prove the value and drive adoption of the Semantic Web paradigm. The first generation of the Web had its Mozilla, Netscape, Internet Explorer, and Apache – and it also had HTML, HTTP, a bunch of good development tools, and a few killer apps and services such as Yahoo! and thousands of popular Web sites. The same things are necessary for the Semantic Web to take off.

And this is where we are today – this all just about to start emerging. There are several companies racing to get this technology, or applications of it, to market in various forms. Within a year or two you will see mass-consumer Semantic Web products and services hit the market, and within 5 years there will be at least a few “killer apps” of the Semantic Web. Ten years from now the Semantic Web will have spread into many of the most popular sites and applications on the Web. Within 20 years all content and applications on the Internet will be integrated with the Semantic Web. This is a sea-change. A big evolutionary step for the Web.

The Semantic Web is an opportunity to redefine, or perhaps to better define, all the content and applications on the Web. That’s a big opportunity. And within it there are many business opportunities and a lot of money to be made. It’s not unlike the opportunity of the first generation of the Web. There are platform opportunities, content opportunities, commerce opportunities, search opportunities, community and social networking opportunities, and collaboration opportunities in this space. There is room for a lot of players to compete and at this point the field is wide open.

The Semantic Web is a blue ocean waiting to be explored. And like any unexplored ocean its also has its share of reefs, pirate islands, hidden treasure, shoals, whirlpools, sea monsters and typhoons. But there are new worlds out there to be discovered, and they exert an irresistible pull on the imagination. This is an exciting frontier – and also one fraught with hard technical and social challenges that have yet to be solved. For early ventures in the Semantic Web arena, it’s not going to be easy, but the intellectual and technological challenges, and the potential financial rewards, glory, and benefit to society, are worth the effort and risk. And this is what all great technological revolutions are made of.

Semantic Web 2.0

Some people who have heard the term “Semantic Web” thrown around too much may think it is a buzzword, and they are right. But it is not just a buzzword – it actually has some substance behind it. That substance hasn’t emerged yet, but it will. Early critiques of the Semantic Web were right – the early vision did not leverage concepts such as folksonomy and user-contributed content at all. But that is largely because when the Semantic Web was originally conceived of Web 2.0 hadn’t happened yet. The early experiments that came out of research labs were geeky, to put it lightly, and impractical, but they are already being followed up by more pragmatic, user-friendly approaches.

Today’s Semantic Web – what we might call “Semantic Web 2.0” is a kinder, gentler, more social Semantic Web. It combines the best of the original vision with what we have all learned about social software and community in the last 10 years. Although much of this is still in the lab, it is already starting to trickle out. For example, recently Yahoo! started a pilot of the Semantic Web behind their food vertical. Other organizations are experimenting with using Semantic Web technology in parts of their applications, or to store or map data. But that’s just the beginning.

The Google Factor

Entrepreneurs, venture capitalists and technologists are increasingly starting to see these opportunities. Who will be the “Google of the Semantic Web?” – will it be Google itself? That’s doubtful. Like any entrenched incumbent, Google is heavily tied to a particular technology and worldview. And in Google’s case it is anything but semantic today. It would be easier for an upstart to take this position than for Google to port their entire infrastructure and worldview to a Semantic Web way of thinking.

If it is going to be Google it will most likely be by acquisition rather than by internal origination. And this makes more sense anyway – for Google is in a position where they can just wait and buy the winner, at almost any price, rather than competing in the playing field. One thing to note however is that Google has at least one product offering that shows some potential for becoming a key part of the Semantic Web. I am speaking of Google Base, Google's open database which is meant to be a registry for structured data so that it can be found in Google search. But Google Base does not conform to or make use of the many open standards of the Semantic Web community. That may or may not be a good thing, depending on your perspective.

Of course the downside of Google waiting to join the mainstream Semantic Web community until after the winner is announced is very large – once there is a winner it may be too late for Google to beat them. The winner of the Semantic Web race could very well unseat Google. The strategists at Google are probably not yet aware of this but as soon as they see significant traction around a major Semantic Web play it will become of interest to them.

In any case, I think there won't be just one winner, there will be several major Semantic Web companies in the future, focusing on different parts of the opportunity. And you can be sure that if Google gets into the game, every major portal will need to get into this space at some point or risk becoming irrelevant. There will be demand and many acquisitions. In many ways the Semantic Web will not be controlled by just one company -- it will be more like a fabric that connects them all together.

Context is King -- The Nature of Knowledge

It should be clear by now that the Semantic Web is all about enabling software (and people) to work with knowledge more intelligently. But what is knowledge? Knowledge is not just information. It is meaningful information – it is information plus context. For example, if I simply say the word “sem” to you, it is just raw information, it is not knowledge. It probably has no meaning to you other than a particular set of letters that you recognize and a sound you can pronounce, and the mere fact that this information was stated by me.

But if I tell you that “sem” it is the Tibetan word for “mind” then suddenly, “*sem* means mind in Tibetan” to you. If I further tell you that Tibetans have about as many words for “mind” as Eskimos have for “snow,” this is further meaning. This is context, in other words, knowledge, about the sound “*sem*.” The sound is raw information. When it is given context it becomes a word, a word that has meaning, a word that is connected to concepts in your mind – it becomes knowledge. By connecting raw information to context, knowledge is formed.

Once you have acquired a piece of knowledge such as “*sem* means mind in Tibetan,” you may then also form further knowledge about it. For example, you may form the memory, “Nova said that ‘*sem* means mind in Tibetan.’” You might also connect the word “sem” to networks of further concepts you have about Tibet and your understanding of what the word “mind” means.

The mind is the organ of meaning – mind is where meaning is stored, interpreted and created. Meaning is not “out there” in the world, it is purely subjective, it is purely mental. Meaning is almost equivalent to mind in fact. For the two never occur separately. Each of our individual minds has some way of internally representing meaning -- when we read or hear a word that we know, our minds connect that to a network of concepts about it and at that moment it means something to us.

Digging deeper, if you are really curious, or you happen to know Greek, you may also find that a similar sound occurs in the Greek word, *sēmantikós* – which means “having meaning” and in turn is the root of the English word “semantic” which means “pertaining to or arising from meaning.” That’s an odd coincidence! “Sem” occurs in Tibetan word for mind, and the English and Greek words that all relate to the concepts of “meaning” and “mind.” Even stranger is that not only do these words have a similar sound, they have a similar meaning.

With all this knowledge at your disposal, when you then see the term “Semantic Web” you may be able to infer that it has something to do with adding “meaning” to the Web. However, if you were a Tibetan, perhaps you might instead think the term had something to do with adding “mind” to the Web. In either case you would be right!

Discovering New Connections

We've discovered a new connection -- namely that there is an implicit connection between “sem” in Greek, English and Tibetan: they all relate to meaning and mind. It's not a direct, explicit connection – it's not evident unless you dig for it. But it's a useful tidbit of knowledge once it's found. Unlike the direct migration of the sound “sem” from Greek to English, there may not have ever been a direct transfer of this sound from Greek to Sanskrit to Tibetan. But in a strange and unexpected way, they are all connected. This connection wasn't necessarily explicitly stated by anyone before, but was uncovered by exploring our network of concepts and making inferences.

The sequence of thought about “sem” above is quite similar to kind of intellectual reasoning and discovery that the actual Semantic Web seeks to enable software to do automatically. How is this kind of reasoning and discovery enabled? The Semantic Web provides a set of technologies for formally defining the context of information. Just as the Web relies on a standard formal specification for “marking up” information with formatting codes that enable any applications that understand those codes to format the information in the same way, the Semantic Web relies on new standards for “marking up” information with statements about its context – its meaning – that enable any applications to understand, and reason about, the meaning of those statements in the same way.

By applying semantic reasoning agents to large collections of semantically enhanced content, all sorts of new connections may be inferred, leading to new knowledge, unexpected discoveries and useful additional context around content. This kind of reasoning and discovery is already taking place in fields from drug discovery and medical research, to homeland security and intelligence.

The Semantic Web is not the only way to do reasoning, inference and discovery -- but it certainly will make these capabilities easier to build, and more widely accessible. And of course, with this improvement will come new questions about how to assess and explain how various inferences were made, and how to protect privacy as our inferencing capabilities begin to extend across ever more sources of public and private data. I don't have the answers to these questions, but others are working on them and I have confidence that solutions will be arrived at over time.

Smart Data

By marking up information with metadata that formally codifies its context, we can make the data itself "smarter." The data becomes self-describing. When you get a piece of data you also get the necessary metadata for understanding it. For example, if I sent you a document containing the word "sem" in it, I could add markup around that word indicating that it is the word for "mind" in the Tibetan language.

Similarly, a document containing mentions of "Radar Networks" could contain metadata indicating that "Radar Networks" is an Internet company, not a product or a type of radar technology. A document about a person could contain semantic markup indicating that they are residents of a certain city, experts on Italian cooking, and members of a certain profession. All of this could be encoded as metadata in a form that software could easily understand. The data carries more information about its own meaning.

The alternative to smart data would be for software to actually read and understand natural language as well as humans. But that's really hard. To correctly interpret raw natural language, software would have to be developed that knew as much as a human being. But think about how much teaching and learning is required to raise a human being to the point where they can read at an adult level. It is likely that similar training would be necessary to build software that could do that. So far that goal has not been achieved, although some attempts have been made. While decent progress in natural language understanding has been made, most software that can do this is limited around particular vertical domains, and it's brittle -- it doesn't do a good job of making sense of terms and forms of speech that it wasn't trained to parse and make sense of.

Instead of trying to make software a million times smarter than it is today, it is much easier to just encode more metadata about what our information means. That turns out to be less work in the end. And there's an added benefit to this approach -- the meaning exists with the data and travels with it. It is independent of any one software program -- all software can access it. And because the meaning of information is stored with the information itself, rather than in the software, the software doesn't have to be enormous to be smart. It just has to know the basic language for interpreting the semantic metadata it finds on the information it works with.

Smart data enables relatively dumb software to be smarter with less work. That's an immediate benefit. And in the long-term as software actually gets smarter, smart data will make it easier for it to start learning and exploring on its own. So it's a win-win approach. Start with by adding semantic metadata to data, end up with smarter software.

Making Statements About the World

Metadata comes down to making statements about the world in a manner that machines, and perhaps even humans, can understand unambiguously. The same piece of metadata should be interpreted in the same way by different applications and readers.

There are many kinds of statements that can be made about information to provide it with context. For example, you can state a definition such as “person” means “a human being or a legal entity.” You can state an assertion such as “Sue is a human being.” You can state a rule such that “if x is a human being, then x is a person.”

From these statements it can then be inferred that “Sue is a person.” This inference is so obvious to you and me that it seems trivial, but most software today cannot do this. It doesn’t know what a person is, let alone what a name is. But if software could do this, then it could for example, automatically organize documents by the people they are related to, or discover connections between people who were mentioned in a set of documents, or it could find documents about people who were related to particular topics, or it could give you a list of all the people mentioned in a set of documents, or all the documents related to a person.

Of course this is a very basic example. But imagine if your software didn’t just know about people – it knew about most of the common concepts that occur in your life. Your software would then be able to help you work with your documents just about as intelligently as you are able to do by yourself, or perhaps even more intelligently, because you are just one person and you have limited time and energy but your software could work all the time, and in parallel, to help you.

Examples and Benefits

How could the existence of the Semantic Web and all the semantic metadata that defines it be really useful to everyone in the near-term?

Well, for example, the problem of email spam would finally be cured: your software would be able to look at a message and know whether it was meaningful and/or relevant to you or not.

Similarly, you would never have to file anything by hand again. Your software could automate all filing and information organization tasks for you because it would understand your information and your interests. It would be able to figure out when to file something in a single folder, multiple folders, or new ones. It would organize everything -- documents, photos, contacts, bookmarks, notes, products, music, video, data records -- and it would do it even better and more consistently than you could on your own. Your software wouldn't just organize stuff, it would turn it into knowledge by connecting it to more context. It could do this not just for individuals, but for groups, organizations and entire communities.

Another example: search would be vastly better: you could search conversationally by typing in everyday natural language and you would get precisely what you asked for, or even what you needed but didn’t know how to ask for correctly, and nothing else. Your

search engine could even ask you questions to help you narrow what you want. You would finally be able to converse with software in ordinary speech and it would understand you.

The process of discovery would be easier too. You could have software agent that worked as your personal recommendation agent. It would constantly be looking in all the places you read or participate in for things that are relevant to your past, present and potential future interests and needs. It could then alert you in a contextually sensitive way, knowing how to reach you and how urgently to mark things. As you gave it feedback it could learn and do a better job over time.

Going even further with this, semantically-aware software – software that is aware of context, software that understands knowledge – isn't just for helping you with your information, it can also help to enrich and facilitate, and even partially automate, your communication and commerce (when you want it to). So for example, your software could help you with your email. It would be able to recommend responses to messages for you, or automate the process. It would be able to enrich your messaging and discussions by automatically cross-linking what you are speaking about with related messages, discussions, documents, Web sites, subject categories, people, organizations, places, events, etc.

Shopping and marketplaces would also become better – you could search precisely for any kind of product, with any specific attributes, and find it anywhere on the Web, in any store. You could post classified ads and automatically get relevant matches according to your priorities, from all over the Web, or only from specific places and parties that match your criteria for who you trust. You could also easily invent a new custom data structure for posting classified ads for a new kind of product or service and publish it to the Web in a format that other Web services and applications could immediately mine and index without having to necessarily integrate with your software or data schema directly.

You could publish an entire database to the Web and other applications and services could immediately start to integrate your data with their data, without having to migrate your schema or their own. You could merge data from different data sources together to create new data sources without having to ever touch or look at an actual database schema.

Bumps on the Road

The above examples illustrate the potential of the Semantic Web today, but the reality on the ground is that the technology is still in the early phases of evolution. Even for experienced software engineers and Web developers, it is difficult to apply in practice. The main obstacles are twofold:

(1) The Tools Problem:

There are very few commercial-grade tools for doing anything with the Semantic Web today – Most of the tools for building semantically-aware applications, or for adding semantics to information are still in the research phase and were designed for expert computer scientists who specialize in knowledge representation, artificial intelligence, and machine learning.

These tools require a large learning curve to work with and they don't generally support large-scale applications – they were designed mainly to test theories and frameworks, not to actually apply them. But if the Semantic Web is ever going to become mainstream, it has to be made easier to apply – it has to be made more productive and accessible for ordinary software and content developers.

Fortunately, the tools problem is already on the verge of being solved. Companies such as my own venture, Radar Networks, are developing the next generation of tools for building Semantic Web applications and Semantic Web sites. These tools will hide most of the complexity, enabling ordinary mortals to build applications and content that leverage the power of semantics without needing PhD's in knowledge representation.

(2) The Ontology Problem:

The Semantic Web provides frameworks for defining systems of formally defined concepts called “ontologies,” that can then be used to connect information to context in an unambiguous way. Without ontologies, there really can be no semantics. The ontologies ARE the semantics, they define the meanings that are so essential for connecting information to context.

But there are still few widely used or standardized ontologies. And getting people to agree on common ontologies is not generally easy. Everyone has their own way of describing things, their own worldview, and let's face it nobody wants to use somebody else's worldview instead of their own. Furthermore, the world is very complex and to adequately describe all the knowledge that comprises what is thought of as “common sense” would require a very large ontology (and in fact, such an ontology exists – it's called Cyc and it is so large and complex that only experts can really use it today).

Even to describe the knowledge of just a single vertical domain, such as medicine, is extremely challenging. To make matters worse, the tools for authoring ontologies are still very hard to use – one has to understand the OWL language and difficult, buggy ontology

authoring tools in order to use them. Domain experts who are non-technical and not trained in formal reasoning or knowledge representation may find the process of designing ontologies frustrating using current tools. What is needed are commercial quality tools for building ontologies that hide the underlying complexity so that people can just pour their knowledge into them as easily as they speak. That's still a ways off, but not far off. Perhaps ten years at the most.

Of course the difficulty of defining ontologies would be irrelevant if the necessary ontologies already existed. Perhaps experts could define them and then everyone else could just use them? There are numerous ontologies already in existence, both on the general level as well as about specific verticals. However in my own opinion, having looked at many of them, I still haven't found one that has the right balance of coverage of the necessary concepts most applications need, and accessibility and ease-of-use by non-experts. That kind of balance is a requirement for any ontology to really go mainstream.

Furthermore, regarding the present crop of ontologies, what is still lacking is standardization. Ontologists have not agreed on which ontologies to use. As a result it's anybody's guess which ontology to use when writing a semantic application and thus there is a high degree of ontology diversity today. Diversity is good, but too much diversity is chaos.

Applications that use different ontologies about the same things don't automatically interoperate unless their ontologies have been integrated. This is similar to the problem of database integration in the enterprise. In order to interoperate, different applications that use different data schemas for records about the same things, have to be mapped to each other somehow – either at the application-level or the data-level. This mapping can be direct or through some form of middleware.

Ontologies can be used as a form of semantic middleware, enabling applications to be mapped at the data-level instead of the applications-level. Ontologies can also be used to map applications at the applications level, by making ontologies of Web services and capabilities, by the way. This is an area in which a lot of research is presently taking place.

The OWL language can express mappings between concepts in different ontologies. But if there are many ontologies, and many of them partially overlap, it is a non-trivial task to actually make the mappings between their concepts.

Even though concept A in ontology one and concept B in ontology two may have the same names, and even some of the same properties, in the context of the rest of the concepts in their respective ontologies they may imply very different meanings. So simply mapping them as equivalent on the basis of their names is not adequate, their connections to all the other concepts in their respective ontologies have to be considered as well. It quickly becomes complex. There are some potential ways to automate the construction of mappings between ontologies however – but they are still experimental. Today, integrating ontologies requires the help of expert ontologists, and to be honest,

I'm not sure even the experts have it figured out. It's more of an art than a science at this point.

Darwinian Selection of Ontologies

All that is needed for mainstream adoption to begin is for a large body of mainstream content to become semantically tagged and accessible. This will cause whatever ontology is behind that content to become popular.

When developers see that there is significant content and traction around a particular ontology, they will use that ontology for their own applications about similar concepts, or at least they will do the work of mapping their own ontology to it, and in this way the world will converge in a Darwinian fashion around a few main ontologies over time.

These main ontologies will then be worth the time and effort necessary to integrate them on a semantic level, resulting in a cohesive Semantic Web. We may in fact see Darwinian natural selection take place not just at the ontology level, but at the level of pieces of ontologies.

A certain ontology may do a good job of defining what a person is, while another may do a good job of defining what a company is. These definitions may be used for a lot of content, and gradually they will become common parts of an emergent meta-ontology comprised of the most-popular pieces from thousands of ontologies. This could be great or it could be a total mess. Nobody knows yet. It's a subject for further research.

Making Sense of Ontologies

Since ontologies are so important, it is helpful to actually understand what an ontology really is, and what it looks like. An ontology is a system of formally defined related concepts. For example, a simple ontology is this set of statements such as this:

A human is a living thing.

A person is a human.

A person may have a first name.

A person may have a last name.

A person must have one and only one date of birth.

A person must have a gender.

A person may be socially related to another person.

A friendship is a kind of social relationship.

A romantic relationship is a kind of friendship.

A marriage is a kind of romantic relationship.

A person may be in a marriage with only one other person at a time.

A person may be employed by an employer.

An employer may be a person or an organization.

An organization is a group of people.

An organization may have a product or a service.

A company is a type organization.

We've just built a simple ontology about a few concepts: humans, living things, persons, names, social relationships, marriages, employment, employers, organizations, groups, products and services. Within this system of concepts there is particular logic, some constraints, and some structure. It may or may not correspond to your worldview, but it is a worldview that is unambiguously defined, can be communicated, and is internally logically consistent, and that is what is important.

The Semantic Web approach provides an open-standard language, OWL, for defining ontologies. OWL also provides for a way to define instances of ontologies. Instances are assertions within the worldview that a given ontology provides. In other words OWL provides a means to make statements that connect information to the ontology so that software can understand its meaning unambiguously. For example, below is a set of statements based on the above ontology:

There exists a person x.

Person x has a first name "Sue"

Person x has a last name "Smith"

Person x has a full name "Sue Smith"

Sue Smith was born on June 1, 2005

Sue Smith has a gender: female

Sue Smith has a friend: Jane, who is another person.

Sue Smith is married to: Bob, another person.

Sue Smith is employed by Acme, Inc, a company.

Acme Inc. has a product, Widget 2.0.

The set of statements above, plus the ontology they are connected to, collectively comprise a knowledge base that, if represented formally in the OWL markup language, could be understood by any application that speaks OWL in the precise manner that it was intended to be understood.

Making Metadata

The OWL language provides a way to markup any information such as a data record, an email message or a Web page with metadata in the form of statements that link particular words or phrases to concepts in the ontology. When software applications that understand OWL encounter the information they can then reference the ontology and figure out exactly what the information means – or at least what the ontology says that it means.

But something has to add these semantic metadata statements to the information – and if it doesn't add them or adds the wrong ones, then software applications that look at the information will get the wrong idea. And this is another challenge – how will all this metadata get created and added into content? People certainly aren't going to add it all by hand!

Fortunately there are many ways to make this easier. The best approach is to automate it using special software that goes through information, analyzes the meaning and adds semantic metadata automatically. This works today, but the software has to be trained or provided with rules and that takes some time. It also doesn't scale cost-effectively to vast data-sets.

Alternatively, individuals can be provided with ways to add semantics themselves as they author information. When you post your resume in a semantically-aware job board, you could fill out a form about each of your past jobs, and the job board would connect that data to appropriate semantic concepts in an underlying employment ontology. As an end-user you would just fill out a form like you are used to doing; under-the-hood the job board would add the semantics for you.

Another approach is to leverage communities to get the semantics. We already see communities that are adding basic metadata “tags” to photos, news articles and maps. Already a few simple types of tags are being used pseudo-semantically: subject tags and geographical tags. These are primitive forms of semantic metadata. Although they are not expressed in OWL or connected to formal ontologies, they are at least semantically typed with prefixes or by being entered into fields or specific namespaces that define their types.

Tagging by Example

There may also be another solution to the problem of how to add semantics to content in the not to distant future. Once a suitable amount of content has been marked up with semantic metadata, it may be possible, through purely statistical forms of machine learning, for software to begin to learn how to do a pretty good job of marking up new content with semantic metadata.

For example, if the string "Nova Spivack" is often marked up with semantic metadata stating that it indicates a person, and not just any person but a specific person that is abstractly represented in a knowledge base somewhere, then when software applications encounter a new non-semantically enhanced document containing strings such as "Nova Spivack" or "Spivack, Nova" they can make a reasonably good guess that this indicates that same specific person, and they can add the necessary semantic metadata to that effect automatically.

As more and more semantic metadata is added to the Web and made accessible it constitutes a statistical training set that can be learned and generalized from. Although humans may need to jump-start the process with some manually semantic tagging, it might not be long before software could assist them and eventually do all the tagging for them. Only in special cases would software need to ask a human for assistance -- for example when totally new terms or expressions were encountered for the first several times.

The technology for doing this learning already exists -- and actually it's not very different from how search engines like Google measure the community sentiment around web pages. Each time something is semantically tagged with a certain meaning that constitutes a "vote" for it having that meaning. The meaning that gets the most votes wins. It's an elegant, Darwinian, emergent approach to learning how to automatically tag the Web.

One thing is certain, if communities were able to tag things with more types of tags, and these tags were connected to ontologies and knowledge bases, that would result in a lot of semantic metadata being added to content in a completely bottom-up, grassroots manner, and this in turn would enable this process to start to become automated or at least machine-augmented.

Getting the Process Started

But making the user experience of semantic tagging easy (and immediately beneficial) enough that regular people will do it, is a challenge that has yet to be solved. However, it will be solved shortly. It has to be. And many companies and researchers know this and are working on it right now. This does have to be solved to get the process of jump-starting the Semantic Web started.

I believe that the Tools Problem – the lack of commercial grade tools for building semantic applications – is essentially solved already (although the products have not hit the market yet; they will within a few years at most). The Ontology Problem is further from being solved. I think the way this problem will be solved is through a few “killer apps” that result in the building up of a large amount of content around particular ontologies within particular online services.

Where might we see this content initially arising? In my opinion it will most likely be within vertical communities of interest, communities of practice, and communities of purpose. Within such communities there is a need to create a common body of knowledge and to make that knowledge more accessible, connected and useful.

The Semantic Web can really improve the quality of knowledge and user-experience within these domains. Because they are communities, not just static content services, these organizations are driven by user-contributed content -- users play a key role in building content and tagging it. We already see this process starting to take place in communities such as Flickr, del.icio.us, the Wikipedia and Digg. We know that communities of people do tag content, and consume tagged content, if it is easy and beneficial enough for them to do so.

In the near future we may see miniature Semantic Webs arising around particular places, topics and subject areas, projects, and other organizations. Or perhaps, like almost every form of new media in recent times, we may see early adoption of the Semantic Web around online porn -- what might be called "the *semantic web*." Whether you like it or not, it is a fact that pornography was one of the biggest drivers of early mainstream adoption of personal video technology, CD-ROMs, and also of the Internet and the Web. But I think it probably is not necessary this time around.

While, I'm sure that the so-called "*semantic web*" could become better from the Semantic Web, it isn't going to be the primary driver of adoption of the Semantic Web. That's probably a good thing -- the world can just skip over that phase of development and benefit from this technology with both hands so to speak.

The World Wide Database

In some ways one could think of the Semantic Web as “the world wide database” – it does for the meaning of data records what the Web did for the formatting documents. But that’s just the beginning. It actually turns documents into richer data records. It turns unstructured data into structured data. All data becomes structured data in fact. The structure is not merely defined structurally, but it is defined semantically.

In other words, it’s not merely that for example, a data record or document can be defined in such a way as to specify that it contains a certain field of data with a certain label at a certain location – it defines what that field of data actually means in an unambiguous, machine understandable way. If all you want is a Web of data, XML is good enough. But

if you want to make that data interoperable and machine understandable then you need RDF and OWL – the Semantic Web.

Like any database, the Semantic Web, or rather the myriad mini-semantic-webs that will comprise it, have to overcome the challenge of data integration. Ontologies provide a better way to describe and map data, but the data still has to be described and mapped, and this does take some work. It's not a magic bullet.

The Semantic Web makes it easier to integrate data, but it doesn't completely remove the data integration problem altogether. I think the eventual solution to this problem will combine technology and community folksonomy oriented approaches.

The Semantic Web in Historical Context

Let's transition now and zoom out to see the bigger picture. The Semantic Web provides technologies for representing and sharing knowledge in new ways. In particular, it makes knowledge more accessible to software, and thus to other people. Another way of saying this is that it liberates knowledge from particular human minds and organizations – it provides a way to make knowledge explicit, in a standardized format that any application can understand. This is quite significant. Let's put this in historical perspective.

Before the invention of the printing press, there were two ways to spread knowledge – one was orally, the other was in some symbolic form such as art or written manuscripts. The oral transmission of knowledge had limited range and a high error-rate, and the only way to learn something was to meet someone who knew it and get them to tell you. The other option, symbolic communication through art and writing, provided a means to communicate knowledge independently of particular people – but it was only feasible to produce a few copies of any given artwork or manuscript because they had to be copied by hand. So the transmission of knowledge was limited to small groups or at least small audiences. Basically, the only way to get access to this knowledge was to be one of the lucky few who could acquire one of its rare physical copies.

The invention of the printing press changed this – for the first time knowledge could be rapidly and cost-effectively mass-produced and mass-distributed. Printing made it possible to share knowledge with ever-larger audiences. This enabled a huge transformation for human knowledge, society, government, technology – really every area of human life was transformed by this innovation.

The World Wide Web made the replication and distribution of knowledge even easier – With the Web you don't even have to physically print or distribute knowledge anymore, the cost of distribution is effectively zero, and everyone has instant access to everything from anywhere, anytime. That's a lot better than having to lug around a stack of physical books. Everyone potentially has whatever knowledge they need with no physical barriers. This has been another huge transformation for humanity – and it has affected every area of human life. Like the printing press, the Web fundamentally changed the economics of knowledge.

The Semantic Web is the next big step in this process – it will make all the knowledge of the human race accessible to software. For the first time, non-human things (software applications) will be able to start working with human knowledge to do things (for humans) on their own. This is a big leap – a leap like the emergence of a new species, or the symbiosis of two existing species into a new form of life.

The printing press and the Web changed the economics of replicating, distributing and accessing knowledge. The Semantic Web changes the economics of processing knowledge. Unlike the printing press and the Web, the Semantic Web enables knowledge to be processed by non-human things.

In other words, humans don't have to do all the thinking on their own, they can be assisted by software. Of course we humans have to at least first create the software (until we someday learn to create software that is smart enough to create software too), and we have to create the ontologies necessary for the software to actually understand anything (until we learn to create software that is smart enough to create ontologies too), and we have to add the semantic metadata to our content in various ways (until our software is smart enough to do this for us, which it almost is already). But once we do the initial work of making the ontologies and software, and adding semantic metadata, the system starts to pick up speed on its own, and over time the amount of work we humans have to do to make it all function decreases. Eventually, once the system has encoded enough knowledge and intelligence, it starts to function without needing much help, and when it does need our help, it will simply ask us and learn from our answers.

This may sound like science-fiction today, but in fact a lot of this is already built and working in the lab. The big hurdle is figuring out how to get this technology to mass-market. That is probably as hard as inventing the technology in the first place. But I'm confident that someone will solve it eventually.

Once this happens the economics of processing knowledge will truly be different than it is today. Instead of needing an actual real-live expert, the knowledge of that expert will be accessible to software that can act as their proxy – and anyone will be able to access this virtual expert, anywhere, anytime. It will be like the Web – but instead of just information being accessible, the combined knowledge and expertise of all of humanity will also be accessible, and not just to people but also to software applications.

The Question of Consciousness

The Semantic Web literally enables humans to share their knowledge with each other and with machines. It enables the virtualization of human knowledge and intelligence. With respect to machines, in doing this, it will lend machines “minds” in a certain sense – namely in that they will at least be able to correctly interpret the meaning of information and replicate the expertise of experts.

But will these machine-minds be conscious? Will they be aware of the meanings they interpret, or will they just be automatons that are simply following instructions without

any awareness of the meanings they are processing? I doubt that software will ever be conscious, because from what I can tell consciousness -- or what might be called the sentient awareness of awareness itself as well as other things that are sensed -- is an immaterial phenomena that is as fundamental as space, time and energy -- or perhaps even more fundamental. But this is just my personal opinion after having searched for consciousness through every means possible for decades. It just cannot be found to be something, yet it is definitely and undeniably taking place.

Consciousness can be exemplified through the analogy of space (but unlike space, consciousness has this property of being aware, it's not a mere lifeless void). We all agree space is there, but nobody can actually point to it somewhere, and nobody can synthesize space. Space is immaterial and fundamental. It is primordial. So is electricity. Nobody really knows what electricity is ultimately, but if you build the right kind of circuit you can channel it and we've learned a lot about how to do that.

Perhaps we may figure out how to channel consciousness like we channel electricity with some sort of synthetic device someday, but I think that is highly unlikely. I think if you really want to create consciousness it's much easier and more effective to just have children. That's something ordinary mortals can do today with the technology they were born with. Of course when you have children you don't really "create" their consciousness, it seems to be there on its own. We don't really know what it is or where it comes from, or when it arises there. We know very little about consciousness today. Considering that it is the most fundamental human experience of all, it is actually surprising how little we know about it!

In any case, until we truly delve far more deeply into the nature of the mind, consciousness will be barely understood or recognized, let alone explained or synthesized by anyone. In many eastern civilizations there are multi-thousand year traditions that focus quite precisely on the nature of consciousness. The major religions have all universally concluded that consciousness is beyond the reach of science, beyond the reach of concepts, beyond the mind entirely. All those smart people analyzing consciousness for so long, and with such precision, and so many methods of inquiry, may have a point worth listening to.

Whether or not machines will ever actually "know" or be capable of being conscious of that meaning or expertise is a big debate, but at least we can all agree that they will be able to interpret the meaning of information and rules if given the right instructions. Without having to be conscious, software will be able to process semantics quite well -- this has already been proven. It's working today.

While consciousness is and may always be a mystery that we cannot synthesize -- the ability for software to follow instructions is an established fact. In its most reduced form, the Semantic Web just makes it possible to provide richer kinds of instructions. There's no magic to it. Just a lot of details. In fact, to play on a famous line, "it's semantics all the way down."

The Semantic Web does not require that we make conscious software. It just provides a way to make slightly more intelligent software. There's a big difference. Intelligence is simply a form of information processing, for the most part. It does not require consciousness -- the actual awareness of what is going on -- which is something else altogether.

While highly intelligent software may need to sense its environment and its own internal state and reason about these, it does not actually have to be conscious to do this. These operations are for the most part simple procedures applied vast numbers of time and in complex patterns. Nowhere in them is there any consciousness nor does consciousness suddenly emerge when suitable levels of complexity are reached.

Consciousness is something quite special and mysterious. And fortunately for humans, it is not necessary for the creation of more intelligent software, nor is it a byproduct of the creation of more intelligent software, in my opinion.

The Intelligence of the Web

So the real point of the Semantic Web is that it enables the Web to become more intelligent. At first this may seem like a rather outlandish statement, but in fact the Web is already becoming intelligent, even without the Semantic Web.

Although the intelligence of the Web is not very evident at first glance, nonetheless it can be found if you look for it. This intelligence doesn't exist across the entire Web yet, it only exists in islands that are few and far between compared to the vast amount of information on the Web as a whole. But these islands are growing, and more are appearing every year, and they are starting to connect together. And as this happens the collective intelligence of the Web is increasing.

Perhaps the premier example of an "island of intelligence" is the Wikipedia, but there are many others: The Open Directory, portals such as Yahoo and Google, vertical content providers such as CNET and WebMD, commerce communities such as Craigslist and Amazon, content oriented communities such as LiveJournal, Slashdot, Flickr and Digg and of course the millions of discussion boards scattered around the Web, and social communities such as MySpace and Facebook. There are also large numbers of private islands of intelligence on the Web within enterprises -- for example the many online knowledge and collaboration portals that exist within businesses, non-profits, and governments.

What makes these islands "intelligent" is that they are places where people (and sometimes applications as well) are able to interact with each other to help grow and evolve collections of knowledge. When you look at them close-up they appear to be just like any other Web site, but when you look at what they are doing as a whole -- these services are *thinking*. They are learning, self-organizing, sensing their environments, interpreting, reasoning, understanding, introspecting, and building knowledge. These are the activities of minds, of intelligent systems.

The intelligence of a system such as the Wikipedia exists on several levels – the individuals who author and edit it are intelligent, the groups that help to manage it are intelligent, and the community as a whole – which is constantly growing, changing, and learning – is intelligent.

Flickr and Digg also exhibit intelligence. Flickr's growing system of tags is the beginnings of something resembling a collective visual sense organ on the Web. Images are perceived, stored, interpreted, and connected to concepts and other images. This is what the human visual system does. Similarly, Digg is a community that collectively detects, focuses attention on, and interprets current news. It's not unlike a primitive collective analogue to the human facility for situational awareness.

There are many other examples of collective intelligence emerging on the Web. The Semantic Web will add one more form of intelligent actor to the mix – intelligent applications. In the future, after the Wikipedia is connected to the Semantic Web, as well as humans, it will be authored and edited by smart applications that constantly look for new information, new connections, and new inferences to add to it.

Although the knowledge on the Web today is still mostly organized within different islands of intelligence, these islands are starting to reach out and connect together. They are forming trade-routes, connecting their economies, and learning each other's languages and cultures. The next-step will be for these islands of knowledge to begin to share not just content and services, but also their knowledge -- what they know about their content and services. The Semantic Web will make this possible, by providing an open format for the representation and exchange of knowledge and expertise.

When applications integrate their content using the Semantic Web they will also be able to integrate their context, their knowledge – this will make the content much more useful and the integration much deeper. For example, when an application imports photos from another application it will also be able to import semantic metadata about the meaning and connections of those photos. Everything that the community and application know about the photos in the service that provides the content (the photos) can be shared with the service that receives the content. Better yet, there will be no need for custom application integration in order for this to happen: as long as both services conform to the open standards of the Semantic Web the knowledge is instantly portable and reusable.

Freeing Intelligence from Silos

Today much of the real value of the Web (and in the world) is still locked away in the minds of individuals, the cultures of groups and organizations, and application-specific data-silos. The emerging Semantic Web will begin to unlock the intelligence in these silos by making the knowledge and expertise they represent more accessible and understandable.

It will free knowledge and expertise from the narrow confines of individual minds, groups and organizations, and applications, and make them not only more interoperable,

but more portable. It will be possible for example for a person or an application to share everything they know about a subject of interest as easily as we share documents today. In essence the Semantic Web provides a common language (or at least a common set of languages) for sharing knowledge and intelligence as easily as we share content today.

The Semantic Web also provides standards for searching and reasoning more intelligently. The SPARQL query language enables any application to ask for knowledge from any other application that speaks SPARQL. Instead of mere keyword search, this enables semantic search. Applications can search for specific types of things that have particular attributes and relationships to other things.

In addition, standards such as SWRL provide formalisms for representing and sharing axioms, or rules, as well. Rules are a particular kind of knowledge – and there is a lot of it to represent and share, for example procedural knowledge, and logical structures about the world. An ontology provides a means to describe the basic entities, their attributes and relations, but rules enable you to also make logical assertions and inferences about them. Without going into a lot of detail about rules and how they work here, the important point to realize is that they are also included in the framework. All forms of knowledge can be represented by the Semantic Web.

Zooming Way, Waaaay Out

So far in this article, I've spent a lot of time talking about plumbing – the pipes, fluids, valves, fixtures, specifications and tools of the Semantic Web. I've also spent some time on illustrations of how it might be useful in the very near future to individuals, groups and organizations. But where is it heading after this? What is the long-term potential of this and what might it mean for the human race on a historical time-scale?

For those of you who would prefer not to speculate, stop reading here. For the rest of you, I believe that the true significance of the Semantic Web, on a long-term timescale is that it provides an infrastructure that will enable the evolution of increasingly sophisticated forms of collective intelligence. Ultimately this will result in the Web itself becoming more and more intelligent, until one day the entire human species together with all of its software and knowledge will function as something like a single worldwide distributed mind – a global mind.

Just the like the mind of a single human individual, the global mind will be very chaotic, yet out of that chaos will emerge cohesive patterns of thought and decision. Just like in an individual human mind, there will be feedback between different levels of order – from individuals to groups to systems of groups and back down from systems of groups to groups to individuals. Because of these feedback loops the system will adapt to its environment, and to its own internal state.

The coming global mind will collectively exhibit forms of cognition and behavior that are the signs of higher-forms of intelligence. It will form and react to concepts about its “self” – just like an individual human mind. It will learn and introspect and explore the

universe. The thoughts it thinks may sometimes be too big for any one person to understand or even recognize them – they will be comprised of shifting patterns of millions of pieces of knowledge.

The Role of Humanity

Every person on the Internet will be a part of the global mind. And collectively they will function as its consciousness. I do not believe some new form of consciousness will suddenly emerge when the Web passes some threshold of complexity. I believe that humanity IS the consciousness of the Web and until and unless we ever find a way to connect other lifeforms to the Web, or we build conscious machines, humans will be the only form of consciousness of the Web.

When I say that humans will function as the consciousness of the Web I mean that we will be the things in the system that know. The knowledge of the Semantic Web is what is known, but what knows that knowledge has to be something other than knowledge. A thought is knowledge, but what knows that thought is not knowledge, it is consciousness, whatever that is. We can figure out how to enable machines to represent and use knowledge, but we don't know how to make them conscious, and we don't have to. Because we are already conscious.

As we've discussed earlier in this article, we don't need conscious machines, we just need more intelligent machines. Intelligence – at least basic forms of it – does not require consciousness. It may be the case that the very highest forms of intelligence require or are capable of consciousness. This may mean that software will never achieve the highest levels of intelligence and probably guarantees that humans (and other conscious things) will always play a special role in the world; a role that no computer system will be able to compete with. We provide the consciousness to the system. There may be all sorts of other intelligent, non-conscious software applications and communities on the Web; in fact there already are, with varying degrees of intelligence. But individual humans, and groups of humans, will be the only consciousness on the Web.

The Collective Self

Although the software of the Semantic Web will not be conscious we can say that system as a whole contains or is conscious to the extent that human consciousnesses are part of it. And like most conscious entities, it may also start to be self-conscious.

If the Web ever becomes a global mind as I am predicting, will it have a “self?” Will there be a part of the Web that functions as its central self-representation? Perhaps someone will build something like that someday, or perhaps it will evolve. Perhaps it will function by collecting reports from applications and people in real-time – a giant collective *zeitgeist*.

In the early days of the Web portals such as Yahoo! provided this function -- they were almost real-time maps of the Web and what was happening. Today making such a map is nearly impossible, but services such as Google Zeitgeist at least attempt to provide approximations of it. Perhaps through random sampling it can be done on a broader scale.

My guess is that the global mind will need a self-representation at some point. All forms of higher intelligence seem to have one. It's necessary for understanding, learning and planning. It may evolve at first as a bunch of competing self-representations within particular services or subsystems within the collective. Eventually they will converge or at least narrow down to just a few major perspectives. There may also be millions of minor perspectives that can be drilled down into for particular viewpoints from these top-level "portals."

The collective self, will function much like the individual self – as a mirror of sorts. Its function is simply to reflect. As soon as it exists the entire system will make a shift to a greater form of intelligence – because for the first time it will be able to see itself, to measure itself, as a whole. It is at this phase transition when the first truly global collective self-mirroring function evolves, that we can say that the transition from a bunch of cooperating intelligent parts to a new intelligent whole in its own right has taken place.

I think that the collective self, even if it converges on a few major perspectives that group and summarize millions of minor perspectives, will be community-driven and highly decentralized. At least I hope so – because the self-concept is the most important part of any mind and it should be designed in a way that protects it from being manipulated for nefarious ends. At least I hope that is how it is designed.

Programming the Global Mind

On the other hand, there are times when a little bit of adjustment or guidance is warranted – just as in the case of an individual mind, the collective self doesn't merely reflect, it effectively guides the interpretation of the past and present, and planning for the future.

One way to change the direction of the collective mind, is to change what is appearing in the mirror of the collective self. This is a form of programming on a vast scale – When this programming is dishonest or used for negative purposes it is called "propaganda," but there are cases where it can be done for beneficial purposes as well. An example of this today is public service advertising and educational public television programming. All forms of mass-media today are in fact collective social programming. When you realize this it is not surprising that our present culture is violent and messed up – just look at our mass-media!

In terms of the global mind, ideally one would hope that it would be able to learn and improve over time. One would hope that it would not have the collective equivalent of psycho-social disorders. To facilitate this, just like any form of higher intelligence, it may need to be taught, and even parented a bit. It also may need a form of therapy now and

then. These functions could be provided by the people who participate in it. Again, I believe that humans serve a vital and irreplaceable role in this process.

How It All Might Unfold

Now how is this all going to unfold? I believe that there are a number of key evolutionary steps that Semantic Web will go through as the Web evolves towards a true global mind:

1. Representing individual knowledge. The first step is to make individuals' knowledge accessible to themselves. As individuals become inundated with increasing amounts of information, they will need better ways of managing it, keeping track of it, and re-using it. They will (or already do) need "personal knowledge management."

2. Connecting individual knowledge. Next, once individual knowledge is represented, it becomes possible to start connecting it and sharing it across individuals. This stage could be called "interpersonal knowledge management."

3. Representing group knowledge. Groups of individuals also need ways of collectively representing their knowledge, making sense of it, and growing it over time. Wikis and community portals are just the beginning. The Semantic Web will take "group minds" to the next level -- it will make the collective knowledge of groups far richer and more re-usable.

4. Connecting group knowledge. This step is analogous to connecting individual knowledge. Here, groups become able to connect their knowledge together to form larger collectives, and it becomes possible to more easily access and share knowledge between different groups in very different areas of interest.

5. Representing the knowledge of the entire Web. This stage -- what might be called "the global mind" -- is still in the distant future, but at this point in the future we will begin to be able to view, search, and navigate the knowledge of the entire Web as a whole. The distinction here is that instead of a collection of interoperating but separate intelligent applications, individuals and groups, the entire Web itself will begin to function as one cohesive intelligent system. The crucial step that enables this to happen is the formation of a collective self-representation. This enables the system to see itself as a whole for the first time.

How it May be Organized

I believe the global mind will be organized mainly in the form of bottom-up and lateral, distributed emergent computation and community -- but it will be facilitated by certain key top-down services that help to organize and make sense of it as a whole. I think this future Web will be highly distributed, but will have certain large services within it as well -- much like the human brain itself, which is organized into functional sub-systems for processes like vision, hearing, language, planning, memory, learning, etc.

As the Web gets more complex there will come a day when nobody understands it anymore – after that point we will probably learn more about how the Web is organized by learning about the human mind and brain – they will be quite similar in my opinion. Likewise we will probably learn a tremendous amount about the functioning of the human brain and mind by observing how the Web functions, grows and evolves over time, because they really are quite similar in at least an abstract sense.

The internet and its software and content is like a brain, and the state of its software and the content is like its mind. The people on the Internet are like its consciousness. Although these are just analogies, they are actually useful, at least in helping us to envision and understand this complex system. As the field of general systems theory has shown us in the past, systems at very different levels of scale tend to share the same basic characteristics and obey the same basic laws of behavior. Not only that, but evolution tends to converge on similar solutions for similar problems. So these analogies may be more than just rough approximations, they may be quite accurate in fact.

The future global brain will require tremendous computing and storage resources -- far beyond even what Google provides today. Fortunately as Moore's Law advances the cost of computing and storage will eventually be low enough to do this cost-effectively. However even with much cheaper and more powerful computing resources it will still have to be a distributed system. I doubt that there will be any central node because quite simply no central solution will be able to keep up with all the distributed change taking place. Highly distributed problems require distributed solutions and that is probably what will eventually emerge on the future Web.

Someday perhaps it will be more like a peer-to-peer network, comprised of applications and people who function sort of like the neurons in the human brain. Perhaps they will be connected and organized by higher-level super-peers or super-nodes which bring things together, make sense of what is going on and coordinate mass collective activities. But even these higher-level services will probably have to be highly distributed as well. It really will be difficult to draw boundaries between parts of this system, they will all be connected as an integral whole.

In fact it may look very much like a grid computing architecture – in which all the services are dynamically distributed across all the nodes such that at any one time any node might be working on a variety of tasks for different services. My guess is that because this is the simplest, most fault-tolerant, and most efficient way to do mass computation, it is probably what will evolve here on Earth.

The Ecology of Mind

Where we are today in this evolutionary process is perhaps equivalent to the rise of early forms of hominids. Perhaps *Australopithecus* or *Cro-Magnon*, or maybe the first *Homo Sapiens*. Compared to early man, the global mind is like the rise of 21st century megacities. A lot of evolution has to happen to get there. But it probably will happen, unless humanity self-destructs first, which I sincerely hope we somehow manage to avoid. And

this brings me to a final point. This vision of the future global mind is highly technological; however I don't think we'll ever accomplish it without a new focus on ecology.

Ecology probably conjures up images of hippies and biologists, or maybe hippies who are biologists, or at least organic farmers, for most people, but in fact it is really the science of living systems and how they work. And any system that includes living things is a living system. This means that the Web is a living system and the global mind will be a living system too. As a living system, the Web is an ecosystem and is also connected to other ecosystems. In short, ecology is absolutely essential to making sense of the Web, let alone helping to grow and evolve it.

In many ways the Semantic Web and the collective minds, and the global mind, that it enables, can be seen as an ecosystem of people, applications, information and knowledge. This ecosystem is very complex, much like natural ecosystems in the physical world. An ecosystem isn't built, it's grown, and evolved. And similarly the Semantic Web, and the coming global mind, will not really be built, they will be grown and evolved. The people and organizations that end up playing a leading role in this process will be the ones that understand and adapt to the ecology most effectively.

In my opinion ecology is going to be the most important science and discipline of the 21st century – it is the science of healthy systems. What nature teaches us about complex systems can be applied to every kind of system – and especially the systems we are evolving on the Web. In order to ever have a hope of evolving a global mind, and all the wonderful levels of species-level collective intelligence that it will enable, we have to not destroy the planet before we get there. Ecology is the science that can save us, not the Semantic Web (although perhaps by improving collective intelligence, it can help).

Ecology is essentially the science of community – whether biological, technological or social. And community is a key part of the Semantic Web at every level: communities of software, communities of people, and communities of groups. In the end the global mind is the ultimate human community. It is the reward we get for finally learning how to live together in peace and balance with our environment.

The Necessity of Sustainability

The point of this discussion of the relevance of ecology to the future of the Web, and my vision for the global mind, is that I think that it is clear that if the global mind ever emerges it will not be in a world that is anything like what we might imagine. It won't be like the Borg in Star Trek, it won't be like living inside of a machine. Humans won't be relegated to the roles of slaves or drones. Robots won't be doing all the work. The entire world won't be coated with silicon. We won't all live in a virtual reality. It won't be one of these technological dystopias.

In fact, I think the global mind can only come to pass in a much greener, more organic, healthier, more balanced and sustainable world. Because it will take a long time for the

global mind to emerge, if humanity doesn't figure out how to create that sort of a world, it will wipe itself out sooner or later, but certainly long before the global mind really happens. Not only that, but the global mind will be smart by definition, and hopefully this intelligence will extend to helping humanity manage its resources, civilizations and relationships to the natural environment.

The Smart Environment

The global mind also needs a global body so to speak. It's not going to be an isolated homunculus floating in a vat of liquid that replaces the physical world! It will be a smart environment that ubiquitously integrates with our physical world. We won't have to sit in front of computers or deliberately logon to the network to interact with the global mind. It will be everywhere.

The global mind will be physically integrated into furniture, houses, vehicles, devices, artworks, and even the natural environment. It will sense the state of the world and different ecosystems in real-time and alert humans and applications to emerging threats. It will also be able to allocate resources intelligently to compensate for natural disasters, storms, and environmental damage – much in the way that the air traffic control systems allocates and manages airplane traffic. It won't do it all on its own, humans and organizations will be a key part of the process.

Someday the global mind may even be physically integrated into our bodies and brains, even down the level of our DNA. It may in fact learn how to cure diseases and improve the design of the human body, extending our lives, sensory capabilities, and cognitive abilities. We may be able to interact with it by thought alone. At that point it will become indistinguishable from a limited form of omniscience, and everyone may have access to it. Although it will only extend to wherever humanity has a presence in the universe, within that boundary it will know everything there is to know, and everyone will be able to know any of it they are interested in.

Enabling a Better World

By enabling greater forms of collective intelligence to emerge we really are helping to make a better world, a world that learns and hopefully understands itself well enough to find a way to survive. We're building something that someday will be wonderful – far greater than any of us can imagine. We're helping to make the species and the whole planet more intelligent. We're building the tools for the future of human community. And that future community, if it ever arrives, will be better, more self-aware, more sustainable than the one we live in today.

I should also mention that knowledge is power, and power can be used for good or evil. The Semantic Web makes knowledge more accessible. This puts more power in the hands of the many, not just the few. As long as we stick to this vision -- we stick to

making knowledge open and accessible, using open standards, in as distributed a fashion as we can devise, then the potential power of the Semantic Web will be protected against being coopted or controlled by the few at the expense of the many. This is where technologists really have to be socially responsible when making development decisions. It's important that we build a more open world, not a less open world. It's important that we build a world where knowledge, integration and unification are balanced with respect for privacy, individuality, diversity and freedom of opinion.

But I am not particularly worried that the Semantic Web and the future global mind will be the ultimate evil – I don't think it is likely that we will end up with a system of total control dominated by evil masterminds with powerful Semantic Web computer systems to do their dirty work. Statistically speaking, criminal empires don't last very long because they are run by criminals who tend to be very short-sighted and who also surround themselves with other criminals who eventually unseat them, or they self-destruct. It's possible that the Semantic Web, like any other technology, may be used by the bad guys to spy on citizens, manipulate the world, and do evil things. But only in the short-term.

In the long-term either our civilization will get tired of endless successions of criminal empires and realize that the only way to actually survive as a species is to invent a form of government that is immune to being taken over by evil people and organizations, or it will self-destruct. Either way, that is a hurdle we have to cross before the global mind that I envision can ever come about. Many civilizations came before ours, and it is likely that ours will not be the last one on this planet. It may in fact be the case that a different form of civilization is necessary for the global mind to emerge, and is the natural byproduct of the emergence of the global mind.

We know that the global mind cannot emerge anytime soon, and therefore, if it ever emerges then by definition it must be in the context of a civilization that has learned to become sustainable. A long-term sustainable civilization is a non-evil civilization. And that is why I think it is a safe bet to be so optimistic about the long-term future of this trend.