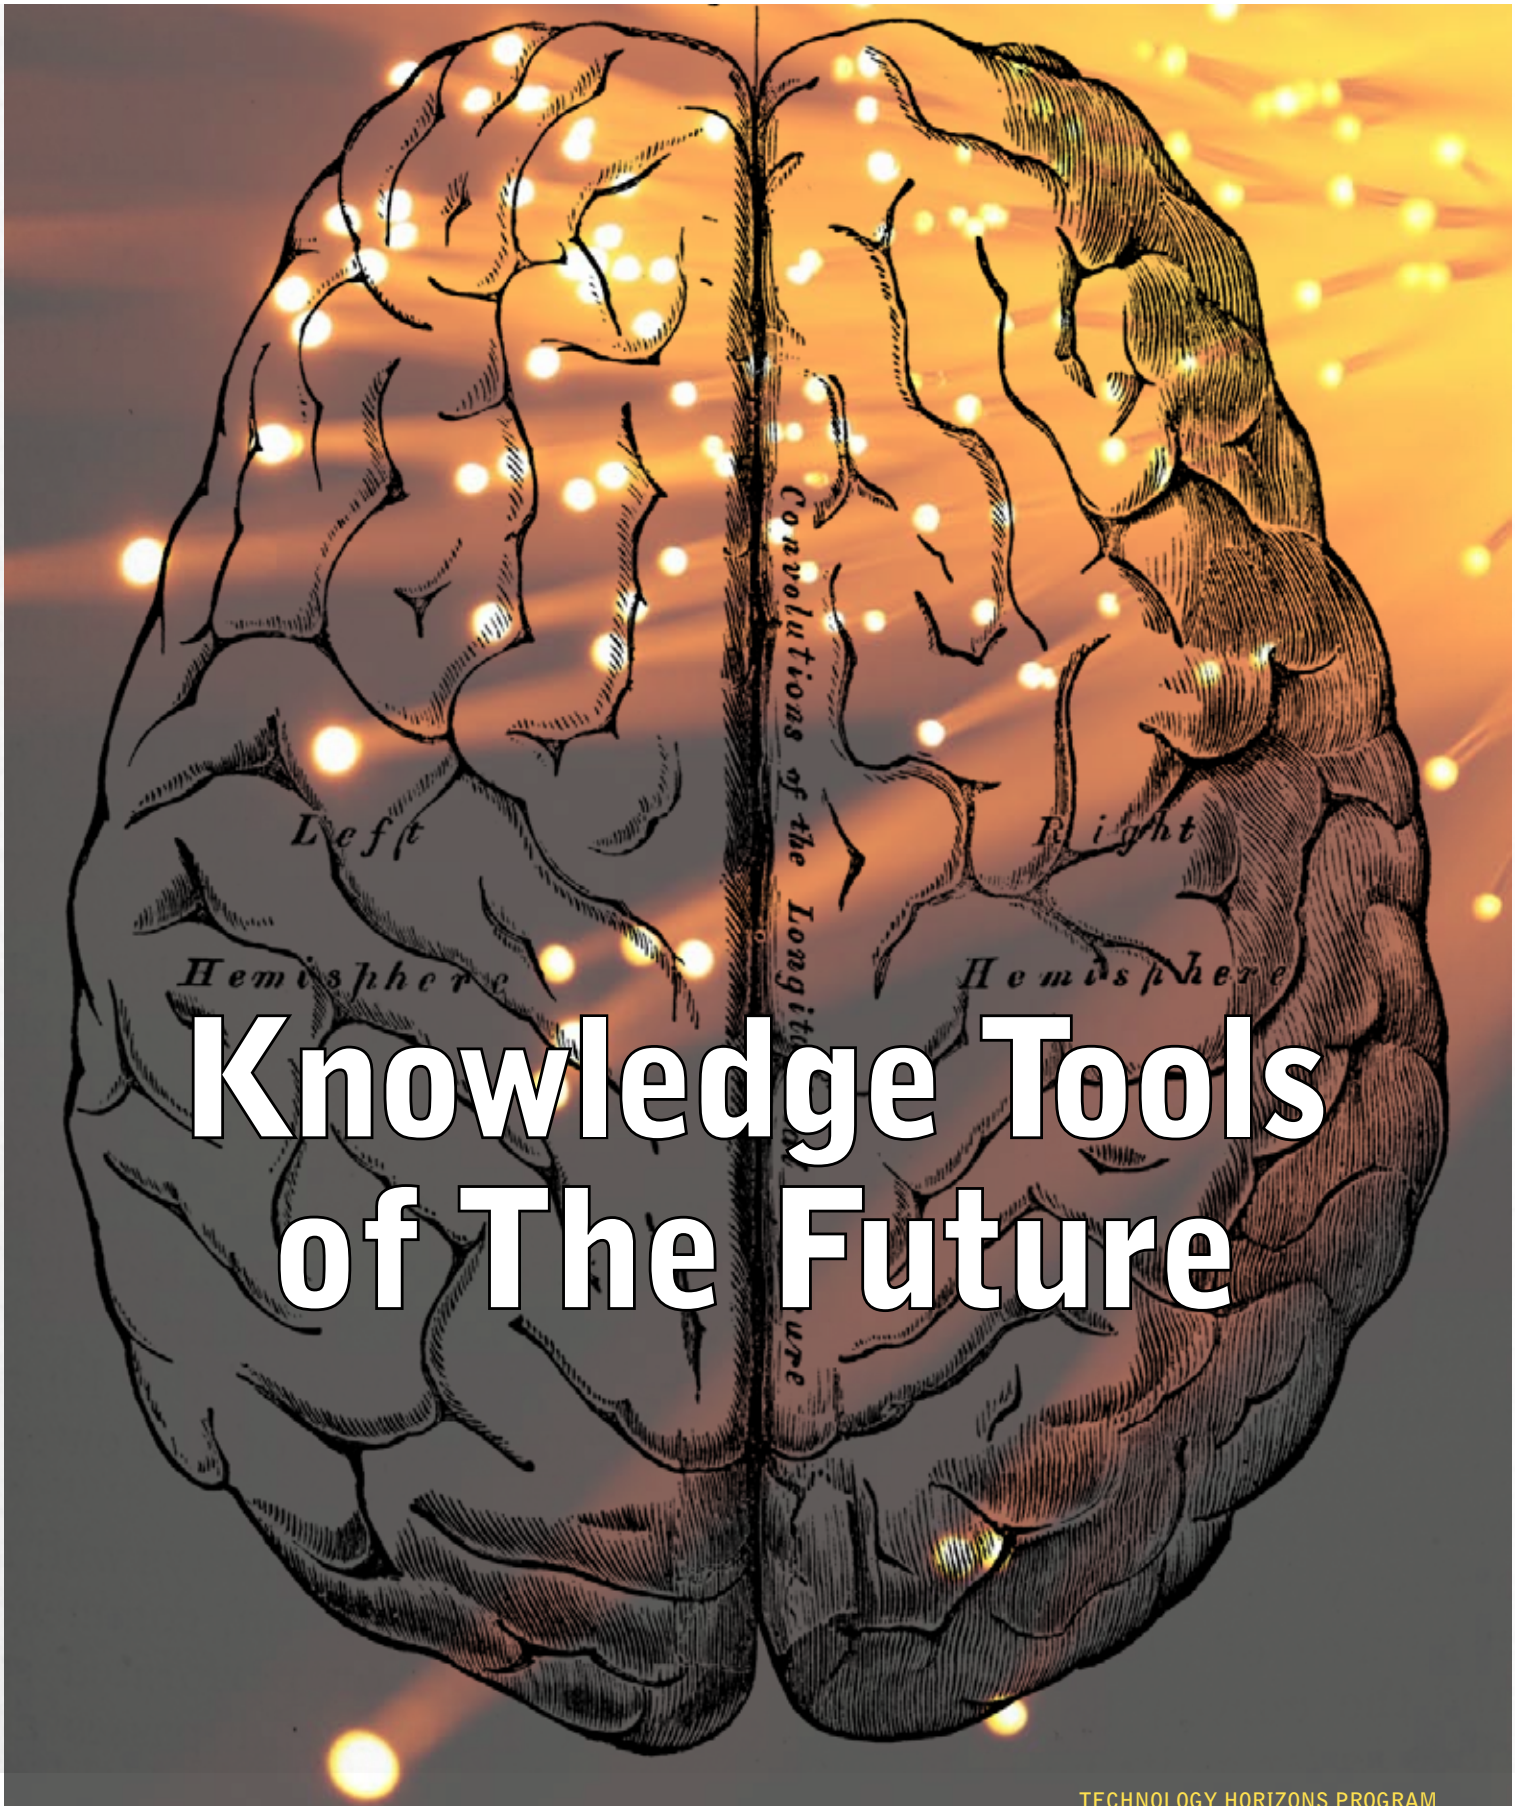

Knowledge Tools of The Future

INSTITUTE FOR THE FUTURE

TECHNOLOGY HORIZONS PROGRAM
November 2008
124 University Avenue, 2nd Floor | Palo Alto, California 94301
650.854.6322 **www.iftf.org**
SR-1179

about the ...

TECHNOLOGY HORIZONS PROGRAM

The Technology Horizons Program combines a deep understanding of technology and societal forces to identify and evaluate discontinuities and innovations in the next three to ten years. We help organizations develop insights and strategic tools to better position themselves for the future. Our approach to technology forecasting is unique—we put humans in the middle of our forecasts. Understanding humans as consumers, workers, householders, and community members allows IFTF to help companies look beyond technical feasibility to identify the value in new technologies, forecast adoption and diffusion patterns, and discover new market opportunities and threats. For more information about the Technology Horizons Program, contact Sean Ness at sness@iftf.org, (650) 233-9517.

INSTITUTE FOR THE FUTURE

The Institute for the Future is an independent, nonprofit strategic research group with more than 40 years of forecasting experience. The core of our work is identifying emerging trends and discontinuities that will transform global society and the global marketplace. We provide our members with insights into business strategy, design process, innovation, and social dilemmas. Our research generates the foresight needed to create insights that lead to action. Our research spans a broad territory of deeply transformative trends, from health and health care to technology, the workplace, and human identity. The Institute for the Future is located in Palo Alto, California.

Author: Alex Pang

Peer Reviewers: Zoe Finkel, Marina Gorbis, Charles Grossel

Editor: Lisa Mumbach

Art Direction and Design: Robin Bogott, Robin Weiss

CONTENTS

INTRODUCTION

The global economy is becoming increasingly knowledge-driven and these days knowledge management is a critical challenge for many organizations. Huge complex systems have been built to store and organize information, monitor consumer data, track resources, process payroll, maintain centralized databases, and standardize organizational practices across larger and larger global companies. The more information and the larger the company, the bigger and more complex the systems.

Will there come a time when the machines take over, when humans become redundant? A generation ago we might have said yes, that the sheer power of artificial intelligence (AI) would soon leave the human brain far behind. But it turns out it's not that simple. Computers are brilliant at handling routine tasks, enforcing standardization and workflows, and connecting workers around the world. But for companies trying to differentiate themselves based on innovation rather than efficiency and lower cost, offloading tasks onto computers, or relying on knowledge management tools, won't guarantee success. The high-margin, cutting edge of the knowledge economy is less about knowledge management than about knowledge creation—creativity and innovation. Creativity is fast becoming the competitive advantage for leading-edge companies.

As creativity becomes more important, one of the most powerful tools is still the human brain, with its ability to hold tacit knowledge—to synthesize information, see patterns, derive insight, and create something new. This ability has yet to be programmed into a computer system. What's more, we are discovering that it's not just the lone human brain working on its own that creates the best innovation these days, but the human brain in collaboration with other human brains, sometimes with many thousands or millions of others, in social networks enabled by the Internet. In other words, there's a social aspect to knowledge, creativity, and innovation that we are just learning to tap. It is this social aspect of knowledge that the new knowledge tools are designed to leverage.

WHAT ARE THE NEW KNOWLEDGE TOOLS?

They don't consist of a single device or system but an array of devices, systems, methodologies, and services sometimes called the "intelligent web." The new tools are applications that exploit things like semantic Web functions, microformats, natural language searching, data-mining, machine learning, and recommendation agents to provide a more productive and intuitive experience for the user. In other words, the new knowledge tools aren't meant to replace humans, they are meant to enable humans to do what they do best—creativity and innovation—without having to do the heavy lifting of brute information processing. We call these new tools "lightweight" because they are designed to be used on existing platforms, require little additional infrastructure, and don't force individuals or organizations to radically change how they work.

WHAT DOES THIS MEAN FOR ORGANIZATIONS?

Organizations are in the middle of a paradigm shift from machine-heavy knowledge management tools designed to maximize efficiency and standardize organizational practices to technically lightweight, human-centered instruments that facilitate creativity and collaboration. It is this human creativity that will differentiate businesses in the future.

Today's generation of knowledge tools—interrelational databases like Freebase and DBPedia, social networks like OpenSocial, information accessing tools like Snapshots—are flexible and relatively easy for individuals and groups to learn, and thus can serve as "outboard" brains. The result is a kind of human–machine symbiosis in which processing-heavy tasks are offloaded onto software, leaving users to collaborate more freely with each other in search of insight, creativity, and experience.

Even as this new generation of knowledge tools evolves, traditional knowledge management will continue to matter, just as agriculture and manufacturing still have a place in service economies. Companies will continue to track resources, process payroll, maintain centralized databases, and manage IT infrastructures. But the new leading edge will not be organized mainly around management, but discovery; using systems to augment human imagination and creativity.

AS WE MAY (YET) THINK: The Origin of Knowledge Tools

In 1945, Vannevar Bush published an article in *Atlantic Monthly* titled, "As We May Think." A former MIT professor and computer pioneer, Bush had spent the war directing the Organization for Scientific Research and Development (OSRD), a government body that sponsored research that produced radar, jet aircraft, and the atomic bomb. OSRD had done a remarkable job of linking basic science and academic scientists to the war effort. Bush was worried that in the future, such fever-paced innovation would be stifled by the sheer volume of information—articles, reports, and papers—scientists and engineers would have to manage. This mattered, Bush believed, because scientific innovation had won World War II, and would continue to be the foundation of future victories.

The solution, he argued, was to develop technologies to help scientists keep track of all this information. Bush envisioned a future in which a yet-to-be-invented device enabled scientists to do science more effectively, to keep up with colleagues' work, and to develop new insights and innovations. Increasing specialization wouldn't come at the expense of breadth; innovation wouldn't be pushed aside as scientists struggled to keep up to date. Bush called this device the Memex. In an era when computers were the size of buildings, the Memex would be an intimate, desk-sized device.

The Memex was never built, but the *idea* of a system that recorded a person's mental journeys and the unique connections they made inspired innovations like the graphical user interface, the personal computer, hypertext publishing, and the Internet. Sixty years later, the dream of creating tools to help us focus our attention, sharpen our memories, boost our creativity, and see connections between disparate things—in short, to help us think—remains as vivid as ever. In an economy driven by innovation and knowledge, the need for such tools has become more and more important. A generation ago, there was some speculation that artificial brains and "strong AI" would provide these kinds of tools; indeed, some feared they'd do their work too well, and make humans obsolete. Even a decade ago, huge knowledge management systems promised to give companies a competitive advantage by crystallizing their knowledge in massive centralized databases.

The Memex was never built, but the *idea* of a system that recorded a person's mental journeys and the unique connections they made inspired innovations like the graphical user interface, the personal computer, hypertext publishing, and the Internet.

Today, we view intelligence and knowledge work in a different way. We have a greater respect for human intelligence and the way it handles ambiguity, sees patterns, and creates connections. We also recognize that much valuable knowledge is tacit—knowledge known by an individual that is difficult to communicate to the wider organization. We can't teach tacit knowledge or describe it, but we can learn it and do it. Finally, we see knowledge as more of a social construct: it's not just locked up in individual brains, but shared by groups, some of the groups quite large. In many ways, these groups don't even know what they know. The challenge is how to distill that knowledge and make it useful.

Our tools are starting to reflect these insights. We see it in systems that combine the analytical and processing capabilities of computers with the tacit intelligence of humans. We also see it in services such as social networking sites like OpenSocial and interrelational databases like Freebase that socialize personal knowledge management and support group collaboration.

Ultimately, these emerging lightweight knowledge tools will be used in different ways by disparate collaborators and in the process will reshape traditional knowledge chains. In their purest form, knowledge chains consist of a clear progression of discrete steps—creation, publication, organization, filtering, and consumption—each produced, more or less, by its own set of people, tools, and systems. These kinds of knowledge *chains* are being replaced by knowledge *ecologies*, in which all of the steps involved in knowledge creation are happening all at once, and any given user, with the help of emerging knowledge tools, may be at once the creator and the consumer of knowledge, and indeed the publisher and organizer as well. Such knowledge ecologies allow us to move from tool to tool and weave together different parts of the previously linear knowledge process in new and interesting ways.

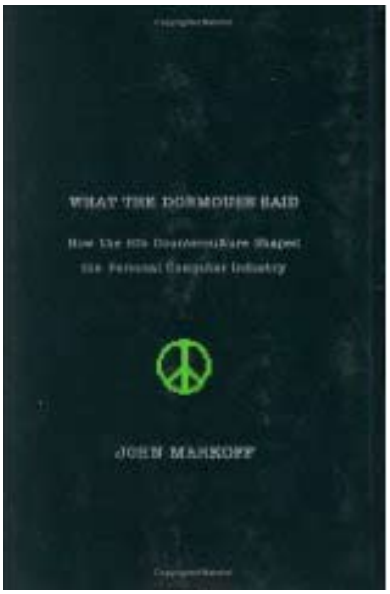

THE TRIUMPH OF LIGHTWEIGHT SYSTEMS:
From Complexity to Simplicity

In his now-classic book *What the Dormouse Said*[1], John Markoff describes three very different Silicon Valley groups that helped create the personal computer: Stanford University–based computer scientists; Stanford Research Institute (SRI) scientists using computers to augment human intelligence; and a loose federation of techies in the Mid-Peninsula Open University. Each had a different vision of how computers could be used. Stanford scientists wanted to create artificial intelligences that would eventually rival humans. At SRI, the Augmentation Research Center built complex systems accessed through novel means, such as mice and graphical interfaces. The Open University dreamed of creating simple computers that regular people could use. Researchers from all three eventually came together to create the first generation of personal computers.

Since then, AI and specialized expert systems have been used in fields as diverse as medical diagnostics, image analysis, and finance, while the SRI vision continues to shape the lives of power users who are not necessarily programmers but want to use the best, fastest, and most powerful technologies. As a result, computers and information technology are as pervasive in the workplace and home as the Stanford and SRI groups imagined. But it's the Open University's vision of computers—improvisational, simple, and user-directed—that has brought us closest to creating systems that complement human intelligence.

Sometimes called Web 2.0 or the intelligent Internet, many of the most successful services aren't very technically sophisticated: a social knowledge tool like del.icio.us is notable for its simplicity as well as its success. Indeed, one of the signal trends in knowledge tools is that simplicity is beating sophistication. Why do lightweight approaches succeed? Because they're built on existing systems and they don't require users to be computer geniuses to make them work. It doesn't require a year to learn the new system.

The rising popularity of StumbleUpon, for example, a channel-surfing approach to finding interesting content on the Internet, could be attributed to the fact that channel-surfing is a very familiar pattern of behavior for a majority of people in the Western world. This observation becomes important when we look to improve methods of accessing and producing knowledge. If you consider how dependent we have become on our current knowledge tools, it is clear that improvements are more likely to be adopted if presented in familiar forms.

Lightweight, social knowledge tools are popular and successful because they have three virtues:

- They're simple, which makes them easy to use, and quick to evolve.
- They're sociable, which allows them to be used by a diverse range of users and communities.
- They have symbiotic relationships with their users.

Since the knowledge tools of the future will be built with these features, let's explore them in greater detail.

INTENTIONAL SIMPLICITY: Good Enough and Easy to Use

Lightweight social knowledge systems offer two significant virtues related to their simplicity. First, both users and programmers find them much easier to work with than more complex systems. Second, their simplicity allows them to be used in a variety of contexts by different groups of users with different levels of technical know-how.

Even though "lightweight" approaches often provide results that are not perfect, the results they provide are good enough. Consider the example of search engines. Users can often do a good job of sorting through short lists, selecting the most relevant result from an ambiguous search, or adding terms to refine a search whose results are too broad. Even better, people can bring context to results. It's easier to select the right search item yourself from a list of items than to explain to a system exactly what you want, or to learn how to use a complex system well enough to construct extremely complex queries. It may be worth spending the time required to go from good to great if you're dealing with legal, patent, chemical, or similar databases. For most work, however, especially work with constantly evolving parameters, it's better to construct a system that will behave predictably for general use.

This helps explain why some technology companies have succeeded with relatively lightweight learning systems, while complex approaches have faltered. Google's PageRank algorithm is relatively simple—it ignores the prepositions and conjunctions of the keywords entered. Incorporating this kind of extra information into its rankings would have made the system only marginally more efficient. Similarly, Amazon doesn't build advanced predictive models for each customer; it calculates correlations between purchases of the products. In contrast, the search company PowerSet has an ambitious program designed to match the semantic meaning of search queries with the semantic meaning of sentences in Web documents, using natural language processing. It sounds impressive, but it also means that the system has evolved slowly, and been hard to apply to a wide variety of data sets: it currently only works on Wikipedia.

Lurking behind the success of these lightweight models is the obvious point that some things are easy for humans to do but difficult or even impossible for computers to do at a reliable level of quality. Lightweight approaches work in part because they exploit uniquely human capabilities. Lightweight knowledge tools let humans deal with context and ambiguity, leaving computers to deal with processing and scale.

Lightweight approaches work in part because humans generally don't need a lot of

data to make good decisions. In fact, humans tend to throw away a lot of information when making choices, and focus on just a couple of points or rules. Psychologists at the University of Oregon have shown empirical support for this theory of a "mental spam filter." Just as successful lightweight knowledge tools focus on a handful of criteria when, say, searching for relevant documents or related books, successful human decision-making hinges on knowing what to pay attention to rather than paying attention to everything. This suggests that the best knowledge systems of the future will continue to help us focus on the right things, in simple ways, and help us share these with collaborators.

SOCIABILITY: Enabling Diversity and Encouraging Community

Lightweight systems also integrate more easily into existing workflows and, perhaps most importantly, across different professional groups. SRI's NLS (oNLine System) required six months training to master. In some professions, familiarity with complex systems is obligatory. For lawyers, mastering Lexis-Nexis is part of the job; likewise, librarians must know the Online Computer Library Center (OCLC) and other electronic catalogs. For people in jobs that require bringing together multidisciplinary perspectives or working in diverse teams, however, it's not worth the time and resources it takes to learn the more complex systems. With lightweight systems, users carry the burden of figuring out how to integrate the tools into their daily work, but the payoff is greater flexibility and utility. Information tools are more likely to be used when they don't require people to redesign their lives around them.

Since lightweight knowledge tools don't require people to change their work habits, and don't require steep learning curves to master, they're more likely to be adopted and used by diverse groups. This makes it easier for people with different jobs, skill sets, and world views to use the same information in a variety of social and work contexts; while the same information is likely to be seen by a larger number of people. A large user base arithmetically increases the amount of information each user has access to in social network services, and it makes each user's informal network exponentially more complex and valuable. As the networks expand, they help to filter information, and to identify new things from novel sources.

Knowledge services with diverse bases can also help users identify other people with similar interests but usefully different perspectives. This is a significant benefit for users in network-centric industries such as product design or software development, where success depends on the ability to build collaborations, and useful expertise may reside both inside and outside a given organization. Traditionally, professionals practiced in the equivalent of villages: whether you were a professor of theoretical physics, a patent attorney, or a cardiovascular surgeon, you were connected to everyone else in your field through common graduate training, professional societies, and annual conferences. Indeed, part of what characterized professionals was their visibility to other professionals: for them, social knowledge tools were redundant.

In other fields, success depends not just on an ability to gather together people with common skills, but also to gather people with different skills. Product design, for example, may bring together engineers, psychologists, ethnographers, and artists. For such groups, social knowledge tools are extremely valuable. They're designed to share information in a technically light but socially rich manner, to facilitate connections to other people as well as to new information, and to support the formations of new communities of practice. In 2007, Google introduced OpenSocial, an application programming interface (API) allowing developers to create applications that access the information available through social networks and update multiple feeds across platforms instead of through each network individually.

This same principle helps power online idea markets, ranging from Amazon's Mechanical Turk to Innocentive. Mechanical Turk is a crowdsourcing service, in which some users post small tasks and others complete them—choosing the more attractive of two colors, or identifying faces in a picture, for example—for small amounts of money. Innocentive allows companies to pose more sophisticated scientific questions, which users then bid to answer. In both cases, a good reason for the success of these services is that they can draw on a wide range of people around the world, across demographic groups and scientific disciplines, including professional laboratories, lone scientists, and even talented hobbyists.

SYMBIOSIS: Changing the Definition of Machine–Human Cooperation

Lightweight knowledge tools have another important feature that separates them from more complex tools: they work and evolve in symbiosis with their users. Lightweight knowledge systems and humans train each other.

Probably the best-known symbiotic lightweight knowledge tools are collaborative filtering and recommendation systems. Developed in the early 1990s by MIT professor Pattie Maes, these systems aggregate user behavior (ratings, songs listened to, or purchases) across a service in order to build a model of relatedness between the items. Commercial services—Netflix, Amazon, and iTunes, among others—as well as knowledge management or social networking systems such as Last.fm, MyStrands, and Digg, all offer collaborative filtering. The Amazon recommendations system, for example, looks at users' purchases, and tells them about other items purchased by people who've ordered the same items. The more they're used, the more adept these systems become at making useful recommendations: experience makes them more adept at anticipating the tastes of their users.

Because of their symbiotic relationship, users sometimes find new and unexpected ways to use technologies. They develop new habits and capabilities. This, in turn, changes the definitions of flexibility, simplicity, and sociability.

Even in easy-to-use systems that we might call intuitively designed, we adapt our thinking to fit the tool, and often forget we have done so. Indeed, the definition of "intuitive" evolves over time. What seems difficult to one group of users will feel natural to another group years later. For example, the first Apple computer mouse, introduced in 1981, had one button because mice with more than one button easily confused novice users. In contrast, the "Mighty Mouse," introduced in 2005, has two buttons on top, a scroll ball, and buttons on the side—relatively complex compared to the original. (In fact, they're not even real buttons, but touch-sensitive regions that react when pressed—a design innovation that could only work if users were extremely familiar with the concept of and placement of buttons.) But for kids (and adults) who grew up with XBOX, Game Boy, and PlayStation controllers with dozens of buttons, and write lengthy text messages with punctuation and emoticons on numerical keypads, a device with four buttons and a scroll ball looks easy to use.

These kinds of behavioral changes are most easily visible in power users as by their very nature power users push the limits and change the capabilities of their technologies. But significant changes can also occur among general users, working with what seem to be fairly routine programs. Google, for example, has trained its users to search using keywords and exact phrases instead of full sentences—and that, in turn, may be reducing the appeal of more complex natural-language query systems. Users can become psychologically invested in both their sources of information and the tools they use to filter them. It is not a metaphorical stretch to argue that computers, networks, and our methods of accessing information online have become "outboard" brains—at least the part of the brain that catalogs and stores information.

One example of a system that's coevolving with users by incorporating more sophisticated analytical tools—and pushing the boundaries of simplicity—is Twine, a product by the start-up Radar Networks. Twine helps users organize, share, and discover information on the Web. Twine extracts semantic information, like the names of people, places, and organizations, from documents that are saved by users, and helps users see the connections among them. Users can share all their stored information with colleagues or the public, and specify when and how their social network data can be shared between sites and companies.

As the new tools come to enable cooperation, collaboration, and organization on a massive scale, the way knowledge is created will change irrevocably. Traditional knowledge chains controlled by gatekeeper institutions and organizations such as publishers or universities will give way to knowledge ecologies controlled by no one in particular but with access for all.

BREAKING TRADITION

From Knowledge Chains to Knowledge Ecologies

Information scientists and managers have traditionally thought of knowledge production as a chain made up of discrete links carried out by different groups: creation, publication, organization, filtering, and consumption. Of course, these are all descriptions of idealized activities that usually aren't completely discrete, nor are they always carried out by different people. In many creative enterprises, creators and consumers are the same, or tasks may combine several links at the same time.

With the emergence of the new lightweight knowledge tools, however, the metaphor of the knowledge chain is even less appropriate. Indeed, the diversity and nature of lightweight tools and the human–machine cooperation they enable, are creating a completely new knowledge landscape, one that requires subtlety, insight, and creativity to navigate. These days, knowledge is created by breaking the traditional chain and refashioning it in the image of a web, or more accurately, a knowledge ecology. In this kind of new knowledge ecology, at any stage—creation, publication, organization, filtering, or consumption—alterations can be made by anyone at any time, with numerous interconnections and reiterations between stages and users. Each element of the ecology is interconnected with every other element and on and on.

Historically, one industry has been organized in ways that closely follow the traditional knowledge chain model: publishing. Until now, the publishing industry has been highly formalized, with well-defined workflows, processes, and products, in which creation, publication, organization, filtering, and consumption are all discrete steps. The way lightweight knowledge tools have begun to change the nature of publishing illustrates the movement from knowledge chains to knowledge ecologies.

CREATION AND PUBLICATION: Growth of Decentralized Peer Production

We have recently witnessed a surge in decentralized peer production online, which has been well documented in the mainstream media as the "rise of the amateur." The phrase "peer production" refers to content that is submitted and edited by a community of users, often outside traditional gatekeeping institutions like publishing houses, companies, and universities. It is important to recognize that many online communities are actually engaging in peer *aggregation*: that is, they are organizing knowledge that has already passed through more traditional processes.

Peer-produced and peer-aggregated knowledge has become a trusted source in the past decade. Regardless of what one thinks about the reliability of Wikipedia, Wikipedia pages frequently dominate the search results of Google queries. This high rank reflects not necessarily the accuracy of the content, but a democratic process in which Internet users are visiting and linking to Wikipedia more than other sites. Companies have also made use of this new model of knowledge production. Newspapers, magazines, and television stations are now offering spaces for their audiences to share relevant information by means of comments or customized user blogs and submission pages. Indeed, many companies have online knowledge bases that are solely populated by the contributions of their consumer/users.

These processes will not displace traditional forms of knowledge production but will add to the diversity of information, and to the amount of information that must be filtered and vetted. Kevin Kelly puts it this way in his book *Out of Control*[2]: "The hive mind produces the raw material that smart design can work on." This smart design includes knowing when and how to pluck relevant content from raw material, and dedicating human and computer resources to the task. We will see the disappearance of top-down editorial control; instead, editors will serve as filters that help to highlight peer-produced and peer-aggregated content for larger audiences.

ORGANIZATION AND FILTERING: Integrating Machine and Human Judgments

Traditional media and new media will become even more interdependent, where each one contributes equally to the development of knowledge. Content created by a traditional source, such as a newspaper or scientific journal, is already finding its way onto blogs, where it is rebroadcast to loyal niche audiences. Similarly, successful ideas incubated on blogs and spread through social networks are being reported in traditional media more often. This relationship will become even more symbiotic as the new knowledge tools make it easier for users to be creators/publishers/consumers of knowledge all at the same time. An ecology of blogs, which repost and link to each other's content, provides another social filter for new information. An individual's collection of favorite blogs works like an adjustable aperture, letting in the appropriate amount and kind of information for a given task or problem.

The challenge, and opportunity for competitive advantage for organizations, is in tracking and making sense of the movement of information throughout these larger-scale knowledge ecologies. One method that has emerged is to combine mass quantities of small human judgments along with the automated weightings of those judgments, as in Google's PageRank algorithm. In addition, platforms like Amazon's Mechanical Turk open up possibilities for other businesses to aggregate human judgments at significantly lower costs than hiring full-time data annotators. The technological driver behind this trend will continue to be the lowered cost and ready availability of digital storage and processing power. These tools will move out from companies that focus entirely on technology to all companies, as organizations see the benefit of tracking, filtering, ranking, and organizing information on their own terms.

CONSUMPTION: Intertwined with Creation and Filtering

It used to be that the consumption of information was a somewhat passive activity. One read a book or an article in a newspaper, magazine, or scholarly journal, and if one had something to say to the author, one had to write a letter to the editor or publisher without the guarantee of a response. In today's growing knowledge ecology, however, one can respond to or even interact with information in real time. The act of consumption has become a feedback signal that determines access to other areas of the knowledge ecology.

Until only recently, an individual's interaction with a bit of knowledge left no mark on that knowledge. Reading a page of a newspaper, product specification, or a textbook had no effect on that page. In the last decade, however, we have witnessed the creation of massive databases storing our interactions with information on the Internet. Every click, every friend added on a social network, every post, comment or tag is stored, aggregated, and eventually processed by algorithms. The majority of the time we are not conscious of the creation and analysis of this data. Yet, the application of this attention and interaction data will powerfully shape our access to knowledge in the future.

Again, the driver for this change has been the decrease in the cost of storage and processing, which enables companies like Google to power large machine-learning algorithms to determine the most relevant pages for each search query. But companies outside the high-tech industry can reap these benefits as well. Organizations can keep track of how audiences respond to different options by running what are essentially large experiments online and in the real world. As they incorporate real-time feedback from larger audiences, such knowledge systems will become more predictive, but predictive in the sense that these systems will be better at guessing what kind of content the crowds of users are most interested in seeing and sharing. (We do not expect to see predictive systems that can reliably foresee changes that rely on variables external to the world of human judgments, like the availability of real-world resources and capital.) As these systems become more predictive, they will both enlarge the knowledge ecology and broaden an individual user's access to it.

FROM KNOWLEDGE MANAGEMENT TO CREATIVITY

Given that one of the defining features of the knowledge tools of the future is the co-evolution of tools and users, the future of knowledge cannot be derived simply by extrapolating technology trends. If knowledge tools co-evolve with their users, we need to start by thinking about how people are likely to use such tools and in what contexts, and how the tools and users will change each other.

A society whose leading economic edge is creative work rather than traditional knowledge management or manufacturing is one in which social knowledge tools will be essential for creating and sustaining economic value. Traditional computer databases and analytical tools are good at doing the kinds of things that characterize traditional knowledge management and information processing. Indeed, the first users of computers in business were in fields requiring massive amounts of sorting, tracking, and processing: things like billing, payroll, insurance claims, banking, and flight reservations. Computers have been exceptionally good at handling these kinds of activities.

In contrast, creative work requires attention to social and cultural context, the ability to make unexpected combinations, and a drive to focus on meaning rather than means. It is, in other words, a fundamentally human activity. The knowledge that informs innovative activity isn't just the product of combinatorics; it's defined by the creation of new insight out of diverse pieces of information. Creative knowledge is more like a verb than a noun: it is activity, not information; and it requires people—usually groups of people and often diverse groups—to bring it about. Using machines to take on the burden of processing tasks frees up time and energy for the human partner to provide insight and creative thinking that will in turn help the machines better predict and provide the information a user needs.

Simple, sociable, and symbiotic: these are the cornerstones of lightweight knowledge tools, and are the characteristics that will define their importance in the future. Competitive advantage in the most advanced service economies won't come from marginal gains in information management, but from the inspired leaps of creative work. George Mason University professor of public policy Christopher Hill argues in his article *The Post Scientific Society*[3], that "the United States is on the threshold of a new era in the development of advanced societies," one that he calls the "post-scientific society." Hill defines a post-scientific society as "a society in which cutting-edge success depends not on specialization, but on integration—on synthesis, design, creativity, and imagination." As he explains,

> A post-scientific society will have several key characteristics, the most important of which is that innovation leading to wealth generation and productivity growth will be based principally not on world leadership in fundamental research in the natural sciences and engineering, but on world-leading mastery of the creative powers of, and the basic sciences of, individual human beings, their societies, and their cultures. …
>
> In the post-scientific society, the creation of wealth and jobs based on innovation and new ideas will tend to draw less on the natural sciences and engineering and more on the organizational and social sciences, on the arts, on new business processes, and on meeting consumer needs based on niche production of specialized products and services in which interesting design and appeal to individual tastes matter more than low cost or radical new technologies. …
>
> [For businesses operating in post-scientific societies,] success will arise in part from the disciplined search for useful new knowledge that, regardless of its origins, can be integrated with intimate knowledge of cultures and consumer preferences. Networks of highly creative individuals and collaborating firms will devise and produce complex new systems that meet human needs in unexpectedly new and responsive ways.

Some of Hill's argument is not entirely convincing. Science is hardly the mechanistic and uncreative activity that Hill implies, and advanced economies are likely to become "post-scientific" in the same way they're "post-agricultural" or "post-industrial." Today these sectors don't make up the bulk of the economy or employ the largest numbers of people, but they're still indispensable. But Hill's larger point about post-scientific society is still worth thinking about.

In such societies, organizations operating at the cutting edge must move past knowledge management tools (which are largely designed to maximize the efficient management and communication of knowledge) to seek out new tools that encourage individual creativity, support collaboration among diverse groups, and support groups in creating new, shared meanings and knowledge—and to do so more or less transparently. These are the lightweight knowledge tools of the future.

We can already see attempts to create infrastructures to support this kind of post-scientific work. Consider the latest trends in office design, which feature heterogeneous spaces to stimulate creativity and support both collaborative work and personal concentration; crossroads spaces that encourage informal communication and serendipitous meetings between people; and flexible and temporary spaces that support nomadic workers and ad hoc groups. A century ago, the cutting-edge office space was a skyscraper, with corporate functions organized by floor. In an age where information-intensive industries competed on the basis of scale and efficiency, these orderly machine-like offices separated by function made perfect sense. In a post-scientific society in which innovation will be driven by creativity and insight, a more jumbled, cross-functional office space that supports such creativity and allows the symbiosis of humans and machines as they create and leverage a growing knowledge ecology is the sensible choice.

What could the future of knowledge tools be? If they continue to be simple, social, and symbiotic, how might they evolve in the future? How will they coevolve with users, and the groups that users form?

Knowledge tools have been influenced by a combination of social and economic trends, new technologies, and user needs and practices. Let's look first at those trends, then see how they might play off each other.

At the global scale, leading enterprises—particularly those based mainly in the post-scientific economies of North America and Europe—will continue to thrive based on their ability to support and harness innovation and creativity. Western-based multinational companies have already been squeezed out of the low end of some markets by Asian companies with lower labor and materials costs; as those competitors move from the low-cost, low-margin space up the value chain, it will become even more important for older companies to play the creative card. (Smaller companies will face even more pressure to maintain a strong creative pace.) To do this, enterprises will need to be able to draw on a variety of perspectives, in order to provide employees with tools that are flexible, easy to use, and simple enough to use without changing the way they normally work.

This will also become more important for nonprofits and organizations such as relief and development agencies. As they face newer, larger, or more extreme challenges in disaster areas, megacities, threatened ecosystems, or regions under political and economic transition, it will be necessary to develop more flexible and creative responses to immediate challenges that don't threaten the long-term viability of the communities they seek to help. (Smart relief organizations already avoid importing massive quantities of food into regions that still have functioning roads and farms, for fear of undermining those institutions and creating other, long-term problems.)

Technological changes will also alter the ways people and groups interact with information, and with each other. We will continue to see the proliferation of richer and more powerful mobile devices; the growth of new interface technologies, and spread of interfaces into new use contexts; and the expansion of always-on social and information services. This growth of ubiquitous and intimate computing will make our interactions with information more intimate, more oblique, and more immediate. The next generation of knowledge and creative tools won't live on desktops and be accessed from workplaces; in broadband and wireless-rich regions, they'll be as ever-present and accessible as cell phone signals.

This new interaction paradigm will be used by a generation of workers (NetGems) accustomed to living—and documenting their lives—online. For NetGens, knowledge tools won't be things attached to workplaces and jobs; they'll morph into extensions of themselves, and will be used casually and constantly. As Ross Mayfield puts it,

> NetGens think of the computer as a door, not a box. … With increasing mobility they tap groups for what they need to get done no matter where they are and make where they are matter. They Google, Flickr, Blog, contribute to Wikipedia, Socialtext it, Meetup, post, subscribe, feed, annotate and above all share. In other words, the Web is increasingly less about places and other nouns, but verbs.[4]

Finally, while the current generation of simple, social, and symbiotic tools has knit together people, groups, and knowledge, we'll see the first stirrings of systems that connect to physical objects as well—or, as Bruce Sterling puts it, "the physical instantiations of objects that have digital lives."[5] Rapid prototyping and the growth of what Sterling playfully calls "blogjects"—products with unique digital IDs, whose production, use histories, functional extensions, and hacks are recorded, shared, and incorporated into their descendants—will start to bring Web 2.0 and collective creative tools off the screen and the world of ideas, and into the world of things as well.

AS WE MAY THINK ... AND CREATE

What will these trends produce? The continued pressure for organizations to support and exploit creativity; the growth of technologies that are more intimate and supportive of users; the emergence of a generation of workers who are digital natives; and the growth of technologies that will move ideas between information and object and back; all will play out in three big ways:

- **The growth of creative co-presence.** Today's creative and collaborative tools are ones that users open, play with, then close. They're a discrete, defined presence in users' lives. The tools of the future will be more like artists' sketchbooks: always-at-hand tools that people will use to think with, and to record ideas. But because they're also shareable, they can change group processes as well. Imagine, for example, an infinitely-large digital whiteboard or notebook that everyone in a group can access, write on, annotate, or import ideas into. In such an environment, the concept of the "brainstorming session"—a discrete time set aside to be creative—will seem quaint and old-fashioned.

- **The evolution of co-cognition.** To date, computer companies have sought to create tools that are either easy to use, or promise dramatic increases in productivity. For the next generation of creative workers, these two design aims will merge in pursuit of a new goal. Ease of use and productivity will be the means, enhancing creativity and cognitive power will be the end. Imagine a generation of creatives carrying or wearing more computing power than sits on your desk today. Imagine them living with a constant background sense of being connected to family and friends; working and playing in smart mobs; pooling experiences and knowledge with trusted humans and virtual agents; and experiencing the Internet as a deep, abiding presence, sometimes

on the edge of their awareness, sometimes in the center, but always there. After a time, their abilities to organize and act collectively will recede into the background of consciousness. At this point, smart mobs, creative partners and collaborators, and the tools that enable and sustain them, become tools that quietly extend the abilities of humans, shaping our thoughts but rarely thought about.

- **The growth of the creative factory.** The growing ability of creatives to move between the world of ideas and the world of things, between the digital concept and its physical realization, will be important for two reasons. First, the ability to rapidly prototype and improve goods will shorten development cycles and time-to-market. Second, more powerfully, sites of physical production are often tremendously creative spaces. They tend to be overshadowed by the hip worlds of design studios and branding shops, or the arcane universes of R&D centers and incubators, but critical innovations—everything from the transistor to semiconductors to the computer mouse—evolved at the intersection of manufacturing and R&D or design. Extending Web 2.0–enabled creative ecologies down onto the factory floor will open up access to new sources of creativity, help ground interesting ideas in the realities of production, and further speed development of new products.

IMPLICATIONS

IMPLICATIONS FOR ORGANIZATIONS

How can organizations most effectively integrate lightweight knowledge tools into their operations to give them a competitive advantage? And how can they think more strategically and proactively about these tools, in order to be more effective as they integrate future generations of tools—and new people who are accustomed to using them?

To take advantage of the opportunities presented by lightweight knowledge tools, it will be necessary for organizations to:

- **Learn which knowledge processes can be automated, which are best left to humans, and how human analysis can be incorporated as feedback into automated processes.** This requires thinking not in terms of IT systems, but in terms of knowledge ecologies—understanding how people, groups, old media like paper and whiteboards and new digital media all work together to complete tasks and create knowledge.

- **Map knowledge processes across the organization.** Such an exercise will reveal some knowledge tasks that can be optimized through automation, others that can be improved with tools, and yet others that can be improved through redesign of existing knowledge ecologies. But organizations have to be careful not to rely too much on automation: process automation can unintentionally make work less creative by enforcing formal routines on complex activities or promoting a misplaced faith in the accuracy or comprehensiveness of the automated systems.

- **Develop strategies for harnessing the ideas and knowledge of individuals in the organization if such a system is not already in place.** Ideally, an internal knowledge system will be simple for all types of users, it will enable social sharing, and it will be able to be integrated into current work practices. The objective should be to improve existing processes rather than implement new systems that will require users and groups to learn new—and potentially less effective—ways of working.

- **Develop strategies for engaging with ideas, attitudes, and data-aggregation efforts from outside the company and explore data-aggregation efforts that might be interested in your company's information.** Interacting with people outside the company is no longer only a matter of marketing and rhetoric: lead users on the Internet expect to interact with a company by means of such activities as spinning their own stories, mapping locations, comparing the waste trail of a product or otherwise examining the supply chain themselves. And because regular people aren't going to centralize their comments for you, companies will have to do the work to aggregate and analyze consumer responses, or hire outside services to do this for them. For example, organizations could crawl Internet traffic or company-hosted forums to detect patterns in customer feedback. In a world where consumer groups develop powerful social networks to scrutinize companies who refuse to practice transparency, opening up your data shows you have nothing to hide.

- **Recruit new hires who are experienced with the new knowledge ecology.** Such people might include information and library scientists, computer scientists with data mining and machine learning backgrounds, and social media designers.

- **Experiment with lightweight knowledge tools available today.** The table below presents several methodologies, tools, and services with which companies should become familiar.

Amazon Mechanical Turk	A marketplace for people to post or complete small tasks for small amounts of money; a crowdsource-enabling service.
DBPedia	DBPedia is an attempt by a community to make a publicly available database of all the structure-able information in Wikipedia; similar to Freebase (see below).
del.icio.us	A "folksonomy" that stores bookmarks for personal gain while contributing to a common good: an aggregate, evolving taxonomy for the Web.
Freebase	A product of the start-up Metaweb, Freebase aims to be a Wikipedia for free, structured data—a centralized, inter-relational database with information about anything in the world.
Microformats	A Web-based, decentralized data formatting approach to allow regular people to embed information into their Web pages so that machines can read it.
Mozilla Weave	A Mozilla service to provide a continuous experience with the Internet across any computer. Currently stores bookmarks, history, saved form information, passwords. Could be mined for attention data. Could become the user's arbiter for all kinds of profile data.
OpenSocial	A set of common interfaces to social network data, developed by Google in 2007 and supported by sites such as Hi5.com, Viadeo, MySpace, Friendster, orkut, Yahoo!, LinkedIn, Salesforce, and more. Allows users to specify when and how their social network data can be shared between sites and companies.
OpenCalais	A free toolkit offered by Thomson Reuters to scrape semantic information from news stories, including topics, companies, countries, people, technologies, products, and more.
PageRank	The algorithm that helps Google rank search results. PageRank assesses the authority of a hyperlinked document by aggregating the authority of pages linking to it.
Powerset	A search company seeking to match the semantic meaning of search queries with the semantic meaning of sentences in Web documents, using natural language processing. Currently, Powerset only indexes Wikipedia.
Resource Description Framework (RDF)	A W3C specification for modeling information as "subject : predicate : object"—for example, "The United States : has as a capital : Washington, D.C." RDF is used in many bottom-up data-creation efforts, including Creative Commons, FOAF (friend of a friend), and MusicBrainz, an open-source, community-contributed database for music metadata.
Snap Shots	A tool to overlay pop-up information over traditional links on a page. Rolling your mouse over an actor's name, for instance, could bring up a bio and picture from the Internet Movie Database. Rolling over a place name brings up a Google map.
TechMeme	A company that scrapes news sites and blogs. Clusters articles into story groups, which are ranked by importance and timeliness.
Twine	A product by start-up Radar Networks that helps users organize, share and discover information on the Web. Twine attempts to extract semantic information (for instance, people, places, and organizations) from documents that are saved by users. Users can share their stored information with colleagues or the public.
X2 and Signals	An integral part of IFTF's research methodology, signals are events, observations, or developments suggestive of a larger trend or shift. Researchers, experts, and public contributors can view, rate, combine, or hypothesize on the impact of signals on online platforms and in workshops. In the case of X2, members of the science community are invited to contribute signals related to the future of science.
Yahoo! SearchMonkey	A tool that simplifies creation of mashups of Yahoo search results with structured information or imagery from any source.

WILD CARD: AUTOMATING CREATIVITY

So far, we've described a future in which electronic and automated systems support creative work, but don't do anything themselves that could be called creative. But what if machines could be designed to be creative? So here's a wild card: automated systems develop the capacity to be creative. (A wild card is an event that has a small likelihood of occurring, but if it does occur, will have a large impact.) To make this particularly interesting, let's focus not on something high-tech or computationally intensive, but something that people see as an intensely personal, and thus human, form of expression: music.

It makes good sense for music to be at the cutting edge of synthetic creativity. Music is a global commodity, it is cheap to distribute and consume, and it is easy to customize for local markets. A large amount of music is not purchased by individuals, nor is it meant to be listened to carefully: many recordings are bought by retailers or game companies, for use as aural background. Despite romantic notions of music as an expression of genius or deep emotions, most music is to the works of Beethoven as industrial painting is to Rembrandt. What's more, music has been evolving into a digital medium for a long time. In the 1960s, composers and musicians first experimented with electronic instruments; twenty years later, music recording and packaging was digitized; twenty years after that, legal and illegal music distribution went digital. It might be only a matter of time before music creation goes digital as well.

IMAGINE:

The synthetic composition process (or "evolved music," as industry advocates call it) begins when software agents create combinations of melodies essentially at random. Melodies that fail to pass basic musicological tests, or are similar enough to copyrighted music to risk legal action, are deleted. The remainder are compared against various databases of songs (recent hits, classics, etc.), and their potential for being a popular hit or part of a long tail market is evaluated. (A long tail market refers to selling a low volume of many different products as contrasted to a blockbuster approach, which aims to sell a huge volume of a few products.)

The melodies that survive are then combined with drum lines and backing instruments. This stage generates a very large number of demos, which are tested against another set of artificial agents, modeling the reactions of music fans. A unique combination of information, generated by recent Amazon purchases, BitTorrent downloads, iTunes sales, and psychological profiles, determines each artificial fan's reactions. Those artificial fans can review and rate thousands of songs per minute.

Demos that pass the gauntlet of artificial fans are then tested on humans. Some are made available as ringtones or posted to music Web sites; in these venues, their popularity can be evaluated explicitly. Other songs are placed in the backgrounds of electronic games, where popularity agents measure the impact they have on player performance. Songs that show promise of influencing psychological or emotional states are tested in shopping malls, subway lines, even prisons and hospitals.

Finally, the demos rated most highly with human audiences are licensed to professional composers and lyricists, who refine the songs and rerecord them with human musicians and singers. Within the music industry, a few composers are known for their ability to work with synthetic compositions; others are rumored to occasionally use evolved music as a basis for their own compositions, but they deny it.

Synthetic music suggests how the kinds of lightweight knowledge tools favored by today's generation might ultimately feed a new kind of artificial intelligence. Software agents, musicological filters, and synthetic fans would all be relatively modest pieces of software. They would succeed not because they're smart, but because they're easy to run in massive numbers. The system would not seek to produce music rivaling Mozart. Over 99% of its melodies and demos would be discarded, but the process is cheap enough to support this kind of profligacy.

It's not a process that mimics the psychology of creative humans. It's computationally intensive, brute-force creativity—a kind of automated brainstorming. It takes simplicity, sociability, and symbiosis—the principles that characterize lightweight knowledge tools—but harnesses them for new purposes. And music would be a relatively complicated medium for such tools: it would probably be considerably easier to create systems to automate creativity in branding, product naming, or logo design. Automated agents wouldn't compete with design giants like Paul Kahn or Milton Glaser; but they would have the advantages of being relatively cheap and prolific.

END NOTES

1. John Markoff, *What the Dormouse Said: How the 60s Counterculture Shaped the Personal Computer,* (New York, New York: Penguin Group, 2005).

2. Kevin Kelly, *Out of Control: The New Biology of Machines, Social Systems, & the Economic World,* (New York, New York, Perseus Group, 1994).

3. "The Post Scientific Society," Issues in Science and Technology, Fall 2007, pp. 78-84 (based on "The American Innovation System in the Post-Scientific Society," paper prepared for the project on Advancing Innovation, Enhancing the Economy, American Institute of Contemporary German Studies, Johns Hopkins University, June 2007, 16 pages).

4. Mayfield, Ross. Ross Mayfield Weblog: Market's Technology & Musings [Internet]. Palo, Alto: RossMayfield. c2003 - [cited October 13, 2008]. Available at http://ross.typepad.com/blog/2005/08/web_of_verbs.html

5. Bruce Sterling. *Shaping Things (Mediaworks Pamphlets)*. (Boston, Massachusetts: The MIT Press. C)