

Improving Our Ability to Improve: A Call for Investment in a New Future

Douglas C. Engelbart
The Bootstrap Alliance
April 2002 (AUGMENT,133320,)

Keynote Address, World Library Summit, April 23 - 26, 2002, Singapore

Summary. In the past fifty years we have seen enormous growth in computing capability – computing is everywhere and has impacted nearly everything. In this talk, Dr. Douglas Engelbart, who pioneered much of what we now take for granted as interactive computing, examines the forces that have shaped this growth. He argues that our criteria for investment in innovation are, in fact, short-sighted and focused on the wrong things. He proposes, instead, investment in an improvement infrastructure that can result in sustained, radical innovation capable of changing computing and expanding the kinds of problems that we can address through computing. In this talk, Dr. Engelbart describes both the processes that we need to put in place and the capabilities that we must support in order to stimulate this higher rate of innovation. The talk closes with a call to action for this World Library Summit audience, since this is a group that has both a stake in innovation and the ability to shape its direction.

Good news and bad news

1

The development of new computing technologies over the past fifty years – in hardware and software – has provided stunningly important changes in the way we work and in the way we solve problems.

I need to get this assertion out in front of you early in this talk, because most of the rest of what I have to say might cause you to think that I have lost track of this progress or that I don't appreciate it. So, let me get it said – we have made enormous strides since the early 1950s, when I first began thinking seriously about ways to use computers to address important social problems. It has truly been a remarkable fifty years.

At my first job at NACA, the forerunner of NASA, right out of engineering school, there was no vision at all of electronic computers. In fact the term "computers" referred to the room full of women sitting at desks using desk calculators to process the wind tunnel data. This was in the late 40's. After I'd earned my doctorate at UC Berkeley and worked there as an acting assistant professor, I applied to Stanford to help develop and teach computer design courses. They told me "Since computing is only a service activity, we don't contemplate ever having courses in

computer design be part of our academic program." When I was pursuing Federal funding for IT projects at SRI in the early 60's, there was some real question about whether there would be the programming talent to work on computing applications of this complexity in Palo Alto California. This certainly inhibited our chances for getting research support.

Later in my research, when I thought about using computers to manipulate symbols and language, rather than to do calculations on numbers, most people thought I was really pretty far gone. I was even advised to stay away from jobs in academia. It seemed clear to everyone else at the time that nobody would ever take seriously the idea of using computers in direct, immediate interaction with people. The idea of interactive computing – well, it seemed simply ludicrous to most sensible people.

So, we have made truly tremendous progress. We are able to solve problems, ranging from weather forecasting to disaster relief to creating new materials to even cracking the very genetic code that makes us humans – problems that we could not have even contemplated without the availability of cheap, widely available, highly functional computers and software. It has been a marvelous 50 years to be in this business.

But that is not what I am going to talk to you about. Not out of lack of appreciation – even a sense of wonder – over what computer technologists have developed – but because I can see that we are not yet really making good progress toward realizing the really substantial payoff that is possible. That payoff will come when we make better use of computers to bring communities of people together and to augment the very human skills that people bring to bear on difficult problems.

In this talk I want to talk to you about that big payoff, think a bit with you about what is getting in the way of our making better progress, and enlist you in an effort to redirect our focus. This audience is made up of the kinds of people who really can change the focus so that we can set at least of part of our development efforts on the right course.

The rewards of focusing on the right course are great. I hope to show you that they can be yours.

The vision: The payoff

2

Before talking and thinking with you about why we keep heading off in the wrong direction, I need to quickly sketch out what I see as the goal – the way to get the significant payoff from using computers to augment what people can do. This vision of success has not changed much for me over fifty years – it has gotten more precise and detailed – but it is pointed at the same potential that I saw in the early 1950s (*ref. 1*). It is based on a very simple idea, which is that when problems are really difficult and complex – problems like addressing hunger, containing terrorism, or helping an economy grow more quickly – the solutions come from the insights and capabilities of people working together. So, it is not the computer, working alone, that produces a solution. It is the combination of people, augmented by computers.

The key word here is "augment." The reason I was interested in interactive computing, even before we knew what that might mean, arose from this conviction that we would be able to solve really difficult problems only through using computers to extend the capability of people to collect information, create knowledge, manipulate and share it, and then to put that knowledge to work. Just as the tractor extends the human's ability to work the earth, and planes extend our ability to move, so does the computer extend our ability to process and use knowledge. And that knowledge production is a group activity, not an individual one. Computers most radically and usefully extend our capabilities when they extend our ability to collaborate to solve problems beyond the compass of any single human mind.

I have found, over the years, that this idea of "augmenting" capability needs clarification for many people. It is useful to contrast "augmentation" with "automation." Automation is what most people have in mind when they think about using computers. Automation is what is going on when we use computers to figure up and print telephone bills or to keep track of bank charges. It is also what is going on when we think about "artificial intelligence." Even though printing up phone bills and AI seem very different, they both share this assumption that the computer stands over there, apart from the human, doing its thing. That is not, in my mind, how we use computers to solve tough problems. We have the opportunity to harness their unique capabilities to provide us new and more effective ways to use our mind and senses – so that the computer truly becomes a way of extending our capabilities.

The shovel is a tool, and so is a bulldozer. Neither works on its own, "automating" the task of digging. But both tools augment our ability to dig. And the one that provides the greatest augmentation, not surprisingly, takes the most training and experience in order to use it really effectively.

For fifty years I have been working to build the computing equivalents of bulldozers. With the added constraint that, in serious knowledge works, it is almost always through the collaborative work of a team, rather than a lone operator, that we make progress.

Now, there is a lot more to say about this vision than that – and I will speak to some of it later in this talk. But, in starting our consideration of how to change course in order to get a bigger payoff from our investment in computing, this focus on augmenting the ability of groups to solve problems is the right starting point.

Evidence of trouble

3

Because we are so accustomed to thinking in terms of the enormous progress and change surrounding computing, it is important to take at least a few moments to look at the evidence that, when it comes to these broader, social and group-centered dimensions of computing, the picture looks quite different.

Difficulty in doing important collaborative work. As one example, my organization, the Bootstrap Alliance, works in loose collaboration with a number of other organizations to help

them develop better ways to improve their ability to learn and to use knowledge – in short, we work with organizations to help them improve their ability to improve.

One organization that we work with is the Global Disaster Information Network – or "GDIN" – which is, itself, a consortium of regional and local disaster response organizations. Organizations that respond to disasters are tremendous examples of organizations that must learn to adapt and use new information quickly. Disasters are, by their nature, unplanned and surprising. Responding requires rapid access to weather information, geographical and mapping information, information about local resources, local communications, the availability of outside resources and organizations – sometimes even about the location of buried mines and unexploded munitions. And, because disaster response involves so many people, from many organizations and jurisdictions, it is critically important to share all of this information about resources and capabilities, as well as information about response status, planned next steps, and so on.

Computers and, in particular, the Internet, clearly play a key role in the efforts to coordinate such disaster response and to improve the ability to improve over the lifecycle of a disaster response effort. But what is striking, as GDIN grapples with these issues, is how difficult it is to harness all the wonderful capability of the systems that we have today in GDIN's effort to improve its ability to improve disaster response. It turns out that it is simply very difficult to share information across systems – where "sharing" means both the ability to find the right information, when it is needed, as well as the ability to use it across systems.

Even harder is the ability to use the computer networks to monitor and reflect status. Anyone that regularly uses e-mail can readily imagine how the chaotic flow of messages between the different people and organizations during a disaster falls far short of creating the information framework that is required for an effectively coordinated response. Make no mistake about it, GDIN and its member disaster response organizations find computers to be very useful – but it is even more striking how the capabilities offered by today's personal productivity and publishing systems are mismatched to the needs of these organizations as they work to coordinate effective response flexibly and quickly.

Difficulties with knowledge governance. As another example of our still relatively primitive ability to deal with information exchange among groups, consider the chaotic and increasingly frightening direction of new laws regarding knowledge governance – most notably reflected in laws regarding copyright. Because it is generally technically advanced, one might think that my country, the United States, would be representative of leading edge capability to deal with knowledge governance and knowledge sharing. But, instead, we are passing increasingly draconian laws to protect the economic value of copies of information. In the US, we are even contemplating laws that would require hardware manufacturers to take steps to encrypt and protect copies (*ref. 2*)

We are doing this while entering a digital era in which the marginal cost of a copy is zero – at a time where the very meaning and significance of the notion of "copy" has changed. It is as if we are trying to erect dikes, using laws, to keep the future from flooding in.

The immediate effect of all this is to enable a dramatic shift in control to the owners of information, away from the users of information (*ref. 3*) – a strategy which will almost certainly fail in the long run and that has confusing and probably damaging economic consequences in the short run.

The most modest conclusion that one might draw from watching the U.S. attempt to deal with knowledge governance in a digital age is that the legislators have a weak understanding of the issues and are responding to the enormous political power of the companies with vested interest in old ways of using information. Looking somewhat more deeply, it seems quite clear that we are ill-prepared to come to terms with an environment in which the social value of knowledge emerges from collaborative use of it. The entire idea of value emerging from sharing, collaboration, and use of knowledge – as opposed to treating knowledge as a scarce resource that should be owned and protected – is anathema to the 20th century knowledge owners, who are fighting hard to protect their turf.

Structural roots of the problem

4

One possible response to my examples is to say, "Doug, be patient. These are new problems and hard problems and it takes time to solve them. We will have better tools and better laws over time. Just wait."

An off-handed response might be that I have been trying to be patient for fifty years. But a much more important, meaningful response is that patience has nothing to do with it. These problems are not due to lack of time, but are instead due to structural factors that create a systematic bias against the improvement of what I call "Collective IQ."

The good news is that, if we can see and understand the bias, we have the opportunity to change it. If we can see how some of the basic assumptions that we bring to the development of computing technologies lead us away from improvement in our ability to solve problems collectively, we can reexamine those assumptions and chart a different course.

Oxymoron: "Market Intelligence" – One of the strongly held beliefs within the United States is that the best way to choose between competing technologies and options for investment is to "let the market decide." In my country we share a mystical, almost religious kind of faith in the efficacy of this approach, growing from Adam Smith's idea of an "invisible hand" controlling markets and turning selfish interest into general good. The "market" assumes the dimensions of faceless, impersonal deity, punishing economically inefficient solutions and rewarding the economically fit. We believe in the wisdom of the market and belief that it represents a collective intelligence that surpasses the understanding of us poor mortal players in the market's great plan.

One of the nice things about getting outside the U.S. – giving a talk here, in Singapore, for example – is that it is a little easier to see what an odd belief this is. It is one of the strange quirks of the U.S. culture.

In any case, it is quite clear that whatever it is that the market "knows," its knowledge is fundamentally conservative in that it only values what is available today. Markets are, in particular, notoriously poor judges of value for things that are not currently being bought and sold. In other words, markets do a bad job at assessing the value of innovation when that innovation is so new that it will actually rearrange the structure of the markets.

This is well understood by people doing market research. Decades ago, when Hewlett Packard was first coming up with the idea of a desktop laser printer – before anyone had experience with such devices and before there was even software available for desktop publishing – market studies of the potential use and penetration for desktop laser printing came up with a very strange answer: people simply did not yet have enough experience with the devices to be able to understand their value. The same thing happened to companies, ten to fifteen years ago, when they did market studies about the potential value and use of digital cameras.

Perhaps the best study of this systematic and very basic conflict between markets and certain kind of innovations is Clayton Christensen's classic and very valuable book, *The Innovator's Dilemma* (ref. 4). Probably most of you are familiar with Christensen's thesis (if you haven't read the book, you should), but, briefly stated, it is that one kind of innovation – Christensen calls it "continuous innovation" – emerges when companies do a good job of staying close to their customers and, in general, "listening to the market." This is the kind of innovation that produces better versions of the kinds of products that are already in the market. If we were all riding tricycles, continuous innovation would lead to more efficient, more comfortable, and perhaps more affordable tricycles.

But it would never, ever produce a bicycle. To do that, you need a different kind of innovation – one that usually, at the outset, results in products that do not make sense to the existing market and that it therefore cannot value. Christensen calls this "discontinuous innovation."

Discontinuous innovation is much riskier, in that it is much less predictable, than continuous innovation. It disrupts markets. It threatens the positions of market leaders because, as leaders, they need to "listen" to the existing market and existing customers and keep building improved versions of the old technology, rather than take advantage of the new innovation. It is this power to create great change that makes discontinuous innovation so valuable over the long run. It is how we step outside the existing paradigm to create something that is really new.

In the past fifty years of history of computing, the one really striking example of discontinuous innovation – the kind where the market's "intelligence" approached an IQ of zero – was early generation of World Wide Web software – and in particular, the Mosaic web browser. There were, as the Web first emerged, numerous companies selling highly functional electronic page viewers – viewers that could jump around in electronic books, follow different kinds of hyperlinks, display vector graphics, and do many other things that early web browsers could not do. The companies in this early electronic publishing business were actually able to sell these "electronic readers" for as much as US \$50 a "seat" – meaning that, when selling electronic document viewers to big companies with many users, this was big business.

Then, along came the Web and Mosaic – a free Web browser that was much less functional than these proprietary offers. But it was *free*! And, more important, it could do something else that these other viewers could not do – it provided access to information anywhere in the world on the Web. As a result, over the next few years, everything changed. We actually did get closer to the goal of computers assisting with collaborative work.

But the key point of the story is that, at first, the "market intelligence" saw no value in web browsers at all. In fact, the market leader, Microsoft, initially started off in the direction of building its own proprietary viewer and network – because that is what market intelligence suggested would work. Fortunately for Microsoft's shareholders, Bill Gates realized that he was facing a discontinuity, and threw the company into a sudden and aggressive campaign to change course.

Despite the Web, despite the example of Mosaic, despite all the work that Christensen has done to teach us about discontinuous innovation, most companies still act as if they believe that the market is intelligent – and, to be sure, this approach really does often work, in the short term. So we are saddled with a systematic, built-in bias against thinking outside the box. And that bias gets in the way of solving hard problems, such as building high performance tools that help groups of people collaborate more effectively.

In a little bit, I will explain how we can overcome such systematic bias and open the doors to the very substantial rewards from continued, productive discontinuous innovation. There is huge opportunity here – and it is an opportunity that will be most available to emerging economies rather than to the incumbents. But, before turning to solutions, I need to tell you about another dimension of systematic bias that is getting in the way of our making important progress in finding new ways to use computers.

The seductive, destructive appeal of "ease of use." A second powerful, systematic bias that leads computing technology development away from grappling with serious issues of collaboration – the kind of thing, for example, that would really make a difference to disaster response organizations – is the belief that "ease of use" is somehow equated with better products.

Going back to my tricycle/bicycle analogy, it is clear that for an unskilled user, the tricycle is much easier to use. But, as we know, the payoff from investing in learning to ride on two wheels is enormous.

We seem to lose sight of this very basic distinction between "ease of use" and "performance" when we evaluate computing systems. For example, just a few weeks ago, in early March, I was invited to participate in a set of discussions, held at IBM's Almaden Labs, that looked at new research and technology associated with knowledge management and retrieval. One thing that was clearly evident in these presentations was that the first source of bias – the tendency to look solely to the invisible hand and intelligence of the market for guidance, was in full gear. Most of the presenters were looking to build a better tricycle, following the market to the next stage of continuous innovation, rather than stepping outside the box to consider something really new.

But there was another bias, even in the more innovative work – and that bias had to do with deciding to set aside technology and user interactions that were "too difficult" for users to learn. I was particularly disappointed to learn, for example, that one of the principal websites offering knowledge retrieval on the web had concluded that a number of potentially more powerful searching tools should not be offered because user testing discovered that they were not easy to use.

Here in Singapore, I see a lot of people wind surfing. I am sure that there are beginner boards and sails, just as in kayaking there are beamy, forgiving boats that are good for beginners, and in tennis there are powerful racquets that make it easy for beginners to wallop the ball even with a short swing. But someone who wants real performance in wind surfing, to have control in difficult conditions, does not want a beginning board. Someone who wants a responsive kayak, that will perform well in following seas and surf, does not want a beginner's boat. A serious tennis player with a powerful swing does not want a beginner's racquet.

Why do we assume that, in computing, ease of use – particularly ease of use by people with little training – is desirable for anyone other than a beginner? Sure, I understand that the big money for a company making surfboards, tennis racquets, skis, golf clubs, and what-have-you is always in the low end of the market, serving the weekend amateur. And surely the same thing is true in computing. That is not surprising. What is surprising is that, in serious discussions with serious computer/human factors experts, who are presumably trying to address hard problems of knowledge use and collaboration, ease of use keeps emerging as a key design consideration.

Doesn't anyone ever aspire to serious amateur or pro status in knowledge work?

Restoring balance

5

I need to remind you of what I said at the beginning of this talk: we have made huge strides forward in computing. It is a wonderful thing to have a large, mass market for equipment that has brought the cost of computing hardware and software down to the point where truly staggering computing capability is available for a few thousand – even a few hundred – dollars. It has been a marvelous fifty years. But I want to alert you to two very important facts:

1. We are still not able to address critically important problems – particularly if those problems demand high performance ability to collect and share knowledge across groups of people.
2. This inability is not an accident, but emerges from values and approaches that are "designed into" our approach to addressing innovation in computing.

These facts are critical for institutions and individuals who are interested in improving our ability to improve. I am pretty sure that this includes everyone in this audience. The important realization – and the message of this talk – is that these institutions and individuals can take big steps forward simply by systematically addressing the biases that are pushing innovation toward lowest common denominator solutions and toward simple continuation down roads that we

already understand. This does not mean that we should stop building easy to use applications that represent continuous innovation. What it does mean is that we *also* need to find ways to address the harder problems and to stimulate more discontinuous innovation.

This focus on new, discontinuous innovation is particularly important for the majority of people and nations in the world who are building emerging economies. It is the developing nations that have the most to gain from developing new ways to share knowledge and to stimulate improvement.

Moving from "invisible hand" to strategy

6

The good news is that it is possible to build an infrastructure that supports discontinuous innovation. There is no need at all to depend on mystical, invisible hands and the oracular pronouncements hidden within the marketplace. The alternative is conscious investment in an improvement infrastructure to support new, discontinuous innovation (*ref. 5*).

This is something that individual organizations can do – it is also something that local governments, nations, and regional alliances of nations can do. All that is necessary is an understanding of how to structure that conscious investment.

ABCs of improvement infrastructure. The key to developing an effective improvement infrastructure is the realization that, within any organization, there is a division of attention between the part of the organization that is concerned with the organization's primary activity – I will call this the "A" activity – and the part of the organization concerned with improving the capability to perform this A-level function. I refer to these improvement efforts as "B" activities. The two different levels of activity are illustrated in Figure 1.

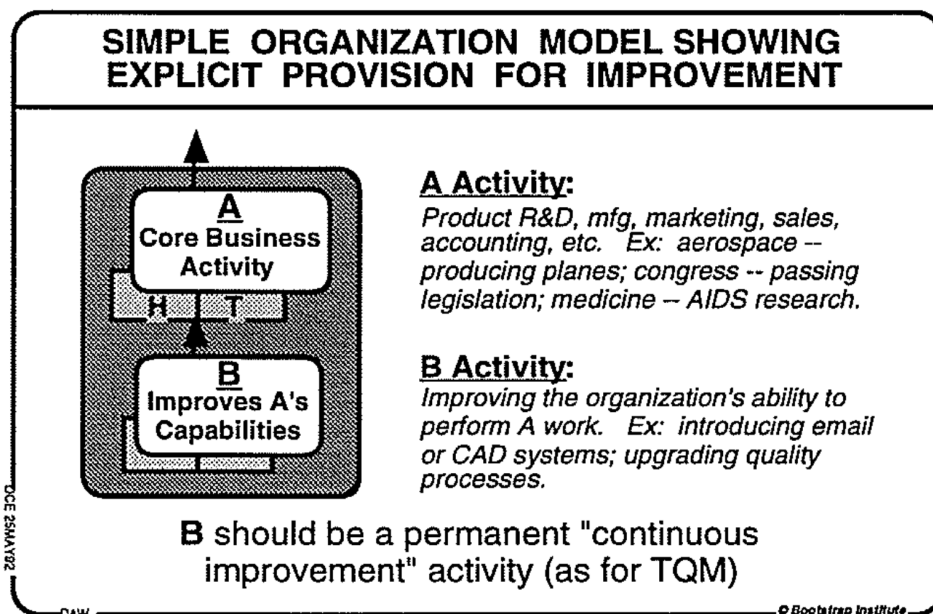


Figure 1. Infrastructure fundamentals: A and B Activities (Engelbart, 1962, 1992)

The investment made in B activities is recaptured, along with an aggressive internal rate of return, through improved productivity in the A activity. If investments in R&D, IT infrastructure, and other dimensions of the B activity are effective, the rate of return for a dollar invested in the B activity will be higher than for a dollar invested in the A activity.

Clearly, there are limits to how far a company can pursue an investment and growth strategy based on type B activities – at some point the marginal returns for new investment begin to fall off. This leads to a question: How can we maximize the return from investment in B activities, maximizing the improvement that they enable?

Put another way, we are asking how we improve our ability to improve. This question suggests that we really need to think in terms of yet another level of activity – I call it the "C" activity – that focuses specifically on the matter of accelerating the rate of improvement. Figure 2 shows what I mean.

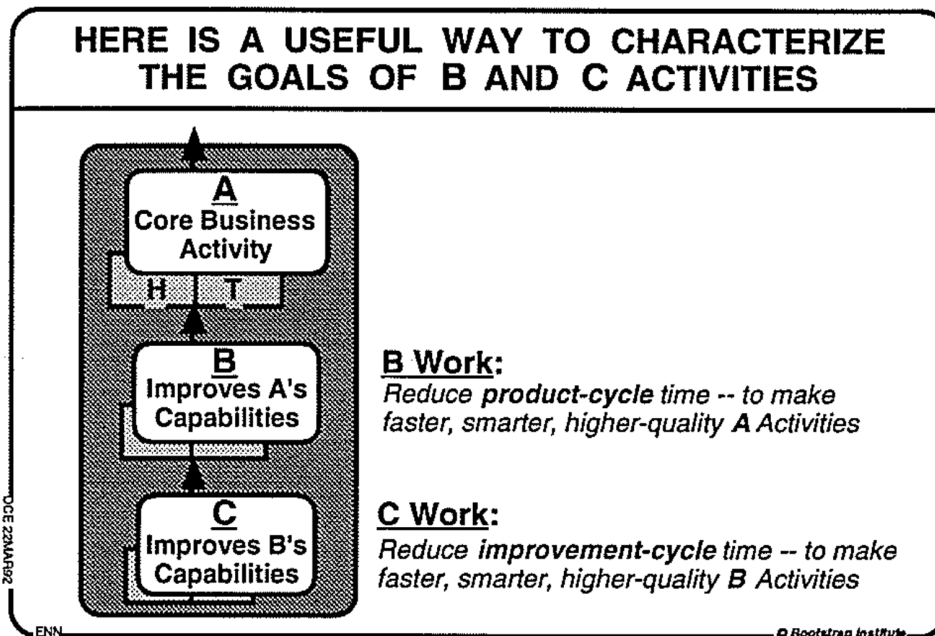


Figure 2. Introducing "C" level activity to improve the ability to improve (Engelbart, 1962)

Clearly, investment in type C activities is potentially highly leveraged. The right investments here will be multiplied in returns in increased B level productivity – in the ability to improve – which will be multiplied again in returns in productivity in the organization's primary activity. It is a way of getting a kind of compound return on investment in innovation.

The highly leveraged nature of investment in type C activities make this kind of investment in innovation particularly appropriate for governments, public service institutions such as libraries, and broad consortia of different companies and agencies across an entire industry. The reason for this is not only that a small investment here can make a big difference – though that certainly is an important consideration – but also because the investment in C activities is typically pre-

competitive. It is investment that can be shared even among competitors in an industry because it is, essentially, investment in creating a better playing field. Perhaps the classic recent example of such investment in the U.S. is the relatively small investment that the Department of Defense made in what eventually became the Internet.

Another example, looking to the private sector, is the investment that companies made in improving product and process quality as they joined in the quality movement. What was particularly important about this investment was that, when it came to ISO 9000 compliance and other quality programs and measures, companies – even competing companies – joined together in industry consortia to develop benchmarks and standards. They even shared knowledge about quality programs. What emerged from this collaborative activity at the C level was significant gain for individual companies at the B and A levels. When you are operating at the C level, collaboration can produce much larger returns than competition.

Investing wisely in improvement

7

Let's keep our bigger goal in mind: we want to correct the current bias, emerging from over-reliance on market forces and the related obsession with ease of use that get in the way of developing better computing tools. We want to do this so that we can use computers to augment the capabilities of entire groups of people as they share knowledge and work together on truly difficult problems. The proposal that I am placing on the table is to correct that bias by making relatively small, but highly leveraged investments in efforts to improve our ability to improve – in what I have called type C activities.

The proposal is attractive not only for quantitative reasons – because it can produce a lot of change with a relatively small investment – but also for qualitative reasons: This kind of investment is best able to support disruptive innovation – the kind of innovation that is necessary to embrace a new, knowledge-centered society. The acceleration in movement away from economic systems based on manufacturing and toward systems based on knowledge needs to be reflected in accelerated change in our ways of working with each other. This is the kind of change that we can embrace by focusing on type C activity and on improvement of our ability to improve.

Given all of that, what do we need to do? If, say, Singapore wants to use this kind of investment as a focus for its development activity, where does it concentrate its attention? If the organizations participating in this World Library Summit want to support and stimulate this kind of investment, where do they begin?

The answer to such questions has two different, but complementary dimensions. The first dimension has to do with process: How do you operate and set expectations in a way that is consistent with productive type C activity? The second dimension has to do with actual tools and techniques.

Process considerations

8

Making an investment in type C activity is not the same as investing in research into new materials or in an ERP system to provide better control over inventory and accounting. Those kinds of investments have very specific objectives and tend to proceed in a straight line from specification to final delivery. Sure, we know that there are usually surprises and unplanned side trips, but that is not the initial expectation. B level investments are supposed to be predictable. Nobody, for example, would think of installing two ERP systems – say, SAP and Peoplesoft – to discover which is better. In B-level investment, you make the design decisions up front and then implement the design.

That is not the way it works with C-level investments. Here, you typically do, in fact, pursue multiple paths concurrently. At the C level we are trying to understand how improvement really happens, so that we can improve our ability to improve. This means having different groups exploring different paths to the same goal. As they explore, they constantly exchange information about what they are learning. The goal is to maximize overall progress by exchanging important information as the different groups proceed. What this means, in practice, is that the dialog between the people working toward pursuit of the goal is often just as important as the end result of the research. Often, it is what the team learns in the course of the exploration that ultimately opens up breakthrough results.

Another difference between innovation at the C level and innovation that is more focused on specific results is that, at the C level, context is tremendously important. We are not trying to solve a specific problem, but, instead, are reaching for insight into a broad class of activities and opportunities for improvement. That means attending to external information as well as to the specifics of the particular work at hand. In fact, in my own work, I have routinely found that when I seem to reach a dead end in my pursuit of a problem, the key is usually to move up a level of abstraction, to look at the more general case.

Note that this is directly counter to the typical approach to solving focused, B-level problems, where you typically keep narrowing the problem in order to make it more tractable. In our work on improving improvement, the breakthroughs come from the other direction – from taking on an even bigger problem.

So, the teams working at the C-level are working in parallel, sharing information with each other, and also tying what they find to external factors and bigger problems. Put more simply, C-level work requires investment integration – a concerted effort to tie the pieces together.

That is, by the way, the reason that the teams that I was leading at SRI were developing ways to connect information with hyperlinks, and doing this more than two decades before it was happening on the web. Hyperlinks were quite literally a critical part of our ability to keep track of what we were doing.

Thinking back to our research at SRI leads me to another key feature of development work at the C level: You have to apply what you discover. That is the way that you reach out and snatch a bit of the future and bring it back to the present: You grab it and use it.

At the C-level, then, the approach focuses on:

- Concurrent development
- Integration across the different concurrent activities though continuous dialog and through constant cross checking with external information
- Application of the knowledge that is gained, as a way of not only testing it, but also as a way to understand its nature and its ability to support improvement.

As a mnemonic device to help pull together these key features of the C-level process, you can take "Concurrent Development," "Integration," and "Application of Knowledge" and put them together in the term "CoDIAK." For me, this invented word has become my shorthand for the most important characteristics of the C-level discovery activity. Figure 3 illustrates the way that the CoDIAK process builds on continuous, dynamic integration of information so that the members of the improvement team can learn from each other and move forward.

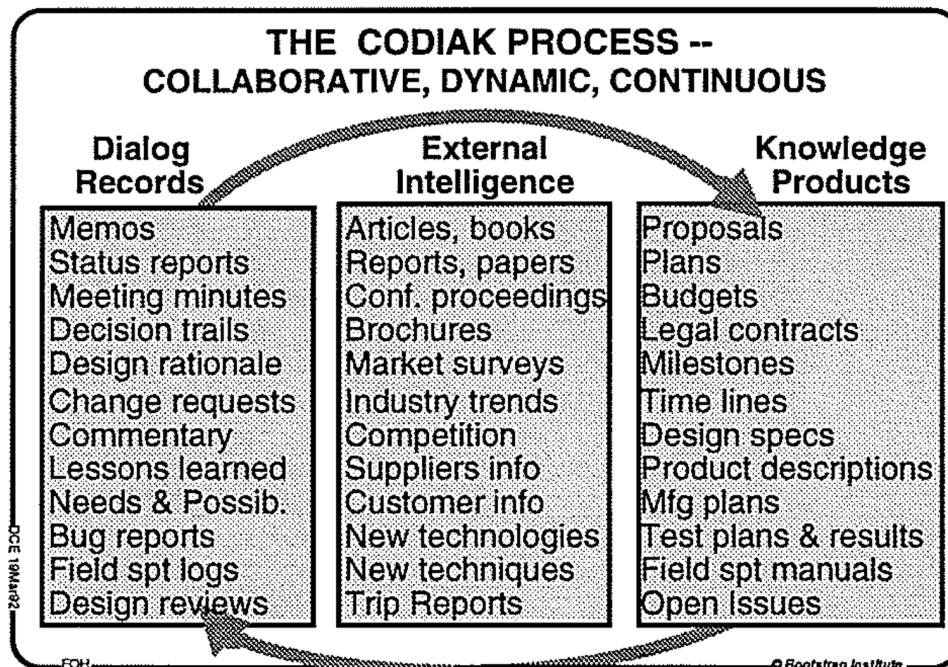


Figure 3. Key elements of the CoDIAK process (Engelbart, 1992)

Investment in tools and techniques

9

Returning once again to the main theme of this talk, my goal is to help the people at this World Library Summit see how they, through their governments and institutions, can make a highly leveraged investment in a different kind of innovation, one that will open up new opportunities and capabilities in computing. Part of what is needed is a new approach to the process of innovation – that is what CoDIAK is all about. But pursuit of CoDIAK requires, in itself, some technical infrastructure to support the concurrent development and continual integration of dialog, external information, and new knowledge. If this sounds somewhat recursive to you, like the snake renewing itself by swallowing its own tail, be assured that the recursion is not an accident. As I just said, one of the key principles in CoDIAK is the application and use of what you learn. That recursive, reflective application gets going right at the outset. So, what do we need to get started?

One of the most important things that we need is a place to keep and share the information that we collect – the dialog, the external information, the things that we learn. I call this the "Dynamic Knowledge Repository," or DKR. It is more than a database, and more than a simple collection of Internet web sites. It doesn't have to be all in one place – it can certainly be distributed across the different people and organizations that are collaborating on improving improvement – but it does need to be accessible to everyone – for reading, for writing, and for making new connections.

The DKR is a wonderful example of the kind of investment that you can start making at the C level, with modest means, that will pay dividends back as it moves up the line to the B and the A levels. This is exactly what I mean when I talk about "bootstrapping." It is a very American term – the image is of someone able to perform the wonderful, impossible trick of pulling himself up by pulling up on his own bootstraps – but the idea is one that we put into practice every time that we "boot up" a computer. A small bit of code in a permanent read only memory knows how to go out to the disk to get more instructions, that in turn know how to do to even more things, such as getting even more instructions. Eventually, this process of using successive steps to lead to ever bigger steps, building on each other, get the whole machine up and running. You start small, and keep leveraging what you know at each stage to solve a bigger and bigger problem.

This is precisely the kind of outcome that can come from investment in building a DKR at the C level. What you learn there can be used to improve work at the C level, which in turn improves ability at the B level, which then translates into new capability at the primary, A level of the organization.

Another key, early investment is in the development of tools to provide access to the knowledge in the DKR for all classes of users, from beginners to professional knowledge workers expecting high performance. This "hyperscope" – that is my term for it – allows everyone to contribute and use the information in the DKR according to his or her ability. It avoids the problem of making everyone, even the pros, play with the same, over-powered tennis racquets that are helpful for beginners.

Tied to the hyperscope is the ability to provide different views of the knowledge in the DKR – and I do mean "views" – stressing the "visual" sense of the term. Moving away from words on a page, we need to be able to analyze an argument – or the results of a meeting – visually. We need to move beyond understanding the computer as some kind of fancy printing machine and begin to use it to analyze and manipulate the symbolic content of our work, extending our own capabilities. We already do this in specialized cases; one of the most spectacular recent examples was the use of high-performance computing in the analysis of the sequences that make up the human genome. Now we need to extend that to the more general class of problems that groups of people encounter as they work together, try to understand each other, and reach collaboratively for decisions.

Another critical focus area for tool and technology development centers on the way that humans interact with computers. We have come far since the early 1950s, when my colleagues felt that such interaction was not worth thinking about. And, as most of you know, it was in the course of trying to broaden the bandwidth of the connection between humans and computers, incorporating both visual and motor skill dimensions, that I developed my most famous invention, the computer "mouse."

There is so much more to be done here – I feel that we have just scratched the surface. Figure 4 provides you with an overview of this very fertile field and opportunity for breakthrough innovation.

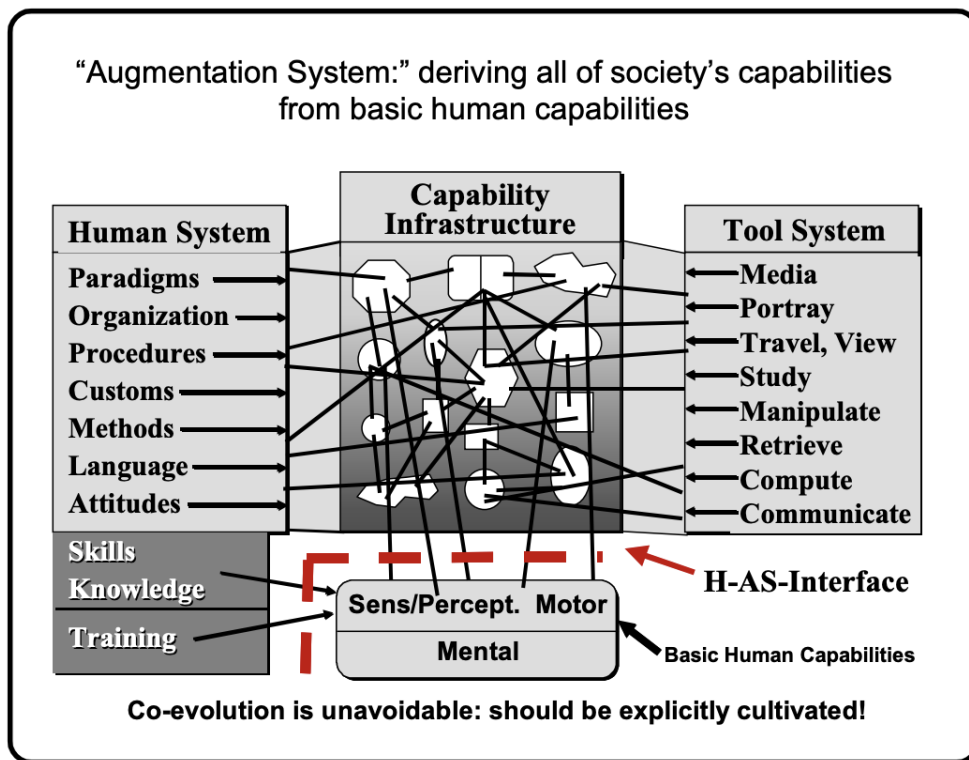


Figure 4. The Human-Augmentation System Interface (Engelbart, 1962, 1992)

The Capability Infrastructure – which is the thing in the middle of this picture and is what we are talking about improving when we are working at the C level of innovation – combines inputs from both the tool system and the human system. The tool system – that's the contribution from the computer, provides access to different media, gives us different ways to portray information, and so on. The human system brings its rich store of paradigms, information captured in customs, and so on. The more static parts of this collection can be added directly into the Capability Infrastructure through construction of ontologies and other artifacts.

The human system, as the part of this framework that is best at learning, also brings the opportunity to develop new skills, benefit from training, and to assimilate and create new knowledge. These dynamic elements are the "magic dust" that makes the whole system capable of innovation and of solving complex problems. These are what make an "augmentation system" different than a mere automation system.

These valuable, dynamic, human inputs must of course come into the system through the human's motor and perceptual capabilities. It is the boundary between these human capabilities and the rest of the infrastructure – represented by the heavy, broken line in this figure and labeled "H-AS Interface" – that, in a very real sense, defines the scope of the capabilities of this augmentation system. If this interface is low-bandwidth and able to pass only a small amount of what the human knows and can do – and what the machine can portray – then the entire system tends to be more "automation" than "augmentation," since the computer and the human are being kept apart by this low-fidelity, limited interface.

If, on the other hand, this interface can operate at high speed and capture great nuance – perhaps even extending to changes in facial expression, heart rate, or fine motor responses, then we greatly increase the potential to integrate the human capabilities directly into the overall system, which means that we can then feed them back, amplify them, and use them.

When you begin to conceive of the human-system interface in this way, the whole notion of "ease of use" – this matter that we are now so obsessed with – appears, as it should, as merely a single and, in the grand scheme of things, not terribly important dimension within a much richer structure. The key to building a more powerful capability infrastructure lies in expanding the channels and modes of communication – not simplifying them.

This is very powerful, exciting stuff. If we begin to act on *this* notion of our relation, as humans, to these amazing machines that we have created, we really begin to open up new opportunities for growth and problem solving.

I could go on and on. And will be happy to for anyone with the patience and interest in understanding what is possible. Just call me up, arrange an audience where, working together, we can make a difference, and I will be there.

The point here, for this talk, for this audience, is that the commitment to the CoDIAK process leads to very specific directions for investments in technology development – the kinds of investments that your companies, agencies, institutions, and governments can make. And the

reason for making them is to open the doors to new kinds of innovation – giving you the power to address much harder, but potentially much richer kinds of problems.

Your involvement matters

10

Well, we are reaching the point in this talk where I should soon be thanking you for your attention and asking whether you have any questions. But, before I go there, I want to tell you again why this matters so much, with the hope of securing your commitment to help in moving us out of the dangerous, disappointing, narrow path that we seem to be stuck with following.

The feature of humans that makes us most human – that most clearly differentiates us from every other life form on Earth – is not our opposable thumb, and not even our use of tools. It is our ability to create and use symbols. The ability to look at the world, turn what we see into abstractions, and to then operate on those abstractions, rather than on the physical world itself, is an utterly astounding, beautiful thing, just taken all by itself. We manifest this ability to work with symbols in wonderful, beautiful ways, through music, through art, through our buildings and through our language – but the fundamental act of symbol making and symbol using is beautiful in itself.

Consider, as a simple, but very powerful example, our invention of the negative – our ability to deal with what something is *not*, just as easily as we deal with what it *is*. There is no "not," no negative, in nature, outside of the human mind. But we invented it, we use it daily, and divide up the world with it. It is an amazing creation, and one that is quintessentially human.

The thing that amazed me – even humbled me – about the digital computer when I first encountered it over fifty years ago – was that, in the computer, I saw that we have a tool that does not just move earth or bend steel, but we have a tool that actually can manipulate symbols and, even more importantly, portray symbols in new ways, so that we can interact with them and learn. We have a tool that radically extends our capabilities in the very area that makes us most human, and most powerful.

There is a native American myth about the coyote, a native dog of the American prairies – how the coyote incurred the wrath of the gods by bringing fire down from heaven for the use of mankind, making man more powerful than the gods ever intended. My sense is that computer science has brought us a gift of even greater power, the ability to amplify and extend our ability to manipulate symbols.

It seems to me that the established sources of power and wealth understand, in some dim way, that the new power that the computer has brought from the heavens is dangerous to the existing structure of ownership and wealth in that, like fire, it has the power to transform and to make things new.

I must say that, despite the cynicism that comes with fifty years of professional life as a computer scientist, inventor, and observer of the ways of power, I am absolutely stunned at the ferocious strength of the efforts of the American music industry, entertainment industry, and other

established interests to resist the new ability that the coyote in the computer has brought from the heavens. I am even more surprised by the ability of these established interests to pass laws that promise punishment to those who would experiment and learn to use the new fire.

As the recipient of my country's National Medal of Technology, I am committed to raising these issues and questions within my own country, but I am also canny enough to understand that, in the short term, it is the nations with emerging economies that are most likely to understand the critical importance and enormous value in learning to use this new kind of fire.

We need to become better at being humans. Learning to use symbols and knowledge in new ways, across groups, across cultures, is a powerful, valuable, and very human goal. And it is also one that is obtainable, if we only begin to open our minds to full, complete use of computers to augment our most human of capabilities.

The Bootstrap Alliance

11

I come to this conference representing my own small organization, the Bootstrap Alliance. We don't sell a product or anything else. But we do offer an opportunity for you to be actively engaged with other people and other institutions that are interested in understanding how to use this new fire that has been brought down from the heavens.

More specifically, the Bootstrap Alliance is an improvement community that is made up of other improvement communities – we are focused on improving the ability to improve, and on helping other groups that share those interests do a better job of it. We exist to help C-level organizations do a better job of being C-level organizations. Our approach to this, not surprisingly, is based on concurrent development, integration, and application of knowledge across those different pioneering communities.

If you are interested in investing in the kind of critically important, highly leveraged mechanisms for change that I talk about here – in using the fire brought down from heaven – please come up and talk to me or e-mail me at doug@bootstrap.org. We have a lot of work to do together, and no time at all to be patient.

Acknowledgements

12

I would like to recognize the assistance I had from Bill Zoellick in preparing this paper, particularly for his contributions concerning recent copyright activity and regarding the interaction of markets and innovation.

References

13

1. Engelbart, Douglas C. "[Augmenting Human Intellect: A Conceptual Framework.](#)" Summary Report, Stanford Research Institute, on Contract AF 49(638)-1024, October 1962.
2. Zoellick, Bill. [CyberRegs](#). Addison-Wesley, Boston, USA, 2002.
3. American Library Association. 26 October, 2000. "[New Digital Copyright Rules Seen As a Defeat for Library Users and the American Public.](#)" Washington Office News Release.
4. Christensen, Clayton. [The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail](#). Harvard Business School Press, Boston, USA, 1997.
5. Engelbart, Douglas C. "[Toward High-Performance Organizations: A Strategic Role for Groupware,](#)" in Proceedings of the GroupWare '92 Conference, San Jose, CA, August 3-5, 1992, ([AUGMENT,132811,](#))