

Computers and Democracy

David Gelernter
Dept Computer Science

“Computers and democracy” is a big, vague, hot-air-balloon of a topic. I’ll try to bring it down to earth; I’m not sure whether I’ll succeed.

American attitudes toward democracy and computers are curiously similar: we depend on both, believe in both, pride ourselves on both; but we don’t necessarily care much about the details of either. The first role of any well-functioning democratic government or desktop computer is to annoy us as little as possible. Traditionally we have not admired, in this country, people who are too eager to be part of government; we are not too crazy about computer experts either.

Computers are likely to affect American democracy mainly by working on American culture—on the atmosphere in which democracy lives. Networked computers at their best are like spring in the air—among human beings, spring doesn’t make things happen, but it gives them an excuse to happen. It gives us an opening. Computers can do that too. They can put us in the mood for change. They can give us an excuse to get good things done.

But the good things won’t happen automatically. Many of us approach technology deferentially. We are content to sit back, see what happens, let matters take their course. Many of us have the idea that software and computers ought to be left to technical experts. Some of us, for that matter, feel the same way about government: that mere people aren’t equipped to decide.

But computers and the internet have a shaping effect, nowadays, on American culture, American language and the texture of daily life. Their influence will increase. Their gravitational field will continue to grow.

To design software and help guide technology evolution takes wisdom, experience and common sense—qualities that are no easier to find among technical experts than anywhere else. Americans ought to take technology in hand, ought to make computers their business (not shrug them off and let the experts decide) the same way they make government their business: not because they are ideally qualified to call the shots; only because they are better qualified than anyone else. Computers could be a great thing for American culture and therefore for American democracy and for life in general. They could also bring out the worst in us, leave us richer in money and worse off in spirit than we have ever been before. The outcome will depend on whether we sit back and let the technical experts run the show, or assert our right and our duty to shape the evolution of this astonishingly, unsettlingly important technology.

First let’s consider the mechanics of democracy: what are the likely consequences of using computers as voting machines, and how should we go about it? Then, the quality of our information supply; finally, the role computers play in creating a national cultural ambiance.

People have been thinking for a while about using computers as voting machines. Proposals come in many shapes and sizes, but let’s take a simple and radical one for concreteness. Let’s say you can vote using any net-connected computer during official voting hours, at home or in any other place.

Let’s forget about the details: we would have to make sure that only qualified people vote, that no one votes more than once, that each voter sees the right ballot for his area. We would have to make sure that accumulating vote totals aren’t messed-with or snooped-upon. We would have to

make sure that net-connected computers are no less accessible than today's voting machines, ballot boxes and so forth. We would have to decide what kind of software to use and how the interface should work. (Presumably each state would make its own decisions.)

We would have to solve other problems that computer-voting wouldn't create but would bring to the fore. For example: if we are allowed to vote using computers in our homes, many of us will spend a good part of every election day within a few steps of a voting machine. Inevitably we will get to wondering whether we can change our votes, and then change them back. We will need a policy and a way to carry it out.

But let's assume that we have solved all these problems. Some are hard, but we could probably solve them all within a year or two if we felt like it. The real question is: is this a good idea?

At first it sounds like a rotten idea; at least it did at first to me, for several reasons.

The most important thing about the process of an election is not that it should be absolutely accurate but that it should be, as far as possible, transparent. Tampering with the results should be difficult and should SEEM difficult. We ought to be able to picture clearly in our minds the procedure that herds millions of scampering little votes into a majestic grand total. If we can picture it, we can believe in it. And today on the whole we CAN picture it; the elections we hold nowadays score fairly high in transparency. A paper card or ballot is a concrete thing. So is a voting machine. Personally I've always liked the solid little crunch that a voting-machine lever makes when you swing it down, and the ringing thrunk when you open the curtain and register your votes. To me these have always been the actual sounds of democracy. I suspect other people feel this way too. It's easy to picture the process that connects our tangible voting acts to the final total. Maybe our picture is incomplete, even wrong; and of course computers are heavily implicated today in vote counting. But at least the voting act itself is easy to understand, and the mental picture exists and is a comforting thing to have.

Computer voting might be far more accurate but less transparent, and I'm not sure that's a good trade. Every programmer knows not to fix what isn't broken. Every programmer also knows that software is hard to get right and prone to break—and everyone else knows it too. I think we COULD build voting software that, even if it is prone to error and fraud, is more reliable overall than the system we use today. And I think we could sell it to the American public. But our elections would become less transparent; your vote would amount to a handful of bits in a complex series of software structures, and very few people can picture how software works. Inevitably we would pay a price in public confidence: we would have less in the bank when disputed elections DID come up, as they are bound to occasionally no matter how we do our voting. The real story of our last election was the crisis that did not happen. If we change our way of voting, our first responsibility will be to do no worse than we are doing today.

At least one other problem with computer voting comes to mind. Voting shouldn't be too hard, but it probably shouldn't be too easy either. Making people leave their homes and go to some public place and maybe even wait on line is no unbearable imposition. But it does quietly underline the importance and the dignity of an election. It seems right and fitting for voting to be done in a public place—usually with flags on hand, with your neighbors standing around, with the feel of a special occasion. Voting is a little (just a little) like getting married. Other things being equal, the nation has an interest in its citizens marrying and voting. Once they've made up their minds to go through with it, we don't want to talk them out of it. But we don't want to make it TOO easy either; we want to make it clear that these are not casual acts and that deliberation is called for. In theory you could get married in front of a PC, too—you would go to the wedding website of your choice and download some appropriate music and the necessary forms. But this is a bad idea.

Nonetheless it almost certainly WILL happen, and computer voting is inevitable too. When we do switch to computer voting, we ought to stipulate something like this: whatever company or agency builds the software is required to be able to explain how it works to any voter in five minutes—not in comprehensive technical detail, but in a fundamentally accurate way. It's a difficult requirement, and admittedly subjective. But it would put pressure in the right direction.

And on second thought computer voting might, despite its disadvantages and even dangers, be good for American democracy, if we play our cards right. It might turn the public's lack of interest in politics into a force to improve politics. By making it easy to get the minimum information you need to vote at the moment you need it and not a moment sooner, it could help make political campaigns as we know them disappear.

Why? Because, first, the social dynamics of computer-based elections would probably allow voters, and even encourage them, to take a few minutes when they cast their ballots. To take their time. Even if you were voting on a computer in a public place, computer habits are different from voting-machine or ballot-box habits. People expect to sit down and think when they use a computer.

So, we might allow each candidate to put a short statement on file, maybe 500 words; you could read as many of these statements as you felt like before casting your vote. You'd click the "show me the statement" button next to the name you were considering, or something like that. Our election laws might have to be changed to allow it, but they might be worth changing.

What's potentially useful here is not the mere existence of these statements—there are plenty of such statements floating around today; what's interesting is the way they would make themselves available at the right moment to be worth reading, and in a setting where reading them seem like a reasonable thing to do. Voters who dump campaign mailings in the trash might nonetheless read these. Voters who ignore the leaflets thrust in their faces by campaign workers hovering hopefully outside the polling-place perimeters might read these statements. There's a big difference between information you ask for and information that is flung at you the way popcorn is thrown at elephants by zoo visitors. When you ask for information your attitude is different, and so the information itself can be presented differently. Most Americans aren't sufficiently interested in politics to ask for information ordinarily—but we are a practical people, so we say, and voters who would never ordinarily solicit this kind of information might conceivably ask for it if they were face-to-face with the ballot AND the asking were easy.

This sort of last minute, point-of-sale message would inevitably carry weight, and—more important—give voters an excuse to do what most of them want to do anyway: ignore the campaign entirely. Block it out, with the thought in mind that the facts will be available when they are needed. The whole plan conflicts with the idea some people have of voters solemnly pondering their election-day choices for months beforehand, but I doubt whether this view of American voters is terribly realistic. In any case we might gradually increase the amount of information that's available at voting time—starting with the 500-word statement, which would no doubt turn into a video snatch, we could add a 2000-worder for voters who ask for more, and then a 5000-word statement and so on. Meanwhile the whole rest of the campaign might grow increasingly vestigial.

The difference between information you ask for and information that is thrown at you affects all of American politics. We turn our politicians into court jesters and professional party crashers, begging for small handouts, mugging and cavorting for attention. And then naturally we despise them. Putting a bit of information in the right place, at the right time, in the right way, won't in itself cause our existing campaign system to collapse. But it might give it an EXCUSE to collapse.

In fact computer voting is probably our best shot at giving it an excuse to collapse. I don't think we will ever have publicly-financed campaigns in this country, to any greater extent than we have them (at the presidential level) today, and I don't think we should. People don't enjoy paying for things they're used to getting for free and don't want anyway. Given the state of American election campaigns, forcing the public to pay for them would be like forcing it to pay for junkmail or obscene phone calls. What candidates desperately need from us is not our money but our attention. They only want money to BUY attention. Donating a little attention free of charge would be the perfect fix—and if we could do it strictly on our own terms, we probably would. Neither the campaigners nor the campaignees like today's system, and it will fall apart if we give it a chance.

There's also a general point here about the likely evolution of technology. It's obvious that the value of a piece of information depends on when you get it. Information that is worthless at the wrong moment can be invaluable at the right moment. Using computers to deliver information at the right time in the right way is likely to be a big topic of technology-and-culture research over the next decade.

For example: we've been working for several years on a research software system, now also a commercial project, called "lifestreams." The idea is to maintain automatically a comprehensive time-ordered journal or diary or scrapbook of your electronic life.

Every document, image, email, Web bookmark, news update or any other piece of information you receive or create gets dropped at the end of a constantly-growing, time-ordered stream. So far we've mainly concentrated on the past and present of this stream. At the head of the stream you see the latest stuff; as new documents arrive, old ones move back. So the stream flows, the way time flows and life flows. The stream has a future too. You store your plans and appointments in the future. They flow toward the present. When they reach "now," they hop over the "now line" and shuffle off into the past.

But we can also use the stream's future as a staging ground for information that we would like to deliver exactly when it's needed. If you're traveling (for example) and will need information about your plane ticket around 3 in the afternoon, about ground transportation two hours later, then information about your hotel, then some phone numbers to call, then some background information for a meeting the next morning, and so on, we can lay out all these documents in sequence in the stream's future (like clothes some mythical valet might lay out on some theoretical bed); each piece of information cruises forward towards "now," and presents itself when its big moment arrives.

So changing campaign politics might be, in certain ways, one facet of a larger problem of technology and culture. The "information age" is all right as far as it goes (it's been underway for at least a century if it exists at all); but what we really need is an "information at the right time" age.

What does democracy need? Citizens who are well informed and thoughtful, and feel responsible to the community. Computers and democracy intersect as soon as we talk about information: computers and networks supply us with some of our information and, more important, they handle nearly all of it. Nearly all the words that are published, and nearly all the prepared words that are spoken, pass through computer word processors, over networks and out of printers.

Orwell writes in his famous essay on politics and the English language: "If thought corrupts language, language can also corrupt thought." It's no accident that the 20th century's most notorious tyrannies expressed themselves in opaque, muddy, swollen language. Crisp clear clean information—sparkling if possible—is a necessity of democratic life. A democratic society had better be an articulate society.

Are computers good or bad for prose? In a sense, the word processor and the cheap printer are the best things that ever happened to writing. To write well is to revise obsessively. Word processors are perfect tools for trying out the look, sound and feel of a sentence before you commit yourself. (Writers don't like to commit themselves to anything as important as a sentence until they are absolutely sure.) Computer printers let you put your draft on paper, change it, print

it again, change it again and on and on with the neurotic persistence of a child working a yo-yo until your editor goes crazy.

Email brought the personal letter back from near death, admittedly in a strange new shape. During the '70s and '80s, phone calls kept growing insidiously cheaper, and then the phones themselves started to get cheaper, and smaller, and wireless—it was a nightmare out of Hitchcock; it seemed clear that before long, we would be crunching cellphones underfoot like cockroaches. And meanwhile the US mails were nothing to write home about; and by around 1990, the personal letter was gasping for breath. The rise of email has put it back in play.

And yet computers can also be bad for writing. Email resurrected the personal letter. It might also be fostering a culture in which our prose is so feeble we have to load it up with little smiley and frowney faces, because otherwise no one can tell whether we're kidding. My sons are often told to write school papers based on research using the Web. The great thing about this form of research is that, for the first time in history, you can write a research paper in less time than it takes to read one. And of course there is no quality control on the Web; there are high-quality sites alongside junk sites, where junk facts are expressed in junk prose.

Once again the technology is neither good nor bad in itself; it's what we make of it. But if I had to guess, I'd predict that the new technology of writing could yield something remarkable for American culture—by giving us an opening, a good excuse, to do what we want to do anyway.

Today we have by far the best writing-tools that ever existed. We are also only a generation or two removed from the greatest outpouring of brilliant writing in American history—an extraordinary eruption that lit up the skies from the 1920s into the '60s or early '70s. There are always a few brilliant writers around; you can find them today—guaranteed. But the interesting thing about the great age of 20th-century American prose is how the formidable parts added up to a remarkable whole, how dozens of sharply defined and distinctive individual voices added up to a national style that caught up the whole country and carried it forward—these authors wrote poems and art stories AND bestsellers and popular magazine pieces. Such writers of vivid, crew-cut English as Anderson and Gertrude Stein and Hemingway leading naturally to the great essays and journalism of White and Mitchell, Liebling and Ernie Pyle, and on to a new crop of novelists and essayists who continued in their own voices the same vibrant, funny, often belligerent national narrative—Mailer and Irwin Shaw, Bellow and Updike and Roth and Eudora Welty and Tom Wolfe, Cynthia Ozick, James Salter. I could go on. This remarkable outpouring ran out of steam in the '60s or maybe the early '70s. But that wasn't so long ago—the engine is still warm, and it would start easily if there were something to kick it back to life. Maybe our unprecedented new writing tools will do that, together with email and its way of making it easy for writers and editors to work together.

Computers wouldn't cause this big cultural event; all the great writing and editing technology in the world never manufactured a single decent sentence. But it might give us an opening or occasion or excuse to get back on track and do great things.

America has many things on its mind; we are not a monomaniac nation. But it seems to me likely that, more or less continuously since the 1930s and occasionally during earlier periods, some Big Theme or other has preoccupied us. The Big Theme is the spare topic in every lagging national conversation. It's the pattern on the national wallpaper. During the '30s the Depression was the big theme. Then the War and its immediate aftermath took over; then, for the forty-odd years between the Berlin airlift and the collapse of the Soviet Union, it was the Cold War. Today, our big theme is computers and the internet.

Identifying a single big national theme is a thought experiment, an exercise in deliberate oversimplification. But such exercises serve a purpose. If you are trying to draw a human figure, for

example, you will have many complex details to deal with. But you'll also have the problem of finding the one characteristic line of action or tension or flow that makes the drawing come to life. That one line is not the figure, or anything close to the figure; it's a radical abstraction. But it's a necessary abstraction. It doesn't suppress the details; it makes the details emerge more clearly.

Granted it seems crazy to propose an analogy between the Cold War and the desktop computer, in the role of national Big Theme—they come from far-apart categories of existence. And I'll admit that such a claim can only be arbitrary and subjective. But I can't shake the strange feeling, and it is strange, that the role the Cold War played in my own childhood is being played in my children's lives by the internet and computers. When I was a young boy President Kennedy, the intrepid Cold Warrior extraordinary, was the official hero of young boys—not because of the Cold War per se; because of PT-109, a story of heroism in the War against Japan that every child knew. Of course we admired other people too; war heroes, sports heroes, Hollywood heroes, TV and pop-music stars, and all Presidents of the United States ex officio. Probably we admired presidents most. Today we're told that Bill Gates is the man children admire most. It seems impossible—but it fits.

Any analogy between the Cold War and computers can only be a sort of half-analogy, something like a half-rhyme or an unresolved dissonance. The half-analogy is inherently awkward and it's not an officially sanctioned rhetorical device, but nonetheless it can be useful.

So consider this half-analogy between the Cold War then and computers now. The most important function of a Big Theme is to answer questions automatically. The big theme is an answering machine with the prerecorded message built-in. If it helps our position in the Cold War, it's good. If it has to do with computers or the internet, it's good.

This is a valuable service at times. It keeps you from going over the same ground repeatedly. It allows you to put a national consensus to work and get things done.

But it can also be dangerous. The Vietnam War for example was a watershed catastrophe in American life. We fought in Vietnam for many reasons, but ultimately the Cold War made us do it—our global struggle with communism and with the Soviet Union. Whether you believe that it was noble or evil for us to have gone into Vietnam, it's hard to doubt that we entered the War without thinking carefully about what we were doing. During its early years, as we edged out calmly into those huge rolling breakers step by step—breakers that eventually swept away so much that was good and valuable in American life, above all our endless confidence in ourselves and our future—there was relatively little discussion or argument about Vietnam. Not much discussion of our goals; not much discussion about whether our methods were in line with our goals. The arguments came later. In a sense, they came too late.

Ideology is a sedative—makes you accept propositions without thinking, tells you the answers without making you solve the problems; it's dangerous to operate a nation when you're under the influence. Today computers are the big theme, and they are a new sort of ideology -- an ideology without ideas. Computers are an ideology the way Bill Gates is a hero. American culture is in peculiar shape.

Today's pre-recorded message is: if it has to do with computers or the internet, it's good. Karl Marx made the famous pronouncement that history repeats itself, the first time as tragedy and then as farce. (No doubt this epigram is what he'll ultimately be remembered for.) Probably there is no national tragedy lurking in our fixation on computers. So far there has only been farce: foolish, expensive mistakes.

Take the collapse of the dot-coms. People poured money into companies that had no logical reason for existing, on the theory that commercial reality could be suspended like alternate-side-of-the-street parking rules so long as your business-plan centered on the Web. The Big Theme made people turn off their brains.

Take the Clinton administration's aggressive promotion of internet in the schools. Most Republicans supported the policy in principle even if they opposed the tax rise that became part of it. Our schools were evidently failing to teach reading, writing, history and arithmetic— but not

many people ever claimed that our students suffered from information underload, or excessively long attention spans, or not enough in-school distractions. I don't claim there was nothing to be said in favor of putting the internet in schools, but there were definitely things to be said against it, on the grounds that at best it was a distraction and at worst, an attractive nuisance—and almost no one said them. Another reminder that we have two political parties in this country so they can disagree, not congratulate each other on bipartisanship. Especially on domestic policy, political parties are supposed to disagree. If they don't, there's probably something wrong with them.

Computers didn't ask to be appointed our national Big Theme. The situation developed by accident and is nobody's fault. Nonetheless— as big themes go, this a bad one. It makes us painfully credulous. We need a loyal opposition in this area and we don't have one. We are setting ourselves up for bigger business failures, and possibly in the long run for an anti-technology backlash. More important, "computers and the internet" makes an unworthy big theme for a great nation, because it is morally, spiritually and intellectually vacuous. It's not that computation is a vacuous topic in itself; not at all. Developing the science and art of computation was one of the twentieth century's greatest achievements. It's just that we choose to treat the subject as if it WERE vacuous. More than a generation after most universities got themselves computer science departments, only a handful of people in and out of universities have any idea what computer science IS. The public mostly doesn't care, and computer scientists mostly don't care whether the public cares or not.

A nation doesn't consciously choose its themes, and it would be foolish to suggest what the next one should be. But we can make predictions if not suggestions. National big themes are likely to grow out of previous themes. In a certain limited sense, modern computer technology did grow out of the Cold War. American military research was the driving force behind the original Arpanet of the late '60s, which became the internet in the '80s; military research funding was a factor in the emergence of nearly every part of modern computing. Of course there were many other contributing factors, too -- freedom of thought, free teaching and research, free markets, and the nation's tendency to attract many of the world's smartest people to its corporations and universities. Ultimately these are all facets of American democracy itself.

No one can predict future themes with any confidence, but if we had to guess, we should probably guess that the next one will in some sense grow out of the current one—unless it has to do with some new and unforeseen national emergency. (In other words, with a resurgent Russian Empire.) But we know that the next theme will be something concrete. People have proposed community, civility, spirituality as important topics for the national agenda, but they wouldn't even make good titles for Directed Studies courses.

One possible new theme might have to do with the emergence of new internet-based cultural institutions. During my lifetime, thousands of major new businesses and corporations have been created in this country, but not a single important new university. No important newspapers have been born, although many have died; some first-rate new magazines have been created, but only one or two swing real cultural weight. Lots of publishing houses have bit the dust. I could go on: it's been an extraordinarily creative era for business institutions, and a dead period for cultural institutions.

Creating new cultural institutions is important because existing successful ones have an obligation to continue succeeding. They are conservative by nature. Yale's most important obligation is to go on being Yale. If technology means that a university might become a radically different kind of institution, Yale is the wrong place to run the experiment. Radical experiments make sense at institutions that have little to lose. That's why new institutions are so indispensable to American culture; they are the right settings for radical experiments.

Of course technology will change existing successful institutions too -- will change Yale, for example, although the changes had better be in character or they won't have a chance and won't deserve one. Universities like Yale are working hard to deliver courses online. It's an interesting, important project. We might also consider the related but different project of putting the university itself online.

We might use the structure I mentioned before, a so-called "lifestream." It would be crazy to suggest that this structure is the only way forward as we build new institutions and cyberfy old ones— not at all; but it gives us a concrete basis for discussion.

What we might do, briefly, is build an electronic campus in the form of a flowing, time-ordered stream of information. Every new document or plan that the university generates, each announcement, newsletter, menu, photo, departmental memo and whatnot is added to the head of the stream—which might actually be more of river; the flow is fast, the volume high...

Onscreen, the stream looks like a receding parade of dominoes, one domino per document. You can see everything, or focus on topics that interest you and make the rest temporarily disappear.

When you look at the head of the stream, you see the latest postings. The stream is in constant motion, like a news-ticker. New things are added; old ones move back. The stream has a future, too; the university's calendar and plans are stored there, and they roll steadily toward the present. So this electronic campus takes the form of a stream with a past, present and future; the future flows into the present, the present flows into the past. New documents are posted in the present and flow into the past; or, if they are plans or calendar items, they're posted in the future, flow into the present and on into the past.

It would be natural for course material, lecture notes, videotaped presentations, concerts, performances, conversations of all kinds to join the stream. And naturally the stream would extend back to 1701 -- all the documents, drawings, photos, films and so forth that document Yale's history would be filed in the stream in proper time order.

Each member of the community has a private stream view, which includes the whole public stream and all your private documents shuffled in— emails, files, memos, Web bookmarks, photos and so on. When a new document shows up on your stream, it might be email for you, a university posting for everyone, or anything in between.

Everything is on the stream. We need simple, powerful ways of finding and focussing on what you want, but I'll skip the details. (A staggering amount of information joins the stream every day, but if you focus on "physics department" or "art gallery" or "football games" or "recent email from schwartz," everything else disappears and you have a manageable substream.) When a freshman arrives, we set him up with his own view of the stream. Your undergraduate years correspond to a four-year-long stretch. When you graduate, you keep your stream account and you can tune in the university whenever you like.

How does the electronic campus compare to the real one? Its disadvantages are obvious. On the real campus people can talk face to face, play football and Beethoven and look each other in the eye. But the electronic campus has certain advantages too. You can be part of it no matter where you are. Everything takes place in historical context: you can rewind the electronic campus to last term or last decade or last century. Everyone can watch the passing scene from a front row seat.

Streams like this are up and running at a few small institutions, for example at the New Haven technology company that built the commercial version of the software, Mirror Worlds Technologies. One of the stream's big advantages at Mirror Worlds Tech is that when you're tuned in to the stream, you feel that you're right at the center of the action even if you are nowhere near the company's offices. One of stream's big disadvantages is exactly the same: you feel you're right at the center of the action, but you aren't. You've lost face-to-faceness, and that's important. You gain something in return; it can be a reasonable trade so long as you don't kid yourself about the cost.

Before long Yale will be in the good position of having both kinds of campus, real and electronic. Similar techniques make it easy to imagine that we could create new universities whose faculty and students are spread out all over the world. These new institutions would exist mainly online. Their members would meet occasionally, but they're unlikely to have anything like an actual campus. Do the potential gains make up for the big losses? I don't know. But I'm sure we'll

have a chance to find out. This experiment will be tried. We're seeing some tentative first steps already.

Everyone knows that it's nearly impossible to start a newspaper nowadays. Of course it might be easy to start an electronic one; but on the whole, Web news services have been disappointing. They're good at providing focussed, specialized news, but not at creating any sort of document that people turn to for the pleasure of browsing and reading. We know one thing for sure about electronic newspapers: when a successful one is created, it won't look anything like an online version of the Times or Wall Street Journal. The organization and design of today's newspapers reflects the character of printing on paper. A successful online newspaper will come in a new shape and use new organizing principles that don't exist on paper. It might be a stream, for example, where you can look at the future to find out what's coming up, and the present to find out what's happening, and you can reach smoothly into the past as far as you'd like. Or it might take some completely different shape. In any case its shape will be new, not a mere translation of paper into a new medium.

Suppose electronic newspapers were structured as streams. You could take the equivalent of the the New York Times stream and the Wall Street Journal stream and shuffle them together the way you'd interleave two decks of cards, and read both newspapers simultaneously. You could follow a story backwards step by step to the beginning if you wanted to. You could make everything disappear but the graphs and pictures. You could read a newspaper once a day, once a week or once a month and it wouldn't make any difference; whenever you tuned it in, you could take up exactly where you left off. The stream is an easy structure to navigate—you can go forward, go back, look for this or that—and so you could listen to a newspaper while you drove to work, by having your car computer read out the newspaper stream. And so on. Whether or not these possibilities are useful remains to be seen. The point is that we can experiment. Newspapers are important to democracy. Nowadays they aren't doing terribly well. Computers will give them a chance to change.

Some of our new institutions won't have any exact analogs in today's cultural universe. Take the following obvious stream-based service, or institution, or whatever you want to call it: New York City is the center of the publishing industry, but it's radically under-served in the book review department. There are only two weekly book reviews for the general public; loads of new books are barely glanced at, or ignored altogether. In fact there are all sorts of New York cultural events that get cursory notice or none. I have an artist friend who worked for years on a big New York gallery show which finally opened not long ago, and everyone waited anxiously for the reviews, if any— and then last week, a review at last; she and her show got blown away in a single paragraph in the New York Times. It shouldn't be. In this allegedly information-rich world, the arts are starved for attention.

Inevitably someone will start up a New York culture-stream service: it will publish news and reviews but will also be a ticket and shopping and calendar and archive operation. You'd tune in the culture stream; at the head of the stream you would see late-breaking news and reviews. You could watch everything or focus on books, music, art, lower Manhattan, whatever. If you read a review for a show you want to see, a seating-plan for the theater is on the stream, so you can look it up; you buy your tickets by clicking the right boxes. The tickets are delivered to your stream, and a note is dropped into the stream's future reminding you to go see the play when the date rolls around.

We're already seeing some tentative movement in this direction; we'll see a lot more before long. These culture streams and related electronic institutions have no analogs in the pre-Web world, but they could easily become more influential than any newspaper or TV or radio outlet that exists today.

Computer technology in itself accomplishes nothing. If we're going to have an American cultural renaissance, computers won't supply the ideas or the will, let alone the people or the money. Computers won't change our election campaigns or improve the quality of our information supply or create any institutions. But they might give us openings, if we choose to go through them.

In conclusion, I'll repeat a few claims and then summarize my basic suggestion.

Information is no good in itself. It becomes valuable when it shows up at the right place, in the right way, at the right time. The technology world hasn't gotten around to taking this problem too seriously yet, but it will.

We could make transparency a goal of software design if we wanted to. We could require that every significant piece of software be able to explain not only how one operates it but how it's built, how it works. Nowadays few people care. As software becomes increasingly central to the workings of society, MORE people will care; at least one hopes they will.

The quality of our language is central to the quality of our democracy. By making it easier than it's ever been to generate written words, computers could be the best or the worst thing that ever happened to our capacity to express ourselves. The outcome depends on whether or not we believe that good writing is important. Today on the whole we don't. But computers could give us an excuse to change our minds.

We each have a fixed number of hours to dispose of, and we spend some of them dealing with information. Other things being equal, the more bytes of information we look at, the less time on average we invest in each. It's good to know a lot, to have seen a lot, been exposed to a lot. It's also good to study important things in depth. I think we need to re-calibrate the all-important balance between depth and breadth in American education and maybe in society at large.

Thanks in part to the Web and to loads of other electronic information sources, our balance is temporarily out of whack. I sometimes think about early American households, where books were rare and you might find a Bible, maybe some Christian or patriotic pamphlets, occasionally the works of Shakespeare. It wasn't an atmosphere that encouraged studiousness, but when a child WAS studious—when he was Abraham Lincoln, for example—he put all his energy into a relative handful of bytes. But he learned them well and thought about them deeply. His thinking was deep instead of broad. Would his education have turned out better if he'd had the Web to surf? I doubt it.

Deep reading, close reading, careful and meticulous reading is the basis of education. We can't let the Web or anything else make us forget how to do it.

Finally, it's been typical of American society in recent generations to have one big basic theme. If that's so, our big basic theme today is computers and the Net. We can do better. It's time for a new theme.

Over the years—in conclusion—we've come to accept responsibility for the air we breathe, the water we drink, the schools we operate, the culture at large. We could accept responsibility also for the computers we live with. We could reject bad software, send it back the kitchen; think carefully about how good software ought to behave and what it should do, and about where the field as a whole ought to go. If we do that, if we choose technological democracy, then (chances are) we will realize the great possibilities of computers and the net.

If we don't—if we sit back and leave computers to technologists and the industry—then society as a whole might easily come to resemble a piece of commercial software. Might grow steadily more complicated, with more and more fancy features that fewer and fewer people understand; with an underlying structure that is opaque and mysterious. The gap between normal citizens and expert users would grow larger and more ominous all the time. Gradually the spirit of civil society might come to resemble the spirit of the Microsoft paperclip—an onscreen cartoon

character, for those of you who never met it, that Microsoft uses to express its cheerful, charming, easy-going, good-natured contempt for us poor child-like fools who use its software.

But it doesn't have to work that way. Computers are one of the most powerful tools we have ever invented, recently. The choices on this ballot proposition are clear. What to do with this amazing tool? (A) We use it. (B) We sit back and let it use us.

I vote for (A). Thank you.