

**HECHT:** Let me introduce myself for those handful of you who don't know me. I'm Bill Hecht, class of '61, your senior employee, and for the past 20 years, and I suppose in perpetuity, Executive Vice President of the Alumni Association.

I want to welcome you to Saturday classes. Those of you who are my age and older remember Saturday classes. Some of us are age will look around and notice that a lot of the younger classes are still missing, and they don't have Saturday classes today. And therefore, some of them haven't figured out that we are going to have class this Saturday.

We provided you with a real, genuine fire hose yesterday evening, and I hope some of you stayed dry. I didn't. This is the intellectual part of the fire hose. Before I introduce this morning's program, a number of alumni have buttonholed me and asked about Mrs. Vest's health. For those of you who don't know, Mrs. Vest suffered a massive coronary about two months ago. The miraculous thing is she's made a remarkable recovery.

And you will note that Chuck and Becky were far less visible at this event and then they normally are. She's doing quite well, and we hope to see her at lunch, but I can't guarantee that. It is, in the words of Chuck Vest, a miraculous recovery. So to those of you who've sent your good wishes, and I know many alumni have, thank you very much.

Today's technology programs are put together by a committee. And one of the first things a committee does, because we're here at MIT, is to call on a distinguished faculty member to help us. And this year it was fairly easy to decide what to do because we had Bill Mitchell, who's the Dean of Architecture, a Professor of Architecture and Media Arts and Sciences, who has come to us in the last several years and brought an enormous amount of excitement not only to the School of Architecture and Planning, but to the whole campus with his ideas, his vision, and his sense of excitement about how physical spaces are just as important in the digital age as they were in the analog age and in the Stone Age and in the Bronze Age and all of those ages that have gone before.

Today's program, as you know, has an interesting and catchy title. It's "The Future of Atoms in an Age of Bits." I'm always encouraged that atoms still have a place since I seem to acquire a few more every week. And I haven't been able to figure out how to do something about that. My wife teases me occasionally that I'm just becoming a real candidate for cloning, and someday there will be enough of me actually to go around.

Without further ado, let me introduce Bill Mitchell who will introduce the speakers in today's program and his own set of remarks. Bill?

[APPLAUSE]

**MITCHELL:** Well, good morning. I'm delighted to be here, and very happy to introduce a, I think, very exciting group of speakers for you this morning. Let me quickly sketch out the program for you-- what we're going to do-- and then I'm going to jump right into a presentation.

I'm going to lead off, and I'm going to talk about the future of all cities in general terms. I'm going to try and lay out a general framework for what's really happening to cities? What's really happening to architecture? What's really happening to our physical environment as the digital revolution unfolds all around us?

I'll be followed by Yossi Sheffy, who's Professor of Civil and Environmental Engineering and Director of the Center for Transportation Studies. And he's going to talk about-- well, the title is, "Transportation Auctions and Exchanges," but from talking to Yossi just before we started, I think his talk a little wider than that.

After the break, we'll have a talk from Rodney Brooks, who's the Fujitsu Professor of Computer Science and Engineering and Director of the Artificial Intelligence Laboratory, and a real pioneer in some very, very exciting areas of artificial intelligence.

And then finally, we'll have Roz Picard, who's an Associate Professor in the Media Lab, the program in Media Arts and Sciences. And she'll talk about the emotionally smart machine.

We are going to organize things in this fashion. We'll have the first two talks, then we'll take a break. And then we'll have the second two presentations. And then following those second two presentations at about 11:30, if we stay on time, and I hope we will, we'll have a question and answer session.

And the way we're going to do this to make it work in this auditorium is that we'll handout cards for you to write questions on, and we'll ask you to write your questions out, hand them to somebody who will be very visible at the time. And we'll pick up those questions and take it from there.

So let me jump in and-- good, the technology worked. I got my first slide up here. I'm going to talk about a subject I've called e-topia, digital telecommunications and the future of cities. Now, in architecture and urban design, the term utopia has been around for a long time.

Architects and urban designers have often dreamt about the ideal cities, ideal futures, and the term the utopia from the Greek follows nicely to describe all of that. Or if you want to take a less optimistic view, sometimes people speak about the possible dystopias of the future.

What I really want to talk about today is our electronically mediated future. So I've taken the Greek root that generates utopia and dystopia and put an "e" front of it as we put an "e" in front of many things today. And I'll talk about e-topia. And I'll leave open the question about whether the future that we face is a future filled with remarkable positive possibilities, or whether the downsides are going to dominate. I'll try and frame the questions for you and ask you to think about where some of this leads.

I'm going to base what I talk about on three recent books. I wrote a book several years ago called *City of Bits-- Space, Place and the Infobahn* that looked very early on in the evolving internet era, looked at what might be happening to cities.

We in the School of Architecture and Planning then moved on and did some very interesting work on the question of high technology and low-income communities. What does this set of developments mean for social equity, for the gap between the haves and the have nots? Is digital telecommunications technology increasing that gap? Is it reducing that gap? What are the potentials for the future?

And most recently, from the MIT Press, I did a book called *e-topia* that pulls a number of these ideas together with the subtitle *Urban Life, Jim-- But Not as We Know It*. You have no idea how difficult it was to get to a subtitle like that through a respectable academic press, but we succeeded.

So in order to open up this topic, let me go right back to some basics-- really this is an urban design 101-- and say something about the relationship between networks and urban structure. As cities have evolved over the centuries, you see that they've acquired increasingly numerous and sophisticated networks to tie them all together, networks of tracks and roads, street systems, road systems, networks of pipes, networks of wires, more recently wireless channels, joining together patterns of production locations, storage locations, switching locations of various kinds, and consumption locations of various kinds.

Now, these networks have never been distinct in their functions, but as we'll see, through today, they start producing all sorts of complex, joint effects. And historically, we've seen that when new networks are overlaid on existing cities, existing urban patterns, they change the distribution of resources and opportunities within the urban fabric. They restructure the urban fabric, and this often creates winners and losers, and some very, very important social equity issues. And I'll come back to these a little bit later on.

Now, digital telecommunications infrastructure is the latest overlay that's transforming cities in this way. It's an infrastructure physically that consists of copper channels, glass channels, and wireless channels of various kinds, and then a collection, as you all know, of systems of servers at various locations, switching nodes, and incredibly numerous end user consumption points.

I want to say something about-- just to characterize what's going on, I'll try to illustrate to you how old and new infrastructures often overlay. It happened a week or so ago. I was driving down the El Camino Real in Palo Alto.

And as you probably know, this is one of the oldest pieces of urban infrastructure on the North American continent. It's the old road that tied together the Spanish missions and settlements way back in the 18th century and made a connection up the coast, as you see here. It was a channel for pedestrian traffic, for horse traffic, for horse drawn vehicles.

As I was driving along there the other day, I noticed a scar down the center of the road. And I drove a little further along, and I found a backhoe continuing that scar down the road. And you could probably guess what happened.

As I took a closer look at it, I saw that in fact, what's happening here is that right down this ancient right of way, this ancient piece of infrastructure, there's a new, very high speed fiber optic backbone being laid. So rather poetically, in exactly the same location that the first infrastructure in this area was laid down, the latest infrastructure is being located.

Let me show you something else about this just to set the scene a little bit. This is a famous monument of modern architecture that you see on the slide here right now. It's in the City of Chandigarh in the North of India, a capital city that was designed by the architect Le Corbusier to be the capital of the Punjab.

And at the heart of his design there's a large open piazza which was intended to function as a face-to-face meeting place, a place where people could come together, exchange ideas, really be something like the ancient Greek Agora. And the center of this was this famous Open Hand Monument grasping at the sky and symbolizing that notion.

Now, it turns out to be the case that this piazza never worked partly because it was too big, partly because it was too hot, partly because it's too difficult to get to. But it represents this ancient idea of the way that people can come together face to face to constitute their community.

Now, down the southern end of Chandigarh, exactly the opposite end of the city from this monumental center, is another monument grasping at the sky, I noticed the last time I was there. Of course, this one actually catches something, unlike the first one.

And this is a satellite dish, a new piece of-- it's a satellite earth station. This is a new piece of infrastructure, and it's making, if you like, a digital oasis in this area. It's a location that it connects this point on the surface of the Earth to the international digital telecommunications network, the global digital telecommunications network.

There's a little switching station there-- microwave tower behind it. And everything within line of sight of that microwave tower is digitally connected to the outside world. So this is a center of community in a very different sense from the center of community symbolized by the open hand monument.

Now, what I want to do quickly is go through some questions relating to what's really going on in cities as a result of this. How is new telecommunications technology changing urban patterns? Is it simply substitutional? Are we merely replacing familiar means by e-this and cyber-that? Or do these new ways of doing things add up to a much more profound structural change?

And I'd like to suggest to you that in fact the changes are fundamental, and that we cannot plan for future urban patterns, and indeed we cannot plan the campus of the future without taking account of cyberspace as well as physical space.

How might we think about this? Well, let's go back to the idea of networks and what they do. Essentially, networks of various kinds that tie cities together support interactions among activities, and they provide the "glue" that hold systems of production and exchange, cities, and regions together.

They have a couple of different functions. Transportation and telecommunications technologies support dispersed interaction among activities. Storage technologies of various kinds, beginning with writing, support asynchronous interaction.

Means you don't have to be in the same time and the same place in order to interact. You can leave a message somewhere, and then somebody can pick it up later. And the great thing about packet switching, the internet, and the World Wide Web is they really combine the two, so we have the possibility of interactions both dispersed and asynchronous.

If you lay out the sorts of possibilities that result from this, you find that something very interesting emerges-- a whole spectrum of different sorts of possibilities begin to develop. I apologize for showing you a rather boring table, but this is useful as a way of beginning to understand what's going on.

If you think about firstly the possibilities for dispersal of interactions, you can have local face-to-face interaction, so the traditional kind of urban pattern, the idea that you do things face-to-face. It's the idea behind the Greek Agora, behind the traditional village, behind the traditional nine to five workplace, and so on.

If you introduce some simple forms all communication at a distance, things like church bells, minarets, smoke signals, semaphores, and so on, you can begin to get partially dispersed interaction. Add a capacity to harness the electromagnetic spectrum for communication, and you get all of the possibilities that we've seen since the telegraph-- the telephone, radio, television, and so on. So there are whole bunch of possibilities for synchronous interaction that begin to develop.

Now, if we start to introduce storage technology, if you think about ephemeral storage like Post-it notes, which don't last very long, local semi synchronous interaction-- you leave a Post-it note somewhere and you get a message communicated that way at a particular location. You can begin to disperse that sort of interaction by introducing some simple transportation technology, so things like pedestrian and bicycle messenger systems disperse things a bit further.

And then if you introduce large scale, high speed transportation systems, you get the whole business of mail systems, these sorts of things. And then with telecommunications, voicemail and email become techniques for semi-synchronous dispersed interaction.

And then finally, if you really make large scale use of storage technology, if you keep things around for a long time, for example, you get the idea of the non-circulating library, the old fashioned manuscript library, for example. Or old-fashioned database systems, where you actually had to go to the computer to get access.

Disperse that a bit further through computer networking, you get things like intranets. And disperse a over wider area, we get the whole structure of the World Wide Web, supporting Dot-com enterprises and so on.

So what's going on these days is that our cities, our campuses, our buildings are held together by a combination of all of these sorts of things, but of course, we're seeing a significant shift towards that bottom left-hand corner. There's a kind of shift across the diagonal going on.

Now, the key question for architects and urban design generally is, how do you find some combination of land use pattern, transportation, and telecommunication capabilities that effectively supports the life of a community? And as technologies change, you end up with different sorts of patterns.

Now, let me show you an interesting example of how a very unusual pattern developed out of a couple of new technologies. This is in the 1920s and the 1930s in the Outback of Australia where, firstly, the land use pattern was a pattern of a very geographically dispersed population. People scattered over thousands of square miles working in the cattle industry and the mining industry and very widely separated from each other.

There are a couple of interesting new technologies, however, at that point. There was a thing called the pedal-powered short-wave radio that enabled telecommunication very effectively over those distances. And air transportation using light airplanes was becoming increasingly feasible. So it was possible to put together some systems that held that community together in a very unusual new way.

Here's the technology. This is a pedal-powered generator of that era. I've actually pedaled one of these things. It's terrific. Keeps you very fit as you're communicating, of course. Here's somebody back in the '20s using one of these things. This is the internet of its era, if you like. This is like sitting at your personal computer and logging into AOL.

Except what's going on here is the pedal generator, an old valve short-wave radio, and this is so early here, he's actually using a Morse key. And this is hooked up to a piece of wire that's hooked up to a eucalyptus tree, and this will allow you to communicate.

And this supported a couple of very interesting systems. It supported a thing called the Royal Flying Doctor Service that enabled you remotely to summon medical advice, or if you're a bit more seriously ill, a medical practitioner would jump in a light plane from someplace like that Alice Springs or Broken Hill or Cloncurry and come out and get you, and take you back to where you could get more effective medical service.

So a combination of telecommunication and transportation-- two new technologies at that point-- two new kinds of telecommunication and transportation technologies providing a new way to deliver medical service. And it also provided a new way to deliver education.

We hear a lot of talk today about remote education and dispersed campuses, and so on. Well, this happened here back in the '20s. And there was a thing called the School of the Air that still operates. And you see the wonderful logo here with the transmission tower and the slogan, "Parted but united."

The kids would assemble in their school rooms every morning, but their school rooms were in their houses dispersed across the Outback. The teachers would be in places like Broken Hill, and everybody would communicate via the radio, and you got a dispersed classroom.

This is today. Actually, it still operates. This is the School of the Air in Broken Hill, except today the technology is a little different. You see a satellite dish there, and very sophisticated telecommunications is used these days.

Let me give you one more example of this kind of new way of holding a city together or a community together by combining technologies. This is something that you'll see in Singapore if you're around there these days-- a system for putting together a road transportation system with some new telecommunications capabilities in order to achieve efficiencies in this case.

This is exactly the opposite situation. This is where you have a very condensed population, very high density, a great demand on the roads. Important issue is making efficient use of available road capacity. So a system is being put together that, at one level, is a sophisticated kind of taxation, an infrastructure funding mechanism.

At another level, it's a demand management and resource allocation system that can give you considerable efficiencies. And it can be used in other ways too, which may not be so positive. A system like this can be used as a kind of surveillance system. It can also be used as a sort of insidious access control system that really starts to control the way people can move around the urban fabric.

And this is what it's like. On the dashboard of every automobile, there's a device like this, a little electronic box that has a debit card in it and a readout showing the amount of money in your debit card in Singapore dollars. And as you drive around, you see these electronic road pricing rates posted. And these can-- depending on the sophistication of the system, and I'm not quite sure how sophisticated this one is in operation, you can vary these rates according to time of day, according to congestion, according to the type of vehicle, and so on.

So what this is showing, it's going to cost you more to drive a large vehicle down this street than it's going to cost you drive a small vehicle down this street. As you drive around then, as you start to move down various streets, you see these kinds of devices that have these kind of booms across the roadway.

And as you drive under one of these things, it sucks the money out of your debit card. And so you may or may not think this is a good thing, but you get the idea of how this is a very, very different way of organizing the flow of traffic. And Yossi Sheffi, I'm sure, will talk about some of these sorts of things probably in some more detail.

And then you can put this together with other things. Many of you are familiar with these kinds of automobile navigation systems now that depend on a GPS satellite system, and keep you located in the city. Well, you can put these together with these road pricing systems, for example, and you can ask these things to find the cheapest route from one location to another, and so on.

Once again, I want to emphasize what's going on here is very sophisticated and complex interaction between a particular pattern of land occupation, pattern of land use, and some modes of transportation, and some telecommunication capabilities. And all of this is being put together in really a very complex system to support the interactions within a community.

Now, I want to take a few minutes to talk about what this really means to the future of cities as we start to introduce more and more sophisticated electronic capabilities, things like the World Wide Web and the internet, things like increasingly intelligent automobiles connected to large scale digital infrastructure systems-- all of the sorts of things that we hear so much about today.

Essentially, the effect of all of these things is to start to loosen spatial and temporal linkages among activities. There's less need for everything to be clustered together spatially because many things can be done at a distance. Not everything, but many things can be done at a distance. And there's less need to coordinate activities because of the capacity for asynchronous interaction.

And the consequences of this on cities are very complex. It turns out that when you loosen these linkages, some components of activities want to decentralize to achieve wider coverage to get to larger markets and so on. Other components want to centralize to achieve economies of scale, to get knowledge spillover effects and these sorts of things.

Yet other components of activities form desirable new alliances. For example, if you can do your work at a distance, maybe you move your residential location to a scenic location or a location that offers particular recreational advantages and so on.

Yet other components of activities begin to float freely so that they can respond flexibly to dynamic conditions. So if the cost of labor is important to an enterprise, for example, if it's very sensitive to that, maybe what happens is activities move around chasing the labor markets, move offshore to places where it's particularly attractive and so on.

And the consequence of this is something that I've called fragmentation and recombination of familiar building types and urban patterns. The ways that things have clustered together in the past begin to break up and new sorts of patterns begin to crystallize out. You can think of it as a bit like a chemical reaction in which on some bonds get broken, some new ones get formed, and a new kind of structure begins to develop that functions in a different kind of way.

And this is a consequence of some of these locational connections among activities being broken but others remaining. For example, you might go to Amazon.com and buy books at a distance. But if you want your hair cut, then you probably want to do that face to face. I can imagine Rod Brooks developing some technology that would enable that to happen at a distance, but it's not a common everyday thing at the moment anyway.

So let me just analyze a couple of these examples of fragmentation and recombination, and then I'm going to draw it all together and leave you with some questions. Here's a traditional building type, a much loved traditional building type, the traditional bookstore. And what this does is combine several functions at one side.

It combines an advertising function by virtue of the banner out front in the display window. It's a place for storage of books. It's a place for browsing by actually physically walking among the books. It's a point of transaction that happens right at the front counter there. And the back office work related to this enterprise literally takes place in the back office back behind all of the books.

And it's all wrapped up in a neat little architectural package. It's all in a box, if you like, because the traditional way of relating all of these activities is through adjacency and through the possibility of direct, face-to-face interaction.

Now, of course, what happens when something like Amazon comes along is that some of these functions are virtualized and decentralized. So the browsing and purchasing functions are virtualized and decentralized, and they fragment, they blow apart, and they recombine with domestic space-- the space for this recombines with domestic space, with workspace, with these kinds of locations. So that aspect of what was accommodated by the bookstore decentralizes.

At exactly the same time, however, there's a radical centralization of the distribution function. You end up with large, national distribution centers where millions and millions of books are kept, highly automated so you get economies of scale, very large so you can keep in stock larger numbers of titles than is possible in that little physical bookstore.

And these, of course, are at different locations. These are located at national airline hubs. And so a very different kind of spatial pattern develops in which one aspect of the activity is decentralized, another aspect is highly centralized, and the back office work, because of business-to-business e-commerce capabilities, is able to float freely to wherever appropriately skilled labor is available at an attractive price. So some of it centralizes, some of it decentralizes, and some of it mobilizes.

And there are direct and immediate consequences of this sort of thing for cities. Retail space, for example, has changed over the years, firstly, as a result of changes in transportation and the development of the suburban mall. More recently as a result of these new patterns of commerce.

And you see, for example, the kind of thing that I'm showing here. This happens to be in downtown Vancouver. And it's literally a white elephant. This is a large Eaton's department store, totally unoccupied and totally derelict as a consequence of the sorts of shifts in retail patterns that have taken place firstly through changes in transportation, secondly as a result of changes more recently in telecommunication.

Let me go very quickly through a couple more examples. You all know what's happened to banking, I think. It used to be that you had to go to a branch bank to transact your business. It was a place to synchronous, face-to-face interaction, literally across the counter with the teller. The bank building was a very important civic kind of structure, often celebrated as this one on Massachusetts Avenue.

The architecture carry some symbolic freight. It was a place-- the architecture, this solid, cut stone, classical architecture of course is intended to symbolize the permanence, the solidity, the power, the respectability of this institution in the community. So that's how that worked.

Then along came automated teller machines. And firstly, these were asynchronous. They allowed 24-hour-a-day, seven-day-a-week service. You didn't have to come face to face with a teller in banking hours. And we've got this fragmentation and recombination once again.

These sorts of things ended up in all kinds of locations-- in supermarkets, in airports, in student union buildings, in gambling casinos-- any place where people needed cash, in fact. So a classic example of this fragmentation and recombination process.

And then more recently, the development of electronic home banking systems has further virtualized and decentralized customer interactions. And this has had all sorts of consequences, in fact, revealed by the slide because this is a slide from a few months ago showing Bank Boston providing this kind of service, and now it's Fleet providing this kind of service.

The consequence of this has been direct and profound for cities. We've seen branch banks being shut down in the thousands all over the world. This is just a clipping from a little while ago from *London Guardian*, for example, about hundreds more Barclays branch banks being shut down.

And we've reached a kind of situation-- it's very interesting-- where a different combination of use of physical space, use of telecommunications, and use of transportation has developed. Some of the branch bank looks something like this. Banking organizations now have electronic fronts and bricks and mortar backs. The main street facade mostly is replaced by a screen logo on an ATM machine or a website.

Often very much on the downside, there's a loss of local, physical presence in communities. And you have very practical questions-- what do you do with these old buildings? And my observation is that Starbucks often takes over the sites. You've probably seen this in many locations. And then the back office facilities, once again, are off chasing inexpensive labor in some offshore location.

New living patterns. I don't have time to go through this in detail, but one of the things we're starting to see is a recombination of the home and the workplace as a consequence of the possibility of telecommuting. As a result of business-to-consumer e-commerce, the home becomes an intensified delivery site.

We're seeing in many contexts the development of 24-hour live/work neighborhoods where people don't commute away, but in fact live and work in the same small-scale neighborhood. Small-scale local environments with global connections is a very interesting pattern starting to develop, for example, in the area around South Park in San Francisco. You can see this developing in a very interesting sort of way.

We're also seeing the reinvention of some traditional building types, if you like, pre-industrial patterns becoming viable once again in the post-industrial year. For example, right throughout Southeast Asia, the traditional packing of the shophouse, which we here in Singapore, is a very, very common traditional type where you have workspace downstairs, living space upstairs.

And this is also common in Europe in a slightly different pattern. A very typical European street, this one in Geneva, shows workspace on the ground floor and then apartments on the upper floors. This sort of fine-grained combination of living and workspace is a traditional pattern.

In the industrial era, we saw something different. We saw the development of bedroom suburbs, industrial and commercial zones in different parts of the city, commuting back and forth among these things and so on. What we're seeing now is a reinvention of the shophouse among other things.

This again is in Palo Alto. Sorry for the slightly fuzzy photograph. But this shows reinvented, high-tech electronic shophouses. On the ground floor here-- and this happens to be on El Camino Real incidentally-- on the ground floor, there are workspace all occupied by little internet start-ups, and on the upper floors there are apartments where typically the folks who work these and start-ups live.

It's pretty disgusting actually. All full of futons and underwear on the floor and kids occupying the space 24 hours a day. But it works. It works. And it's very interesting to see this kind of reinvention going on.

Finally, what's happening to public space? Well, turns out that in the digital era, many traditional attractions that pull people to public space no longer work in the way they used to. We're seeing a shift to dispersed from more focused interaction in many contexts. We're seeing a shift of many activities from public space to private space.

And this has some interesting and difficult consequences. I'll show you one example of how this kind of thing works. Here's a very important public space in Hong Kong-- one of the racecourses in Hong Kong. You may know if you know Hong Kong that the racecourse is a very important place. People love to go to the races.

And it functions very much as a public space because many people show up there. It's a place to network. It's a place to do deals. It's a place to be known as a prominent member of the community and so on. And this is all held together by adjacencies.

Traditionally at a racecourse, you had to be in the stands to see the races run. The bookies were there on the rails to take your money. All of these sorts of things were both synchronous and tightly focused at a particular location.

This enabled, as many of you know, in this particular context, The Hong Kong Jockey Club to be a very rich and powerful organization. But now what's starting to happen is, for example, a development of these sorts of devices. This is an online betting device that enables you to plug in to a telephone outlet anywhere, place your bets any time up to the running of the race from any location.

You can listen to the race on the radio, or you can watch it on television, or you may not even bother. And then, of course, what happens is that immediately the race is run. Your account is electronically credited or debited. Inevitably, in my case, debited.

And then here's a clipping from the *South China Morning Post* just showing a slightly more up-to-date version of this kind of thing, and nicely packaged as you can see-- and I think this is rather good-- packaged in a little pouch with your cell phone and a place to keep your business cards. It's a rather elegant kind of summary of the way that things are moving.

Now, what does this mean socially? Well, what it means is that if this happens on a large scale, the old modes of social interaction structured by synchronous assembly in particular urban locations begins to break down. And this nice *New Yorker* cartoon begins to illustrate quite nicely this sort of mechanism.

If you can't read it, this is a typical *New Yorker* barroom cartoon. And the customer is saying to the bartender, "To avoid the tedium of this endless socializing, Eddie, I've decided beginning Monday, to obtain all my future booze from Amazon.com" And you can begin to understand the sort of social consequences that can flow from this.

So to summarize and wrap up, you can trace these sorts of processes of fragmentation and recombination of the urban fabric in any kind of field that you want to look at. You could look at it in retailing. You could look at it in the pattern of housing. You can look at it in the workplace. You can look at it in education. You can look at it in medical services.

Any kind of traditional building type or urban pattern, you can begin to see the changes that are taking place as a consequence of this massive shift of activity across the diagonal of this table from the traditional way of doing things, which was entirely face to face, synchronous, based on physical assembly, across to a context where not everything happens electronically, at a distance, asynchronously-- far from it-- but a much larger proportion happens that way, loosening the traditional bonds among activities and allowing new sorts of patterns to develop.

This affects markets. It affects organizations. It affects communities. We see this process of fragmentation and recombination of familiar building types and urban patterns, and this leaves us with some very interesting questions to face. There are some design questions, and this particularly affects architects and urban designers. How do you organize systems of physical and virtual places connected by transportation telecommunication links that operate in some of the new ways I've described?

There are some interesting empirical research questions. How do these systems actually work out in practice when you put them in place? What the consequences? How do people actually use these things? Does telecommuting, for example, reduce traffic, reduce number of trips, or does it restructure, or is the complementarity effect? Lots of interesting questions like this.

There are some real estate questions. What new opportunities start to be created? What opportunities traditionally have been there become history? And then the questions I want to leave you with are some very fundamental questions of what this means to social equity within our communities.

There are some short-term issues that follow from the impossibility of deploying new infrastructure everywhere at once. We have to dig a lot of ditches and erect a lot of towers and so on. So in the short-term, there's an inevitable digital divide. Some people are going to get new technology before other people. So there's a difficult, complex, very, very important question of managing the transition to new patterns and new technologies so that very difficult questions of inequity are not created on too large a scale.

There's a longer-term question that the new infrastructure very heavily mediated by digital communications is potentially ubiquitous, potentially provides a very efficient mechanism for distribution of access to economic opportunity and basic services. That doesn't happen automatically as a consequence of the availability of the technology. But with the right sorts of social policies, with the right sorts of design interventions, those are very attractive possibilities that can begin to be realized in the longer term.

And then in the very long term, we face an issue that almost certainly education, to a level that enables people to operate effectively with these new capabilities, education rather than physical access to the infrastructure and to the technology is probably the key issue. So it's a moment of dramatic and sweeping change that raises some very profound questions, and I encourage you to think about these questions and raise some questions for us to discuss later in the morning. Thank you.

And let me now introduce Yossi Sheffi. Yossi?

**SHEFFI:**

Good morning, everybody. I'll start right from the beginning. Two steps back. Good Morning, everybody. The crowd is alive. Okay. We'll talk a little bit about lots of stuff. And I'll try to talk fast to get us-- not back on time. This is hopeless. But get close to back on time. Because originally, I thought I had 40 minutes for the presentation. I came in at 30. Now I have less, but I'll try to make it to work.

The outline which we'll cover is as follows. Ah, it works. We'll talk about the internet and electronic commerce in general, some business models on the internet, business-to-consumer issues, e-retailing, transportation issues related to this, and intelligent transportation systems and congestion all in one. This is actually a long agenda. We'll try to go through relatively quickly.

This graph is the obligatory graph. Anybody who talks about the internet has to show something that goes like this. So this is mine. It can be anything. The number of servers. The business-to-business transaction, business-to-consumer transaction-- anything tied to the internet. I did it. We're okay.

In fact, because of this, this is tied many times to network phenomenon and the winner-take-all that happens sometimes in network type of businesses, of which many internet businesses are. And it's the reason to try to get there first. He who gets there first many times wins the battle.

But in general, we're facing with millions of servers and billions of devices and trillions of dollars in e-commerce and a world connected like never before. And the question is, what does this mean? We developed certain competencies now.

For example, we can manage large communities online. We can run huge auctions and match buyers and sellers in a very effective way. And we can communicate very, very effectively large amounts of data through the Web.

Let me very quickly go through three basic business models that sprung up mainly on the internet. These are not totally new business models, but on the internet, you can look at three types of businesses. One is when a buyer is trying to connect to several sellers. This is mainly a business model where it prevalent in business-to-business type transactions.

When one buyer-- for example, when I try to buy something from MIT, I try to buy a box number two pencils, I actually go into a website and hit a box of number two pencils. MIT ties to several providers on the Web, and somehow this box of pencils finds its way to my office. So this is a business model that in fact MIT Procurement is using.

The other one is a seller-side business model, when it's all the catalogs. And you go to LL Bean, J Crew, or industrial catalogs to buy what the seller puts online. The last model is the aggregation model, when you try to aggregate many buyers and sellers together. In the consumer-to-consumer space, the most famous example is, of course, eBay when it brings together to one digital place a lot of buyers and a lot of sellers at the same time the transact, to go through business.

I'm not going to spend too long on this. But just mention that these various business models are in fact tied to all kinds of auctions. Many times, seller-based business models are tied to seller-based auction of people auctioning off capacity or ability to do work.

On the buyer's side, there are many times you'll find so-called reverse auction or RFP, RFQ type bids when buyer put their business online and sellers are bidding against it. And you'll find also double-sided sellers in exchanges just like the New York Stock Exchange.

There are exchanges popping all over the world. The United States now, there are about 800 industrial exchanges operating where people are trading anything from transportation services to chemicals to plastics to paper to paper clips-- what have you.

Let's now go closer to home and look at the effect of e-commerce going from business to the consumer with an emphasis on transportation. The traditional business model before the Web was that stuff got from a factory to a distribution center to a store, and then we all came to the store and picked it up.

Of course, after the Web, it's done a little differently where it's just sitting there, not going anywhere, and the hope is that the little vans will come and drop some stuff right next to us as we do whatever it is that we do, which is kind of nirvana, right? We're just sitting there doing what we're doing, and the stuff comes and just appears.

This, in fact, is a scheme, is kind of a model of various delivery systems that are being used in practice. In practice, actually in some cases like Peapod, they actually do not bypass the store. The delivery is actually from the store. Or they used until lately.

Well, now they're not doing anything at all. But until a few months ago, they were collecting-- they started doing from warehouses. But before that, when you got an order from Peapod, they went to the local Stop and Shop here in Boston, collected your stuff, and brought it to your home.

Sometimes these outfits like PetSmart are tied to distribution centers of retailers, like in this case for Kmart, and they will use exactly this model. They will collect at the distribution center and deliver it to the home. Officedepot.com does not use its own warehouses.

Actually, does not use the same warehouses. The storage, this was a specialized warehouse. eToys is using other parties. Amazon.com is trying many cases to go directly from the supplier to the home. They found out actually it's not working very well, and they have to build their own distribution center.

But the challenges here-- there are several challenges in all these models of delivering. First the all, take Amazon for example. From '96 to '99, these sales went up 100 fold. The losses went up 120 fold. You don't need to be a graduate of MIT to lead the trend-- is not promising.

So the question is, how do you make money in this business? It's a serious question. And a lot of it is tied to the delivery. A lot of it is tied to how we can deliver efficiently. And when you look at the carriers, at UPS, FedEx, and the Postal Service who are doing the bulk of the e-commerce delivery, first of all, they are facing real challenge in the future. Right now, less than 10% of what they deal with is tied to e-commerce, and most of this, by the way, is business-to-business deliveries still, even on e-commerce front.

So if everybody would start really doing home delivery, they'll face real capacity problems. They just can't do it. Right now, FedEx and UPS have very similar cost structure. They make about-- the revenue is about \$29 on a package delivered to business. And they deliver on the average to business, about 25 times per hour. 25 deliveries per hour to business.

To home, the average package cost between \$10 and \$15, and they can go do only five to eight deliveries an hour. Without getting into the cost again, you start realizing there's a problem there. The home deliveries are not making money for this outfit. They are all still riding on the business delivery.

All of them have problems with perishables-- how to deliver perishables. It cannot go through their distribution center. All of them have problems with delivering heavy stuff. When you want to order your sofa online, it's hard to imagine the FedEx guy running with a sofa up the stairs. It just doesn't work, right?

And of course, there's a problem of congestions of all these vans running around the city and all these trucks running around the city and creating a hell of a lot of congestion, which you can see in downtown New York today.

So how do we deal with some of this? Well, there are some alternatives. With perishables, people are dealing with dedicated fleets like webvan is trying to do, put a system of warehouses and their own fleets to deliver perishables. NationStreet is a system for dealing with large packages.

They use commercial deliveries between urban centers, and then they tie to agents that deliver anyway large stuff for stores and for other outfits to deliver the sofa and refrigerator and whatever else they want to do. It's a company in fact right here in the Boston area. The challenge still is to provide good service, and to make it profitable, and to develop the systems to control this, which nobody has actually figured out yet.

The keys to efficient home delivery is first of all, efficient customer receiving. What kills home delivery is that fact that if you come to a business, you know that between 9:00 and 5:00, somebody will be there to take it. When you come to the home, it's a problem.

Because first of all, you spend some time on just getting to the door, and getting past the dog, and trying to get somebody to understand that are not out there with a gun trying to-- that you're just trying to deliver a package. Then half the time, that somebody's not even at home. And you have to do it a second time and a third time.

People are trying to go through it by putting all kind of boxes that you can deliver into little home warehouses, which has its own costs and its own set of liabilities. So it's a real problem. The customer receiving is one of the keys to efficient home delivery.

Some of the other keys are density. About two months ago, two young people from a well known consulting firm decided they wanted to start a system for home delivery. They had all kind of gizmos. The toilet paper would sit on something, some device-- when it goes to a certain level, it will transmit to the store, and transmit to them, and they'll take it and deliver a new batch of toilet paper.

And they did they a whole study of what kind of people would use it-- you know, market research. And they found out all the red-headed, left-handed women all over the place who may use it. And I tried to explain to them that their market segment should be Beacon Street, and their second market segment should be Marlborough Street, and then they go to Commowearth Avenue.

Because this business will rise and fall on the density of the delivery. If you can come to a high rise on Beacon Street and deliver 50 to 70 deliveries in one address, you are making money on this. If you have to go to Newton or to Western, and go half a mile between houses, nobody will pay you the money that it takes to actually do it.

Density. Second, in order to deliver a good level of service, you have to deliver frequently, which again raises the cost. You want to deliver in large sizes, and a lot of these specialized outfits are going to multi-product delivery. So people like Kozmo and people like Webvan, you can order grocery and videotape, and CD, and electronic stuff-- lots of other things. Then the key, of course, you need to have the right technology and to understand a little bit about logistics and optimization and how to do it right.

Let's now go to the congestion question, and in face, Professor Michtell did a great set up for me on this when we deal with the congestion question. With all of this, the question is, will the increased home deliveries be offset by a reduction in trips that we take. So maybe there's not going to be a big congestion question.

What about returns? How do we return the stuff that we get home? Still, in most cases, we have to do something, if at least to take a trip to FedEx or to the Postal Service or to a store or sometimes when we are able to return stuff.

The parcel services are talking about starting to deliver several times a day to get to the immediacy that people start to demand. Well, nobody really knows, but I can just tell you the truck manufacturers are very bullish and the stock is it all time high because they think that there will be a lot more trucks and pick up stuff running around the cities.

The question is also, will suppliers deliver direct? It may be that because the level of service is not that great, many suppliers are thinking about taking the delivery into their own hands, which is actually by and large for society a terrible idea, because you will have starting a lot of deliveries to your home by half empty trucks because nobody will have the density to do it effectively even in an urban area.

So how do we solve the congestion question general? Well, the first idea is eliminate trips. We hope that by telecommuting, by attending broadcasted events, by e-shopping, teleconferencing, maybe people will have to drive less. Of course, there's the opposite effect that in fact Professor Mitchell was talking about, and this is we will have an expanded set of opportunity.

We can now work from anywhere, so maybe we'll drive more, driving longer distances. Maybe we will wait always until the last minute and then somebody will have to speed with half empty truck to get us the stuff, or we will have to go with a trip that was not planned before, not part of a more elaborate trip that achieved several goals. Maybe we'll have to go just to the something right away. So it's not clear that telecommuting really works.

Solution number two is the internet on wheels. This is what the new devices that Ford now-- Ford and GM and Chrysler are now making. They're actually internet on wheels, right? They have four wheels, and a steering wheel and engine, but they're basically a browser on wheels.

And the trucks, by the way, are already marvels of technology. Most commercial trucks-- I'm talking about the over-the-road, heavy semi trailers have computers in the cab. All the bookkeeping is done electronically.

The tax-- they have to pay tax based on how many miles they drive in each state. This is now collected automatically with a global positioning system fed immediately into the truck and into the company. So there are lots of trucks already becoming center of economic activity. And this is an example of a truck driver with his computer.

In future cars, you of course have better operation, ties so AAA, the dealers. You've seen the new-- what is this? Mercedes ad-- the guy there doesn't want to look for direction, presses a button, and just immediately connected to ask for direction. Strange solution, by the way, because you can of course get the direction automatically with a navigation system.

But the idea is to reduce the [INAUDIBLE] of driving. Maybe we'll say, well, we can't eliminate driving, but we can make driving more fun. So from the car, we can do a lot of things. And in fact, Mercedes has a van now that is literally an office on wheels, outfitted like a full-fledged office with full-fledged communication system.

We still need to mitigate condition, right? We still need to do something even if we have office on wheels and people are still running around all over the place. So there's a lot of schemes that came out for traffic management.

In fact, in the central artery in the big, big project in Boston, we are going through-- if you go through the tunnel, the tunnel is totally outfitted with sensors and variable message signs that tell you if something happens in the tunnel and you should avoid the tunnel possibly. It doesn't really help once you're in the tunnel to know that. But maybe knowing why you're standing still for two hours is better than not knowing why you're standing still for two hours.

But there's a lot of money being poured into so-called ITS congestion-- intelligent transportation system-- that's what ITS stands for. I'm talking about many, many billions of dollars that are poured into research on this. On the combination of road sensors and variable messaging and dynamic routing.

This does not really happen today, but the vision is that the sensors will collect stuff from the road. They'll go to some massive central computers that will calculate for every car that will request it, what's the best way to go to the destination today?

There's something common to all of these devices and to all of these schemes. And what is common is they don't work. And they don't work for very fundamental reasons. Adding electronic capacity and adding physical capacity is the same thing. Adding capacity does not work, actually. Because of equilibrium effect. Because once you add more capacity, people will travel more.

There's a famous story of the Long Island Expressway congested on day one. Why did we build it? Because there was the whole effect of anticipation during the years of building and people moved and located in anticipation of building the highway. It doesn't matter if it's electronic capacity or physical capacity, the equilibrium effect will always happen.

So this actually doesn't quite work. And building-- and I bet, if anybody in this room who wants to take it, dollars to donuts, that the new Southeast Expressway underground will be as congested as the elevated one on day one. There is no issue that it will be as congested. Because all the traffic that now goes around the city was, aha, we have another lane. It's equilibrium effect. You can actually prove it mathematically, but it actually happens in reality.

So the question really-- what the solution is is actually-- there is a solution to this. And the solution is the reason that I'm never working on ITS is because for me to work on something that is so expensive and developing in this when I know the solution is there rather than there is kind of hard.

And the solution is [INAUDIBLE] pricing is what Singapore is trying to do, is price the use of the road. And there are several ways of doing congestion pricing. Really simple is to put electronic tolls, for example, around downtown. And now, tolls could be electronic, not only in Singapore, in the United States.

Even right now in the [INAUDIBLE] going to the airport, you see this-- what's it called? [INAUDIBLE]? Whatever this card that you put in. Same thing that they use in Singapore. It's been used in the United States now for 10 years in Dallas, Oklahoma, several other places. Just made it to Boston.

So if you want to change people's behavior, say, going into downtown, you can put some electronic tolls around downtown. And you can get a little more but still simple pricing. If you don't want people to come between 8:00 and 9:00-- there's a lot congestion-- make the price high between 8:00 and 9:00 so people will distribute themselves and come early or later. So you have the distribution of the traffic.

If you really were to use the capability, you can price every piece of real estate. And as the car drives over it, it gets charged just like in Singapore. They do it with the overhangs. But it is the idea. Long term, I think the idea the-- I mean, it's a depressing idea, but I think it's inevitable-- that the idea of driving into the central business district of a large urban area and being expected to pay only for parking or for gas is just not going to happen like this anymore. We will have to start paying for the use of the real estate, the pavement that we are driving on at one point or another.

And this payment does not have to be fixed, of course. It can change by the time of day. Go down Government Center at 2:00 AM in the morning, nobody's charging anything. Go there at 8:00 AM, you pay \$5 every so often.

You can also do it if you really want by dynamic conditions. Because if you really have sensors that do the charging and collect data, then the dreams of economies is coming through. You can charge by externalities. You can charge based on the cost that you impose on others.

So when you're driving downtown, your car increases the congestion on everybody else. You're charged according to how much you impose on the rest of society. And this can be done on a rainy day, you may be imposing more than on a good weather day. When you go to certain-- you want to go through Fenway during a Red Sox, game you get charged more just going through this because you impose congestion on others.

The really Star Wars solution is to do real-time auctions. The real Star Wars solution is to come to an intersection and have my car bit with your car about the right to use the intersection. And in fact, if I'm the third in line, I can give some of my money to the cars in front to make them move faster, right? So you get into a whole game theory aspect.

And it doesn't have to be money. And I always say that in Cambridge [INAUDIBLE] people say, but it's not equitable. The rich will be better off. Well, to answer this, first of all, the rich are better off anyway. And second answer, it doesn't have to be money. The government can, of course, distribute its own units of accessibility. I call them "environs" or "accessirons" or something. And these will be the units that will be traded rather than money. So me and Bill Gates will have a fair shot at the intersection.

If you think about it, it obviates traffic signals and all messaging, and it let's emergency vehicles, of course, go through automatically. Set all the traffic lights. But it actually raises some bigger questions around this. And the bigger question about the use of auctions in general are things like-- you can use auctions for a lot of things, for purchase anything, of course.

But beyond this, to set markets that right now are not set in the most efficient way, and it does not have to be money. It can be any unit that can be tradable. And so you can think about organ transplants, where every citizen has the same number of units to begin with. I'm not talking about the organs. I'm talking about the same currency to begin with, with which to trade. The whole issues like school choice, how to choose, which school that's dividing communities can be done on an auction basis.

You can start thinking about the end of representative democracy, about direct democracy. When people have access to and are able to vote, in fact, on any question. And the technology's not far fetched. I mean, this is something that we are a lot closer to that people realize if we want to go this way-- if we want to go this way.

So let me leave you here. And the I think I'm not going to do the questions. This is usually when I go like this. Let me finish here and I'll see you later.

**MITCHELL:** Well, against all of the odds, we are not too far behind time. So it is now the break. We were supposed to break at 10:10 and reassemble at 10:30. Why don't we give you a little bit more time, so we'll reassemble at 10:35. If you could be back here at 10:35. And we'll take you further to the outer reaches of the technology and what it means.

Let me make a couple of announcements as we get started. Firstly, about submitting questions. Now, the mechanism we should have for submitting questions, of course, should be something along the lines of you just fire the questions wirelessly to me from your palm pilots or something like that. But we don't have that kind of technology.

What we have are students with appropriate T-shirts walking around with little index cards. And if you ask one of these students-- put up your hand and ask one of these students, you can get an index card and write your question on the index card and hand it back to the student, and it will find its way eventually to me. And we'll use those questions in the question and answer session later on. So I guess I should apologize for that very low technology, but it will work, and that's always very important.

Secondly, a couple of people asked me during the break whether the slides and the videos from these presentations would be available. And I checked with the Alumni Association. The answer is yes. You should watch out on the website for the Alumni Association, and you'll find information about obtaining video, slide, whatever records of these presentations you need. So watch out on the website for that.

So let me now introduce the second session, which will take us further to the outer reaches of the technology that's going to affect our future. And our first speaker is Rodney Brooks, as I mentioned before. I provoked Rodney a little bit by talking about not being able to get your hair cut over the internet, and he assures me in the break that that's not too far away and many other things that are going to be very important to us. So let me introduce Rodney.

**BROOKS:**

Thanks, Bill. As I listened to the first two speakers, I realized that if you're a pessimist, you could interpret Bill's talk to be saying that our cities are going to disintegrate. And you get to interpret Yossi's talk to say that out transportation is going to disintegrate. And I realize that when I'm here to tell you is that your self identity is going to disintegrate. But that's only if you're a pessimist.

We, as humankind, have been using machines for the last few thousand years, and the types of our machines have profoundly affected human life, as was shown in the previous two talks. And we've tried to maximize the autonomy of our machines and minimize the human effort and we need to regulate them.

It's much easier to drive an automobile today than it was when they were first built just over 100 years ago. And with computation, it gives new levels of autonomy for our machines. So things that used to require that people or animals in the loop will no longer require us to be in the loop so much.

And with trying to build machines and make them more autonomous, there's also being a fascination with building artificial creatures. On the left here-- well, these creations were by Vaucanson in France in the early 18th century, around 1738 or so. This is a mechanical duck which operated like a duck. It even had a digestive system where it could eat and provide some excrement driven by a cam system.

This is an automaton that could write something, and you could actually program it to write about 40 different characters in a string in the back of it, and then it would write letters. So people have been trying to build things look lifelike.

This century, there's been more capability for that. This is a robot that was built in the late 1940s by Grey Walter in England using two vacuum tubes. It was a tortoise. There was a series of articles in *Scientific American* about this around 1949, 1950.

And more recently with digital electronics, we've been able to build more capable sorts of artificial creatures. This was one built in our lab that used to wander around looking for soda cans and collecting them. This is one that used to give tours of the AI lab. This is a prototype for a Mars explorer. And the ideas from here eventually did go to Mars. And this is a more recent European robot.

In the last few years-- three, four, five years-- people have been trying to build robots with human form. This is one by Honda Motor Corporation. It's not so well known in the US, but in Japan, all of Honda's advertisements on TV feature this robot, which they now have a whole fleet of humanoid robots, which have [INAUDIBLE] operated control, but have human form.

And then these are a couple of robots that we've built here at MIT. This is a robot Cog and this is a robot called Kismet along with its creator, Cynthia Brazil, who just received her PhD for this work yesterday. And I want to show you a little about Kismet and show you some of Cynthia's work with Kismet to give you an idea of where we are currently in building artificial creatures, creatures which are lifelike.

Kismet a doesn't have a body. It has a neck and some shoulder-type motion. It has a face which can give facial expressions. And it has sensors inside its face-- a couple Foveal cameras here and then some wider angle cameras hidden in the nose, a mouth, lips, et cetera. And it can speak and listen to you as it interacts with you.

And our work is based on four principles. We think it's important to embody robots to true social interaction, but we also follow development patents that human children go through, and we integrate lots of sensors together. And what we're trying to do with these robots is get beyond the sort of thing that we have on the Web, which is a very impersonal sort of interaction.

And instead, when I come up to someone here, he looks up at me, he acknowledges that he's seeing me by nodding his head, he even gestured with his hands, and that's the sort of natural interaction that we have evolved over millions of years. And we're trying to understand that interaction by building it into robots and seeing what sort of interactions we can have with robots and what new ways we can interact with technology through that in this natural sort of way. And I see he's nodding there. He's following what I'm saying.

And this is completely unconscious in us. We do it all the time. And it's only when the person we're interacting with sort of breaks the rules about how it works that we notice that they're not interacting in a normal sort of human way. So can we put that on board robots?

So over the last few years, building intelligent robots has sort of agreed-- everyone sort of agreed upon having three levels of primitive actions, skills, and then behaviors. And what Cynthia's been doing has been adding a social level, which requires a fairly complex architecture, and I'm not going to go into any of the details of that.

But what she's been doing is getting the robot to be tuned to the human and the human to be tuned to the robot. And humans are very adaptable, and they can adapt to slight differences in the way the robot acts, the way people act, and tune themselves into interacting with it. And I want to show you some examples of that in the second.

And the competencies that you should be looking for in the videos I'm about to show you are being able to see or direct the visual attention of the robot, have the robot recognize socially communicated reinforcement like the gentleman in the front row who just nodded as I said something. He reinforced that he was understanding what I was saying. And that regulated our conversation.

If he'd given me a puzzled look, I would have put in an extra sentence or two to explain what I meant. The robot is going to communicate its internal state to the human. So it's going to be symmetric. And then they'll have this regulation of social interaction.

So first on the left here, you'll see just the robot looking around for its toy. And you'll be able to see from the way it's moving what it's doing. It's searching clearly here as you watch its eyes. And then it will find the toy. It gives a sort of social signal that it's seen the toy. It gets slightly happier. And we can interpret what's going on. But at the same time, it's observing the world.

So now someone comes in from the right here, and watch the robot as the person comes in. You can see that the robot was watching the person. So as an external observer here, we can understand what the robot is doing, what is on its mind, if you like. Now, that's not always the case with our laptops or our web interfaces. We often can't tell what's happening.

On the right here, we'll see an example of social amplification. These are naive subjects who did not know the robot previously who are sitting in front of the robot.

**SUBJECT:** Am I a little too close to you? I could stand back.

**KISMET:** [INAUDIBLE].

**BROOKS:** So in both cases, the robot signaled when the person was too close. Well, now this kid's experimenting. How close can he get?

**KISMET:** [INAUDIBLE].

**BROOKS:** So in both cases there, the robot gave a social cue that the person was too close to it. Why did it care? Well, it cared because if the person is too close, the baseline of its stereo cameras is too large relative to something that close, and it can't see the face clearly, and it can't tell what's going on.

So it gives a social cue that the person's too close as you would if I came up and stuck my nose two inches in front of your face. And the person reacted naturally to that because the robot had got the right cue that this naive person who didn't know anything about the robot a priori was able to interpret in the natural sort of way.

Now, the robot has an emotional space. It happens to be a three-dimensional space. Three axes here-- valence, arousal, and stance. And at any point in time, the robot is somewhere in this emotional space, and it will be outputting through its facial expression an indication of where it is in that emotional space. But it will also interpret actions from the person in a different way depending on where it is in the space.

So it has emotions of a sort, and it reacts to the world differently based on its emotional state. And it displays its emotional state the world. Let me just show you it displaying its emotional state coupled with speech here.

**KISMET:** Do you really think so? Do you really think so? Do you really think so?

**BROOKS:** So this is saying the same thing at different points in its emotional space.

**KISMET:** Do you really think so? Do you really think so? Do you really think so?

**BROOKS:** At this point in time when these videos were taken, the robot had no understanding of the content its speech there. It was just expressing a string of phonemes. And when it interacts with the person or the person talks to it in the next video, it doesn't understand the content of what the person is saying. But it does produce effect, and it understands the affect of the person speaking.

So in this next video, we've got some relatively naive subjects who were told to praise the robot, prohibit the robot, and it was up to them to see how they could get the robot to understand those sorts of instructions. So here is first praise in different languages. Doesn't matter about the actual words.

**SUBJECT:** [SPEAKING FRENCH]

**SUBJECT:** Nice robot. You're such a cute little robot. [INAUDIBLE].

**SUBJECT:** Good job, Kismet. Very good, Kismet.

**SUBJECT:** Look at my smile.

**BROOKS:** Now, notice how she imitates the robot back and forth here.

**SUBJECT:** Aw, that was cute.

**BROOKS:** Here they were told to get its attention and tell us when they got its attention.

**SUBJECT:** Kismet? Hey, Kismet.

**SUBJECT:** [SPEAKING FRENCH]

**BROOKS:** She knows it got it.

**SUBJECT:** Kismet, do you like the toy?

**BROOKS:** Now, here's someone told to prohibit it. And she's really tough on this robot. She really pushes it into a corner.

**SUBJECT:** No. No. You're not do that. No. It's not appropriate. No. No.

**SUBJECT:** [NON-ENGLISH]

**BROOKS:** Now, watch the person here.

**SUBJECT:** Where'd you put your body?

**BROOKS:** And notice how she imitates the robot there. And lastly, some soothing.

**SUBJECT:** [SPEAKING FRENCH]

**BROOKS:** So the important thing here is that the people are able--

**SUBJECT:** Aw, you're better than Jibo.

**BROOKS:** You're better than Jibo, she just said. The important thing here is that the people are able to interpret how the robot is moving around in its emotional space.

Now, when you couple that together with turn taking on the robot's-- as a behavior, the robot's just babbling, but it can engage. This is actually Cynthia and a graduate student who built the robot, so these are not naive. The one's will be naive.

**SUBJECT:** [INAUDIBLE].

**KISMET:** [INAUDIBLE].

**BROOKS:** So the robot and the people are turn taking.

**KISMET:** [INAUDIBLE].

**SUBJECT:** Really?

**BROOKS:** Notice the way the robot is making and breaking eye contact with the people.

**SUBJECT:** [INAUDIBLE].

**KISMET:** [INAUDIBLE].

**SUBJECT:** [INAUDIBLE].

**KISMET:** [INAUDIBLE].

**BROOKS:** Now, this next subject is naive.

**SUBJECT:** Hi.

**KISMET:** Hello.

**BROOKS:** She figures out turn taking, but she doesn't know what to say to it.

**SUBJECT:** Hello?

**KISMET:** [INAUDIBLE].

**BROOKS:** Notice she knows it's her turn, but what to say?

**SUBJECT:** Hello.

**KISMET:** [INAUDIBLE].

**BROOKS:** Now, here's another naive subject. He talks about 25 minutes. There's a cute cut here.

**KISMET:** [INAUDIBLE].

**SUBJECT:** Do you laugh at all? I laugh a lot. Carol says I laugh a lot.

**KISMET:** [INAUDIBLE].

**SUBJECT:** I try not to laugh at her. Okay. You're adorable. Who are you? What are you?

**KISMET:** [INAUDIBLE].

**BROOKS:** Watch it follow him and where's its eyes walk now.

**SUBJECT:** I want to show you something.

**KISMET:** [INAUDIBLE].

**SUBJECT:** This is a watch that my girlfriend gave me.

**KISMET:** [INAUDIBLE].

**SUBJECT:** Yeah, look it's got a little blue light on it too. I almost lost it this week.

**KISMET:** [INAUDIBLE].

**SUBJECT:** You know what it's like to lose something?

**KISMET:** [INAUDIBLE].

**SUBJECT:** You are amazing.

**KISMET:** [INAUDIBLE].

**BROOKS:** Notice him imitating the robot.

**SUBJECT:** I'm going to try closing my eyes.

**KISMET:** [INAUDIBLE].

**SUBJECT:** Let me try opening my eyes.

**KISMET:** [INAUDIBLE].

**SUBJECT:** Am I too close?

**KISMET:** [INAUDIBLE].

**SUBJECT:** No, stop! Oh, no, no. I got to talk now. No, no, stop! Listen to me. Listen to me. I think we have something going on. Oh, I think there's something here between us.

**KISMET:** [INAUDIBLE].

**SUBJECT:** Stop! You got to let me talk. Kismet, I think we got something going on here.

**KISMET:** We're in love.

**SUBJECT:** You and me. You're amazing. What are you?

**BROOKS:** I think we'll stop there. So here we've got the machine able to elicit from a person the responses they normally give to other people. It's not the responses a person normally gives to a machine.

If that Paperclip on Microsoft Word comes up and asks you some dumb question, you just click it out of existence. You don't care about that paper clip. But here with this physical presence and the ability to have these social interaction cues that we're all used to, it elicits from people the same sorts of cues back.

And we've seen this sort of thing happening in toys. The Tamagotchi-- the Furby did this a little bit. And a toy that I've been involved with, which will be coming out this year from Hasbro at Christmas, is a humanoid doll with facial expressions, but not quite as much interaction as you just saw then.

So we're starting to see in our lives machines which have some of these properties being mass produced. And in the labs, as I just showed you, we've gone a lot further. Where does this ultimately lead? Well, as these machines not only have emotional content, but also have intellectual content coupled with that so they're not just babbling but are saying intelligent things, will we ultimately accept these machines into our lives?

And this has been a theme in science fiction over the last few years. This is from just last year, *The Millennium Man*, where this robot spent 200 years trying to be classified as equivalent to a human. We're just starting to put the pieces together. We're starting to get machines able to interact with us this way. If the progression continues to happen, will we ultimately be faced with these questions of whether to accept machines as ultimately our equals?

What will their ultimate nature be? Will they just be hunks of junk designed by someone else? This was a phrase that was used by John Searle, a philosophy professor at Berkeley that's described Deep Blue when it beat the world chess champion. Ah, it's just a hunk of junk.

Or will they, in principle, be fully fledged animats in the same way living creatures are animats that have genuine emotions, visceral emotions. Will we treat them as beings rather than things? Will we empathize with them? And as my former student, Cynthia Breazeal says, will we treat them as an appliance or as a friend?

So that's a pretty difficult question, but let's ask ourselves whether we are willing to attribute real emotions to machines? Can a robot be afraid? I think most AI researchers these days would say that robots or programs can reason about facts. They can make decisions. They can have goals.

But I think most researchers in AI, they're willing to say robots can act as if they're afraid. They can seem to be afraid. They can simulate having fear. But I'm not sure most people today would be willing to say that robots could really have fear. They just simulate fear.

I talked about the robot Kismet having emotions and being a point in that three dimensional emotional space. Is it the same as having emotions, or is that just a simulation or a model of emotions? I don't think most people today are willing to say that robots are viscerally afraid.

But is that just because of our technological capabilities right now, or is that an absolute that we think machines can never be viscerally afraid-- can never have real emotions? Emotions are reserved for us, us humans. Machines-- they're just cold, hard, calculating engines.

Well, when faced with those sorts of arguments, I reflect on mankind's retreat from specialness over the last few centuries. With Galileo and others, we had to admit that Earth was no longer the center of the universe. Earth wasn't the special location in the universe. It was just a little outpost in an unremarkable galaxy. It wasn't the center.

With Darwin, we had to admit that humans and animals have common ancestors. And that was in the 1860s, and still in many places in the US today, that is not commonly accepted. Over the last few years, we've understood that the mechanisms of DNA-- we see that we're very similar to yeast. We share a lot of genes with the yeast.

With computation, we have come to see human thought as something that fits on machines, and we've had to give up that we're better chess players than the machines. Biochemistry shows that we are collections of tiny machines, little molecules doing their thing together and interacting, and that makes us.

And over the last two or three years, we've seen that human flesh and body plans are subject to technological manipulation because we've seen that happen with Dolly the sheep and in mutations in other animals where we can control how the body plans form. The same could be done in principle to humans-- is not done on the moral basis.

So over time, we've had to retreat from specialness to be pretty much just machines. And in my view, the way we've put the wall around it, we say, well, robots can't really feel emotions. It's only humans. That's what makes us special. But perhaps as we build these robots, and they get better and better at this, we may have to make this jump and retreat from specialness even more.

And it's important to note that the emotional models are only one component of us getting to treat robots as equals to us. The emotional content is often in the eye of the observer. The level of engagement is often in the eye of the observer.

And we saw that in those videotapes with the person interacting with the robot there. The woman who was saying just, "Hello, hello" wasn't really interacting, but the guy who was following her really was interacting. And the observer is a component of the dynamics of the behavior of the robot.

So as our systems become more complex and the engagement for more people is longer term, the illusion that they're not artificial will be shattered less often. Now, is that enough to make them equal to us?

And I like this quote from Sherry Turkle when she came to my lab a few years ago. And this is from her book *Life on the Screen*. She came into my lab and Cog noticed her after it entered the room. And she found herself competing with another visitor for its attention. She felt sure that Cog's eyes had caught her own.

And this left her shaken, not because Cog was that fantastic, but for years, she'd sort of been brushing off my robotic creatures as not being real creatures. But she had found herself-- get rid of those quotation marks for little bit-- only for a few minutes. Or probably in her case, it was only about 30 seconds. So she's continually skeptical about the research project, but she had behaved for an instant of time as in the presence of another being.

When we put naive subjects in front of Kismet, some of them will sit there for 15 or 20 or 25 minutes and interact as though it's another being. So we've gone from almost nothing to 25 minutes in the last seven years. Maybe we'll never get beyond 25 minutes.

But as our technology increases, maybe we'll be able to get those interactions to go on for hours and days and weeks. And then what is the status of those machines? Are they just machines, or are we continuing to treat them in human like ways?

So when I talk about this, and especially when I talk about it to reporters, reporters always say, but will these machines, when they get really smart, will they want to take over from us? And this has been the subject of a lot of science fiction. Will the robots want to take over from us?

And I used to sort of brush that off and say, oh, we don't need to worry about that. That's far in the future. But recently, I realized that the answer is no. The robots will never take over from us. And it was sort of driven home to me about six months ago when I was waiting for an elevator in the AI Lab, and the doors opened and out walked Hugh Herr, who's a researcher in our lab-- also has an appointment at Harvard Medical School.

And from here up, he was pretty much human. From here down, he was totally robot. He happens to be a double amputee, and in our Leg Lab at the Artificial Intelligence Lab, we've been building prosthetic legs, which are intelligent robots attached to people.

So he was half robot, half person. And it wasn't like it was in the lab. It wasn't like it was a nicely dressed up leg. There were computers hanging on the sides, cable harnesses hanging out. And this person is walking along, and he's clearly robot down here and clearly person up there.

Now, that got me thinking a little more that ultimately-- there's Hugh with one of the legs on-- one of the prototypes. That got me thinking a little more that although in the last Millennium we came to rely on our machines, in the new Millennium, we're going to become our machines. And I'll show you how that's happening in just a second in a deeper way than just attaching artificial limbs.

So my answer to reporters now is, no, we don't need to fear the machines because we, the man-machines, will always be a step ahead of the machine-machines because we're going to be adopting the best technology into our bodies.

Why would we do this? Well, leg and arm prostheses are easy to understand. You want better. You don't want an artificial arm on an amputee that can't do much. You want capabilities with fingers, and so you want to build a really good robot like The Six Million Dollar Man.

But also now there are tens of thousands of people walking around with cochlear implants. These are people who've become deaf and now have implants in their ears, which take in the signal electronically and then directly connect to their nerve system at about six different frequency bands and give them enough hearing to understand speech-- not understand music, but to understand speech. And there are tens of thousands of people with those implanted-- direct electronic to neural connections. A lot of that work's being done at MIT, by the way.

Also, people have started to experiment with retinal implants. Professor John Wyatt here at MIT is one of those people-- number of people in other places-- where for people who have become blind through macular degeneration, the idea is to try and put essentially a video chip in the eye to replace the front end of the retina and then connect up to the neural system.

And they have done clinical trials where they have put chips in blind people's eyes. Not permanently-- only for about 24 hours at a time-- and had connections bonded to the person's visual system, and people have been able to not see, but detect light changes and those sorts of things.

So it's a long way off from giving sight, but hearing works-- hearing speech works. Retinal implants-- it looks like it will work in the next few years. So there's good clinical reasons to be connecting electronics to our neural systems.

Electrodes in the thalamus has now become a common treatment for Parkinson's disease. So good clinical reasons are driving our digital technology and our robotic technology to become interfaced to people.

There've been other experiments in animals with nerve fibers growing through chips. And we have a long-term project in the AI Lab with one of our faculty members, Steve Massaquoi, who's also an MD who wants to solve some of the tremor problems that Parkinson's disease patients have by doing connections from the cerebellum to the muscles directly and bypass some of the nerve fibers.

So getting digital machinery into the control system of a person to counter diseases. That's not there yet. But it's certainly things that people are writing research proposals about and getting research funding for.

And then other people are trying to put intracranial implants of neural circuits into patients with deficiencies in their brain to try and recover some lost functions. Again, this is very experimental, but you sort of see the trend from things that are connecting electronics and digital electronics to the nerve system into more and more complex things driven by clinical desires to help people who have medical problems.

In the United Kingdom, pets these days have locator chips implanted under their skin, and there's even actually a couple of academics who have made some names for themselves by doing this also in the United Kingdom. This is so you can locate and identify your pets. Maybe this will ultimately be used by some governments to label people.

But then from a more medical point of view, there's the idea of guardian angels where you have implanted in your body something that constantly measures your vital signs, your temperature, your blood sugar level, et cetera. Because it turns out, everyone has a unique signature of what is normal for them, and by monitoring that and then alerting via wireless networking some machine in your home when things are going out of parameters, that can tell you when your coming down with something.

So there's clinical reasons that you might want this guardian angel embedded in your body. And people are actively working on this, and so I think we can expect to see that happen in a few years.

Beyond that, if we get these intracranial implants and neural circuits working and can actually get some understanding across the digital neural divide, which is somewhat more general than where we're currently able to fully do it with cochlear implants and partially with retinal, I think there will be a real push to have connections to the Net as it becomes more and more pervasive.

I tell my kids, okay, so you want to rebel against me by putting a stub through your tongue. Well, your kids are going to rebel against you by getting a wireless internet connection directly in their head.

Now, in the Olympics we don't let athletes who have been taking steroids compete. So maybe when these first get implanted, we'll stop kids with them from taking their SATs or something like that. But it's no longer a good enough excuse that your eyesight is poor not to take the SATs. Get some glasses or go down to the mall and get some surgery on your eyes like the other kids are doing.

So I think before very long, if this technology works out, we'll make it compulsory. There'll be the iSATs or the eSATs or something where you have to have an internet connection in order to take them. So that digital technology is getting into our flesh for medical reasons. But there's another wave coming a slightly different direction.

We've had 50 years of molecular biology, which have produced analysis tools to let us know what's going on inside cells. Those analysis tools are now being turned around to be engineering tools. We've seen gross level genetic interventions in cloning and genetic therapies. Some of the work down at the AI Lab right now is we have digital control over living cells where we have software that compiles down to DNA which gets inserted into living cells, and then that hijacks the mechanisms of the living cells to do some computation and control at the molecular level of the proteins being formed by that cell, which then go across the cell boundary and communicate with other cells, which then receive messages.

And they have their own compiled piece of software in the DNA form inside them, and they do some other computations. So we're now able-- this is in E coli, not in mammalian cells-- but in E coli, we're able to control digitally the processes inside living cells. Hugh Herr, who I mentioned before, is also building cultured muscle cells as part of robots.

So biotechnology is now also getting us to the point where we're able to control and intervene in what was our pristine bodies before. So with all this, will we ultimately be immortal with all this technology? I think more and more of our bodies will be able to be changed and repaired in better ways than we've currently been able to do.

We won't have to go in with a chisel and a hammer and bash away at a thigh bone when we want to insert a new hip joint as we currently do. It will get a little technologically more precise the way we do things.

We'll be able to change more part of our bodies. We'll be able to repair more parts of our bodies. And we might ultimately become something like a B52. You know, the B52s that are the flying around now have the same serial numbers as the ones that were flying 40 years ago, but the engines have been changed.

The flaps have been changed. All the pieces get changed out from underneath them, they're still the same B52. In the next 100 years, I don't think we're going to be able to replace everything in the body, but lots of pieces will be replaceable.

So in the foreseeable future, and foreseeable, I think, is the 50-- well, is 100 years foreseeable? In the foreseeable future, we'll not become immortal. I don't think we're going to be able to download ourselves into digital domain and completely get rid of our bodies. I think we're going to have these bodies which have components replaced.

But I think we will live longer and longer. I think that we may become more intellectually powerful than in the past, and we have seen that even with the tools that are slightly outside of ourselves becoming more intellectually powerful over the history of time. And there's a curve that goes like that on this intellectual power.

But as we get into these changes, we may not be the same species anymore. As we have more and more genetic interventions and as we have more and more of these intellectual interventions because our intellectual societal interactions are part of what makes the same species, I think we will become something else.

And we will be as alien in 100 years from now and 200 years from now to today's people as us, the year 2000 people, would be with all our paraphernalia to someone from 1,000 years or 2,000 years ago. So our species is going to change. And the robots will not takeover, but we will become partially robots. I think I'll hand over to the next speaker. Thank you.

**MITCHELL:** As Rodney was talking about these strange new creatures with a certain amount of emotional intelligence, difficult to understand, maybe want to take over from us, seem a little bit alien, I was reminded of the way I felt when I first realized I was the parent of a teenager.

With that, let me introduce some Rosalind Picard, one of our very interesting and exciting young faculty members from the Media Lab who will speak about the emotionally smart machine. Rosalind?

**PICARD:** Good morning. I'm waiting for my computer to wake up. No reflection on Rod's great talk. Actually, I tend to often get plagued by technological problems these days with talks, which actually aids the subject matter of my talk, which is about how people deal with frustration and technology. Except this is not one of our experiments right now.

Rod talked about the Kismet project which-- I think it's really time for us to loosen up here. Where you saw Kismet really having a lot of fun with people, he also mentioned the Paperclip, which I think is well known by lot of you, and how we don't think of that as a social because it's not embodied. But I'm going to tell you about some surprises with people's behavior that may make you think differently about even the technology in your laptop, palmtop, and soon-to-be your jacket and shoes and so forth.

First, just briefly, a little bit more about emotion. Sometimes people, when they hear us talk about emotion, they think, what is it? If those AI guys just decided that they couldn't make any more progress in the really important areas-- intelligence and reasoning and so forth. Why are they interested in emotion? Well, it turns out there've been a lot of really surprising findings lately about the role of emotion in intelligence. Not just best sellers like Dan Goldman's book *Emotional Intelligence*, but also top sellers like Damasio's book *Descartes' Error* by a-- he's a neuroscientist.

And the findings suggest that even when you are engaged in the most rational and intelligent of tasks like perceiving information or making decisions in a very reasonable way, that the other parts of the brain, such as the limbic system, these lower level structures, are actually engaged in this as well.

So what we have thought to be the higher level rational functions in the cortex, the visual cortex, the auditory cortex and so forth, actually involve signaling that happens in the limbic system even before we engage the cortex. This will be familiar to you if you think of the example of fear.

You're hiking through the woods, and all of a sudden, you jump out of the way only to realize moments later that that was a stick and not a snake. Your limbic system recognized the snake, got your body to respond before the visual cortex kicked in and said, oh, it's just a stick.

Now, we all know, of course, that people who are emotional in the sense of some outburst of emotion are not being very reasonable. What's surprising though is it's also been found that in people with a certain kind of brain damage, if you actually have too little emotion, it's also the case that you don't act in such a reasonable way.

This is where Data has led a lot of people astray-- the *Star Trek* character who has the emotion chip, the android. And when Data is reading poetry or wanting to learn about romance, he clicks on his emotion chip and can engage in those things. But when he's trying to be the rational scientist, he clicks off his emotion chip.

Well, that's really misleading. If you had an emotion chip, if we replaced that part of your body, and you turned it off, and if it was functioning the way that emotions function in a human brain now, what we would find is you actually would cease to become rational when it was turned off, especially when it comes to social and personally significant interactions. So what we're finding is that emotion isn't just important for being emotional in poetry and social situations, but also for rational and even scientific reasoning.

Now, what may surprise you is that you're not just social when you're confronted with a character that has a human body and face, but when you're confronted even with something like a TV set or a regular non-humanoid-like computer, we default to a lot of social forms of interaction.

Here, the work of Reeves and Nass at that other institution on the West Coast, Stanford, is very interesting. And this is a picture of a man eating dinner with his TV set. What they have shown-- and initially we thought this was just Stanford computer science students this worked on-- what they've shown, for example, is that when you interact with a computer, let's say a computer gives you a presentation, the computer has no face or voice or body, it just simply presents you with some information.

And at the end says, please rate how this computer did. And let's say you really like it, so you give it a seven. Then you go over to another computer, and it says, please rate how that computer did. Well, it turns out that Stanford computer science students are a little nicer to the computer face to face than they are behind its back. Behind its back, they give it a six.

Now, they deny doing this when it happens. And in fact, Reeves and Nass have run dozens of these kinds of experiments where they take a classical human-human interaction, take out the human, put in the computer, and ask if the results of the interaction still hold such as that you tend to be slightly nicer face to face. And the results still hold.

Now, we have found that this kind of thing even works on the East Coast with MIT and others. But this has very important implications for how we design not just robotic interactions, but all of the interactions and the technology that are starting to fill our cities and our automobiles and our natural spaces.

So it's important then to look very closely at the subtle cues that we read in human-human interaction because those appear to carry over to human-computer interaction. So let's say everybody's putting Microsoft Windows on your dashboard these days. All the car companies we're talking to talk about a frightening situation-- fear.

So suppose that your automatic navigation system is going to give you help. When should it do it? Well, we can look at things like when a person is offering you help. And suppose that they just accidentally don't do it at a good time. Well, what do you do? Well, your natural social cues kick in, and you will subtly send them a message that that was bad timing.

Well, if the person reads those affective cues of disliking and rejection, then they will note that that behavior they did was inappropriate and modify the interaction the next time, hopefully-- if they're smart, if they're emotionally savvy. What Reeves and Nass' theory predicts is that you can just take out the word "human" in a situation that works, put in the word "computer," and then the same prediction should hold.

If that navigation system or that machine in your kitchen kicks in, turns out people express their feelings at it. In fact, people are very expressive at these machines. I was surprised to read-- you guys may have heard about the guy in Colorado who got so frustrated with this machine he actually shot it three times through the monitor and two times through the hard drive and was taken away by the police.

A surprising statistic-- even kids, users under the age 25 in a very large survey done in the UK, 25% of them admitted to having physically kicked their machine out of frustration with it. So we express emotion, maybe not the best of ways, but we do it quite naturally to machines.

So the emotionally savvy machine should see this little bit of feedback that is positive or negative, how intense it is, and that should go into its learning system. It should recognize the state of what it's doing and use that to adjust its behavior.

Now, Rod showed a picture of Kismet's emotions lying in a space, the arousal, valence, stance space. The first two dimensions of that are two of the most common forms of describing emotion. By the way, emotion theorists don't really know what emotion is. They disagree on definitions of it. There's even a paper that goes through over 100 definitions of it. But nonetheless, we can begin to describe certain aspects of it and use those in designing technology.

The arousal space has to do-- you saw Kismet when he was tired at the bottom, excited at the top. When your about now, end of the morning, your blood sugar is low, you're probably all pretty low on this arousal space.

Valence is positive or negative. If you're liking something, you're over on the right. If you're disliking it, you're over on the left. You can imagine the tremendous applications of sensing just these two aspects of emotion in consumer feedback and product feedback and certainly in advertising and those areas.

One place where my student Matt Norwood built this in-- you may have heard of information appliances-- computing sneaking into everything you interact with. You can simply give valence feedback to our coffee machine in the Media Lab. If you liked the cappuccino it gave you-- good job, Mr. Java. Thumbs up. If you don't like it-- darn this machine. All I got was three cups of frothy milk when I asked for decaf. You can hit the thumbs down.

The machine keeps track of what it's doing and associates that with the kind of user feedback so that it can, in this case, let the service people know the points of greatest irritation. So that's all associated with interstate diagrams and so forth.

Similarly, we have worked somewhat with regular interfaces. You may have heard of IBM's emotion mouse that senses some physiological signals. What actually seems to be a little more appropriate to do from the mouse is to sense pressure changes. As you're moving things around, you naturally tend to pull towards things you like, push away from things you don't, bash on things.

We've characterized four basic clusters of sort of bashing clicking patterns among users who are frustrated to sort of gage the intensity of their interaction. So the thumbs, up thumbs down gives you a quick and dirty measure of valence. Something like the intensity of hitting something or the repeatedness of that helps to give a clue to the arousal level.

There are many other things we can measure. And in specific spaces, we're not just interested in arousal and valence, but in emotions like confusion, frustration, stress, interest, boredom. One wearable computing system that was developed by my graduate student, Jocelyn Scheirer, shown here.

She is not usually this upset. She's furrowing her brow trying to make a face of confusion. Here she's wearing the glasses she designed in stealth mode So that I can't see what she's expressing with her face, but the computer can.

And here's an early version of them with electromyograph sensors in the brow sensing corrugator and other muscle activity above the brow. I'll play a short video of these glasses being used to communicate the affect of expression directly to the computer and possibly through the computer to a lecturer.

This is a mock lecture situation. The lecture you will hear is deliberately confusing. But there was no rehearsal or anything for the students. They just plopped down and listened, Raul Fernandez on your right is just wearing the sensor, and Jocelyn is wearing glasses. And you'll see the bar graph in front of Jocelyn go up when she furrows her brow in confusion, and Raul's go up when he furrows his brow in confusion.

**PROFESSOR:** Every college graduate should have adequate writing skills. A writing exam should be given before students receive their college diploma. No matter what field one enters, unambiguous writing is essential to not avoid understanding.

Since misunderstanding cannot occur with unambiguous, the university misunderstands its purpose when it graduates students who do not understand the importance of not ambiguously communicating their understanding to others. In order to graduate those students who have unambiguous writing skills, they should have to take a test which would unambiguously assess the level of ambiguity present in their writing.

**PICARD:** Kind of reminds you of some lectures at MIT. Now, that was a very simple sensing coupled with some signal processing to clean up the signal, and with a little bit more signal processing, slightly different flavor of the sensor, we can discriminate upward and downward expressions-- upwards that tend to be indicative of interest or openness, downward that tend to be indicative of confusion and this sadness or disapproval.

What we can also do with computer vision is look at the entire system of facial actions that's happening and begin to recognize whole patterns of facial expressions. I chose not to show you that here because what people often forget is that there are a lot of applications where we've tried to bring our technology into the real world.

We've deployed it among real, truly frustrated users around MIT. And they often do not want a camera monitoring their behavior. In fact, none of them so far have checked box saying they would like a camera.

When however there's a sensor that only measures a tiny aspect of what they're doing, they are a little more comfortable with that initially. So we're recognizing that sometimes less is more when it comes to sensing. Maybe a camera seems less obtrusive, but ultimately psychologically, it can be even more obtrusive.

We've also made some very exciting progress in wearable systems, systems we've sewn into clothing, blood volume pressure [INAUDIBLE] respiration sensor sewn into a sports bra, skin connectivity sensors sewn into gloves and shoes that, sensing for physiological signals and a bunch of features of these and using new pattern recognition tools-- we've been developing new coupled with old-- we're up to a discrimination rate of greater than 81% on a group of eight emotions expressed by one person over many weeks of data, which is very exciting.

It says that the computer you wear could really get to know things about your body and get to know something I think that's particularly important. If this-- get to know your response to it. If your computer is just in your briefcase, and you only bring it out a few hours a day, it's okay if it kind of irritates because you could just put it back. If however, your computer is in every room of your house and in your clothing, it's really important that it not irritate you. So it's very important that it begin to read your responses to its behavior.

We're doing this in transportation systems. There's been a lot of work, of course, on pilots over the years. We're looking now at the average driver and stress that he or she goes under. We've run experiments with drivers in the Boston area. We don't have to have these million dollar simulators and put people under accidents and stuff like that. We just let them drive around Boston. And it's amazing the range of stress we get.

Given different driving conditions like resting in the garage, driving through the city of Cambridge, avoiding pedestrians, bicyclists, traffic lights, and so forth, driving through the traffic lights around here, highway driving, non-rush hour, we're up to 96% recognition accuracy discriminating those conditions and slightly lower than that when you're correlating the results with their self-report level of stress.

Self-report level of stress is ultimately very hard to pin down. People have different perceptions of their own emotions and different willingness in terms of reporting them as well. This is work with Jennifer Healey who just completed her doctoral thesis.

Now, perhaps the most common emotion we see around machines that we really want them to recognize is frustration. And here we see that it's not so easy to recognize. In the upper left, the user is simply tensing his facial muscles. In the upper right, she's shaking her hair. In the lower left, she's shaking her hands. And then in the center, Jonathan Klein is actually having physical contact with the machine.

We see many ways of expressing frustration, as alluded to earlier, from yelling at it to physically interacting with it. And what we would like to do is move beyond just recognizing it to actually responding to it in an appropriate way and parallel evolve both aspects of the system.

Now, I won't go into detail on this. The basic idea is to figure out what happens in a human-human interaction, take out the human, put in the computer. If it's a good thing to do in human-human interaction to acknowledge the frustration, and that's true only in certain situations, then we will try to, in those same situations, do that in the computer-human situation.

So we built a system that practices those skills. This is Jonathan Klein's thesis work. And what it does-- the system he built us what's called here an emotion-support agent. Basically responds with a little bit of active listening, empathy, and sympathy.

We compared it to two control conditions-- an agent that just ignored the user's emotions, which is what most systems do now, and an agent that basically ignored the emotions, but then at the end said, is there anything frustrating that happened, oh, and boy, did they report lots of things as frustrating? That was the vent condition. But the computer did not respond to what they reported. Just let them reported.

We frustrated two groups of users. They did not know they were going to be frustrated, of course. We debriefed them on that afterwards. If you tell a subject that we're going to try to frustrate you, they won't get frustrated. So we told them they were going to come in and test play this cool new computer graphics game.

So they came in and they played this game, and in the middle of the game, we had things go wrong. It would freeze up-- internet delays. And we thought, well, that's not going to really frustrate people. So what we have to do is get them personally invested in this.

We didn't want people who are just great game players. So what we did is got people with minimal game playing experience, but who we told this was a test of their intelligence. And since these were MIT students, they are really invested themselves in it.

So the clock's racing ahead. They're trying to solve these puzzles. And we have things freeze up on them. They got very frustrated down here. After interacting with one of these three conditions, 12, 11, 11 people-- each of these conditions-- they were invited to go back to the game that had caused their frustration. They had to play it for three more minutes. After three minutes, the quit button came on, and we measured their behavior after that three minutes. Did they hang in there and play with it, or did they get up and leave?

Now, the human-human situation prediction is that if this were human and a human frustrated you, and this were some other interaction, and it reduced your frustration, then when you had to go back to that original human and interact with them, if you were still very frustrated, you would minimize your interaction with them. If however, you were feeling better, you would be more inclined to hang out longer.

What did we find? We found a very significant difference. The people who had received the emotionally savvy interaction, which by the way, took a lot less time to do than it takes me to explain it-- it's very quick. Just a couple minutes. They stayed over here significantly longer than the controls. And that was true across the low frustration and high frustration conditions.

Now, this should raise a lot of questions in your mind. Not just issues Rod brought up about machines that might feel-- if a computer's expressing empathy to you, I mean, does it really feel? I've had one person say to me, gee, I'll accept it's sympathy for my pain once it can feel pain. I want that robot to feel my pain.

Well, we all know that it is possible for a husband to show appropriate empathy to his wife during labor even though presumably he will never feel the pain of childbirth and so forth. There are many cases where we can appropriately address somebody's frustration without actually feeling that pain. So I'm not sure that's a requirement.

I'm going to whip through these because there are a lot of them here. But there's the issue of-- Rod addressed the specialness, the thought of us becoming less important, the thought that if you could-- in the Lascaux cave paintings, when they duplicated them so that more people could see them, somehow the originals became less special because of this ability to duplicate it in some way.

So to the extent that we can duplicate something in machine, does that make it less special? Now, I actually differ with Rod on the extent to which we're duplicating these things in machines. But it's a matter of extent.

So there are several things here. One factor that's particularly disturbing and was brought to my attention by a very large computer maker whom you all know-- some of their executives, former researchers, were Sloan Fellows visiting. And when they heard about our work with the emotion support agent, they said, gee, we're not surprised by that.

We ran a very large survey of our customers, and the customers who had used our technology-- and compared the customers who had used our technology and had problems with it and gotten emotionally smart support to the customers who had used our technology and not had any problems with it. Which of those two groups was most likely to buy our product again?

Turned out the group who had had problems with their technology and gotten good support was significantly more likely to buy their brand again than the people who had had no problems with the machines. Now, that should disturb you.

The short take-home message is frustrate your customers, handle it well, and they'll be more likely to buy your product. So was that then their strategy? Really deliberately frustrate people? Well, we never got around to answering that question. But you can guess the real answer is, these days you have to get a product out there so fast that there's no time to really get all the kinks out of it.

So when it's only 80% ready, we, the customers, bear the brunt, the frustration, the stress, the resulting health consequences of that increased stress. And we don't see the price-performance curve include the amount of stress that it adds to us.

One more one I'll mention briefly since yesterday was commencement day, and that's-- many years ago, I don't know if you were at the commencement when Lee Iacocca spoke. But on that beautiful June afternoon in our lovely courtyard, he took the wind out of the sails, if you will. Well, this one refers to the wind out of the sails.

Everybody is there celebrating-- and you know how wonderful it is to be there celebrating this great occasion of commencement-- and he yelled at the graduates not congratulations, but you must get angry. And they're like, huh? Wait a minute. This is a joyful day. Does he realize where he is?

And he said it again. You must get angry. Well, for a long time, people talked about this. And what they realized the message was, was if you really want to make a great change in the world, the negative emotion of anger is a very strong motivator.

In other words, someone who just always tries to alleviate all those negative emotions, get rid of the customer's frustration and so forth, is also potentially robbing you of the goad you need to make positive change. So negative valence emotions do not necessarily mean that they are a bad thing. We know of many cases where the strong negative emotions lead to much better change overall.

There are lots of keys-- ways to address this, and I'll give you a URL and book pointer at the end to where you can find out much more about these. We are concerned about these issues, but we do also think that they're addressable. So we welcome dialog on all of this as we go.

I hear a lot-- people sometimes say, gee, you guys talk about making things that think, making machines emotionally savvy. What about making people that think and people emotionally savvy? And in fact, the number insurer of physicians in the Boston area said that they had done a large study of which physicians who had had a particular thing go wrong, which ones were most likely to get sued?

Both groups had the same kind of problem go wrong. These guys were likely to get sued. These guys weren't. What was the big difference? Well, the big difference seemed to have as a very large component of it the ability to have this appropriate rapport, empathy, emotional skills communicated to the patient. So they put this very large dollar figure on this problem.

Some of you who have autistic kids-- and by the way, people at places like MIT are much more likely to have autistic kids than people at more liberal arts schools and so forth-- may know that autistic kids suffer-- it's very broad class of disorder, but there's a strong tendency among them to suffer from an inability to recognize and respond appropriately to emotion.

So we've taken some of our technology and built a system that works with kids. This is the ASQ system developed by Kathy Blocher that showed clips of various emotional scenarios to the autistic kids. And in trials with this at the Dan Marino Center down in Florida, the six kids who lasted through the lengthy trials showed several significant improvements in their ability to recognize emotion within the context of the computer environment. And there appears to be some extension of that into the home environment. But of course, it takes a long time to really make solid claims about that.

So I'm wrapping up here. I'll just say briefly, I've mentioned some issues in sensing signals. We've been building new sensors-- mice, pressure sensors, new tangible interfaces, wearable interfaces that gather information from you and then run that through pattern recognition signal processing to try to infer things about your state.

So this is different from really knowing how you feel. Your feelings truly involve thoughts as well as physical expressions. So we can't really recognize what you're feeling. We just recognize expressions of it. And then I've given you one example of a system that tries to reduce frustration, respond to those emotions, and just mention that we're engaged in work now, trying to help teach affective skills.

Now, in all of this work where we mention emotion, it's just extremely important to remind ourselves that with emotion, there is a time that you should pay careful attention to it. There's also a time to ignore it. There's a time to express emotion and a time to suppress it and so forth.

And every time, however, we need to balance. And it's the balance that's been missing from computers. I am not arguing that we should make computers these emotionally gushy, goofy, friendly machines and so forth. What we're simply trying to do is to move to a balance that is not present today because of an almost total lack of attention to this topic. So what we're trying to move toward is machines that appropriately express, respond to, and show respect for human emotions.

So I'll close by leaving you with the URL where I invite you to come and download papers, find more information. There's a book on this that addresses a whole spectrum of topics related to this as well published by MIT Press. And we'll welcome your questions now. Thank you.

**MITCHELL:** Well, thank you, everybody. And we still do have quite a reasonable amount of time for questions. And I promise I will get you out of here in time for lunch. Our low-tech question collection system worked extremely well, and so I have a very large stack of questions here. And I'm going to ask the first one. And I'm going to toss it out to our speakers. And meanwhile, I'm going to look through the rest of the questions to see.

So I'm going to toss this one to all of you. Anybody can pick it up. And the question is how will all these technologies transform the third world?

**BROOKS:**

I'll start with that. We've seen over the last few years, one, to me at least, surprising thing that these technologies have done is we've managed to shift intellectual work to the third world. So a lot of computer software development is now done in India and some in China because of availability of networks.

I actually think that over the next few years, we're going to see more shift of labor for things like the hair cutting. Or maybe not the hair cutting exactly, but things which apparently require a person to be there and do something physical, but they don't actually have to be that intellectual.

So you can imagine security guards could be in a third world country for some installation here in Boston where there are cameras sitting around, watching at night or watching during the day. And when they detect an anomaly of some sort but they really need a human to make the judgment of whether this is a serious anomaly or whether it's just a piece of paper blowing through the scene, that decision could be made by someone in the third world country who gets shipped those images with a latency of a second or something. They look at it, they make the decision, and they're providing valuable work here in Boston which is not highly skilled. So that's just the edge of it.

Then as the robots become more physically able to do things with manipulation as they are being pushed by surgery robot research, we'll also get physical work where the people can tele-operate in a supervisory role. Robots doing things like cleaning bathrooms in a remote location. That's the hardest thing to do, to get people to actually get in the bathrooms and clean them in commercial cleaning.

That can be remote, and as soon as that can be remote, that can be around the other side of the world. So I think this physical work will migrate around the world and labor markets don't have to be restricted to where we are physically.

**PICARD:**

As you may know, the MIT Media Lab's opening a whole new center, the OKAWA Center for Future Children addressing in large effort how to have a positive effort on the third world. And actually, I would flip that around too. The third world can have a very positive impact on us as well. That involves things like instead of sending computers, send a printer and the capability for them to print to their own computers-- new technology, penny PCs, PCs that cost a penny each.

On another note, when it comes to actually using these things, there's always a barrier of language. And there's work in our labs such as that by Deb Roy on computers that would learn the language as you interact with them in ways that somewhat imitate the way we think children learn. In fact, Kismet's next step is to do that as well.

And one of the really important things to note about the Kismet project there is that it's paying attention to a key signal that humans seem to use in learning, and it's a universal signal across all cultures, and that's the affect. Is the child pleasing the person who's teaching them the English, the caregiver, or not?

And if the computer is going to learn from you, it's really got to have this feedback in a natural and social way so you don't have to first go through this huge learning curve to interact with it. If it defaults to learning what you approve and disapprove of, then it can adjust its learning loops to better adjust to your language and your culture.

**SHEFFI:**

Just to add, I'm in fact quite optimistic about the use of technology in general because in the third world, to advance the third world, we're seeing already with, for example, the use of cell phone, that one gets beyond the need to put expensive infrastructure in the ground in order to get communication.

Once we start having wireless broadband, it will bring opportunities, I think, for education and communication to the third world that are simply prohibitive now because of the cost of the infrastructure that has to go into the ground.

And if we can leapfrog a generation through the use of the wireless technology and the inexpensive devices, I think we can get to a world that is much more connected and much more educated. And personally, I feel that there a big hope there for balancing the current social and economic inequities in the world in general.

**MITCHELL:**

There are a couple of questions here that are directed towards me. Let me jump in and answer these questions. Then I'm going to ask Professor Sheffi to answer a question, but he needs his reading glasses to read the question.

One of the questions here is a question-- who's concerned about the social dynamics of work-at-home-based culture? And the questioner points out, there is value in personal interactions at the office around the coffee pot and so on.

And a second question very closely connected to this asks, isn't the absence of face-to-face transactions reducing the quality of personal interactions, de-energizing social interaction and so on? These questions come up a lot. And let me sketch a couple of interesting answers to these because some of the answers are counter intuitive.

The way to think about these questions is to recognize that one technology doesn't usually separate-- a new technology doesn't usually simply substitute for an older technology. You get typically some substitution effects, but much more often, complementarity effects and very interesting kinds of interactions.

So a lot of these sorts of questions, very legitimate questions, arise from the sense that probably what happens is that maybe email substitutes for face-to-face interaction, and clearly there's something lost in that and so on, but it turns out to be much more complicated.

Let me do a little thought experiment here-- a little survey among the audience. Let me ask all of you to think for a moment about what's the single most common use of email? I assume you mostly use email. Of your own uses of email, what's the single thing you use email for most?

Now, if your like most people I ask this question to, it's for arranging face-to-face meetings. That's typically what's the highest use of email. And the reason for that is you use very low cost, asynchronous, very convenient interaction to arrange the highest cost, of the most high-quality interaction that you have, which is face to face. So two things, in fact, end up working together in a way that's a little bit counter intuitive, but is not simply a substitution effect.

What happens in the workplace is something like this-- that through electronic telecommunication, we get more flexibility in the workplace, more flexibility in where we do work, but that flexibility doesn't end up producing one simple result. So typically what happens, for example, is the that telecommuting capabilities give you the capability to do private work in many different locations.

You can get access to databases-- all of these kinds of things. So that becomes very flexible. The sort of work that in a typical office building might be done a Dilbert-like cubicle alone no longer has to be done in that sort of context. You can do it anywhere.

But a lot of things that really do require that face-to-face interaction-- delicate negotiations, creative problem solving in a group and so on-- all of these kinds of things still end up happening very much in face-to-face context. So we're seeing a shift in the functions of office space.

And this is very, very clear if you look at the way people are building office space these days. Much less emphasis on private workspace. That can happen anywhere, and happens on the road, it happens at home, it happens in offices that are not private offices, but you just occupy for a short amount of time when you need it and so on. Much less emphasis on that in the centralized workplace and much more emphasis, in fact, on the social functions, the place to hang out, the meeting rooms, the place where you can come together in a congenial atmosphere to work as a group.

So I think what we're seeing generally is a restructuring of the workplace and indeed the recognition that face-to-face interaction remains very important, but happens in different patterns in different ways in different locations. We're very unlikely to see a scenario, to summarize, in which everybody spends all their time at home in darkened rooms in their underwear typing email messages to each other. That's not likely to be the way that it goes.

**SHEFFI:**

Well, I'm not sure about that, just to add to this. And I say I'm not sure about that just because I look at my son. My son happens to just finish his freshman year at MIT. Makes me very proud. But the way throughout his teenage years that he was interacting with his-- and he came to me and had some problem and says, Joe suggested that I do this and that. Who is Joe? Joe is my friend. Who is your friend? You never brought him over. Well, he lives in the Philippines. I mean, I never saw him.

The way he automatically and naturally was treating people that he never saw in his life, using and ICQ and all these methods to communicate with him internally, creating his own community, I looked at him, I said, as I said, this is a person from another space. So I'm not sure. I still have doubts about where we are going with all of this.

**MITCHELL:**

Well, one of the ways to think about this is what to developers is a kind of economy of presence, if you like, where we have the whole spectrum of communication possibilities ranging in that table you remember I showed from a face-to-face synchronous communication, very high emotional intensity, very high cost, very high opportunity cost typically, all the way down to the most attenuated thing, which is a distributed, asynchronous interaction-- all the possibilities around it.

And essentially, what we do in our lives is-- say we have a certain amount of interaction capacity-- it's expanding with the new technologies. But we allocate our activities among those different cells of the matrix, if you like, depending on our personality, that particular situation.

So we make various trade-offs, and it's a very complex kind of allocation decision, I think, that we make, or set of allocation decisions we make. And it's changing. The way people do this is changing. And teenagers probably are different. I've observed this with my teenager.

Anyway, let me pass it on. There's a question.

**SHEFFI:**

Oh, you want me to read the question. Okay, it's a question directed to me. Let me read the question first. As an economics major at MIT, I find the pricing solution elegant. As a person-- which means there's a distinction there between MIT graduate and a person-- I am very put off by the notion of everything being measured in price. I don't want to live in that way. This is the psychological cost-- underlined psychological cost-- there is a psychological cost to this kind of solution. Looking for other approaches.

Well, it's really a question of outlook on life. It's a question of-- it goes quite deep. In 1978, I was in the United States. I came to the US in 1975. Some of you may have detected this is not a North Boston accent. Came in '75. '78 we had the embargo on gas lines from the Iran embargo.

And I delivered a paper in the Transportation Research Board meeting in Washington DC-- suggested that gas stations would be allowed each to have two prices at the pump-- one, which will be government controlled and another they'll be able to charge whatever they damn please. And let the users-- those who want to wait in line will wait, those who want to pay will pay.

And this being in Washington, I was tarred and feathered, and got there with a still alive [INAUDIBLE], but just hardly. There is an element in society, in Western society probably, that likes the fact that we are all waiting in line together. We are all waiting in line-- regardless of how-- I'm not putting it down-- regardless of how bad it is, but we are all in it together.

Rather than try to come up with a more economically viable alternative and saying, the rich are rich, and they are really not like the rest of us. And this is just another element of if they may be able to get through an intersection faster-- and it doesn't have to be rich, of course, as I mentioned before. It can be done in any other unit. It does not have to be monetary unit, but clearly whatever unit we'll be using, that they will develop a market in this unit that will involve money as well.

So it's not a matter of-- the other solutions are staring us in the face in every element of life, whether we want it to do-- when we're talking about transportation, we like the system that somebody allocates the preset green light regardless of what the traffic is, or we want to be able to do it more dynamically, and I realize that the people may be willing to pay.

And being able to extract it-- and I did not study economics-- but it still seems to be a total waste and, of course, dead weight loss in terms of economics that there are people who are willing to pay in terms of money, and we extract the payment in terms of waiting time or other inconvenience.

And the reason that it's so wasteful is because money can be put to great purposes. This whole thing can be revenue neutral or money can be put in building shelters or helping the homeless or solving a lot of other of society's ill. And this goes to waste when people have to wait through waiting time. So I don't know. I'm still hooked on my solution. That's the answer.

**BROOKS:**

The question here for me-- why do you believe that humans will always be better than machines at designing other machines? I want to address that in two ways. Patrick Winston-- he teaches the undergraduate artificial intelligence course here at MIT.

And he always tells a story that when he was young in Illinois, he had a raccoon that was very dexterous. And no matter how he tried to lock up the refrigerator so the raccoon couldn't get into it, the raccoon always managed to break in and steal the food. And he said, but it never occurred to him to ask himself the question whether that raccoon was smart enough to build a robot raccoon.

So maybe someone comes from Alpha Centauri and looks and says, oh, look at those cute little humans. Look at them running around. Oh, that one thinks he can build an artificial human. They're not smart enough to do that. So it may be arrogance on our part to even think that we can try to build machines that are as good as human machines because we're just below the threshold where we can do that.

So one approach to that is to try and build artificial evolutionary methods and have big, crunching computers go through and do brute force guided search through evolutionary techniques to evolve machines that are smart. And there's been a lot of progress in that technique over the last 10 years, especially as machines have gotten larger.

So indeed, in fact, some of these evolutionary techniques can design circuits that no human is capable of designing-- much better circuits than humans are capable of. So I don't think that people can design machines better than machines can. In fact, machines are crucial to designing the next generation of machines. So then the question is, well, would these smart robots be able to do things quicker than us?

But I think my point was that we will incorporate the best parts of the machines into ourselves, so we'll just slightly keep ahead of them. Our species will change and eventually the question will drift away about who's machine, who's people ultimately when we go down a few hundred years from now.

**PICARD:**

There's two questions on here which Rod and I will tackle together. The first is, it appropriate that toys for children display human attributes? That is, is it psychologically beneficial to the child? It's a very large question. I have experience with this in two areas.

One is in the case of children who have certain impairments with social emotional skills, and there the psychologists now we are working with are extremely excited about this. The kids already migrate to the computers.

Adult autistics who are high functioning actually love sending email and interacting on the Web. They say-- and this should be a warning sign to us-- they say that it levels the playing field for them to communicate via the internet.

The toy effort-- my former student Jonathan Klein has been working with Rod's company, iRobot, and he mentioned the Hasbro doll coming out. And there have been other efforts like that. Rod mentioned one, the Furby, and we've done one with Tigger and Microsoft's got the Barney.

And I understand that there's a lot of psychological interest in this, and people like Sherry Turkle are studying the effects of children putting their little Furbies to bed. Is this really different than what they do with their Teddy bear and their doll that doesn't have all these things? The answer there to the best of my knowledge is we still really don't know. It's sort of uncharted waters. And if you want to add to that?

**BROOKS:**

No, I think I agree with you.

**PICARD:** Yeah, okay. The second one specifically asked if females were more communicative with Kismet than males. And I'll just mention that in our work when we've looked at are people nicer to the computer face to face, or do they respond more to the emotionally savvy agent as a function of their gender, we have actually found no significant gender differences. We expected to find some, but they weren't there in the real data.

However, it is known that females as a whole are better than males as a whole, statistically speaking, at recognizing the intended communication of emotion. So if somebody intentionally stands up here and communicates a bunch of basic emotions, the women tend to score a bit higher than the men at recognizing what they intended to communicate.

**BROOKS:** Explicitly with Kismet, we haven't tried to measure that in any way. You may have noticed, and maybe this is why the question came, that all the subjects who were trying to praise Kismet or prohibit Kismet were women, it is true that getting the affect out of voices is different for male voices and female voices.

The two students who were doing this work both happened to be women, so they decided to only train the system with women because they actually had to have two separate models, and they didn't want to build two.

**MITCHELL:** The question says, these discussions of telecommuting and e-universities seem to neglect our need for socialization which is provided for most of us by work and school. Will people really settle for life in a cocoon, however well equipped? Let me ask how others would like to? Would you like to? Would you like me to?

**SHEFFI:** Why don't you start?

**MITCHELL:** Why don't I start? Okay. What I'd suggest here is that the way you really need to think about these things-- and we really need to shift the way one thinks about these things-- often things like telecommuting and any kind of distributed thing thought of in terms of individual-to-individual interactions. And in fact, there are many, many more-- and that's useful. That's very useful, and one aspect that is certainly important.

But there are many other aspects. One of the things we found most fruitful in distributed education, for example, is connecting one social setting to another social setting rather than connecting an individual to individual. So you get a kind of richness out of the interconnection of social settings.

For example, I've spent a lot of time over the last few years teaching design studios in which students at MIT physically in the design studio in a very face-to-face design setting where there's a lot of social stuff going on, I can assure you, link up electronically to design studios on the other side of the world. We've done it with Japan, Portugal, Australia-- lots of different places.

And the students in these settings work jointly on design problems. And that turns out to be an extraordinarily effective and useful thing to do because what you're able to do is essentially put together the social capital and the social dynamic of a couple of very different groups. It's very interesting issues start to arise related to some of the other things we've been talking about.

For example, if you put together the two groups of designers to talk about an issue in different parts of the world, you discover, because of cultural differences, they frame the issues differently, they interact with each other in different ways, the discourse unfolds in very, very different fashion.

So for example, MIT students working on a design problem with each other are extraordinarily aggressive verbally in the way they interact back and forth. They won't say that's just a bad idea. They'll say that's the stupidest idea I've ever heard in my life. And the discussion goes in that kind of framework.

Now, a group of Japanese students doing the same kind of thing interact in a very, very different kind of way, as many of you will recognize. The cultural conventions are that if you want to say no to something, you say yes unenthusiastically as a way of communicating. And a lot more emphasis-- and these are cliches, but they're roughly true-- a lot more emphasis on building consensus and so on.

Now, if you're adept culturally in moving back and forth from one context to another, you know how to adjust your style of interaction depending on geographic context. So if you're smart, you don't behave in Japan the same way that you behave in the United States.

However, when you create an electronic space of interaction, which is neither Japan nor the United States, or nor MIT specifically, and you have to work together in that space, it's nobody's particular territory, so it's not immediately evident whose cultural convention should prevail for framing a question and working through a set of interactions.

It turns out to be educationally extraordinarily fruitful because what the students have to do is figure out on the fly and negotiate on the fly among themselves and with each other how they're going to interact with each other, how they're going to deal with this, how they're going to deal with the cultural issues on so on, that foregrounds those issues and teaches them something very, very important.

**BROOKS:**

Yeah, I agree with what you said, but I wanted to take a slightly different tack on it. I think this is an issue that's facing MIT. In Course 6 at least, we're starting to experiment with putting lectures for our big courses on the Web. We still have recitations with a faculty member, but why give the same lecture which is not interactive every term when you can just record it on the Web.

Now, should we make that material available to the rest of the world? And if people in the Philippines or somewhere come and go through all the course material, have they had the same experience getting an MIT degree that way as they would coming here?

And I think we'd probably all agree, no, it's important for being here on campus and having those interactions and learning to say, that's the most ridiculous thing I've ever heard, Bill, because that has given all of you some abilities to go out and make really tough technical decisions and be able to push technology along.

But what is that component that you get by being here? And what components can you get without being here? And how do you still get that component by being here? And I think there's a lot of questions for us at MIT to answer about how we want to distribute our education and the role of campus itself.

**SHEFFI:**

One case study about this-- I also think that this discussion has a lot to do with MIT, with the future of the residential universities in general. I did my undergrad degree in Israel, which is a very different experiment. There it's a very different experience. There's a lot less social interaction. People go to classes, some of them literally even a few years ago were doing it remotely.

And the main reason is that people go to Israel. In Israel, they go to college after at least three years of service in the armed forces which serves-- it's age 18 to 21-- which serves the same function of getting away from home, getting some social interaction norms established, growing up. So when they get to the university, they look at it simply as knowledge acquisition rather than the whole experience of going through a university.

And we have the residential university in the United States offers, of course, both in one package. And it is not clear that in the future many people will not opt for getting the knowledge acquisition through efficient Web-based mechanism and signing up for two-year camp that may be MIT doing other-- getting the social growing and the social interaction, but somehow being able to deliver those separately.

**MITCHELL:**

So you never thought of MIT as the place to learn social interaction, did you? So we're running out of time, so let me wrap up with just a comment that I think may be an interesting way to summarize some things.

If you look to MIT's future and ask where we're placing our bets about the future of education, the future of the campus and so on, I think you can say we're diversifying our bets, and we're trying to do two things simultaneously, and expecting that they're going to work in a complimentary way.

We're doing some very exciting research and some very significant investment in educational technology in new ways of teaching and learning, in new strategies for connecting the physical setting of MIT to a wider world and enriching that setting by distributing some of our activities and so on. So that's one direction.

At exactly the same time, we're going into the largest construction program that MIT has ever had in its history, and we're intensifying the physical campus in a way that I think is going to be enormously exciting. Working very hard on intensifying the social character, the opportunities for interaction and interconnection in a direct, face-to-face way in the physical campus located right here in Cambridge, Massachusetts.

We have a new master plan being done which is extraordinarily good. Some magnificent new buildings-- a building by Frank Gehry that will accommodate the lab for artificial intelligence, the lab for computer science, linguistics and philosophy, and a group of other things. [INAUDIBLE] extraordinary building done by Frank Gehry.

Fumihiko Magi is doing a new building for the Media Lab that this is all structured around some really exciting ideas about how you make working groups in face-to-face interaction in very effective ways. Steve Hall-- wonderful, really avant-garde New York architect, is doing a new undergraduate residence that we think is going to be enormously exciting in revitalizing the campus. We're doing a new central athletic facility.

So I think we're going to get both in the immediate future. We're going to get a very intense, active, socially vibrant campus, and it's going to be wired to the world and able to do things that we were never able to do in the past. So thank you. Lunch time.