

The Design *of* Future Things

Donald A. Norman

BASIC

BOOKS

A MEMBER OF THE PERSEUS BOOKS GROUP
NEW YORK

Copyright © 2007 by Donald A. Norman
Published by Basic Books,
A Member of the Perseus Books Group

All rights reserved. Printed in the United States of America. No part of this book may be reproduced in any manner whatsoever without written permission except in the case of brief quotations embodied in critical articles and reviews. For information, address Basic Books, 387 Park Avenue South, New York, NY 10016-8810.

Books published by Basic Books are available at special discounts for bulk purchases in the United States by corporations, institutions, and other organizations. For more information, please contact the Special Markets Department at the Perseus Books Group, 2300 Chestnut Street, Suite 200, Philadelphia, PA 19103, or call (800) 255-1514, or e-mail special.markets@perseusbooks.com.

Designed by Timm Bryson
Set in 11.5 point Minion

Library of Congress Cataloging-in-Publication Data

Norman, Donald A.

The design of future things / Donald A. Norman.

p. cm.

Includes bibliographical references and index.

ISBN-13: 978-0-465-00227-6

ISBN-10: 0-465-00227-7

1. Design, Industrial—Psychological aspects. 2. Human engineering. I. Title.

TS171.4.N668 2007

745.2—dc22

2007035377

10 9 8 7 6 5 4 3 2 1

Cautious Cars and Cantankerous Kitchens

How Machines Take Control

I'm driving my car through the winding mountain roads between my home and the Pacific Ocean. Sharp curves drop off steeply amidst the towering redwood trees and vistas of the San Francisco Bay on one side and the Pacific Ocean on the other. It's a wonderful drive, the car responding effortlessly to the challenge, negotiating sharp turns with grace. At least, that's how I am feeling. But then I notice that my wife is tense: she's scared. Her feet are braced against the floor, her shoulders hunched, her arms against the dashboard. "What's the matter?" I ask, "Calm down, I know what I'm doing."

Now imagine another scenario. I'm driving on the same winding, mountain road, and I notice that my car is tense: it's scared. The seats straighten, the seat belts tighten, and the dashboard starts beeping at me. I notice the brakes are being applied automatically. "Oops," I think, "I'd better slow down."

Do you think the idea of a frightened automobile fanciful? Let me assure you, it is not. This behavior already exists on some luxury automobiles—and more is being planned. Stray out of

your lane, and some cars balk: beeping, perhaps vibrating the wheel or the seat or flashing lights in the side mirrors. Automobile companies are experimenting with partial correction, helping the driver steer the car back into its own lane. Turn signals were designed to tell other drivers that you are going to turn or switch lanes, but now they are how you tell your own car that you really do wish to turn or change lanes: “Hey, don’t try to stop me,” they say to your car. “I’m doing this on purpose.”

I was once a member of a panel of consultants advising a major automobile manufacturer. I described how I would respond differently to my wife than my car. “How come?” asked fellow panelist Sherry Turkle, an MIT professor and an authority on the relationship between people and technology. “How come you listen to your car more than your wife?”

How come, indeed. Sure, I can make up rational explanations, but they will miss the point. As we start giving the objects around us more initiative, more intelligence, and more emotion and personality, we now have to worry about how we interact with our machines.

Why do I appear to pay more attention to my car than to my wife? The answer is complex, but in the end, it comes down to communication. When my wife complains, I can ask her why, then either agree with her or try to reassure her. I can also modify my driving so that she is not so disturbed by it. But I can’t have a conversation with my car: all the communication is one way.

“Do you like your new car?” I asked Tom, who was driving me to the airport after a lengthy meeting. “How do you like the navigation system?”

“I love the car,” said Tom, “but I never use the navigation system. I don’t like it: I like to decide what route I will take. It doesn’t give me any say.”

Machines have less power than humans, so they have more authority. Contradictory? Yes, but, oh, so true. Consider who has more power in a business negotiation. If you want to make the strongest possible deal, who should you send to the bargaining table, the CEO or someone at a lower level? The answer is counterintuitive: quite often, the lower-level employee can make the better deal. Why? Because no matter how powerful the opposing arguments, the weak representative cannot close the deal. Even in the face of persuasive arguments, he or she can only say, “I’m sorry, but I can’t give you an answer until I consult with my boss,” only to come back the next day and say, “I’m sorry, but I couldn’t convince my boss.” A powerful negotiator, on the other hand, might be convinced and accept the offer, even if later, there was regret.

Successful negotiators understand this bargaining ploy and won’t let their opponents get away with it. When I discussed this with a friend, a successful lawyer, she laughed at me. “Hey,” she said, “if the other side tried that on me, I’d call them on it. I won’t let them play that game with me.” Machines do play this game on us, and we don’t have any way of refusing. When the machine intervenes, we have no alternatives except to let it take over: “It’s this or nothing,” they are saying, where “nothing” is not an option.

Consider Tom’s predicament. He asks his car’s navigation system for directions, and it provides them. Sounds simple.

Human-machine interaction: a nice dialogue. But notice Tom's lament: "It doesn't give me any say." Designers of advanced technology are proud of the "communication capabilities" they have built into their systems. But closer analysis shows this to be a misnomer: there is no communication, none of the back-and-forth discussion that characterizes true dialogue. Instead, we have two monologues. We issue commands to the machine, and it, in turn, commands us. Two monologues do not make a dialogue.

In this particular case, Tom does have a choice. If he turns the navigation system off, the car still functions, so because his navigation system doesn't give him enough say over the route, he simply doesn't use it. But other systems do not provide this option: the only way to avoid them is not to use the car. The problem is that these systems can be of great value. Flawed though they may be, they can save lives. The question, then, is how we can change the way we interact with our machines to take better advantage of their strengths and virtues, while at the same time eliminating their annoying and sometimes dangerous actions.

As our technology becomes more powerful, its failure in terms of collaboration and communication becomes ever more critical. Collaboration means synchronizing one's activities, as well as explaining and giving reasons. It means having trust, which can only be formed through experience and understanding. With automatic, so-called intelligent devices, trust is sometimes conferred undeservedly—or withheld, equally undeservedly. Tom decided not to trust his navigational system's instructions, but in some instances, rejecting technology can cause harm. For example, what if Tom turned off his car's antiskid brakes or the stability control? Many drivers be-

lieve they can control the car better than these automatic controls. But antiskid and stability systems actually perform far better than all but the most expert professional drivers. They have saved many lives. But how does the driver know which systems can be trusted?

Designers tend to focus on the technology, attempting to automate whatever possible for safety and convenience. Their goal is complete automation, except where this is not yet possible because of technical limitations or cost concerns. These limitations, however, mean that the tasks can only be partially automated, so the person must always monitor the action and take over whenever the machine can no longer perform properly. Whenever a task is only partially automated, it is essential that each party, human and machine, know what the other is doing and what is intended.

Two Monologues Do Not Make a Dialogue

SOCRATES: You know, Phaedrus, that's the strange thing about writing. . . . they seem to talk to you as if they were intelligent, but if you ask them anything about what they say, from a desire to be instructed, they go on telling you just the same thing forever.

—Plato: Collected Dialogues, 1961.

Two thousand years ago, Socrates argued that the book would destroy people's ability to reason. He believed in dialogue, in conversation and debate. But with a book, there is no debate: the written word cannot answer back. Today, the book is such a

symbol of learning and knowledge that we laugh at this argument. But takè it seriously for a moment. Despite Socrates's claims, writing does instruct because we do not need to debate its content with the author. Instead, we debate and discuss with one another, in the classroom, with discussion groups, and if the work is important enough, through all the media at our disposal. Nonetheless, Socrates's point is valid: a technology that gives no opportunity for discussion, explanation, or debate is a poor technology.

As a business executive and as a chair of university departments, I learned that the process of making a decision is often more important than the decision itself. When a person makes decisions without explanation or consultation, people neither trust nor like the result, even if it is the identical course of action they would have taken after discussion and debate. Many business leaders ask, "Why waste time with meetings when the end result will be the same?" But the end result is not the same, for although the decision itself is identical, the way it will be carried out and executed and, perhaps most importantly, the way it will be handled if things do not go as planned will be very different with a collaborating, understanding team than with one that is just following orders.

Tom dislikes his navigation system, even though he agrees that at times it would be useful. But he has no way to interact with the system to tailor it to his needs. Even if he can make some high-level choices—"fastest," "shortest," "most scenic," or "avoid toll road"—he can't discuss with the system why a particular route is chosen. He can't know why the system thinks route A is better than route B. Does it take into account the long

traffic signals and the large number of stop signs? And what if two routes barely differ, perhaps by just a minute out of an hour's journey? He isn't given alternatives that he might well prefer despite a slight cost in time. The system's methods remain hidden so that even if Tom were tempted to trust it, the silence and secrecy promotes distrust, just as top-down business decisions made without collaboration are distrusted.

What if navigation systems were able to discuss the route with the driver? What if they presented alternative routes, displaying them both as paths on a map and as a table showing the distance, estimated driving time, and cost, allowing the driver to choose? Some navigation systems do this, so that the drive from a city in California's Napa Valley to Palo Alto might be presented like this:

FROM ST. HELENA, CA TO PALO ALTO, CA

	DISTANCE	ESTIMATED TIME	ROUTE	TOLLS
1	94.5 Miles	1 Hour 46 Minutes	Via Dumbarton Bridge	\$4
2	98.3 Miles	1 Hour 50 Minutes	Via San Francisco Bay Bridge	\$4
3	103.6 Miles	2 Hours 10 Minutes	Via Golden Gate Bridge	\$5

This is a clear improvement, but it still isn't a conversation. The system says, "Here are three choices: select one." I can't ask for details or seek some modification. I am familiar with all these routes, so I happen to know that the fastest, shortest, cheapest route is also the least scenic, and the most scenic route is not even offered. But what about the driver who is not so

knowledgeable? We would never settle for such limited engagement with a human driver. The fact that navigation systems offering drivers even this limited choice of routes are considered a huge improvement over existing systems demonstrates how bad the others are, how far we still have to go.

If my car decides an accident is imminent and straightens the seat or applies the brakes, I am not asked or consulted; nor am I even told why. Is the car necessarily more accurate because, after all, it is a mechanical, electronic technology that does precise arithmetic without error? No, actually it's not. The arithmetic may be correct, but before doing the computation, it must make assumptions about the road, the other traffic, and the capabilities of the driver. Professional drivers will sometimes turn off automatic equipment because they know the automation will not allow them to deploy their skills. That is, they will turn off whatever they are permitted to turn off: many modern cars are so authoritarian that they do not even allow this choice.

Don't think that these behaviors are restricted to the automobile. The devices of the future will present the same issues in a wide variety of settings. Automatic banking systems already exist that determine whether you are eligible for a loan. Automated medical systems determine whether you should receive a particular treatment or medication. Future systems will monitor your eating, your reading, your music and television preferences. Some systems will watch where you drive, alerting the insurance company, the rental car agency, or even the police if they decide that you have violated their rules. Other systems monitor for copyright violations, making decisions about what should be permitted. In all these cases, actions are apt to be

taken arbitrarily, with the systems making gross assumptions about your intentions from a limited sample of your behavior.

So-called intelligent systems have become too smug. They think they know what is best for us. Their intelligence, however, is limited. And this limitation is fundamental: there is no way a machine has sufficient knowledge of all the factors that go into human decision making. But this doesn't mean we should reject the assistance of intelligent machines. As machines start to take over more and more, however, they need to be socialized; they need to improve the way they communicate and interact and to recognize their limitations. Only then can they become truly useful. This is a major theme of this book.

When I started writing this book, I thought that the key to socializing machines was to develop better systems for dialogue. But I was wrong. Successful dialogue requires shared knowledge and experiences. It requires appreciation of the environment and context, of the history leading up to the moment, and of the many differing goals and motives of the people involved. I now believe this to be a fundamental limitation of today's technology, one that prevents machines from full, humanlike interaction. It is hard enough to establish this shared, common understanding with people, so how do we expect to be able to develop it with machines?

In order to cooperate usefully with our machines, we need to regard the interaction somewhat as we do interaction with animals. Although both humans and animals are intelligent, we are different species, with different understandings and different capabilities. Similarly, even the most intelligent machine is a different species, with its own set of strengths and weaknesses,

its own set of understandings and capabilities. Sometimes we need to obey the animals or machines; sometimes they need to obey us.

Where Are We Going? Who Is in Charge?

“My car almost got me into an accident,” Jim told me.

“Your car? How could that be?” I asked.

“I was driving down the highway using the adaptive cruise control. You know, the control that keeps my car at a constant speed unless there is a car in front, and then it slows up to keep a safe distance. Well, after awhile, the road got crowded, so my car slowed. Eventually, I came to my exit, so I maneuvered into the right lane and then turned off the highway. By then, I had been using the cruise control for so long, but going so slowly, that I had forgotten about it. But not the car. I guess it said to itself, ‘Hurrah! Finally, there’s no one in front of me,’ and it started to accelerate to full highway speed, even though this was the off-ramp that requires a slow speed. Good thing I was alert and stepped on the brakes in time. Who knows what might have happened.”

We are in the midst of a major change in how we relate to technology. Until recently, people have been in control. We turned the technology on and off, told it which operation to perform, and guided it through its operations. As technology became more powerful and complex, we became less able to understand how it worked, less able to predict its actions. Once computers and microprocessors entered the scene, we often found ourselves lost and confused, annoyed and angered. But

still, we considered ourselves to be in control. No longer. Now, our machines are taking over. They act as if they have intelligence and volition, even though they don't.

Machines monitor us with the best of intentions, of course, in the interest of safety, convenience, or accuracy. When everything works, these smart machines can indeed be helpful, increasing safety, reducing the boredom of tedious tasks, making our lives more convenient, and performing tasks more accurately than we could. It is indeed convenient that the automobile automatically slows when a car darts too closely in front of us, that it shifts gears quietly and smoothly, or, in the home, that our microwave oven knows just when the potatoes are cooked. But what about when the technology fails? What about when it does the wrong thing or fights with us for control? What about when Jim's auto notices that there are no cars in front of it, so it accelerates to highway speed, even though it is no longer on a highway? The same mechanisms that are so helpful when things are normal can decrease safety, decrease comfort, and decrease accuracy when unexpected situations arise. For us, the people involved, it leads to danger and discomfort, frustration and anger.

Today, machines primarily signal their states through alerts and alarms, meaning only when they get into trouble. When a machine fails, a person is required to take over, often with no advance warning and often with insufficient time to react properly. Jim was able to correct his car's behavior in time, but what if he couldn't have? He would have been blamed for causing an accident. Ironically, if the actions of a so-called intelligent device lead to an accident, it will probably be blamed on human error!

The proper way to provide for smooth interaction between people and intelligent devices is to enhance the coordination and cooperation of both parties, people and machines. But those who design these systems often don't understand this. How is a machine to judge what is or is not important, especially when what is important in one situation may not be in another?

I have told the story of Jim and his enthusiastic car to engineers from several automobile companies. Their responses always have two components. First, they blame the driver. Why didn't he turn off the cruise control before exiting? I explain that he had forgotten about it. Then he was a poor driver, is their response. This kind of "blame-and-train" philosophy always makes the blamer, the insurance company, the legislative body, or society feel good: if people make errors, punish them. But it doesn't solve the underlying problem. Poor design, and often poor procedures, poor infrastructure, and poor operating practices, are the true culprits: people are simply the last step in this complex process.

Although the car companies are technically correct that the driver should remember the mode of the car's automation, that is no excuse for poor design. We must design our technologies for the way people actually behave, not the way we would like them to behave. Moreover, the automobile does not help the driver remember. In fact, it seems more designed to help the driver forget! There is hardly any clue as to the state of the cruise control system: the car could do a far better job of reminding the driver of what control it has assumed.

When I say this to engineers, they promptly introduce the second component of their response: "Yes, this is a problem, but don't worry. We will fix it. You're right; the car's navigation sys-

tem should realize that the car is now on the exit road, so it should automatically either disconnect the cruise control or, at least, change its setting to a safe speed.”

This illustrates the fundamental problem. The machine is not intelligent: the intelligence is in the mind of the designer. Designers sit in their offices, attempting to imagine all that might happen to the car and driver, and then devise solutions. But how can the designers determine the appropriate response to something unexpected? When this happens to a person, we can expect creative, imaginative problem solving. Because the “intelligence” in our machines is not in the device but in the heads of the designers so when the unexpected happens, the designer isn’t there to help out, so the machine usually fails.

We know two things about unexpected events: first, they always occur, and second, when they do occur, they are always unexpected.

I once got a third response from an automobile company engineer about Jim’s experience. He sheepishly admitted that the exit lane problem had happened to him, but that there was yet another problem: lane changing. On a busy highway, if a driver decides to change lanes, he or she waits until there is a sufficiently large gap in the traffic in the new lane, then quickly darts over. That usually means that the car is close to those in front and behind. The adaptive cruise control is likely to decide the car is too close to the car in front and therefore brake.

“What’s the problem with that?” I asked. “Yes, it’s annoying, but it sounds safe to me.”

“No,” said the engineer. “It’s dangerous because the driver in back of you didn’t expect you to dart in and then suddenly put on the brakes. If they aren’t paying close attention, they could

run into you from behind. But even if they don't hit you, the driver behind is annoyed with your driving behavior.”

“Maybe,” said the engineer, laughing, “the car should have a special brake light that comes on when the brakes are applied by the automobile itself rather than by the driver, telling the car behind, ‘Hey, don't blame me. The car did it.’”

The engineer was joking, but his comments reveal the tensions between the behavior of people and machines. People take actions for all sorts of reasons, some good, some bad, some considerate, some reckless. Machines are more consistent, evaluating the situation according to the logic and rules programmed into them. But machines have fundamental limitations: they do not sense the world in the same way as people, they lack higher order goals, and they have no way of understanding the goals and motives of the people with whom they must interact. Machines, in other words, are fundamentally different: superior in some ways, especially in speed, power, and consistency, inferior in others, especially in social skills, creativity, and imagination. Machines lack the empathy required to consider how their actions impact those around them. These differences, especially in what we would call social skills and empathy, are the cause of the problems. Moreover, these differences—and therefore these conflicts—are fundamental, not ones that can be quickly fixed by changing the logic here or adding a new sensor there.

As a result, the actions of machines contradict what people would do. In many cases, this is perfectly fine: if my washing machine cleans clothes very differently than I would, I don't care as long as the end result is clean clothes. Machine automation works here because once the washing machine has been loaded and started, it is a closed environment. Once started, the

machine takes over, and as long as I refrain from interfering, everything works smoothly.

But what about environments where both people and machines work together? Or what happens with my washing machine if I change my mind after it has started? How do I tell it to use different setting, and once the washing cycle has started, when will the changes take effect—right away or with the next filling of the machine? Here, the differences between the way machines and people react really matter. Sometimes, it appears that the machine is acting completely arbitrarily, although if the machine could think and talk, I suspect it would explain that from its point of view, the person is the one being arbitrary. To the person, this can be frustrating, a continual battle of wills. To the observer, it can be confusing, for it is never clear who is in charge or why a particular action has been taken. It doesn't really matter whether the machine or the person is correct: it is the mismatch that matters, for this is what gives rise to aggravation, frustration, and, in some cases, damage or injury.

The conflict between human and machine actions is fundamental because machines, whatever their capabilities, simply do not know enough about the environment, the goals and motives of the people, and the special circumstances that invariably surround any set of activities. Machines work very well when they work in controlled environments, where no pesky humans get in the way, where there are no unexpected events, and where everything can be predicted with great accuracy. That's where automation shines.

But even though the machines work well when they have complete control of the environment, even here they don't quite do things the way we would. Consider the "smart" microwave. It

knows just how much power to apply and how long to cook. When it works, it is very nice: you simply have to put in fresh salmon and tell the machine you are cooking fish. Out it comes, cooked to perfection, somewhere between a poached fish and a steamed one, but perfect in its own way. “The Sensor features detect the increasing humidity released during cooking,” says the manual, “[and] the oven automatically adjusts the cooking time to various types and amounts of food.” But notice that it doesn’t determine if the microwave cooks the food in the same way that a person would. A person would test the firmness, look at the color, or perhaps measure the internal temperature. The microwave oven can’t do any of this, so it measures what it can: the humidity. It uses the humidity to infer the cooking level. For fish and vegetables, this seems to work fine, but not for everything. Moreover, the sensing technology is not perfect. If the food comes out undercooked, the manual warns against using the sensor a second time: “Do not use the Sensor features twice in succession on the same food portion—it may result in severely overcooked or burnt food.” So much for the intelligent microwave.

Do these machines aid the home dweller? Yes and no. If machines can be said to have a “voice,” theirs is certainly condescending, offering no hint as to how or why they do what they do, no hint as to what they are doing, no hint as to the amount of doneness, cleanliness, or drying the machine is inferring from its sensing, and no idea of what to do when things don’t work properly. Many people, quite appropriately in my opinion, shun these devices. “Why is it doing this?” interested parties want to know. There is no word from the machines and hardly a word from the manuals.

In research laboratories around the world, scientists are working on even more ways of introducing machine intelligence into our lives. There are experimental homes that sense all the actions of their inhabitants, turning the lights on and off, adjusting the room temperature, even selecting the music. The list of projects in the works is impressive: refrigerators that refuse to let you eat inappropriate foods, tattletale toilets that secretly tell your physician about the state of your body fluids. Refrigerators and toilets may seem an unlikely pairing, but they team up to monitor eating behavior, the one attempting to control what goes into the body, the other measuring and assessing what comes out. We have scolding scales watching over weight. Exercise machines demanding to be used. Even teapots shrilly whistling at us, demanding immediate attention.

As we add more and more smart devices to daily life, our lives are transformed both for good and for bad. This is good when the devices work as promised—and bad when they fail or when they transform productive, creative people into servants continually looking after their machines, getting them out of trouble, repairing them, and maintaining them. This is not the way it was supposed to be, but it certainly is the way it is. Is it too late? Can we do something about it?

The Rise of the Smart Machine

Toward a Natural, Symbiotic Relationship

The hope is that in not too many years, human brains and computing machines will be coupled together very tightly,

and that the resulting partnership will think as no human brain has ever thought.

—J. C. R. Licklider, “Man-Computer Symbiosis,” 1960.

In the 1950s, the psychologist J. C. R. Licklider attempted to determine how people and machines could interact gracefully and harmoniously, or in what he called a “symbiotic relationship,” so that the resulting partnership would enhance our lives. What would it mean to have a graceful symbiosis of people and technology? We need a more natural form of interaction, an interaction that can take place subconsciously, without effort, whereby the communication in both directions is done so naturally, so effortlessly, that the result is a smooth merger of person and machine, jointly performing a task.

There are numerous instances of “natural interaction.” Let me discuss four that demonstrate different kinds of relations: between people and traditional tools, between horse and rider, between driver and automobile, and one involving machine automation, “recommendation” systems that suggest books to read, music to listen to, and films to watch.


Skilled artisans work their materials through their tools, just as musicians relate with their instruments. Whether used by a painter or sculptor, woodworker or musician, their tools and instruments feel like a part of the body. So, craftspeople do not act as if they are using tools but as if they are directly manipulating the items of interest: paint on canvas, sculptured material, wood, or musical sounds. The feel of the materials provides feedback to the person: smooth and resonant here, bumpy or rough there. The interaction is complex but pleasurable. This

symbiotic relationship only occurs when the person is well skilled and the tools are well designed. When it happens, this interaction is positive, pleasurable, and effective.

Think of skilled horseback riders. The rider “reads” the horse, just as the horse can read its rider. Each conveys information to the other about what is ahead. Horses communicate with their riders through body language, gait, readiness to proceed, and their general behavior: wary, skittish, and edgy or eager, lively, and playful. In turn, riders communicate with horses through their body language, the way they sit, the pressures exerted by their knees, feet, and heels, and the signals they communicate with their hands and reins. Riders also communicate ease and mastery or discomfort and unease. This interaction is positive example two. It is of special interest because it is an example of two sentient systems, horse and rider, both intelligent, both interpreting the world and communicating their interpretations to each other.

Example three is similar to the horse and rider, except that now we have a sentient being interacting with a sophisticated, but nonsentient, machine. At its best this is a graceful interaction between the feel of the automobile, the track, and the actions of the driver.

I think of this when I sit beside my son while he drives my highly tuned German sports car at high speed on the racetrack that we have rented for the afternoon. We approach a sharp curve, and I watch as he gently brakes, shifting the car’s weight forward, then turns the steering wheel so that as the front end of the car turns, the rear end, now with reduced weight bearing down, skids, putting the car into a deliberate, controlled skid, known as an “oversteer” condition. As the rear end swings



around, my son straightens the steering wheel and accelerates, shifting the car's weight back to the rear wheels so that we are once again accelerating smoothly down a straightaway with the pleasure of feeling in complete control. All three of us have enjoyed the experience: me, my son, and the car.

Example four, the recommendation system, is very different from the other three for it is slower, less graceful, and more intellectual. Nonetheless, it is an excellent example of a positive interaction between people and complex systems, primarily because it suggests without controlling, without annoyance: we are free to accept or ignore its recommendations. These systems work in a variety of ways, but all suggest items or activities that you might like by analyzing your past selections or activities, searching for similarities to other items in their databases, and by examining the likes and dislikes of other people whose inter-

ests appear similar to yours. As long as the recommendations are presented in a noninvasive fashion, eliciting your voluntary examination and participation, they can be helpful. Consider the search for a book on one of the internet websites. Being able to read an excerpt and examine the table of contents, index, and reviews helps us decide whether to make a purchase.

Some sites even explain why they have made their recommendations, offering to let people tune their preference settings. I have seen recommendation systems in research laboratories that watch over your activities, so if you are reading or writing, they suggest articles to read by finding items that are similar in content to what is on your display. These systems work well for several reasons. First, they do offer value, for the suggestions are often relevant and useful. Second, they are presented in a nonintrusive manner, off to the side, without distracting you from the primary task but readily available when you are ready. Not all recommendation systems are so effective, for some are intrusive—some seem to violate one's privacy. When done well, they demonstrate that intelligent systems can add pleasure and value to our interactions with machines.

A Caveat

When I ride a horse, it isn't any fun for me or the horse. Smooth, graceful interaction between horse and rider requires considerable skill, which I lack. I don't know what I am doing, and both I and the horse know this. Similarly, I watch drivers who are neither skilled nor confident struggle with their automobiles, and I, as a passenger, do not feel safe. Symbiosis is a wonderful concept, a cooperative, beneficial relationship. But in some cases, as in my

first three examples, it requires considerable effort, training, and skill. In other cases, such as in my fourth example, although no high-level skill or training is required, the designers of these systems must pay careful attention to appropriate modes of social interaction.

After I had posted a draft version of this chapter on my website, I received a letter from a group of researchers who were exploring the metaphor of horse and rider to the control of automobiles and airplanes. The “H-metaphor,” they called it, where “H” stands for “horse.” Scientists at the American National Aeronautics and Space Administration research facilities at Langley, Virginia, were collaborating with scientists at the German Aerospace Center’s Institute for Transportation Systems in Braunschweig, Germany, to understand just how such systems might be built. I visited Braunschweig to learn more about their work (fascinating stuff, to which I return in chapter 3). Riders, it seems, delegate the amount of control they give to the horse: when using “loose reins,” the horse has authority, but under “tight reins,” the rider exerts more control. Skilled riders are in continual negotiation with their horses, adjusting the amount of control they maintain to the circumstances. The American and German scientists are trying to replicate this relationship with human-machine interaction—not only with cars but with houses and appliances.

Symbiosis, in the sense meant by Licklider half a century ago, is a merger of two components, one human, one machine, where the mix is smooth and fruitful, the resulting collaboration exceeding what either is capable of alone. We need to understand how best to accomplish this interaction, how to make it so natural that training and skill are usually not required.

Skittish Horses, Skittish Machines

What would it mean for a car and driver to interact much as a skilled rider interacts with a horse? Suppose a car were to balk or act skittish when getting too close to the cars ahead or when driving at a speed it computed to be dangerous? Suppose the car responded smoothly and gracefully to appropriate commands and sluggishly and reluctantly to others? Would it be possible to devise a car whose physical responsiveness communicated the safety status to the driver?

What about your house? What would it mean to have a skittish house? I can see my vacuum cleaner or stove acting up, wanting to do one thing when I wanted it to do another. But my house? Today companies are poised to transform your home into an automated beast, always looking out for your best interests, providing you with everything you need and desire, even before you know you need or desire it. Many companies are anxious to equip, wire, and control these “smart homes”—homes that control the lighting according to their perception of your moods, that choose what music to play or that direct the television images to move from screen to screen as you wander about the house. All these “smart” and “intelligent” devices pose the question of how we will be able to relate to all this smartness. If we want to learn to ride a horse, we have to practice or, better yet, take lessons. So, do we need to practice how to use our home, to take lessons on getting along with our appliances?

What if we could devise natural means of interaction between people and machines? Could we learn from the way that skilled riders interact with horses? Perhaps. We would need to determine the appropriate behavioral mappings between the

behaviors and states of the horse and rider and those of the car and driver. How would a car indicate nervousness? What is the equivalent for a car to a horse's posture or skittishness? If a horse conveys its emotional state by rearing back and tensing its neck, what might the equivalent be for a car? What if suddenly your car reared back, lowering its rear end while raising the front, perhaps moving the front end left and right?

Natural signals akin to what the horse receives from its rider are actually being explored in research laboratories. Research scientists in the automobile companies are experimenting with measures of emotion and attention, and at least one automobile model sold to the public does have a television camera located on the steering column that watches drivers, deciding whether or not they are paying attention. If the automobile decides that a crash is imminent but the driver is looking elsewhere, it brakes.

Similarly, scientists are hard at work developing smart homes that monitor the inhabitants, assessing their moods and emotions, and adjusting room temperature, lighting, and background music. I've visited several of these experiments and observed the results. At one research facility at a European university, people were asked to play a stressful video game, then allowed to rest afterwards in a special experimental room equipped with comfortable chairs, friendly and aesthetically pleasing furniture, and specially equipped lighting designed to relax the inhabitants. When I tried it, I found it to be a calm and restful environment. The goal of the research was to understand how to develop room environments appropriate to a person's emotional state. Could a home relax its inhabitants automatically when it detected stress? Or perhaps the home could take on a zingy, upbeat mood with

bright lights, lively music, and warm colors when it determined that the inhabitants needed an energy boost.

Thinking for Machines Is Easy; Physical Actions Are Hard; Logic Is Simple, Emotion Difficult

“Follow me,” says Manfred Macx, the hero/narrator of Charles Stross’s science fiction novel *Accelerando*, to his newly purchased luggage. And follow him it does, “his new luggage rolling at his heels” as he turns and walks away.

Many of us grew up with the robots and giant brains of novels, movies, and television, where machines were all-powerful, sometimes clumsy (think of *Star Wars*’ C-3PO), sometimes omniscient (think of *2001*’s HAL), and sometimes indistinguishable from people (think of Rick Deckard, hero of the movie *Blade Runner*: is he human or replicant?). Reality is rather different from fiction: twenty-first century robots can’t conduct any meaningful communication with people; indeed, they are barely capable of walking, and their ability to manipulate real objects in the world is pathetically weak. As a result, most intelligent devices—especially in the home, where costs must be kept down and reliability and ease of use kept up—concentrate on mundane tasks such as making coffee, washing clothes and dishes, controlling lights, heating, and air conditioning, and vacuuming, mopping, and cutting the grass.

If the task is very well specified and the environment under control, then intelligent machines can indeed do a reasonable, informed job. They can sense temperature and moisture, as well as the amount of liquid, clothing, or food, and thus determine

when the laundry is dry or the food is cooked. The latest models of washing machines can even figure out what kind of material is being washed, how large the load is, and how dirty the clothes are, and adjust itself accordingly.

Vacuum cleaners and mops work as long as the pathway is relatively smooth and clear of obstacles, but the luggage that follows its owner in Stross's *Accelerando* is still beyond the capability of affordable machines. Nonetheless, though, this is precisely what a machine might be able to do, for it doesn't require real interaction with people: no communication, no safety-related issues, just follow along. What if someone tried to steal the freewheeling suitcase? It could be programmed to scream loudly at any attempt, and Stross tells us that it has learned the owner's "fingerprints, digital and phenotypic": thieves might be able to steal it, but they wouldn't be able to open it.

But could the luggage really make its way through crowded streets? People have feet, the better to step over and around obstacles, to go up and down stairs and over curbs. The luggage, with its wheels, would behave like a handicapped object, so it would need to seek out curb cuts at street intersections and ramps and elevators to maneuver within buildings. Human wheelchair users are often stymied: the wheeled luggage would be even more frustrated. And beyond curbs and stairs, navigating through city traffic would likely defeat its visual processing systems. Its ability to track its owner, avoid obstacles, and find paths navigable by a nonlegged device, while avoiding collisions with automobiles, bicycles, and people, would surely be compromised.

There is an interesting disjunction between the things people and machines find easy and hard. Thinking, which once was

held up as the pinnacle of human achievement, is the area in which machines have made the greatest progress, especially any thinking that requires logic and attention to detail. Physical actions, such as standing, walking, jumping, and avoiding obstacles, are relatively easy for people, but difficult if not impossible for machines. Emotions play an essential role in human and animal behavior, helping us judge what is good or bad, safe or unsafe, while also providing a powerful communication system for conveying feelings and beliefs, reactions and intentions among people. Machine emotions are simplistic.

Despite these limitations many scientists are still striving to create the grand dream of intelligent machines that will communicate effectively with human beings. It is in the nature of research scientists to be optimists, to believe that they are doing the most important activity in the world and, moreover, that they are close to significant breakthroughs. The result is a plethora of news articles, such as this one:

Researchers say robots soon will be able to perform many tasks for people, from child care to driving for the elderly.

Some of the country's leading robotics experts gathered here Saturday at the annual meeting of the American Association for the Advancement of Science to present their latest research and talk about a future rife with robots. . . .

[Y]our future could include: a huggable teddy bear that tutors your kids in Spanish or French; an autonomous car that drives you to work while you nap, eat or prepare your PowerPoint presentation; a Chihuahua-sized pet dinosaur that learns whether you like to cuddle, play or be left

alone; a computer that can move its screen to help your posture or match your task or mood; and a party-bot that greets your guests at the door, introduces them in case you've forgotten their names, and entertains them with music, jokes and finger food.

Many conferences are held to discuss progress on the development of "smart environments." Here is the wording of one invitation among the many that I receive:

Symposium on Affective Smart Environment. Newcastle Upon Tyne, UK.

Ambient Intelligence is an emerging and popular research field with the goal to create "smart" environments that react in an attentive, adaptive and proactive way to the presence and activities of humans, in order to provide the services that inhabitants of these environments request or are presumed to need.

Ambient Intelligence is increasingly affecting our everyday lives: computers are already embedded in numerous everyday objects like TV sets, kitchen appliances, or central heating, and soon they will be networked, with each other. . . . [B]io-sensing will allow devices to perceive the presence and state of users and to understand their needs and goals in order to improve their general living conditions and actual well-being.

Do you trust your house to know what is best for you? Do you want the kitchen to talk to your bathroom scale, or perhaps to have your toilet run an automatic urinalysis, sharing the results

with your medical clinic? And how, anyway, would the kitchen really know what you were eating? How would the kitchen know that the butter, eggs, and cream taken out of the refrigerator were for you, rather than for some other member of the household, or for a visitor, or maybe even for a school project?

Although monitoring eating habits wasn't really possible until recently, we can now attach tiny, barely visible tags on everything: clothes, products, food, items, even people and pets, so everything and everybody can be tracked. These are called radio frequency identification (RFID) tags. No batteries are required because these devices cleverly take their power from the very signal sent to them asking them to state their business, their identification number, and any other tidbits about the person or object they feel like sharing. When all the food in the house is tagged, the house knows what you are eating. RFID tags plus TV cameras, microphones, and other sensors equals "Eat your broccoli," "No more butter," "Do your exercises." Cantankerous kitchens? That's the least of it.

"What if appliances could understand what you need?" asked one group of researchers at the MIT Media Lab. They built a kitchen with sensors everywhere they could put them, television cameras, and pressure gauges on the floor to determine where people were standing. The system, they said, "infers that when a person uses the fridge and then stands in front of the microwave, he/she has a high probability of re-heating food." "KitchenSense," they call it. Here is their description:

KitchenSense is a sensor-rich networked kitchen research platform that uses CommonSense reasoning to simplify control interfaces and augment interaction. The system's

sensor net attempts to interpret people's intentions to create fail-soft support for safe, efficient and aesthetic activity. By considering embedded sensor data together with daily-event knowledge, a centrally-controlled OpenMind system can develop a shared context across various appliances.

If people use the refrigerator and then walk to the microwave oven, they have a "high probability of reheating food." This is highfalutin scientific jargon for guessing. Oh, to be sure, it is a sophisticated guess, but a guess it is. This example makes the point: the "system," meaning the computers in the kitchen, doesn't know anything. It simply makes guesses—statistically plausible guesses based on the designer's observations and hunches. But these computer systems can't know what the person really has in mind.

To be fair, even statistical regularity can be useful. In this particular case, the kitchen doesn't take any action. Rather, it gets ready to act, projecting a likely set of alternative actions on the counter so that if by chance one of them is what you are planning to do, you only have to touch and indicate yes. If the system doesn't anticipate what you had in mind, you can just ignore it—if you can ignore a house that constantly flashes suggestions to you on the counters, walls, and floors.

The system uses CommonSense (any confusion with the English term "common sense" is deliberate). Just as CommonSense is not really a word, the kitchen doesn't actually have any real common sense. It only has as much sense as the designers were able to program into it, which isn't much, given that it can't really know what is going on.

But what if you decide to do something that the house thinks is bad for you, or perhaps simply wrong? “No,” says the house, “that’s not the proper way to cook that. If you do it that way, I can’t be responsible for the result. Here, look at this cookbook. See? Don’t make me say ‘I told you so.’” This scenario has shades of *Minority Report*, the Steven Spielberg movie based upon the great futurist Philip K. Dick’s short story by that name. As the hero, John Anderton, flees from the authorities, he passes through the crowded shopping malls. The advertising signs recognize him, calling him by name, tempting him with offers of clothes and special sale prices just for him. A car advertisement calls out, “It’s not just a car, Mr. Anderton. It’s an environment, designed to soothe and caress the tired soul.” A travel agency entices him: “Stressed out, John Anderton? Need a vacation? Come to Aruba!” Hey, signs, he’s running away from the cops; he isn’t going to stop and buy some clothes.

Minority Report was fiction, but the technology depicted in the movie was designed by clever, imaginative experts who were very careful to depict only plausible technologies and activities. Those active advertising signs are already close to becoming a reality. Billboards in multiple cities recognize owners of BMW’s Mini Cooper automobile by the RFID tags they carry. The Mini Cooper advertisements are harmless, and each driver has volunteered and selected the phrases that will be displayed. But now that this has started, where will it stop? Today, the billboard requires its audience to carry RFID tags, but this is a temporary expedient. Already, researchers are hard at work, using television cameras to view people and automobiles, then to identify them by their gait and facial features or their model, year, color,

and license plate. This is how the City of London keeps track of cars that enter the downtown area. This is how security agencies expect to be able to track suspected terrorists. And this is how advertising agencies will track down potential customers. Will signs in shopping malls offer special bargains for frequent shoppers? Will restaurant menus offer your favorite meals? First in a science fiction story, then in a movie, then on the city streets: look for them at your nearest shops. Actually, you won't have to look: they will be looking for you.

Communicating with Our Machines: We Are Two Different Species

I can imagine it now: it's the middle of the night, but I can't sleep. I quietly get out of bed, careful not to wake up my wife, deciding that as long as I can't sleep, I might as well do some work. But my house detects my movement and cheerfully announces "good morning" as it turns on the lights and starts the radio news station. The noise wakes my wife: "Why are you waking me up so early?" she mumbles.

In this scenario, how could I explain to my house that behavior perfectly appropriate at one time is not so at another? Should I program it according to the time of day? No, sometimes my wife and I need to wake up early, perhaps to catch a morning flight. Or I might have a telephone conference with colleagues in India. For the house to know how to respond appropriately, it would need to understand the context, the reasoning behind the actions. Am I waking up deliberately? Does my wife still want to sleep? Do I really want the radio and the

coffeemaker turned on? For the house to understand the reasons behind my awakening, it would have to know my intentions, but that requires effective communication at a level not possible today or in the near future. For now, automatic, intelligent devices must still be controlled by people. In the worst of cases, this can lead to conflict. In the best of cases, the human+machine forms a symbiotic unit, functioning well. Here, we could say that it is humans who make machines smart.

The technologists will try to reassure us that all technologies start off as weak and underpowered, that eventually their deficits are overcome and they become safe and trustworthy. At one level they are correct. Steam engines and steamships used to explode; they seldom do anymore. Early aircraft crashed frequently. Today, they hardly ever do. Remember Jim's problem with the cruise control that regained speed in an inappropriate location? I am certain that this particular situation can be avoided in future designs by coupling the speed control with the navigation system, or perhaps by developing systems in which the roads themselves transmit the allowable speeds to the cars (hence, no more ability to exceed speed limits), or better yet, by having the car itself determine safe speeds given the road, its curvature, slipperiness, and the presence of other traffic or people.

I am a technologist. I believe in making lives richer and more rewarding through the use of science and technology. But that is not where our present path is taking us. Today we are confronting a new breed of machines with intelligence and autonomy, machines that can indeed take over, for us in many situations. In many cases, they will make our lives more effective,

more fun, and safer. In others, however, they will frustrate us, get in our way, and even increase danger. For the first time, we have machines that are attempting to interact with us socially.

The problems that we face with technology are fundamental. They cannot be overcome by following old pathways. We need a calmer, more reliable, more humane approach. We need augmentation, not automation.