


Our Robots, Ourselves: Robotics and the Myths of Autonomy

David A. Mindell

audiobook / *ebooks / Download PDF / ePub / DOC



DOWNLOAD



+

READ ONLINE

#502386 in eBooks 2015-10-13 2015-10-13 File Name: B00SI02AWK | File size: 65.Mb

David A. Mindell : Our Robots, Ourselves: Robotics and the Myths of Autonomy before purchasing it in order to gage whether or not it would be worth my time, and all praised Our Robots, Ourselves: Robotics and the Myths of Autonomy:

1 of 1 people found the following review helpful. Excellent introduction and case studiesBy Kristian GouldAn elegant and fascinating introduction to some key issues related to "autonomous" transportation systems and remote-operated technology. The case studies are excellent, although they do not go into as much depth as in Mindell's previous book on the Apollo project. This book is highly recommended for students, researchers, practitioners and others interested

in the field of human-machine interaction. It will clearly gain importance with the current debate around use of "autopilot" in cars. 1 of 1 people found the following review helpful. If you think self-driving vehicles are around the corner, you may want to read this book. By Richard E. Berger. Mindell has long worked with semi-autonomous systems and is familiar with their strengths and weaknesses. I thought (like many others) that self-driving vehicles will be common by the end of the 2020's. After reading his book, I have reconsidered that view - I think it will take longer, possibly much longer. I thought this book was a very interesting and informative read. My only complaint was that the message was sometimes obscured by the detailed recounting of some of the projects he has worked on. 0 of 1 people found the following review helpful. Terrific read, Important ideas. By gliterati29. Our Robots Ourselves doesn't just read like a true life adventure story - it will change the way we think about robots as they take a larger place in our lives. Mindell's expertise as a pioneer in robotics since the early years of undersea exploration into current space, aeronautics, automobile and other applications gives him the credibility to argue that full autonomy isn't the best direction but collaboration between people and robots will give us the safest and most efficient use of these modern marvels. And the plain language allows low techies like me to read it and understand the ideas. I'm planning to give this book to every teenaged techie I know to show how interesting and exciting a career in high tech can be.

It is required reading as we seriously engage one of the most important debates of our time. — Sherry Turkle, author of *Reclaiming Conversation: The Power of Talk in a Digital Age*. From drones to Mars rovers — an exploration of the most innovative use of robots today and a provocative argument for the crucial role of humans in our increasingly technological future. In *Our Robots, Ourselves*, David Mindell offers a fascinating behind-the-scenes look at the cutting edge of robotics today, debunking commonly held myths and exploring the rapidly changing relationships between humans and machines. Drawing on firsthand experience, extensive interviews, and the latest research from MIT and elsewhere, Mindell takes us to extreme environments — high atmosphere, deep ocean, and outer space — to reveal where the most advanced robotics already exist. In these environments, scientists use robots to discover new information about ancient civilizations, to map some of the world's largest geological features, and even to "commute" to Mars to conduct daily experiments. But these tools of air, sea, and space also forecast the dangers, ethical quandaries, and unintended consequences of a future in which robotics and automation suffuse our everyday lives. Mindell argues that the stark lines we've drawn between human and not human, manual and automated, aren't helpful for understanding our relationship with robotics. Brilliantly researched and accessibly written, *Our Robots, Ourselves* clarifies misconceptions about the autonomous robot, offering instead a hopeful message about what he calls "rich human presence" at the center of the technological landscape we are now creating. From the Hardcover edition.

"Mindell brings an altogether refreshing perspective to a field that can sometimes get lost in the 'what ifs'." — *Financial Times*. "Neither overly optimistic nor doomy, MIT professor Mindell offers a clear-eyed, reasoned overview of current and potential robotics achievements — and why the machines will always need us." — *Discover*. "Science fiction has become science as we find more and more uses for robots. The MIT professor explores the 'hidden world' of robotics and the controversial relationship between humankind and what it has created." — *Sacramento Bee*. "Mindell's ingenious and profoundly original book will enlighten those who prophesy that robots will soon make us redundant, and challenges us all to think more precisely and creatively about how machines can augment human potential." — David Autor, Professor of Economics, MIT. "Authoritative.... [Mindell] leaves us with a better understanding of what lies ahead for our daily lives." — *Kirkus*. "An expansively researched and enjoyably accessible treatment of robotic automation, recommended for readers of popular science and those with an interest in artificial intelligence and automation." — *Library Journal*. "A careful, measured extrapolation of contemporary technological trends." — *Booklist*. "A lucid, hype-free exploration of how robotic automation really works — in concert with human design, intention, and action." — Ian Bogost, Ivan Allen College Distinguished Chair in Media Studies and Professor of Interactive Computing, Scheller School of Business, Georgia Institute of Technology. "My thanks to the author for bringing scholarship and sanity to a debate which has run off into a magic la-la land in the popular press." — Rodney Brooks, founder Rethink Robotics and iRobot. "A must read for those who aspire to be effective contributors to the robotics of our future." — David R. Scott, Commander, Apollo 15. About the Author: David A. Mindell is Professor of Aeronautics and Astronautics and Distinguished Professor of the History of Engineering and Manufacturing at MIT. He has twenty-five years of experience as an engineer in undersea robotic exploration, as a veteran of more than thirty oceanographic expeditions, and as a pilot and engineer of autonomous aircraft. He is the award-winning author of *Iron Coffin: War Technology and Experience Aboard the USS Monitor* and *Digital Apollo: Human and Machine in Spaceflight*. He founded Humatics Corporation, which creates technologies to render autonomy transparent, safe, and trustworthy by transforming how robots and autonomous systems work in human environments. Excerpt. copy; Reprinted by

permission. All rights reserved. CHAPTER 1 Human, Remote, Autonomous

LATE IN THE NIGHT, HIGH ABOVE THE ATLANTIC OCEAN IN THE LONG, OPEN STRETCH

between Brazil and Africa, an airliner encountered rough weather. Ice clogged the small tubes on the aircraft's nose that detected airspeed and transmitted the data to the computers flying the plane. The computers could have continued flying without the information, but they had been told by their programmers that they could not. The automated, fly-by-wire system gave up, turned itself off, and handed control to the human pilots in the cockpit: thirty-two-year-old Pierre Cedric Bonin and thirty-seven-year-old David Robert. Bonin and Robert, both relaxed and a little fatigued, were caught by surprise, suddenly responsible for hand flying a large airliner at high altitude in bad weather at night. It is a challenging task under the best of circumstances, and one they had not handled recently. Their captain, fifty-eight-year-old Marc Debois, was off duty back in the cabin. They had to waste precious attention to summon him. Even though the aircraft was flying straight and level when the computers tripped off, the pilots struggled to make sense of the bad air data. One man pulled back, the other pushed forward on his control stick. They continued straight and level for about a minute, then lost control. On June 1, 2009, Air France flight 447 spiraled into the ocean, killing more than two hundred passengers and crew. It disappeared below the waves, nearly without a trace. In the global, interconnected system of international aviation, it is unacceptable for an airliner to simply disappear. A massive, coordinated search followed. In just a few days traces of flight 447 were located on the ocean's surface. Finding the bulk of the wreckage, however, and the black box data recorders that held the keys to the accident's causes, required hunting across a vast seafloor, and proved frustratingly slow. More than two years later, two miles deep on the seafloor, nearly beneath the very spot where the airliner hit the ocean, an autonomous underwater vehicle, or AUV, called Remus 6000 glided quietly through the darkness and extreme pressure. Moving at just faster than a human walking pace, the torpedo-shaped robot maintained a precise altitude of about two hundred feet off the bottom, a position at which its ultrasonic scanning sonar returns the sharpest images. As the sonars pinged to about a half mile out either side, the robot collected gigabytes of data from the echoes. The terrain is mountainous, so the seafloor rose quickly. Despite its intelligence, the robot occasionally bumped into the bottom, mostly without injury. Three such robots worked in a coordinated dance: two searched underwater at any given time, while a third one rested on a surface ship in a three-hour pit stop with its human handlers to offload data, charge batteries, and take on new search plans. On the ship, a team of twelve engineers from the Woods Hole Oceanographic Institution, including leader Mike Purcell, who spearheaded the design and development of the searching vehicles, worked in twelve-hour shifts, busy as any pit crew. When a vehicle came to the surface, it took about forty-five minutes for the engineers to download the data it collected into a computer, then an additional half hour to process those data to enable a quick, preliminary scroll-through on a screen. Looking over their shoulders were French and German investigators, and representatives from Air France. The mood was calculating and deliberate, but tense: the stakes were high for French national pride, for the airliner's manufacturer, Airbus, and for the safety of all air travel. Several prior expeditions had tried and failed. In France, Brazil, and around the world, families awaited word. Interpreting sonar data requires subtle judgment not easily left solely to a computer. Purcell and his engineers relied on years of experience. On their screens, they reviewed miles and miles of rocky reflections alternating with smooth bottom. The pattern went on for five days before the monotony broke: a crowd of fragments appeared, then a debris field—a strong signal of human-made artifacts in the ocean desert. Suggestive, but still not definitive. The engineers reprogrammed the vehicles to return to the debris and “fly” back and forth across it, this time close enough that onboard lights and cameras could take pictures from about thirty feet off the bottom. When the vehicles brought the images back to the surface, engineers and investigators recognized the debris and had their answer: they had found the wreckage of flight 447, gravesite of hundreds. Soon, another team returned with a different kind of robot, a remotely operated vehicle (ROV), a heavy-lift vehicle specially designed for deep salvage, connected by a cable to the ship. Using the maps created by the successful search, the ROV located the airliner's black box voice and data recorders and brought them to the surface. The doomed pilots' last minutes were recovered from the ocean, and investigators could now reconstruct the fatal confusion aboard the automated airliner. The ROV then set about the grim task of retrieving human remains. The Air France 447 crash and recovery linked advanced automation and robotics across two extreme environments: the high atmosphere and the deep sea. The aircraft plunged into the ocean because of failures in human interaction with automated systems; the wreckage was then discovered by humans operating remote and autonomous robots. While the words (and their commonly perceived meanings) suggest that automated and autonomous systems are self-acting, in both cases the failure or success of the systems derived not from the machines or the humans operating on their own, but from people and machines operating together. Human pilots struggled to fly an aircraft that had been automated for greater safety and reliability; networks of ships, satellites, and floating buoys helped pinpoint locations; engineers interpreted and acted on data produced by robots. Automated and autonomous vehicles constantly returned to their human makers for information, energy, and guidance. Air France 447 made tragically clear that as we constantly adapt to and reshape our surroundings, we are also remaking ourselves. How could pilots have become so dependent on computers that they flew a perfectly good airliner into the sea? What becomes of the human roles in activities like transportation, exploration, and warfare when more and more of the critical tasks seem to be done by machines? In the extreme view,

some believe that humans are about to become obsolete, that robots are "only one software upgrade away" from full autonomy, as *Scientific American* has recently argued. And they tell us that the robots are coming—coming to more familiar environments. A new concern for the strange and uncertain potentials of artificial intelligence has arisen out of claims that we are on the cusp of superintelligence. Our world is about to be transformed, indeed is already being transformed, by robotics and automation. Start-ups are popping up, drawing on old dreams of smart machines to help us with professional duties, physical labor, and the mundane tasks of daily life. Robots living and working alongside humans in physical, cognitive, and emotional intimacy have emerged as a growing and promising subject of research. Autonomy—the dream that robots will one day act as fully independent agents—remains a source of inspiration, innovation, and concern. The excitement is in the thrill of experimentation; the precise forms of these technologies are far from certain, much less their social, psychological, and cognitive implications. How will our robots change us? In whose image will we make them? In the domain of work, what will become of our traditional roles—scientist, lawyer, doctor, soldier, manager, even driver and sweeper—when the tasks are altered by machines? How will we live and work? We need not speculate: much of this future is with us today, if not in daily life then in extreme environments, where we have been using robotics and automation for decades. In the high atmosphere, the deep ocean, and outer space humans cannot exist on their own. The demands of placing human beings in these dangerous settings have forced the people who work in them to build and adopt robotics and automation earlier than those in other, more familiar realms. Extreme environments press the relationships between people and machines to their limits. They have long been sites of innovation. Here engineers have the freest hand to experiment. Despite the physical isolation, here the technologies' cognitive and social effects first become apparent. Because human lives, expensive equipment, and important missions are at stake, autonomy must always be tempered with safety and reliability. In these environments, the mess and busyness of daily life are temporarily suspended, and we find, set off from the surrounding darkness, brief, dream-like allegories of human life and technology. The social and technological forces at work on an airplane's flight deck, or inside a deep-diving submersible, are not fundamentally different from those in a factory, an office, or an automobile. But in extreme environments they appear in condensed, intense form, and are hence easier to grasp. Every airplane flight is a story, and so is every oceanographic expedition, every space flight, every military mission. Through these stories of specific people and machines we can glean subtle, emerging dynamics. Extreme environments teach us about our near future, when similar technologies might pervade automobiles, health care, education, and other human endeavors. Human-operated, remotely controlled, and autonomous vehicles represent the leading edge of machine and human potential, new forms of presence and experience, while drawing our attention to the perils, ethical implications, and unintended consequences of living with smart machines. We see a future where human knowledge and presence will be more crucial than ever, if in some ways strange and unfamiliar. And these machines are just cool. It's not alone in my lifelong fascination with airplanes, spacecraft, and submarines. Indeed, technological enthusiasm, as much as the search for practical utility, drives the stories that follow. It's no coincidence that similar stories are so often the subject of science fiction—something about people and machines at the limits of their abilities captures the imagination, engages our wonder, and stirs hopes about who we might become. This enthusiasm sometimes reflects a naive faith in the promise of technology. But when mature it is an enthusiasm for basic philosophical and humanistic questions: Who are we? How do we relate to our work and to one another? How do our creations expand our experience? How can we best live in an uncertain world? These questions lurk barely below the surface as we talk to people who build and operate robots and vehicles. Join me as I draw on firsthand experience, extensive interviews, and the latest research from MIT and elsewhere to explore experiences of robotics and automation in the extreme environments of the deep ocean and in aviation (civil and military) and spaceflight. It is not an imagination of the future, but a picture of today: I'll see how people operate with and through robots and autonomy and how their interactions affect their work, their experiences, and their skills and knowledge. Our stories begin where I began, in the deep ocean. Twenty-five years ago, as an engineer designing embedded computers and instruments for deep-ocean robots, I was surprised to find that technologies were changing in unexpected ways the work of oceanography, the ways of doing science, the meaning of being an oceanographer. The realization led to two parallel careers. As a scholar, I study the human implications of machinery, from ironclad warships in the American Civil War to the computers and software that helped the Apollo astronauts land on the moon. As an engineer, I bring that research to bear on present-day projects, building robots and vehicles designed to work in intimate partnership with people. In the stories that follow I appear in some as a participant, in others as an observer, and in still others as both. These years of experience, research, and conversation have convinced me that we need to change the way we think about robots. The language we use for them is more often from twentieth-century science fiction than from the technological lives we lead today. Remotely piloted aircraft, for example, are referred to as "drones," as though they were mindless automata, when actually they are tightly controlled by people. Robots are imagined (and sold) as fully autonomous agents, when even today's modest autonomy is shot through with human imagination. Rather than being threatening automata, the robots we use so variously are embedded, as are we, in social and technical networks. In the pages ahead, we will explore many examples of how we work together with our machines. It's the combinations

that matter. It is time to review what the robots of today actually do, to deepen our understanding of our relationships with these often astoundingly capable human creations. I argue for a deeply researched empirical conclusion: whatever they might do in a laboratory, as robots move closer to environments with human lives and real resources at stake, we tend to add more human approvals and interventions to govern their autonomy. My argument here is not that machines are not intelligent, nor that someday they might not be. Rather, my argument is that such machines are not inhuman. Let us name three mythologies of twentieth-century robotics and automation. First, there is the myth of linear progress, the idea that technology evolves from direct human involvement to remote presence and then to fully autonomous robots. Political scientist Peter W. Singer, a prominent public advocate for autonomous systems, epitomizes this mythology when he writes that “this concept of keeping the human in the loop is already being eroded by both policymakers and the technology itself, which are both rapidly moving toward pushing humans out of the loop.” Yet there is no evidence to suggest that this is a natural evolution, that the “technology itself,” as Singer puts it, does any such thing. In fact there is good evidence that people are moving into deeper intimacy with their machines. We repeatedly find human, remote, and autonomous vehicles evolving together, each affecting the other. Unmanned aircraft, for example, cannot occupy the national airspace without the task of piloting manned aircraft changing too. In another realm, new robotic techniques for servicing spacecraft changed the way human astronauts serviced the Hubble Space Telescope. The most advanced (and difficult) technologies are not those that stand apart from people, but those that are most deeply embedded in, and responsive to, human and social networks. Second is the myth of replacement, the idea that machines take over human jobs, one for one. This myth is a twentieth-century version of what I call the iron horse phenomenon. Railroads were initially imagined to replace horses, but trains proved to be very poor horses. Railroads came into their own when people learned to do entirely new things with them. Human-factors researchers and cognitive scientists find that rarely does automation simply “mechanize” a human task; rather, it tends to make the task more complex, often increasing the workload (or shifting it around). Remotely piloted aircraft do not replicate the missions that manned aircraft carry out; they do new things. Remote robots on Mars do not copy human field science; they and their human partners learn to do a new kind of remote, robotic field science. Finally, we have the myth of full autonomy, the utopian idea that robots, today or in the future, can operate entirely on their own. Yes, automation can certainly take on parts of tasks previously accomplished by humans, and machines do act on their own in response to their environments for certain periods of time. But the machine that operates entirely independently of human direction is a useless machine. Only a rock is truly autonomous (and even a rock was formed and placed by its environment). Automation changes the type of human involvement required and transforms but does not eliminate it. For any apparently autonomous system, we can always find the wrapper of human control that makes it useful and returns meaningful data. In the words of a recent report by the Defense Science Board, “there are no fully autonomous systems just as there are no fully autonomous soldiers, sailors, airmen or Marines.” To move our notions of robotics and automation, and particularly the newer idea of autonomy, into the twenty-first century, we must deeply grasp how human intentions, plans, and assumptions are always built into machines. Every operator, when controlling his or her machine, interacts with designers and programmers who are still present inside it—perhaps through design and coding done many years before. The computers on Air France 447 could have continued to fly the plane even without input from the faulty airspeed data, but they were programmed by people not to. Even if software takes actions that could not have been predicted, it acts within frames and constraints imposed upon it by its creators. How a system is designed, by whom, and for what purpose shapes its abilities and its relationships with the people who use it. My goal is to move beyond these myths and toward a vision of situated autonomy for the twenty-first century. Through the stories that follow, I aim to redefine the public conversation and provide a conceptual map for a new era. As the basis for that map, I will rely throughout the book on human, remote, and autonomous when referring to vehicles and robots. The first substitutes for the awkward “manned,” so you can read “human” as shorthand for “human occupied.” These are of course old and familiar types of vehicles like ships, aircraft, trains, and automobiles, in which people’s bodies travel with the machines. People generally do not consider human-occupied systems to be robots at all, although they do increasingly resemble robots that people sit inside. “Remote,” as in remotely operated vehicles (ROVs), simply makes a statement about where the operator’s body is, in relation to the vehicle. Yet even when the cognitive task of operating a remote system is nearly identical to that of a direct physical operator, great cultural weight is attached to the presence or absence of the body, and the risks it might undergo. In the most salient example, remotely fighting a war from thousands of miles away is a different experience from traditional soldiering. As a cognitive phenomenon, human presence is intertwined with social relationships. Automation is also a twentieth-century idea, and still carries a mechanical sense of machines that step through predefined procedures; “automated” is the term commonly used to describe the computers on airliners, even though they contain modern, sophisticated algorithms. “Autonomy” is the more current buzzword, one that describes one of the top priorities of research for a shrinking Department of Defense. Some clearly distinguish autonomy from automation, but I see the difference as a matter of degree, where autonomy connotes a broader sense of self-determination than simple feedback loops and incorporates a panoply of ideas imported from artificial intelligence and other disciplines.

And of course the idea of the autonomy of individuals and groups pervades current debates in politics, philosophy, medicine, and sociology. This should come as no surprise, as technologists often borrow social ideas to describe their machines. Even within engineering, autonomy means several different things. Autonomy in spacecraft design refers to the onboard processing that takes care of the vehicle (whether an orbiting probe or a mobile robot) as distinct from tasks like mission planning. At the Massachusetts Institute of Technology, where I teach, the curriculum of engineering courses on autonomy covers mostly "path planning"—how to get from here to there in a reasonable amount of time without hitting anything. In other settings autonomy is analogous to intelligence, the ability to make human-like decisions about tasks and situations, or the ability to do things beyond what designers intended or foresaw. Autonomous underwater vehicles (AUVs) are so named because they are untethered, and contrast with remotely operated vehicles (ROVs), which are connected by long cables. Yet AUV engineers recognize that their vehicles are only semiautonomous, as they are only sometimes fully out of touch. The term "autonomous" allows a great deal of leeway; it describes how a vehicle is controlled, which may well change from moment to moment. One recent report introduces the term "increasing autonomy" to describe its essentially relative nature, and to emphasize how "full" autonomy—describing machines that require no human input—will always be out of reach. For our purposes, a working definition of autonomy is: a human-designed means for transforming data sensed from the environment into purposeful plans and actions. Language matters, and it colors debates. But we need not get stuck on it; I will often rely on the language (which is sometimes imprecise) used by the people I study. The weight of this book rests not on definitions but on stories of work: How are people using these systems in the real world, experiencing, exploring, even fighting and killing? What are they actually doing? Focusing on lived experiences of designers and users helps clarify the debates. For example, the word "drone" obscures the essentially human nature of the robots and attributes their ill effects to abstract ideas like "technology" or "automation." When we visit the Predator operators' intimate lairs we will discover that they are not conducting automated warfare—people are still inventing, programming, and operating machines. Much remains to debate about the ethics and policy of remote assassinations carried out by unmanned aircraft with remote operators, or the privacy concerns with similar devices operating in domestic airspace. But these debates are about the nature, location, and timing of human decisions and actions, not about machines that operate autonomously. Hence the issues are not manned versus unmanned, nor human-controlled versus autonomous. The questions at the heart of this book are: Where are the people? Which people are they? What are they doing? When are they doing it? Where are the people? (On a ship . . . in the air . . . inside the machine . . . in an office?) The operator of the Predator drone may be doing something very similar to the pilot of an aircraft—monitoring onboard systems, absorbing data, making decisions, and taking actions. But his or her body is in a different place, perhaps even several thousand miles away from the results of the work. This difference matters. The task is different. The risks are different, as are the politics. People's minds can travel to other places, other countries, other planets. Knowledge through the mind and senses is one kind of knowledge, and knowledge through the body (where you eat, sleep, socialize, and defecate) is another. Which one we privilege at any given time has consequences for those involved. Which people are they? (Pilots . . . engineers . . . scientists . . . unskilled workers . . . managers?) Change the technology and you change the task, and you change the nature of the worker—in fact you change the entire population of people who can operate a system. Becoming an air force pilot takes years of training, and places one at the top of the labor hierarchy. Does operating a remote aircraft require the same skills and traits of character? From which social classes does the task draw its workforce? The rise of automation in commercial-airline cockpits has corresponded to the expanding demographics of the pilot population, both within industrialized countries and around the globe. Is an explorer someone who travels into a dangerous environment, or someone who sits at home behind a computer? Do you have to like living on board a ship to be an oceanographer? Can you explore Mars if you're confined to a wheelchair? Who are the new pilots, explorers, and scientists who work through remote data? What are they doing? (Flying . . . operating . . . interpreting data . . . communicating?) A physical task becomes a visual display, and then a cognitive task. What once required strength now requires attention, patience, quick reactions. Is a pilot mainly moving her hands on the controls to fly the aircraft? Or is she punching key commands into an autopilot or flight computer to program the craft's trajectory? Where exactly is the human judgment she is adding? What is the role of the engineer who programmed her computer, or the airline technician who set it up? When are they doing it? (In real time . . . after some delay . . . months or years earlier?) Flying a traditional airplane takes place in real time—the human inputs come as the events are happening and have immediate results. In a spaceflight scenario, the vehicle might be on Mars (or approaching a distant asteroid), in which case it might take twenty minutes for the vehicle to receive the command, and twenty minutes for the operator to see that the action has occurred. Or we might say that craft is landing "automatically" when actually we can think of it as landing under the control of the programmers who gave it instructions months or years earlier (although we may need to update our notions of "control"). Operating an automated system can be like cooperating with a ghost. These simple questions draw our attention to shifts and reorientations. New forms of human presence and action are not

trivial, nor are they equivalent—a pilot who risks bodily harm above the battlefield has a different cultural identity from one who operates from a remote ground-control station. But the changes are also surprising—the remote operator may feel more present on the battlefield than pilots flying high above it. The scientific data extracted from the moon may be the same, or better, when collected by a remote rover than by a human who is physically present in the environment. But the cultural experience of lunar exploration is different from being there. Let's replace dated mythologies with rich, human pictures of how we actually build and operate robots and automated systems in the real world. The stories that follow are at once technological and humanistic. We shall see human, remote, and autonomous machines as ways to move and reorient human presence and action in time and in space. The essence of the book boils down to this: it is not "manned" versus "unmanned"; that matters, but rather, where are the people? Which people? What are they doing? And when? The last, and most difficult questions, then, are: How does human experience change? And why does it matter?

CHAPTER 2 Sea CRAMPED BUT COMFORTABLE, THE SUBMARINE'S INTERIOR looked like a cross between a commercial airliner and a 1950s camper. Though it was 1997, the ambiance was Cold War diner—switches, glowing tubes, knobs, and handles, green paint, linoleum, and stainless steel appliances. A constant, deafening whoosh reminded me that the very air I breathed came from a machine. The navy crew of ten called instructions and technical lingo to one another as though they were piloting an airplane ("Sierra, this is Victor, mark time complete"). As in an airplane, two pilots sat facing forward, pilot on the left, copilot on the right. The space was so small that the captain's bed lay right on the floor behind them. I stood next to him, leaning across the captain's sleeping body to look over the pilot's shoulders. I was part of a team of engineers, oceanographers, and archaeologists that had joined this submarine, the U.S. Navy's NR-1, and its mother ship, MV Carolyn Chouest, on an expedition to search for ancient shipwrecks in the Mediterranean. The NR-1 was a vestige of an earlier time, formerly dedicated to secret missions against the Soviet Union, now applied to civilian science. Built in the 1960s as an experiment in how to make a small nuclear submarine, it was about 150 feet long and could stay submerged for long periods. In the 1980s it recovered parts of the space shuttle Challenger after its pieces fell into the ocean. Now we were about seventy miles northwest of Sicily, in the Tyrrhenian Sea, searching around a geological feature called Skerki Bank. It looked like nothing but ocean from above, but Skerki Bank harbored two large rocky reefs that jutted up to just below the surface. The nasty ground sowed trouble in the ancient world: its treacherous topography lay right in the main shipping route between Carthage (modern-day Tunisia) and Rome's port of Ostia, gouging the hulls of many a doomed trader. On this day's NR-1 dive, I was in charge of the submarine's search for the remains of these ancient vessels. On the mother ship, just a few hours earlier, I sat down with Robert Ballard, chief scientist and architect of the overall expedition, to plan the search. Ballard, best known as the discoverer of the wreck of the Titanic, is an expert at finding the remains of human disasters on the seafloor. Together we laid out a series of track lines—regular, specific runs on the seafloor to guide NR-1 across a broad area. "Stick to the track lines," Ballard admonished me. "Don't go chasing after every target that appears on the sonar—yours! I'll never finish the survey." After our planning, a quick boat ride delivered me to the sub. I stepped onto its black hull, just a few feet above the sea surface. The sub had a bright red conning tower or "sail" on top, about the height of a grown person. I climbed into a door on the side of the sail and descended through the hull via a narrow ladder. Once inside, a crewman closed the hatch behind me. The sky disappeared with a feeling of finality; I would not get out for days. I stood aside as the crew prepared to dive. Checks, calls, communications; in a flurry of hand-cranked valves the sub began to descend with a gentle downhill pitch. My bunk was on top of the narrow corridor. Surrounded by pipes and brackets, it had only a small opening at the foot end. From there I had to wriggle inside to get my head in place for sleeping. Once in position, I could not turn over. Lying on my back, a bunch of pipes hung right in front of my face, and a few inches behind them was the sub's steel hull. On the other side of that was, well, three thousand feet of water. The first time I slept up there I awoke feeling claustrophobic and had to get down immediately and walk around to relax. The second time, it seemed a little more cozy but still made me nervous. By night three it felt like home. After descending for a few minutes we reached the bottom just outside Skerki Bank, about three thousand feet deep, and began our survey, looking for telltale signs of shipwrecks. On its sides, NR-1 had "side scan" sonars, which could see a few hundred meters out to each side. But NR-1's main feature was the forward-looking sonar. Every couple of seconds, the sonar on the nose of the submarine assaulted the watery space with a ping of high-frequency sound, then collected the echoes and displayed them on a computer screen. The sonar was designed to look up under ice for possible Russian submarines. Mounted on the NR-1, it pointed downward and forward of the sub; it could see a soda can from three thousand yards (and we saw quite a few on the bottom of the Mediterranean). The trouble was, the sonar only displayed "targets"—fuzzy blobs of pixels. To make out what they were, the crew had to laboriously drive over to each target and closely observe it either out the window or with NR-1's many cameras. And NR-1 was quite slow; it could make just a knot or two across the bottom, about a human walking pace. If we detected something three thousand yards out on the sonar, it could take hours to drive over to it and have a look. About an hour after we began our dive, Scott, the lieutenant on navigation watch who also served as the sonar man, noticed something on the display. It was a target not more than a few pixels across, but Scott thought it might be man-made. It

had a denser inner area surrounded, halo-like, by a ring of less dense reflections. This is not what rocks look like on sonar. As we moved past the target, the position and appearance of the blob did not change, though the grazing angle of the sonar changed—another indication of something solid, substantial, and possibly of human origin. Scott recommended that we depart from our track line and approach. It was the start of what was supposed to be a two-day survey and already my leadership was being tested. Just an hour or two had passed since Ballard's admonition to stick to the track lines. But I had to trust the crew. If this departure turned out to be a wild-goose chase, then I'd earn the credibility to turn down future requests. I descended into the NR-1's viewing area, a cramped compartment on the bottom of the hull with small windows. We were traveling about forty feet off the seafloor at a leisurely pace. Outside I saw undifferentiated green, the color a result of NR-1's green lights. Squinting, I could make out the sand, and get a sense of motion only when a ripple or rock slid by to break the visual monotony. As we approached the mysterious target, I prepared to see a pile of rocks. Instead, what emerged out of the green filled me with awe. Ceramic jars from an ancient world, more than a hundred of them, lay about the ocean floor. They were scattered, but in an identifiable pattern in two distinct piles about ten meters apart. This was the site of an ancient shipwreck. Long ago, the wooden hull rotted away, leaving the cargo exposed just as it was stacked in the hull. Lead stocks from two lead anchors clearly identified the bow. The wreck was pristine, untouched and unseen since settling here more than two thousand years ago. As the first person to see it since the day it sank, I was moved by the magnitude of the passage of time, and by the power of physical presence to abridge that time. The U.S. Navy's NR-1 nuclear research submarine above Skerki D, the remains of a first-century BCE shipwreck 800 meters (3,000 feet) deep in the Mediterranean Sea. (COURTESY NATIONAL GEOGRAPHIC SOCIETY) I named our discovery Skerki D, a scientific-sounding way to describe the fourth known shipwreck on Skerki Bank. We announced it to our colleagues on the surface through an underwater telephone, a scratchy, uneven channel that, at its best, garbled voices like an old walkie-talkie. We carefully noted the position and took a lot of pictures. At the end of our survey, after about a day and a half, we planned to return to the surface. Formal, clear language squeaked over the underwater telephone: "Interrogative: what is weather on surface?" A gale was brewing up there, which would have made it unsafe for us to transfer back to the Carolyn Chouest. So we returned to Skerki D and took a few more pictures. NR-1 has wheels, so we just moved off the site a few hundred feet and planted the submarine on the bottom. There we sat, at three thousand feet, waiting for the weather to clear—for nearly two days, watching war movies in the tiny galley. Finally we got word that the weather was clearing, and we ascended as quickly as we had dived. I returned to the Carolyn Chouest feeling serene but excited about our successful hunting, only to find my shipboard colleagues green, a little seasick, and tired from a rough couple of days. We had indeed been in a different world, less than a mile away but straight down. What came next was a natural experiment comparing the emotional power of embodied experience to the cognitive power of remote presence. For I was not a native submariner but a robotics engineer. The amount of time I spent physically on the seafloor was dwarfed by the amount of time I spent remotely there, telepresent through the medium of remote robots and fiber-optic cables. My home technology was the remote robot Jason, built and run by the Deep Submergence Laboratory of the Woods Hole Oceanographic Institution (WHOI). The Volkswagen-sized Jason waited out the rough weather lashed to the deck of the Carolyn Chouest. As soon as the weather cleared and NR-1 got out of the way, we quickly tasked Jason to carry out an intense, computer-controlled survey of the wreck site. We sat on board the ship in a darkened, air-conditioned control room while Jason, connected to the surface via a high-bandwidth fiber-optic cable, descended the depths to the Skerki D site. We watched on video, monitored sensors, and frenetically programmed computers. On this particular dive, seven years of work came together: sensors, precision navigation systems, and computerized controls coordinated to hover Jason above Skerki D and move it at a snail's pace to run precise track lines, just a meter apart, above the wreck site. Sonars and digital cameras bounced sound and light off the wreck, gathering gobs of data and transmitting it to hard drives on the ship. An acoustic navigation system I had built monitored Jason's position with subcentimeter precision several times per second, giving exact location tags to all the data. Then engineers and graduate students set to work, compiling the images into extensive photomosaics of the site and assembling the sonar data into a high-precision topographical map. This map connected navigation, computers, sensors, and data processing into a single pipeline. We had done pieces of this before, but never all together, and never on such an interesting and important site. The robotic survey expanded and quantified what I had seen out the window of NR-1. Where the submarine provided a visceral experience of presence, the robot digitized the seafloor into bits. Then as we pored over the data from the comparative comfort of the surface ship, we explored the virtual site in detail, discovering a great deal about it that was not visible when I was "there." We could now say the wreck site was about twenty meters long by five meters wide, with two distinct piles of ancient jars called amphoras. Many of the amphoras lay in small craters, apparently scoured out just for them by thousands of years of gentle bottom currents. Most of the amphoras were quite varied in appearance, although three identical ones lay, almost as if they had been lashed together, in a single crater. The seafloor, apparently flat to my naked eye peering through the window, actually had a gentle crescent just a few centimeters high that marked the outline of Skerki's ship's hull, buried just below the mud line. When we showed the digital maps to one of the archaeologists on board, he exclaimed, "You've just done in four

hours what I spent seven years doing on the last site I excavated." Yet no scuba-diving archaeologist ever had a map nearly as detailed and precise as our map of Skerki D—in fact, it was the most precise map ever made of the ocean floor, albeit of a tiny square in the vast ocean. The Skerki D survey was the culmination of at least eight years of engineering. We had learned how to digitize the seafloor with ultra-high precision. That would change both what was possible in archaeology and in how we explore human history in the deep sea. We would now learn how to "excavate" an archaeological site without ever touching it. We would now learn how to do a new kind of archaeology focused on the deep water and ancient trade routes that connected civilizations. It would let us ask new questions. But not everyone would welcome the new methods.