Once upon a time, education, industry, and art were integrated in the work of a village artisan. By the time that I went to school, college-bound kids like me had to sit in rather sterile classrooms, while the kids taking up trades got to go to a vocational school that had all the cool stuff—machine tools, welders, electronic test equipment, and the like. At the time, this split seemed vaguely punitive for someone like me. I couldn't understand why an interest in making things was taken as a sign of lesser intelligence. It wasn't until I became an MIT professor and had an official excuse to buy those kinds of tools that I realized the problem predates my high school by quite a bit. The Kellys and Meejins and Shellys and Dalias of the world are correcting a historical error that dates back to the Renaissance.

In the fifteenth and sixteenth centuries, stiff paintings of religious subjects gave way to much more diverse images of breathtaking realism and luminosity. The advances that made this possible were first of all technical, including the introduction of (originally Flemish) oil paints and the development of mathematical techniques for establishing perspective. Oil-based paints dried more slowly, allowing more complex brushstrokes; the intensity of the colors led to thinner layers that better reflected light; and the viscosity of the paints improved their coverage on wood and canvas. At the same time, vision studies led to geometrical solutions for the problem of how to best project a three-dimensional scene onto a two-dimensional surface. These techniques were created by and for artists; Leonardo da Vinci, for example, was continually experimenting with new paints.

Improvements in paints and perspective would not have had the impact they did if not for the simultaneous development of artists to use them. Your average Middle Ages painter worked as an artisan in a guild, with the same (lack of) creative control as a carpenter. An aspiring painter would progress from apprentice to journeyman to master, finally gaining full admission to the guild by producing a "masterpiece." The real mastery of the guilds, however, was of the marketplace: they were very effective monopolists, controlling both the supply of and standards for skilled laborers. The work itself was done to detailed specifications drawn up by, say, a church that wanted an altarpiece illustrated with a particular scene.

The guild system began to break down under both the fragmentation of the crafts into increasingly narrow specialties and their aggregation into workshops that could produce increasingly complex, complete artifacts (and which formed the seeds for the later development of factories). But creative individuals were able to escape from the guild system because of another invention: customers. Artisans became artists when a population emerged that had both the discretionary income and the intellectual interest to acquire art.

Led by merchant families, most notably the Medicis in Florence, and the occasional municipality or pre-Enlightened monarch, a community of patrons began to emerge that bought art from and for individuals rather than (literally) dogmatic institutions. Michelangelo and Leonardo da Vinci started their careers as apprentices but ended up valued for their personal expression rather than their occupational productivity. Leonardo da Vinci ultimately represented just himself. He was not the CEO of a nascent da Vinci Industries, with a fiduciary

responsibility to its shareholders to maximize the brand's value (which is a good thing, otherwise his legacy might instead have been a line of Mona Lisa and Luigi dolls).

As remarkable as these new artists and their new materials and methods for painting were, their subject material was more significant still. Paintings began to represent the world of people rather than that of myths or gods. This human focus both reflected and shaped the defining intellectual movement of the Renaissance: humanism. What liberal arts students now study as the humanities emerged in the Renaissance as an expression of a growing sense of mastery by (selected) people over their world.

In Italy, humanism grew in part out of a fifteenth-century attempt to understand Roman and Greek ruins, both architectural and intellectual, an effort that today would be called reverse engineering. While much ancient knowledge and experience had been lost in the intervening centuries, the existing record of ruins and ancient documents provided a handy template for reconstructing a society that had worked pretty well for the Romans (other than the decline-and-fall issue). This enterprise benefited considerably when the Ottomans conquered Constantinople in 1453, freeing up a ready supply of fleeing Greek scholars. When they sought refuge in Italy they brought with them writings and knowledge that had long been lost to the West.

A second ingredient in the emergence of humanism arose as a reaction against the strictures and scriptures of the prevalent monastic, ecclesiastical seat of advanced education. While nominally still religiously observant, the growing urban mercantile economy and rule by civic authority brought a need for training future secular leaders with relevant skills. The human part of *humanism* comes from *studia humanitatis* ("studies of humanity"), referring to this shift in focus from immortal to mortal subjects, and was associated with a growing interest in how people were reflected in everything from portraiture to pedagogy.

These threads of humanism came together with the dusting off (sometimes quite literally) of a curriculum based around the four-part quadrivium (geometry, arithmetic, astronomy, music) and the three-part trivium (grammar, logic, rhetoric). These Latin terms refer back to four- and three-way road intersections, the latter notable as a place where people would congregate and share knowledge that naturally came to be known as trivial, or trivia. The trivium and the quadrivium together make up the seven "liberal arts." Both of these words warrant comment. "Liberal" in this sense is not the opposite of "conservative"; it referred to the liberation that the study of these subjects was thought to bring. And "art" did not mean just creative expression; it meant much more broadly the mastery that was developed over each of these domains. Liberal arts originally has this rather rousing meaning as a mastery over the means for personal liberation. They're now associated with academic study that is remote from applications, but they emerged in the Renaissance as a humanist pathway to power.

In 1513 Niccolò Machiavelli wrote *The Prince*, the classic (and infamous) guide to governance, on how to use rhetoric to win friends and manipulate people. Renaissance social engineering also gave birth to the concept of utopia, if not the reality. The term first appeared in a book of that name, written by Sir Thomas More in 1516; his utopian vision was very much a humanist paradise, governed by reason and based on a belief in the power of ideas. It was against this backdrop of the growing recognition of the practical importance of language and reasoning that familiarity with the liberal arts emerged as a modern notion of literacy. These skills became an expectation of any active participant in civil society.

Unfortunately, the ability to make things as well as ideas didn't make the cut; that was relegated to the *artes illiberales*, the "illiberal arts," that one pursued for mere economic gain. With art separated from artisans, the remaining fabrication skills were considered just mechanical production. This artificial division led to the invention of unskilled labor in the Industrial Revolution.

As with the revolution in painting in the Renaissance, this transition in industry was triggered in part by advances in materials, in this case

the use of iron and steel, which in turn both led to and benefited from the development of steam power. These developments in materials and power made possible modern machinery, most notably mechanized looms. These could produce much more cloth than traditional artisans could (from 50,000 pieces in 1770 England to 400,000 pieces in 1800), and thus could clothe many more people (from 8.3 million people in 1770 England to 14.2 million in 1821). Newly unemployed craft workers crowded into growing cities to seek employment operating the machines that would replace not only the jobs but also the skills of still more workers. Unintended consequences of this shift included a layer of black smoke covering earth and sky, generated from burning coal in the factories, and the epidemics of cholera, smallpox, typhoid, typhus, and tuberculosis that followed from packing people around the factories.

This new division of labor between people and machines became explicit with Joseph-Marie Jacquard's invention of the programmable loom, first demonstrated in Paris in 1801. He introduced an attachment that could read instructions on punched cards (more reliably than Florida's voters) to control the selection of shuttles containing colored threads and thereby program patterns into fabric.

Because the looms could now follow instructions, their operators no longer needed to. The job of the weaver was reduced to making sure that the loom was supplied with thread and cards. Lyon's silk weavers, threatened by this challenge to their livelihood, rather reasonably destroyed Jacquard's looms. But the looms won; commercial weaving turned from a skilled craft into menial labor.

The invention of industrial automation meant that a single machine could now make many things, but it also meant that a single worker who used to do many things now did only one. Thinking about how to make things had became the business of specialized engineers; the Ecole Polytechnique was set up in France in 1794 to train them, and in Britain there was an unsuccessful attempt to forbid the export of both its engineers and the machines they developed because of the perceived strategic importance of both.

Tellingly, in Britain, where the separation between art and artisans was furthest along, scientific progress suffered. The great acoustics discoveries of the nineteenth century occurred in France and Germany, where there was a lively exchange in workshops that made both musical and scientific instruments, rather than in England, where *handwerk* had become a pejorative term.

From there, the relative meaning of literacy diverged for machines, their designers, and their users. First, the machines. Around 1812, the mathematician Charles Babbage conceived that it would be possible to construct a machine to do the tedious job of calculating mathematical tables, and in 1823 he received government support to build his "Difference Engine." He failed to finish it (Babbage was also a pioneer in bad management), but the Difference Engine did produce one very useful output: the inspiration for the Analytical Engine. Babbage realized that an improved steam-powered engine could follow instructions on Jacquard's punched cards to perform arbitrary mathematical operations, and could change the operation by changing the cards rather than the machine. By the mid-1830s Babbage had failed to complete this new machine as well, limited by the timeless problems of underfunding and mismanagement, and by the available manufacturing technology that could not make parts with the complexity and tolerances that he needed.

Jacquard's punched cards reappeared in 1882 when Herman Hollerith, a lecturer at MIT who had worked for the U.S. Census Bureau as a statistician, sought a way to speed up the hand-tallying of the census. He realized that the holes in the cards could represent abstract information that could be recorded electrically. The result was the Hollerith Electric Tabulating System, which counted the 1890 census in a few months rather than the years that a hand tally would have required. The greater consequence of this work was the launch in 1896 of his Tabulating Machine Company, which in 1924 became IBM, the International Business Machines Corporation.

Information-bearing punched cards made machines more flexible in what they could do, but that didn't change anyone's notion of the nature of people versus machines. That challenge surfaced in an initially obscure paper published by the twenty-four-year-old Alan Turing in Cambridge in 1936. In "On Computable Numbers, with an Application to the Entscheidungsproblem," he tackled one of the greatest outstanding mathematical questions of his day, the *Entscheidungsproblem* ("decision problem") posed by the great mathematician David Hilbert in 1928: can there exist, at least in principle, a definite procedure to decide whether a given mathematical assertion is provable? This is the sort of thing that someone like Turing, employed in the business of proving things, might hope to be possible. His rather shocking answer was that it wasn't. Alonzo Church, who would become Turing's thesis adviser at Princeton, independently published the same conclusion in 1936, but Turing's approach was later considered by everyone (including Church) to be much more clever.

To make the notion of "procedure" explicit, Turing invented an abstract mechanism that he called an LCM, a *logical computing machine* (everyone else just called it a Turing machine). This device had a paper tape that could contain instructions and data, and a head that could move along the tape reading those instructions and interpreting them according to a fixed set of rules, and could then make new entries onto the tape. This sort of machine could follow a procedure to test the truth of a statement, but Turing was able to show that simple questions about the working of the machine itself, such as whether or not it eventually halts when given a particular set of instructions, cannot be answered short of just watching the machine run. This means that it might be possible to automate the solution of a particular problem, but that there cannot be an automated procedure to test when such an approach will succeed or fail.

This dramatic conclusion set a profound limit on what is knowable. The seemingly steady advance of the frontiers of knowledge had halted at a fundamentally unanswerable question: the undecidability of testing a mathematical statement. But Turing's solution contained an even greater consequence: He showed that the particular details of the design of the Turing machine didn't matter, because any one of them can emulate any other one by putting at the head of its tape a set of instructions describing how the other one works. For instance, a Mac Turing machine can use a PC Turing machine tape by starting it off with a PC specification written in Mac language. This insight, now called the Church-Turing thesis, is the key to machine literacy. Any machine that can emulate a Turing machine can solve the same problems as any other, because it is general enough to follow machine translation instructions. This property has since been shown to be shared by systems ranging from DNA molecules to bouncing billiard balls.

Turing's thoughts naturally turned to building such a universal computing device; in 1946 he wrote a "Proposal for Development in the Mathematics Division of an Automatic Computing Engine (ACE)" for the UK's National Physical Laboratory (NPL). Like Babbage, Turing proved himself to be better at proposing machines than building them, but those machines of course did get built by his successors. The story picks up across the Atlantic, where, shades of Babbage's math tables, the U.S. Army funded the construction of an all-electronic machine to be built with vacuum tubes to calculate artillery range tables. The ENIAC (Electronic Numerical Integrator and Computer) was publicly dedicated at the University of Pennsylvania in 1946. Its first calculations weren't range tables, though. They were something much more secret and explosive: mathematical models for the nuclear bomb effort at Los Alamos. Those calculations arrived via John von Neumann; getting him interested in computers was perhaps ENIAC's most important consequence.

Von Neumann is on the short list of the smartest people of the past century; those who knew him might say that he *is* the short list. He was a math wizard at Princeton, where he overlapped with Turing, and he was an influential government adviser. When von Neumann heard about the ENIAC through a chance encounter he planted himself at the University of Pennsylvania, recognizing how much more the ENIAC could do than calculate range tables. It was, after all, the first general-purpose programmable digital electronic computer.

It was also a rather clumsy first general-purpose programmable digital electronic computer. It tipped the scales at a svelte thirty tons, and for maximum operational speed it was programmed by plugboards that took days to rewire. It's charitable to even call it programmable. But his experience with this computer did lead von Neumann to propose to the Army Ordnance Department in 1945 construction of the EDVAC (Electronic Discrete Variable Computer), and in 1946 he elaborated on this idea in a memo, "Preliminary Discussion of the Logical Design of an Electronic Computing Instrument." He made the leap to propose that programs as well as data could be stored electronically, so that the function of a computer could be changed as quickly as its data. He proved to be a better manager than Babbage or Turing; the EDVAC was finished in 1952 (although the first stored-programs computers became operational at the universities of Manchester and Cambridge in 1949).

Having invented the modern architecture of computers, von Neumann then considered what might happen if computers could manipulate the physical world outside of them with the same agility as the digital world inside of them. He conceived of a "universal constructor," with a movable head like a Turing machine, but able to go beyond making marks to actually move material. Because such a thing was beyond the capabilities of the barely functional computers of his day, he studied it in a model world of "cellular automata," which is something like an enormous checkerboard with local rules for how checkers can be added, moved, and removed. Von Neumann used this model to demonstrate that the combination of a universal constructor and a universal computer had a remarkable property: self-reproduction. The computer could direct the constructor to copy both of them, including the program for the copy to make yet another copy of itself. This sounds very much like the essence of life, which is in fact what von Neumann spent the rest of his life studying. I'll return to this idea in "The Future" to look at the profound implications for fabrication of digital self-reproduction.

While von Neumann was thinking about the consequences of connecting a universal computer to machinery to make things, the first general-purpose programmable fabricator was actually being built at MIT. The Whirlwind computer was developed there in the Servomechanism Laboratory, starting in 1945 and demonstrated in 1951. Intended for the operation of flight simulators, the Whirlwind needed to respond to real-time inputs instead of executing programs submitted as batch jobs. To provide instantaneous output from the computer, the Whirlwind introduced displays screens. But if the computer could control a screen, that meant that it could control other things in real time. At the request of the air force, in 1952 the Whirlwind was connected to an industrial milling machine. The mechanical components for increasingly sophisticated aircraft were becoming too difficult for even the most skilled machinists to make. By having the Whirlwind control a milling machine, shapes were limited only by the expressive power of programs rather than by the manual dexterity of people. The machines of Babbage's day weren't up to making computers, but finally computers were capable of making machines. This in turn raised a new question: How could designers tell computers how to make machines?

The answer was the development of a new kind of programming language for doing what became known as computer-aided manufacturing (CAM) with numerically controlled (NC) machines. The first of these, Automatically Programmed Tools (APT), ran on the Whirlwind in 1955 and became available on IBM's 704 computer in 1958. It was a bit like the theoretical programs for a Turing machine that would specify how to move the read/write head along the abstract tape, but now the head was real, it could move in three dimensions, and there was a rotating cutting tool attached to it. APT is still in use, and is in fact one of the oldest active computer languages.

APT was a machine-centric representation: it described steps for the milling machine to follow, not results that a designer wanted. Real computer aid on the design side came from the next major computers at MIT, the TX–0 and TX–2. These were testbeds for computing with transistors instead of vacuum tubes, and sported a "light pen" that allowed the operator to draw directly on the display screen. In 1960 Ivan Sutherland, a precocious student supervised by Claude Shannon (inventor of the theory of information that forms the foundation of digital communications), used the combination of the TX–2 and the light pen to create the seminal "Sketchpad" program. Sketchpad let a designer sketch shapes, which the computer would then turn into precise geometrical figures. It was the first computer-aided design (CAD) program, and remains one of the most expressive ones.

The TX–2 begat Digital Equipment Corporation's PDP (Programmed Data Processor) line of computers, aimed at work groups rather than entire organizations. These brought the cost of computing down from one million dollars to one hundred thousand and then ten thousand dollars, sowing the seeds for truly personal computing and the PCs to come. The consequences of that personalization of computation have been historic. And limited.

Personal computers embody centuries of invention. They now allow a consumer to use a Web browser to buy most any product, from most anywhere. But online shopping is possible only if someone somewhere chooses to make and sell what's wanted. The technology may be new, but the economic model of mass production for mass markets dates back to the origin of the Industrial Revolution.

Unseen behind e-commerce sites are the computers that run industrial processes. Connecting those computers with customers makes possible what Stan Davis calls "mass customization," allowing a production line to, for example, cut clothes to someone's exact measurements, or assemble a car with a particular set of features. But these are still just choices made from among predetermined options. The real expressive power of machines that make things has remained firmly on the manufacturing rather than the consumer side.

Literacy has, if anything, regressed over time to the most minimal meaning of reading and writing words, rather than grown to encompass the expressive tools that have come along since the Renaissance. We're still living with the historical division of the liberal from the illiberal arts, with the belief that the only reasons to study fabrication are for pure art or profane commerce, rather than as a fundamental aspect of personal liberation.

The past few centuries have given us the personalization of expression, consumption, and computation. Now consider what would happen if the physical world outside computers was as malleable as the digital world inside computers. If ordinary people could personalize not just the content of computation but also its physical form. If mass customization lost the "mass" piece and become personal customization, with technology better reflecting the needs and wishes of its users because it's been developed by and for its users. If globalization gets replaced by localization.

The result would be a revolution that contains, rather than replaces, all of the prior revolutions. Industrial production would merge with personal expression, which would merge with digital design, to bring common sense and sensibility to the creation and application of advanced technologies. Just as accumulated experience has found democracy to work better than monarchy, this would be a future based on widespread access to the means for invention rather than one based on technocracy.

That will happen. I can say this so firmly because it is my favorite kind of prediction, one about the present. All of the technologies to personalize fabrication are working in the laboratory, and they are already appearing in very unusual but very real communities of users outside the laboratory. The stories of these technologies, and people, are the subject of the rest of this book.

Hard Ware

If the conception of universal programmed assemblers is one of the peaks of human invention, then the reality of the engineering software used in manufacturing today is somewhere in the pits. This sorry state of affairs can appear to present an insurmountable obstacle to an unwary would-be personal fabricator; however, the same kind of rapid-prototyping approach to hardware can also be applied to create software aimed at the needs of mortals rather than machines.

The intuitive drawing metaphor of Ivan Sutherland's original Sketchpad program has been replaced in state-of-the-art engineering software with obscure "user interfaces." That those words are even used to describe how computers appear and respond to their users indicates the problem: people, like printers, are considered to be a kind of computer peripheral requiring a compatible communication interface. There's been no compelling reason to make engineering software easy to use; these programs have been written by engineers, for engineers, who make a career out of using one of them. Here's what the on-screen controls for one high-end computer-aided design (CAD) package look like: