

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
4 January 2007 (04.01.2007)

PCT

(10) International Publication Number
WO 2007/001962 A2

(51) International Patent Classification:
A61K 48/00 (2006.01) C12N 15/09 (2006.01)
G06F 19/00 (2006.01)

(21) International Application Number:
PCT/US2006/023763

(22) International Filing Date: 19 June 2006 (19.06.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/692,327 20 June 2005 (20.06.2005) US

(71) Applicant (for all designated States except US):
KURZWEIL TECHNOLOGIES, INC. [US/US];
15 Walnut Street, Wellesley Hills, Massachusetts 02481
(US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **KURZWEIL, Raymond, C.** [US/US]; 203 Lake Avenue, Newton, Massachusetts 02161 (US).

(74) Agents: **MALONEY, Denis G.** et al; Fish & Richardson PC, PO Box 1022, Minneapolis, MA 55440-1022 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW

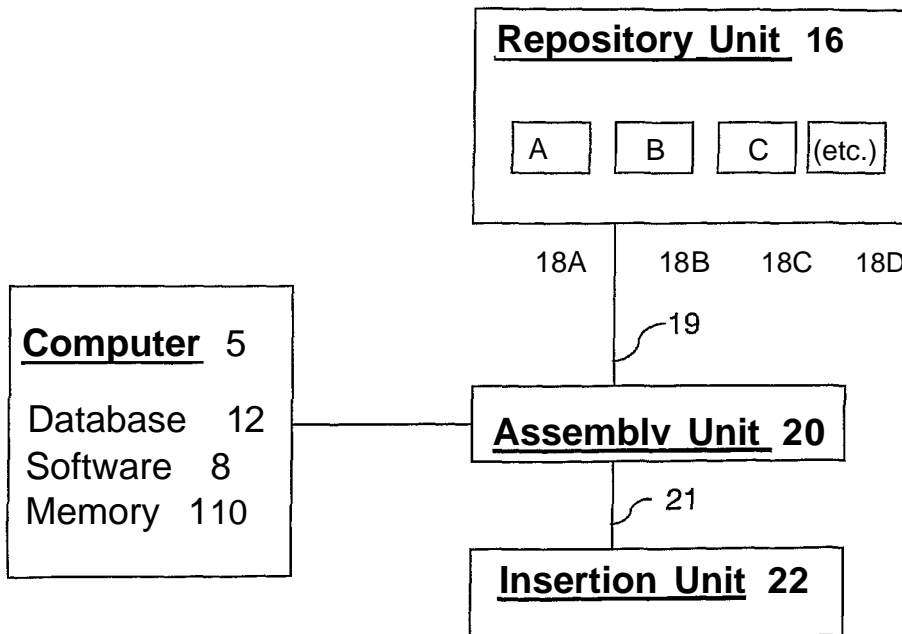
(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: SYSTEMS AND METHODS FOR GENERATING BIOLOGICAL MATERIAL



(57) Abstract: The invention relates to systems and methods for synthesizing biological material.

WO 2007/001962 A2

SYSTEMS AND METHODS FOR GENERATING BIOLOGICAL MATERIAL

CROSS REFERENCE TO RELATED APPLICATION

This application claims the benefit of the filing date of U.S. Provisional
5 Application No. 60/692,327, which was filed on June 20, 2005, the contents of which
(including all appendices) are hereby incorporated by reference to this description.

BACKGROUND

The invention relates to systems and methods for synthesizing biological material.
10 Methods of manipulating biological systems and processes through the use of
cell-free assays and nanotechnology have developed rapidly in recent years. For
example, cell-free systems for directing protein synthesis, RNA transcription, and RNA
splicing are known and more efficient and robust systems are evolving, such those that
include the use of very small sample volumes and that use nanotechnology.
15 Nanotechnology has been applied to the manipulation of cells and cellular processes,
including cell sorting based on the type, size, or function of a cell. Micro-fabricated
fluidic channels have been developed for sizing and sorting DNA molecules. Photonic
pressure has been used to transport cells over the length of defined fluidic channels. Bio-
chips have been developed which have the ability to operate with extremely small sample
20 volumes (on the order of nanoliters) and to perform analyses at much higher rates than
can be achieved by traditional methods. Many of the existing bio-chip and microfluidic
technologies use electrical, mechanical or forces to perform switching within the
microfluidic channels. Certain optical-based technologies describe the use of lasers to
define an optical path having an intensity gradient sufficient to propel the particles along
25 a path but sufficiently weak enough that the particles are not trapped in an axial direction.
Other lasers can interrogate particles to identify predetermined phenotypical
characteristics, and upon recognition of a particular phenotype, can deflect the particles in
along a different specified path. As the field of nanotechnology grows and the
understanding of the mechanics of biological processes evolves, the two fields will
30 continue to merge.

SUMMARY

In some embodiments, a system includes a computer configured to execute instructions for synthesizing biological material and an assembly unit electrically connected to the computer, the assembly unit being configured to synthesize the biological material based on the instructions executed by the computer.

Embodiments can include one or more of the following.

The system can include an insertion unit. The system can include a repository unit. The assembly unit further comprises one or more of an input channel and an output channel. The system can include separate device for a user to communicate wirelessly with the computer. The computer can include one or more of a memory unit, software, and a database. The database can include one or more of DNA sequence, RNA sequence and polypeptide sequence information. The repository unit can include one or more different types of monomelic biological components. The monomelic biological components can include one or more of a nucleotide, amino acid, and tRNA. Each tRNA can be attached to an amino acid.

The assembly unit can include one or more of a polymerase or a ribosome. The assembly unit can include a robot that mimics the activity of a polymerase or a ribosome. The biological material can be a nucleic acid or polypeptide. The biological material can be an RNA.

The insertion unit can be attached to the assembly unit. The system can be coated with a biocompatible material.

In some embodiments, a computer configured to execute instructions for synthesizing biological material, a central unit responsive to execution of the instructions to control an assembly unit, and an assembly unit electrically connected to the central unit, the assembly unit being configured to synthesize the biological material based on the instructions executed by the computer.

Embodiments can include one or more of the following.

The system can include a repository unit. The assembly unit can include one or more of an input channel and an output channel. The computer can be separate from the central unit and the assembly unit. The computer can reside outside of the cell and the central unit and the assembly unit reside inside the cell. The computer can reside within

the central unit, and the central unit and the assembly unit can reside inside the cell. The system can include a separate device for a user to communicate wirelessly with the computer. The computer can include one or more of a transmitter, software, and a database. The central unit can include one or more of a memory, a receiver, an engine,
5 and an antenna. The database can include one or more of DNA sequence, RNA sequence and polypeptide sequence information.

The system can include a repository unit, and the repository unit comprises one or more different types of monomeric biological components. The monomeric biological components can include one or more of a nucleotide, amino acid, and tRNA. Each tRNA
10 can be attached to an amino acid. The assembly unit can include one or more of a polymerase or a ribosome. The assembly unit can include a robot that mimics the activity of a polymerase or a ribosome. The biological material can be a nucleic acid or polypeptide. The biological material can include an RNA. The system can be coated with a biocompatible material.

In some embodiments, a method includes synthesizing a biological material and
15 introducing the biological material into a cell. The biological material is synthesized by a system that includes a computer configured to execute instructions for synthesizing biological material and an assembly unit electrically connected to the computer, the assembly unit being configured to synthesize the biological material based on the
20 instructions executed by the computer.

Embodiments can include one or more of the following.

The system further can include an insertion unit and a repository unit. The system further can include one or more of an input channel and an output channel on the assembly unit. The step of synthesizing the biological material can be initiated by a
25 signal from a user operating a device separated from the system, wherein the user uses the device to send a signal to a receiver located in the computer of the system.

In some embodiments, a method includes synthesizing a biological material and introducing the biological material into a cell. The biological material can be synthesized by a system that includes a computer configured to execute instructions for synthesizing
30 biological material, a central unit responsive to execution of the instructions to control an assembly unit, and an assembly unit electrically connected to the computer, the assembly

unit being configured to synthesize the biological material based on the instructions executed by the computer.

Embodiments can include one or more of the following.

The system can include a repository unit. The method can also include, prior to
5 the synthesizing step, of putting at least the central unit and the assembly unit of the system inside the cell. The system further can include one or more of an input channel and an output channel on the assembly unit. Putting the central unit and the assembly unit of the system into the cell can include electroporation, microinjection, or a lipophilic carrier. Synthesizing the biological material can be initiated by a signal from a user
10 operating a device separated from the system, wherein the user uses the device to send a signal to a receiver located in the central unit of the system. The computer can reside inside the central unit.

In some embodiments, a cell comprising a system includes a central unit responsive to instructions to control an assembly unit and an assembly unit electrically
15 connected to the central unit, the assembly unit being configured to synthesize biological material based on the instructions. The instructions are executed by a computer.

Embodiments can include one or more of the following.

The computer can reside inside the central unit. The computer can reside outside the central unit and outside the cell. The system can include a repository unit attached to
20 the assembly unit. The assembly unit can include one or more of an input channel and an output channel on the assembly unit. The computer can include one or more of a transmitter, software, and a database. The central unit can include one or more of a memory, a receiver, an engine, and an antenna. The repository unit can include one or more different types of monomelic biological components. The monomeric biological
25 components can include one or more of a nucleotide, amino acid, and tRNA. Each tRNA can be attached to an amino acid. The cell can originate from a mammal. The cell can originate from a human, mouse, rat, monkey, dog, cat, or rabbit. The cell can be in a tissue of a human.

In some embodiments, a method of treating a human includes administering a
30 system to the human. The system includes a central unit responsive to instructions to control an assembly unit and an assembly unit electrically connected to the central unit,

the assembly unit being configured to synthesize biological material based on the instructions. The instructions are executed by a computer.

Embodiments can include one or more of the following.

The computer can reside inside the central unit. The computer can reside outside
5 the central unit and outside the cell. The system can include a repository unit attached to the assembly unit. The assembly unit can include one or more of an input channel and an output channel. The computer can include one or more of a transmitter, software, and a database. The central unit can include one or more of a memory, a receiver, an engine, and an antenna. The database can include one or more of DNA sequence information,
10 KNA sequence information, and polypeptide sequence information. The repository unit can include one or more different types of monomeric biological components. The monomeric biological components can include one or more of a nucleotide, amino acid, and tRNA. Each tRNA is attached to an amino acid. The human can have a cancer, a tissue disorder, or a disorder of the nervous system. The system can be administered by
15 tissue graft, microprojectile technology, or by a lipophilic carrier.

In some embodiments, a method of treating a human includes administering a cell comprising. The system includes a central unit responsive to instructions to control an assembly unit and an assembly unit electrically connected to the central unit, the assembly unit being configured to synthesize biological material based on the
20 instructions. The instructions are generated by a computer.

Embodiments can include one or more of the following.

The computer can reside inside the central unit. The computer can reside outside the central unit and outside the human. The system includes a repository unit attached to the assembly unit. The assembly unit of the system can include one or more of an input
25 channel and an output channel on the assembly unit. The computer can include one or more of a transmitter, software, and a database. The central unit can include one or more of a memory, a receiver, an engine and an antenna. The database can include DNA sequence information, RNA sequence information, and polypeptide sequence information. The repository unit can include one or more different types of monomeric
30 biological components. The monomeric biological components can include one or more of a nucleotide, amino acid, and tRNA. Each tRNA can be attached to an amino acid.

The human can have a cancer, a tissue disorder, or a disorder of the nervous system. The cell can be administered by tissue graft.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and
5 advantages of the invention will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

FIG. 1 is a block diagram of a biological material synthesizing system such as a system residing outside of a cell.

10 FIG. 2 is a block diagram of a database.

FIG. 3 is a block diagram of a biological material synthesizing system, such as a system including components that can reside inside a cell.

FIG. 4 is flow chart of a process for generating biological material using the system of FIG. 2.

15

DETAILED DESCRIPTION

Described are systems and methods for synthesizing biological materials.

Referring to FIG. 1, the systems include at least a computer 5 and an assembly unit 20. The computer 5 is configured to generate instructions for the synthesis of
20 biological material, and includes, for example, a database 12, software 8, and memory 110. The assembly unit 20 is electronically connected to the computer 5 and is configured to synthesize the biological material based on the instructions received from the computer. The assembly unit 20 stores machinery for synthesizing the biological material and can be located inside or outside of a cell. When the assembly unit is located
25 outside of the cell, the assembly unit can be physically attached to the computer, and also to a repository unit 16 containing one or more different types of monomeric biological components. The repository unit 16 can be attached to the assembly unit 20 at the site of an input channel on the assembly unit 20. The assembly unit 20 uses the biological components stored in the repository unit for the synthesis of the biological material. The
30 assembly unit 20 can also be attached to an insertion unit 22, such as at the site of an output channel on the assembly unit. Synthesized biological material can pass from the

assembly unit 20 into the cell through the insertion unit 22. When the assembly unit 20 is located inside the cell, the assembly unit can be attached to a central unit 10, which is also located within the cell. The central unit can include the computer 5, or the central unit can be separated from the computer, and the computer can be stored, for example, outside the cell. The central unit can include a memory 110, receiver 14, engine 108 and antenna 102. A computer 5, located outside the cell, can include a database 12, software 8, and a transmitter 6 for transmitting instructions regarding biological synthesis from the computer to the central unit. The synthesized biological material can pass from the assembly unit directly into the cell, such as through an output channel. The biological materials synthesized by the systems featured herein can be multimeric molecules, such as nucleic acids (*e.g.*, DNA or RNA) and polypeptides. A system that synthesizes biological material or a cell containing such a system can be used to treat a human who has a disorder, such as cancer or a neurological disorder.

Systems for synthesizing biological material.

In general the computer 5 provides instructions to the assembly unit 20 indicating how to generate biological material. The assembly unit 20 generates the biological material using monomeric components (*e.g.*, monomeric components 18a-18d) stored in the repository unit 16. In some embodiments, the system operates outside the cell, and further includes an insertion unit 22. When the system functions outside of the cell, the biological material synthesized within the assembly unit passes through the insertion unit 22 and into the cell. The insertion unit resembles microinjection apparatus known in the art. For example, the insertion unit can include an injection pipet with an external diameter of about 1 micrometer (*e.g.*, about 0.6 micrometer, about 0.8 micrometer, about 1.0 micrometer, about 1.2 micrometer, about 1.4 micrometer) and tubing connecting the assembly unit and the injection pipet. The tubing can have an external diameter of about 60-70 micrometers (*e.g.*, 62, 64, 66, 68 micrometers). An insertion unit includes a hollow portion that can hold biological material generated by the assembly unit. The assembly unit can be in fluid communication with the insertion unit such that biological material can flow from the assembly unit to the insertion unit. The insertion unit also includes a tip with an opening disposed at the tip. The tip is configured to pierce the membrane of

cell without permanently damaging the cell. The insertion unit inserts the biological material stored in the hollow portion into the cell by flowing the biological material through the opening in the tip.

Still referring to FIG. 1, in order to generate the biological material the computer 5
5 communicates information with the assembly unit 20. In general, the computer 5 includes software 8 that provides instructions to the assembly unit 20. The software 8 includes a user interface that allows a user to select the type of biological material to be generated by the assembly unit 20. Based on the type of biological material the user selects to generate, the software 8 interfaces with a database 12 to determine an
10 appropriate set of instructions to send to the assembly unit 20.

For example, as shown in FIG. 2, the database 12 can store information regarding cell phenotype 35, the type or unit of biological material to be synthesized 40, the monomeric components 42 of the biological material, and the assembly instructions 44 for the material. As depicted in FIG. 2 the database includes a table with fields
15 representing the stored data, however other types of databases such as relational databases, or flat file systems, etc. could be used. The computer can maintain the information of a human's genetic code, altered to replace deleterious information (*e.g.*, deleterious mutations) with benign or beneficial information. By this approach, cells that are genetically programmed to synthesize a malfunctioning or nonfunctioning
20 polypeptide, or programmed not to synthesize a particular essential polypeptide, can be manipulated to express the properly functioning polypeptide by execution of the program by the computer and using data stored in the database.

Referring to FIG. 3, an assembly unit 20 can include an input channel 19 and an output channel 21. An assembly unit located inside a cell may or may not be attached to
25 a repository. The assembly unit 20 can harvest the biological building blocks (*e.g.*, nucleotides, amino acids) directly from the cell cytoplasm and through the input channel 19. Synthesized biological material is released from the assembly unit and into the cell through the output channel 21.

The biological material synthesizing system can be programmed such that the
30 assembly unit 20 receives instructions for synthesizing one particular type of biological material (*e.g.*, nucleic acid, polypeptide). Alternatively, the assembly unit can receive

instructions for synthesizing more than one type of biological material. For example, the computer can be programmed such that the assembly unit 20 receives instructions for synthesizing all the polypeptides necessary to support cell function. The computer can be programmed and re-programmed by a user using a device located outside of the cell. The device can be a wireless device, such as a second computer or a remote control device.

A repository unit can include, for example, nucleotides, such as deoxyribonucleotides for assembling deoxyribonucleic acid (DNA), or ribonucleotides for assembling ribonucleic acid (RNA), or amino acids for assembling polypeptides. DNA is a polymer of deoxyribonucleotide subunits and RNA is a polymer of ribonucleotide subunits. A nucleotide is a nitrogenous base (*e.g.*, a purine, a pyrimidine), a sugar (*e.g.*, a ribose, a deoxyribose), and one or more phosphate units. Table 1 lists exemplary nucleotides and amino acids that can be included in the repository unit. The repository unit can include a stock of amino acids attached to tRNA. Ribosomes within the assembly unit, or machinery that mimics a ribosome, can use the stock of tRNA to assemble a polypeptide by a mechanism similar to that which occurs *in vivo*.

Table 1. Exemplary monomeric components of the repository unit of a system.

<u>Component Class</u>	<u>Monomeric Component</u>
Deoxynucleotides	Deoxyadenosine 5'-triphosphate Deoxyguanosine 5'-triphosphate Deoxycytodine 5'-triphosphate Deoxythymidine 5'-triphosphate
Ribonucleotides	adenosine 5'-triphosphate guanosine 5'-triphosphate cytodine 5'-triphosphate uridine 5'-triphosphate
Amino Acids	Alanine Asparagine Cysteine Aspartic acid Glutamic acid

	Phenylalanine
	Glycine
	Histidine
	Isoleucine
	Lysine
	Leucine
	Methionine
	Proline
	Glutamine
	Arginine
	Serine
	Threonine
	Valine
	Tryptophan
	Tyrosine

The repository unit 16 can also include carbohydrates for attachment to the polypeptides. Alternatively, necessary carbohydrates can be attached to the polypeptide by the endogenous cellular machinery after the polypeptide passes through the assembly unit output channel.

The assembly unit 20 can include machinery for synthesizing biological material. The assembly unit 20 can include enzymes (*e.g.*, RNA polymerases, ribosomes) to facilitate RNA and polypeptide synthesis. Alternatively, or in addition, the assembly unit 20 can include machinery (*e.g.*, manmade machinery) that mimics the endogenous cellular machinery.

The repository unit 16 and assembly unit 20 can be fluid-filled with salts and buffer to provide an environment that mimics the interior of a cell. Such an environment facilitates the integrity of the molecular structures, including secondary and tertiary polypeptide structures formed as the amino acids are linked together in the assembly unit. The salts can include, for example, potassium, magnesium, calcium, zinc, ammonium or sodium salts, while suitable buffers include MES, Tris, HEPES, MMT and the like. The

fluid can also include serum, such as bovine fetal serum, to provide additional components to support assembly of the biological materials. Other suitable components include reducing agents, such as dithiothreitol (DTT), chelating agents such as EDTA, and polymers such as polyethylene glycols (PEGs). The fluid within a system can also
5 contain antibacterial and antifungal agents to prevent contamination. Appropriate antibacterial agents include but are not limited to streptomycin and penicillin, and fungizone is an example of an appropriate antifungal agent.

A technique for assembling the biological material can be programmed into the computer 5, and the programmable interface and the internal operating circuitry and/or
10 the signal processor, which may be one or more of a microprocessor or nanoprocessor, can be configured to allow adjustable and/or selectable operational configurations of the device to operate in the desired feedback mode or modes. The computer 5 can communicate with the assembly unit directly (FIG. 1) or can communicate wirelessly with a central unit 10 (FIG. 3). The computer and assembly unit can operate outside of
15 the cell (see, e.g., FIG. 1), or the system, or parts of the system, can be introduced into the interior of the cell, where it can assemble the biological material and distribute the material directly into the interior of the cell (see, e.g., FIG. 3).

Referring to FIG. 3, a wireless systems can include a computer 5, a central unit 10, and an assembly unit 20. The computer can reside inside the cell (such as inside a central
20 unit 10), or the computer 5 can reside outside the cell. A computer residing outside the cell can send instructions to the central unit that resides inside the cell. Due to the limited size of the central unit, some or all information processing can occur in a processor located on the computer 5. For example, the database 12 can be stored on the computer 5 and a reduced set of data or instructions (e.g., and executable instruction set) can be
25 wirelessly transmitted from a transmitter 6 on computer 5 to a receiver 14 on the central unit 10. The central unit 10 stores the received instructions or computer executable code in a memory 110. During operation, the central unit 10 executes the instructions stored on the memory 110 and sends instructions to the assembly unit 20 for generation of biological material.

30 The processor can be a microprocessor, a nanoprocessor, or a set of micro-engines. In general, a microprocessor is a computer processor on a microchip. A

nanoprocessor can be a processor having limited memory for executing a reduced set of instructions.

In general, a processor can be implemented in digital electronic circuitry, in computer hardware, firmware, software, or in combinations of these. The processor
5 described herein can be implemented as a computer program product, *e.g.*, a computer program tangibly embodied in an information carrier, *e.g.*, in a machine-readable storage device or in a propagated signal, for execution by, or to control the operation of, a data processing apparatus, *e.g.*, a processing device, a computer, or multiple computers. A computer program can be written in any form of programming language, including
10 compiled, assembled, or interpreted languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program can be deployed to be executed on one computer or on multiple computers at one site or distributed across multiple sites and interconnected by a communication network.

15 The system can be coated in a material that is biocompatible, such as a coating formed from polyurethane, or an amorphous titanium nitride. A material is biocompatible if it can come in contact with at least one part of the body without causing significant health hazards.

Referring to FIG. 4, a process 140 for modulating cell activity is shown. The
20 process 140 includes monitoring 142 at least one condition (*e.g.*, pH, hypotonicity, temperature) in a cell. The information about the cell is processed 144 by the computer, and a transmitter 6 transmits 148 the information to a central unit 10. The monitoring step can include monitoring a condition visually, such as by monitoring a condition of the human, or the monitoring step can include measuring conditions within the cell by
25 sensors located on one or more components of the system. The central unit receives the information and sends instructions to the assembly unit in the cell. The central unit can be located in the cell, or the central unit can reside outside the cell (*e.g.*, included in the computer 5). A wireless device, such as a remote control device, located outside the cell can be used to send instructions to the central unit 10 in the cell to direct the synthesis of
30 particular biological material. The assembly unit executes 152 the instructions and

generates biological material. The assembly unit then deposits 154 the material into the cell.

Uses. The system and cells featured can be used to treat a disorder, such as a proliferative disorder (*e.g.*, a cancer). For example, to eliminate unwanted cells (*e.g.*,
5 tumor cells) in a human, the computer can be programmed to synthesize a protein that is toxic to a cell, such as an α -sarcin polypeptide. The toxic polypeptide can mimic a polypeptide of any origin, such as of mammalian (*e.g.*, human) origin, bacterial origin, or fungal origin. Alternatively, the computer can be programmed to synthesize a double-stranded RNA (dsRNA), such as a short hairpin RNA, that can downregulate gene
10 expression by hybridizing to an endogenous RNA of the human, effectively shutting down translation of the endogenous RNA by the process of RNA interference. The result is the death of the cell and subsequently, the death of the tumor. In another alternative, the computer can be programmed to synthesize a single-stranded antisense RNA or microRNA, that can downregulate gene expression by hybridizing to an endogenous
15 RNA of the human.

The computer can be programmed to synthesize a dsRNA for the downregulation of any gene, and a computer can be programmed according to the features of the target cell. For example, the computer can be programmed to synthesize a dsRNA that targets a Src gene, such as for treating tumors of the colon. A computer programmed to synthesize
20 a dsRNA that targets the RAS gene can be used to treat tumors of the pancreas, colon or lung, and a computer programmed to synthesize a dsRNA that targets the c-MYC gene can be used to treat a neuroblastoma.

A system can be introduced into the unwanted cell by any method described below. For example, microprojectile technology can be used to propel a system into the
25 cells of a tumor mass.

The systems can be delivered to a human through use of a tissue graft. For example, cells can be cultured *in vitro* for use in a tissue graft. Before transferring the tissue to a human, the systems are introduced into the cells of the graft, such as through a liposomal carrier, or by an electrical pulse. The systems delivered to the tissue graft can
30 be programmed to synthesize a therapeutic biological component, *e.g.*, a therapeutic nucleic acid or polypeptide. In particular, a therapeutic polypeptide synthesized by the

system can be secreted by the cells of the tissue graft, where they are taken up by neighboring cells in need of therapeutic polypeptides. The tissue grafts can be applied to diseased or damaged tissue, *e.g.*, to treat burns or diseased organs, such as diseased heart, liver, or kidney tissue.

5 A system can replace a cell's nucleus. For example, the nucleus can be removed, such as by micromanipulation, and a system can then be injected into the cell. The system is programmed with the information needed to synthesize the biological material and to maintain cell growth and survival. For example, adult neural cells can be subjected to an exchange of a nucleus for a system. The system is programmed with
10 information regarding proteins required for neural cell survival and neurite outgrowth. Neural cells carrying a system can be transplanted into patients having a neurological disorder, such as due to genetic disposition or trauma, to replace nerve function. For example, the cells can be transplanted into or near the spinal cord of paraplegic patients to restore function to the central nervous system, and consequently to improve or restore
15 mobility.

 A system can be used to treat a human having a variety of different disorders. As discussed above, a human having a cancer can be treated with a system. For example, the human can have colon cancer, breast cancer, pancreatic cancer, lung cancer, liver cancer, gall bladder cancer, endometrial cancer, a glioblastoma, a squamous cell carcinoma,
20 ovarian cancer, prostate cancer, Ewing Sarcoma, myxoid liposarcoma, leukemia, an adenocarcinoma, and the like.

 A system can be also be used to treat a human who experiences acute or chronic pain, or an autoimmune disorder. Exemplary autoimmune disorders include, but are not limited to, rheumatoid arthritis, systemic lupus erythematosus, Sjogren's syndrome,
25 scleroderma, mixed connective tissue disease, dermatomyositis, polymyositis, Reiter's syndrome or Behcet's disease, type I (Insulin dependent), type II diabetes mellitus, Hashimoto's thyroiditis, Graves' Disease, multiple sclerosis, myasthenia gravis, encephalomyelitis, pemphigus vulgaris, pemphigus vegetans, pemphigus foliaceus, Senear-Usher syndrome, Brazilian pemphigus, psoriasis (*e.g.*, psoriasis vulgaris), atopic
30 dermatitis, inflammatory bowel disease (*e.g.*, ulcerative colitis or Crohn's Disease), a disorder resulting from an organ, tissue, or cell transplant (*e.g.*, a bone marrow

transplant), such as acute or chronic GVHD, or aplastic anemia, endogenous uveitis, nephrotic syndrome, primary biliary cirrhosis, lichen planus, pyoderma gangrenosum, alopecia areata, a Bullous disorder, chronic viral active hepatitis, auto immune chronic active hepatitis, acquired immune deficiency syndrome (AIDS), and the like.

5 A system can be used to treat a human infected with a pathogen, *e.g.*, a virus, bacteria, or fungus. For example, the human can have a virus, such as a hepatitis virus (*e.g.*, Hepatitis A, B, C, D, E, F, G, H), respiratory syncytial virus, Herpes simplex virus, cytomegalovirus, Epstein Barr Virus, Kaposi's Sarcoma-associated Herpes Virus, JC Virus, rhinovirus, myxovirus, coronavirus, West Nile Virus, St. Louis Encephalitis
10 flavivirus, Tick-borne encephalitis virus, Murray Valley encephalitis flavivirus, Simian Virus 40, Human T Cell Lymphotropic Virus, Moloney-Murine Leukemia Virus, encephalomyocarditis virus, measles virus, Varicella zoster virus, yellow fever virus, adenovirus, poliovirus, poxvirus, and the like.

 The human can be infected with a bacteria, such as *Mycobacterium ulcerans*,
15 *Mycobacterium tuberculosis*, *Mycobacterium leprae*, *Staphylococcus aureus*,
Streptococcus pneumoniae, *Streptococcus pyogenes*, *Chlamydia pneumoniae*,
Mycoplasma pneumoniae, and the like.

 A human can be treated with other pharmaceutical compositions, or other therapy regimens, in addition to treatment with a system. The methods can also be used
20 therapeutically or prophylactically.

 A cell containing a system can generate large amounts of a protein that are secreted and harvested for use as therapeutic agents. For this purpose, the system in the cell is programmed to synthesize a particular polypeptide of interest. For example, the system can be programmed to synthesize large quantities of insulin for packaging and
25 marketing for the treatment of diabetes or human growth hormone for treatment of dwarfism in children. The system can be programmed to synthesize any variation of a polypeptide, including variants discovered to have greater efficacy or fewer side effects than naturally occurring polypeptides.

 The compositions and methods provided may also be used, *e.g.*, as a research tool,
30 to examine the function of various proteins and genes *in vitro* in cultured or preserved

dermal tissues and in animals. The system can be applied to examine the function of any gene.

Delivery of a system to a cell.

5 A system can be introduced into a cell by any method, including any method traditionally used to introduce nucleic acids into cells. For example, a system can be introduced into a cell by microinjection, electroporation, by liposomes, or by microprojectile technology.

A system can be delivered to a cell as a component of a membranous molecular assembly, *e.g.*, a liposome or a micelle. As used herein, the term "liposome" refers to a vesicle composed of amphiphilic lipids arranged in at least one bilayer, *e.g.*, one bilayer or a plurality of bilayers. Liposomes include unilamellar and multilamellar vesicles that have a membrane formed from a lipophilic material and an aqueous interior. The aqueous portion contains the system. The lipophilic material isolates the aqueous interior from an aqueous exterior. Liposomes are generally useful for the transfer and delivery of active ingredients (*e.g.*, a system) to the site of action (*e.g.*, to the interior of a cell). Because the liposomal membrane is structurally similar to biological membranes, when liposomes are applied to a tissue, the liposomal bilayer fuses with the bilayer of the cellular membranes. As the merging of the liposome and cell progresses, the internal aqueous contents that include the system are delivered into the cell where the system can synthesize biological components. In some cases the liposomes are also specifically targeted, *e.g.*, to direct the system to a particular cell type (see methods of targeting below).

A liposome containing a system can be prepared by a variety of methods. In one example, the lipid component of a liposome is dissolved in a detergent so that micelles are formed with the lipid component. For example, the lipid component can be an amphipathic cationic lipid or lipid conjugate. The detergent can have a high critical micelle concentration and may be nonionic. Exemplary detergents include cholate, CHAPS, octylglucoside, deoxycholate, and lauroyl sarcosine. Systems are then added to the micelles that include the lipid component. The system can be coated with an anionic material such that the cationic groups on the lipid interact with the system and condense

around the system to form a liposome. After condensation, the detergent is removed, *e.g.*, by dialysis, to yield a liposomal preparation containing the system.

If necessary a carrier compound that assists in condensation can be added during the condensation reaction, *e.g.*, by controlled addition. For example, the carrier
5 compound can be a polymer, such as spermine or spermidine. pH can also adjusted to favor condensation.

One major type of liposomal composition includes phospholipids other than naturally-derived phosphatidylcholine. Neutral liposome compositions, for example, can be formed from dimyristoyl phosphatidylcholine (DMPC) or dipalmitoyl
10 phosphatidylcholine (DPPC). Anionic liposome compositions generally are formed from dimyristoyl phosphatidylglycerol, while anionic fusogenic liposomes are formed primarily from dioleoyl phosphatidylethanolamine (DOPE). Another type of liposomal composition is formed from phosphatidylcholine (PC) such as, for example, soybean PC, and egg PC. Another type is formed from mixtures of phospholipid and/or
15 phosphatidylcholine and/or cholesterol.

Examples of other methods to introduce liposomes into cells *in vitro* and *in vivo* include U.S. Pat. No. 5,283,185; U.S. Pat. No. 5,171,678; WO 94/00569; WO 93/24640; WO 91/16024; Feigner, *J. Biol. Chem.* 269:550, 1994; Nabel, *Proc. Natl. Acad. Sci* 90:11307, 1993; Nabel, *Human Gene Ther.* 3:649, 1992; Gershon, *Biochem.* 32:7143,
20 1993; and Strauss *EMBO J.* 11:417, 1992.

One cationic lipid conjugate includes derivatization of the lipid with cholesterol ("DC-Choi") which has been formulated into liposomes in combination with DOPE (See, Gao, X. and Huang, L., *Biochim. Biophys. Res. Commun.* 179:280, 1991). Lipopolylysine, made by conjugating polylysine to DOPE, has been reported to be
25 effective for transfection in the presence of serum (Zhou, X. *et al*, *Biochim. Biophys. Acta* 1065:8, 1991). For certain cell lines, these liposomes containing conjugated cationic lipids, are said to exhibit lower toxicity and provide more efficient transfection than the DOTMA-containing compositions. Other commercially available cationic lipid products include DMRIE and DMRIE-HP (Vical, La Jolla, California) and Lipofectamine
30 (DOSPA) (Life Technology, Inc., Gaithersburg, Maryland).

Liposomal formulations are particularly suited for topical administration, liposomes present several advantages over other formulations. Such advantages include reduced side effects related to high systemic absorption of the administered drug, increased accumulation of the administered drug at the desired target, and the ability to administer a system to skin cells. In some implementations, liposomes are used for delivering a system to epidermal cells and also to enhance the delivery of the systems into dermal tissues, e.g., into skin. For example, the liposomes can be applied topically. Topical delivery of drugs formulated as liposomes to the skin has been documented (see, e.g., Weiner *et al*, *Journal of Drug Targeting*, 2:405-410, 1992, and du Plessis *et al*, *Antiviral Research*, 18:259-265, 1992; Mannino, R. J. and Fould-Fogerite, S., *Biotechniques* 6:682-690, 1988; Itani, T. *et al*. *Gene* 56:267-276. 1987; Nicolau, C. *et al*. *Meth. Enz.* 149:157-176, 1987; Straubinger, R. M. and Papahadjopoulos, D. *Meth. Enz.* 101:512-527, 1983; Wang, C. Y. and Huang, L., *Proc. Natl. Acad. Sci. USA* 84:7851-7855, 1987).

Non-ionic liposomal systems can also be used to deliver a system to the skin. Non-ionic liposomal formulations include Novasome I (glyceryl dilaurate/cholesterol/polyoxyethylene-10-stearyl ether) and Novasome II (glyceryl distearate/ cholesterol/polyoxyethylene-10-stearyl ether). Such formulations containing the systems are useful for treating a dermatological disorder.

Surfactants find wide application in formulations such as emulsions (including microemulsions) and liposomes (see above). Compositions including a system can include a surfactant. In one embodiment, the system is formulated as an emulsion that includes a surfactant.

If the surfactant molecule is not ionized, it is classified as a nonionic surfactant. Nonionic surfactants include nonionic esters such as ethylene glycol esters, propylene glycol esters, glyceryl esters, polyglyceryl esters, sorbitan esters, sucrose esters, and ethoxylated esters. Nonionic alkanolamides and ethers such as fatty alcohol ethoxylates, propoxylated alcohols, and ethoxylated/propoxylated block polymers are also included in this class. The polyoxyethylene surfactants are the most popular members of the nonionic surfactant class.

If the surfactant molecule carries a negative charge when it is dissolved or dispersed in water, the surfactant is classified as anionic. Anionic surfactants include carboxylates such as soaps, acyl lactylates, acyl amides of amino acids, esters of sulfuric acid such as alkyl sulfates and ethoxylated alkyl sulfates, sulfonates such as alkyl
5 benzene sulfonates, acyl isethionates, acyl taurates and sulfosuccinates, and phosphates. The most important members of the anionic surfactant class are the alkyl sulfates and the soaps.

If the surfactant molecule carries a positive charge when it is dissolved or dispersed in water, the surfactant is classified as cationic. Cationic surfactants include
10 quaternary ammonium salts and ethoxylated amines. The quaternary ammonium salts are the most used members of this class.

If the surfactant molecule has the ability to carry either a positive or negative charge, the surfactant is classified as amphoteric. Amphoteric surfactants include acrylic acid derivatives, substituted alkylamides, N-alkylbetaines and phosphatides.

15 A system can be delivered to a cell as a micellar formulation. In micelles amphipathic molecules are arranged in a spherical structure such that all the hydrophobic portions of the molecules are directed inward, leaving the hydrophilic portions in contact with the surrounding aqueous phase. The converse arrangement exists if the environment is hydrophobic.

20 A mixed micellar formulation suitable for delivery through transdermal membranes may be prepared by combining a system with an alkali metal C8 to C22 alkyl sulphate, and micelle forming compounds. Exemplary micelle forming compounds include lecithin, hyaluronic acid, pharmaceutically acceptable salts of hyaluronic acid, glycolic acid, lactic acid, chamomile extract, cucumber extract, oleic acid, linoleic acid,
25 linolenic acid, monoolein, monooleates, monolaurates, borage oil, evening of primrose oil, menthol, trihydroxy oxo cholanyl glycine and pharmaceutically acceptable salts thereof, glycerin, polyglycerin, lysine, polylysine, triolein, polyoxyethylene ethers and analogues thereof, polidocanol alkyl ethers and analogues thereof, chenodeoxycholate, deoxycholate, and mixtures thereof. The micelle forming compounds may be added at
30 the same time or after addition of the alkali metal alkyl sulphate. Mixed micelles will

form with substantially any kind of mixing of the ingredients but vigorous mixing is preferred in order to provide smaller size micelles.

In one method a first micellar composition is prepared which contains the system and at least the alkali metal alkyl sulphate. The first micellar composition is then mixed
5 with at least three micelle forming compounds to form a mixed micellar composition. In another method, the micellar composition is prepared by mixing the composition containing the system, the alkali metal alkyl sulphate and at least one of the micelle forming compounds, followed by addition of the remaining micelle forming compounds, with vigorous mixing.

10 Phenol and/or m-cresol may be added to the mixed micellar composition to stabilize the formulation and protect against bacterial growth. Alternatively, phenol and/or m-cresol may be added with the micelle forming ingredients. An isotonic agent such as glycerin may also be added after formation of the mixed micellar composition.

The specific concentrations of the essential ingredients can be determined by
15 relatively straightforward experimentation.

In another embodiment, a system may be incorporated into a particle, *e.g.*, a microparticle. Microparticles can be produced by spray-drying, but may also be produced by other methods including lyophilization, evaporation, fluid bed drying,
vacuum drying, or a combination of these techniques.

20 Polymeric particles, *e.g.*, polymeric in microparticles, can be used as a sustained-release reservoir of systems that are taken up by cells only released from the microparticle through biodegradation. The polymeric particles in this embodiment should therefore be large enough to preclude phagocytosis (*e.g.*, larger than 10 μm and preferably larger than 20 μm). Such particles can be produced by the same methods to
25 make smaller particles, but with less vigorous mixing of the first and second emulsions. That is to say, a lower homogenization speed, vortex mixing speed, or sonication setting can be used to obtain particles having a diameter around 100 μm rather than 10 μm . The time of mixing also can be altered.

Larger microparticles can be formulated as a suspension, a powder, or an
30 implantable solid, to be delivered by intramuscular, subcutaneous, intradermal, intravenous, or intraperitoneal injection; via inhalation (intranasal or intrapulmonary);

orally; or by implantation. These particles are useful for delivery of a system when slow release over a relatively long term is desired. The rate of degradation, and consequently of release, varies with the polymeric formulation.

5 Microparticles preferably include pores, voids, hollows, defects or other interstitial spaces that allow the fluid suspension medium to freely permeate or perfuse the particulate boundary. For example, the perforated microstructures can be used to form hollow, porous spray dried microspheres.

10 Polymeric particles containing the systems can be made using a double emulsion technique, for instance. First, the polymer is dissolved in an organic solvent. A preferred polymer is polylactic-co-glycolic acid (PLGA), with a lactic/glycolic acid weight ratio of 65:35, 50:50, or 75:25. Next, systems in aqueous solution are added to the polymer solution and the two are mixed to form a first emulsion. The solutions can be mixed by vortexing or shaking, and in a preferred method, the mixture can be sonicated. Most preferable is any method by which the system receives the least amount of damage while
15 still allowing the formation of an appropriate emulsion. For example, a Vibra-cell model VC-250 sonicator is useful for making polymeric particles.

Targeting. In some embodiments, the system is targeted to a particular cell. For example, a liposome or particle or other structure that includes a system can also include
20 a targeting moiety that recognizes a specific molecule on a target cell. The targeting moiety can be a molecule with a specific affinity for a target cell. Targeting moieties can include antibodies directed against a protein found on the surface of a target cell, or the ligand or a receptor-binding portion of a ligand for a receptor found on the surface of a target cell. For example, the targeting moiety can recognize a cancer-specific antigen
25 (*e.g.*, CA15-3, CA19-9, CEA, or HER2/neu) or a viral antigen, thus delivering the system to a cancer cell or a virus-infected cell. Exemplary targeting moieties include antibodies (such as IgM, IgG, IgA, IgD, and the like, or a functional portion thereof), or ligands for cell surface receptors.

30

Route of Delivery.

A composition that includes a system can be delivered to a human subject by a variety of routes. Exemplary routes include intravenous, topical, nasal, pulmonary, and ocular.

5 The systems can be incorporated into pharmaceutical compositions suitable for administration. Such compositions typically include at least one system and a pharmaceutically acceptable carrier. As used herein the language "pharmaceutically acceptable carrier" is intended to include any and all solvents, dispersion media, coatings, antibacterial and antifungal agents, isotonic and absorption delaying agents, and the like,
10 compatible with pharmaceutical administration. The use of such media and agents for pharmaceutically active substances is well known in the art. Except insofar as any conventional media or agent is incompatible with the system, use thereof in the compositions is contemplated.

15 Pharmaceutical compositions featured in may be administered in a number of ways depending upon whether local or systemic treatment is desired and upon the area to be treated. Administration may be topical (including ophthalmic, intranasal, transdermal), oral or parenteral. Parenteral administration includes intravenous drip, subcutaneous, intraperitoneal or intramuscular injection, or intrathecal or intraventricular administration.

20 The route and site of administration may be chosen to enhance targeting. For example, to target muscle cells, intramuscular injection into the muscles of interest would be a logical choice. Lung cells might be targeted by administering the composition containing the system in aerosol form. The vascular endothelial cells could be targeted by coating a balloon catheter with a composition including the systems and mechanically introducing the composition.

25 Formulations for topical administration may include transdermal patches, ointments, lotions, creams, gels, drops, suppositories, sprays, liquids and powders. Conventional pharmaceutical carriers, aqueous, powder or oily bases, thickeners and the like may be necessary or desirable.

30 Compositions for oral administration include powders or granules, suspensions or solutions in water, syrups, elixirs or non-aqueous media, tablets, capsules, lozenges, or troches. In the case of tablets, carriers that can be used include lactose, sodium citrate

and salts of phosphoric acid. Various disintegrants such as starch, and lubricating agents such as magnesium stearate, sodium lauryl sulfate and talc, are commonly used in tablets. For oral administration in capsule form, useful diluents are lactose and high molecular weight polyethylene glycols. When aqueous suspensions are required for oral use, the
5 nucleic acid compositions can be combined with emulsifying and suspending agents. If desired, certain sweetening and/or flavoring agents can be added.

Compositions for intrathecal or intraventricular administration may include sterile aqueous solutions which may also contain buffers, diluents and other suitable additives.

Formulations for parenteral administration may include sterile aqueous solutions
10 which may also contain buffers, diluents and other suitable additives. Intraventricular injection may be facilitated by an intraventricular catheter, for example, attached to a reservoir. For intravenous use, the total concentration of solutes should be controlled to render the preparation isotonic.

For ocular administration, ointments or droppable liquids may be delivered by
15 ocular delivery systems known to the art such as applicators or eye droppers. Such compositions can include mucomimetics such as hyaluronic acid, chondroitin sulfate, hydroxypropyl methylcellulose or poly(vinyl alcohol), preservatives such as sorbic acid, EDTA or benzylchromium chloride, and the usual quantities of diluents and/or carriers.

Iontophoresis (transfer of ionic solutes through biological membranes under the
20 influence of an electric field) (Lee *et al*, Critical Reviews in Therapeutic Drug Carrier Systems, 1991, p. 163), phonophoresis or sonophoresis (use of ultrasound to enhance the absorption of various therapeutic agents across biological membranes, notably the skin and the cornea) (Lee *et al*, Critical Reviews in Therapeutic Drug Carrier Systems, 1991, p. 166), and optimization of vehicle characteristics relative to dose position and retention
25 at the site of administration (Lee *et al*, Critical Reviews in Therapeutic Drug Carrier Systems, 1991, p. 168) may be useful methods for enhancing the transport of topically applied compositions across skin and mucosal sites.

Other embodiments are within the scope of the following claims.

Chapter Five: GNR: Three Overlapping Revolutions

There are few things of which the present generation is more justly proud than the wonderful improvements which are daily taking place in all sorts of mechanical appliances. . . . But what would happen if technology continued to evolve so much more rapidly than the animal and vegetable kingdoms? Would it displace us in the supremacy of earth? Just as the vegetable kingdom was slowly developed from the mineral, and as in like manner the animal supervened upon the vegetable, so now in these last few ages an entirely new kingdom has sprung up, of which we as yet have only seen what will one day be considered the antediluvian prototypes of the race. . . . We are daily giving [machines] greater power and supplying by all sorts of ingenious contrivances that self-regulating, self-acting power which will be to them what intellect has been to the human race.

—*Samuel Butler, 1863 (four years after publication of Darwin's The Origin of Species,)*

Who will be man's successor? To which the answer is: We are ourselves creating our own successors. Man will become to the machine what the horse and the dog are to man; the conclusion being that machines are, or are becoming, animate.

—*Samuel Butler, 1863 letter, "Darwin among the Machines"*¹

The first half of the twenty-first century will be characterized by three overlapping revolutions—in Genetics, Nanotechnology, and Robotics. These will usher in what I referred to earlier as Epoch Five, the beginning of the Singularity. We are in the early stages of the "G" revolution today. By understanding the information processes underlying life, we are starting to learn to reprogram our biology to achieve the virtual elimination of disease, dramatic expansion of human potential, and radical life extension. Hans Moravec points out, however, that no matter how successfully we fine-tune our DNA-based biology, humans will remain "second-class robots," meaning that biology will never be able to match what we will be able to engineer once we fully understand biology's principles of operation.²

The "N" revolution will enable us to redesign and rebuild—molecule by molecule—our bodies and brain and the world with which we interact, going far beyond the limitations of biology. The most powerful impending revolution is "R": human-level robots with their intelligence derived from our own but redesigned to far exceed human capabilities. R represents the most significant transformation, because intelligence is the most powerful "force" in the universe. Intelligence, if sufficiently advanced, is, well, smart enough to anticipate and overcome any obstacles that stand in its path.

While each revolution will solve the problems from earlier transformations, it will also introduce new perils. G will overcome the age-old difficulties of disease and aging but establish the potential for new bioengineered viral threats. Once N is fully developed we will be able to apply it to protect ourselves from all biological hazards, but it will create the possibility of its own self-replicating dangers, which will be far more powerful than anything biological. We can protect ourselves from these hazards with fully developed R, but what will protect us from pathological intelligence that exceeds our own? I do have a strategy for dealing with these issues, which I discuss at the end of chapter 8. In this chapter, however, we will examine how the Singularity will unfold through these three overlapping revolutions: G, N, and R.

<hl>Genetics: The Intersection of Information and Biology

<epi>It has not escaped our notice that the specific pairing we have postulated immediately suggests a possible copying mechanism for the genetic material.

<&p\s>—*James Watson and Francis Crick?*

<epi>After three billion years of evolution, we have before us the instruction set that carries each of us from the one-cell egg through adulthood to the grave.

<epis>—*Dr. Robert Waterston, International Human Genome Sequencing Consortium*⁴

<tx>Underlying all of the wonders of life and misery of disease are information processes, essentially software programs, that are surprisingly compact. The entire human genome is a sequential binary code containing only about eight hundred million bytes of information. As I mentioned earlier, when its massive redundancies are removed using conventional compression techniques, we are left with only thirty to one hundred million bytes, equivalent to the size of an average contemporary software program.⁵ This code is supported by a set of biochemical machines that translate these linear (one-dimensional) sequences of DNA "letters" into strings of simple building blocks called amino acids, which are in turn folded into three-dimensional proteins, which make up all living creatures from bacteria to humans. (Viruses occupy a niche in between living and nonliving matter but are also comprised of fragments of DNA or RNA.) This machinery is essentially a self-replicating nanoscale replicator that builds the elaborate hierarchy of structures and increasingly complex systems that a living creature comprises.

<bh>Life's Computer


<btx>In the very early stages of evolution information was encoded in the structure of increasingly complex organic molecules based on carbon. After billions of years biology evolved its own computer for storing and manipulating digital data based on the DNA molecule. The chemical structure of the DNA molecule was first described by J. D. Watson and F. H. C. Crick in 1953 as a double helix consisting of a pair of strands of polynucleotides with information encoded at each position by the choice of nucleotides.⁶ We finished transcribing the genetic code at the beginning of this century. We are now beginning to understand the detailed chemistry of the communication and control processes by which DNA commands reproduction through such other complex molecules and cellular structures as messenger RNA (mRNA), transfer RNA (tRNA), and ribosomes.

At the level of information storage the mechanism is surprisingly simple. Supported by a twisting sugar-phosphate backbone, the DNA molecule contains up to several million rungs, each of which is coded with one letter drawn from a four-letter alphabet; each rung is thus coding two bits of data in a one-dimensional digital code. The alphabet consists of the four base pairs: adenine-thymine, thymine-adenine, cytosine-guanine, and guanine-cytosine. The DNA strings in a single cell would measure up to six feet in length if stretched out, but an elaborate packing method coils it to fit into a cell only 1/2500 of an inch across.

Special enzymes can copy the information on each rung by splitting each base pair and assembling two identical DNA molecules by rematching the broken base pairs. Other enzymes actually check the validity of the copy by checking the integrity of the base-pair matching. With these copying and validation steps, this chemical data-processing system makes only about one error in ten billion base-pair replications.⁷ Further redundancy and error-correction codes are built into the digital data itself, so meaningful mutations resulting from base-pair replication errors are rare. Most of the errors resulting from the one-in-one-billion error rate will result in the equivalent of a "parity" error, which can be detected and corrected by other levels of the system, including matching against the corresponding chromosome, which can prevent the incorrect bit from causing any significant damage.⁸ Recent research has shown that the genetic mechanism detects such errors in transcription of the male Y chromosome by matching each Y-

chromosome gene against a copy on the same chromosome.⁹ Once in a long while a transcription error will result in a beneficial change that evolution will come to favor.

In a process technically called translation, another series of chemicals put this elaborate digital program into action by building proteins. It is the protein chains that give each cell its structure, behavior, and intelligence. Special enzymes unwind a region of DNA for building a particular protein. A strand of mRNA is created by copying the exposed sequence of bases. The rRNA essentially has a copy of a portion of the DNA letter sequence. The mRNA travels out of the nucleus and into the cell body. The mRNA codes are then read by a ribosome molecule, which represents the central molecular player in the drama of biological reproduction. One portion of the ribosome acts like a tape-recorder head, "reading" the sequence of data encoded in the mRNA base sequence. The "letters" (bases) are grouped into words of three letters each called codons, with one codon for each of twenty possible amino acids, the basic building blocks of protein. A ribosome reads the codons from the mRNA and then, using tRNA, assembles a protein chain one amino acid at a time.



The notable final step in this process is the folding of the one-dimensional chain of amino acid "beads" into a three-dimensional protein. Simulating this process has not yet been feasible because of the enormous complexity of the interacting forces from all the atoms involved. Supercomputers scheduled to come online around the time of the publication of this book (2005) are expected to have the computational capacity to simulate protein folding, as well as the interaction of one three-dimensional protein with another.

Protein folding, along with cell division, is one of nature's remarkable and intricate dances in the creation and re-creation of life. Specialized "chaperone" molecules protect and guide the amino-acid strands as they assume their precise three-dimensional protein configurations. As many as one third of formed protein molecules are folded improperly. These disfigured proteins must be immediately destroyed or they will rapidly accumulate, disrupting cellular functions on many levels.

Under normal circumstances, as soon as a misfolded protein is formed, it is tagged by a carrier molecule, ubiquitin, and escorted to a specialized proteosome, where it is broken back down into its component amino acids for recycling into new (correctly folded) proteins. As cells age, however, they produce less of the energy needed for optimal function of this mechanism. Accumulations of these misformed proteins aggregate into particles called protofibrils, which are thought to underlie disease processes leading to Alzheimer's disease and other afflictions.¹⁰

The ability to simulate the three-dimensional waltz of atomic-level interactions will greatly accelerate our knowledge of how DNA sequences control life and disease. We will then be in a position to

rapidly simulate drugs that intervene in any of the steps in this process, thereby hastening drug development and the creation of highly targeted drugs that minimize unwanted side effects.

It is the job of the assembled proteins to carry out the functions of the cell, and by extension the organism. A molecule of hemoglobin, for example, which has the job of carrying oxygen from the lungs to body tissues, is created five hundred trillion times each second in the human body. With more than five hundred amino acids in each molecule of hemoglobin, that comes to 1.5×10^{19} (fifteen billion billion) "read" operations every minute by the ribosomes just for the manufacture of hemoglobin.

In some ways the biochemical mechanism of life is remarkably complex and intricate. In other ways it is remarkably simple. Only four base pairs provide the digital storage for all of the complexity of all human life and all other life as we know it. The ribosomes build protein chains by grouping together triplets of base pairs to select sequences from only twenty amino acids. The amino acids themselves are relatively simple, consisting of a carbon atom with its four bonds linked to one hydrogen atom, one amino ($-\text{NH}_2$) group, one carboxylic acid ($-\text{COOH}$) group, and one organic group that is different for each amino acid. The organic group for Alanine, for example, has only four atoms (CH_3-) for a total of 13 atoms. One of the more complex amino acids, arginine (which plays a vital role in the health of the endothelial cells in our arteries) has only seventeen atoms, in its organic group for a total of twenty-six atoms. These twenty simple molecular fragments are the building blocks of all life.

The protein chains then control everything else: the structure of bone cells, the ability of muscle cells to flex and act in concert with other muscle cells, all of the complex biochemical interactions that take place in the bloodstream, and, of course, the structure and functioning of the brain."

Designer Baby Boomers

<tx>Sufficient information already exists today to slow down disease and aging processes to the point that baby boomers like myself can remain in good health until the full blossoming of the biotechnology revolution, which will itself be a bridge to the nanotechnology revolution (see Resources page). In *Fantastic Voyage: Live Long Enough to Live Forever*, which I coauthored with Terry Grossman, M.D., a leading longevity expert, we discuss these three bridges to radical life extension (today's knowledge, biotechnology, and nanotechnology).¹² I wrote there: "Whereas some of my contemporaries may be satisfied to embrace aging gracefully as part of the cycle of life, that is not my view. It may be 'natural,' but I don't see anything positive in losing my mental agility, sensory acuity, physical limberness, sexual desire, or any other human ability. I view disease and death at any age as a calamity, as problems to be overcome."

Bridge one involves aggressively applying the knowledge we now possess to dramatically slow down aging and reverse the most important disease processes, such as heart disease, cancer, type 2 diabetes, and stroke. You can, in effect, reprogram your biochemistry, for we have the knowledge today, if aggressively applied, to overcome our genetic heritage in the vast majority of cases. "It's mostly in your genes" is only true if you take the usual passive attitude toward health and aging.

My own story is instructive. More than twenty years ago I was diagnosed with type 2 diabetes. The conventional treatment made my condition worse, so I approached this health challenge from my perspective as an inventor. I immersed myself in the scientific literature and came up with a unique program that successfully reversed my diabetes. In 1993 I wrote a health book (*The 10% Solution for a Healthy Life*) about this experience, and I continue today to be free of any indication or complication of this disease.¹³

In addition, when I was twenty-two, my father died of heart disease at the age of fifty-eight, and I have inherited his genes predisposing me to this illness. Twenty years ago, despite following the public guidelines of the American Heart Association, my cholesterol was in the high 200s (it should be well below 180), my HDL (high density lipoprotein, the "good" cholesterol) below 30 (it should be above 50), and my homocysteine (a measure of the health of a biochemical process called methylation) was an unhealthy 11 (it should be below 7.5). By following a longevity program that Grossman and I developed, my current cholesterol level is

130, my HDL is 55, my homocysteine is 6.2, my C-reactive Protein (a measure of inflammation in the body) is a very healthy 0.01, and all of my other indexes (for heart disease, diabetes, and other conditions) are at ideal levels.¹⁴

When I was 40, my biological age was around 38. Although I am now 56, a comprehensive test of my biological aging (measuring various sensory sensitivities, lung capacity, reaction times, memory, and related tests) conducted at Grossman's longevity clinic measured my biological age at 40.¹⁵ Although there is not yet a consensus on how to measure biological age, my scores on these tests matched population norms for this age. So according to this set of tests I have not aged very much in the last 16 years, which is confirmed by the many blood tests I take, as well as the way I feel.

These results are not accidental; I have been very aggressive about reprogramming my biochemistry. I take 250 supplements (pills) a day and receive a half-dozen intravenous therapies each week (basically nutritional supplements delivered directly into my bloodstream, thereby bypassing my GI tract). As a result, the metabolic reactions in my body are completely different than they would otherwise be.¹⁶ Approaching this as an engineer, I measure dozens of levels of nutrients (such as vitamins, minerals, and fats), hormones, and metabolic by-products in my blood and other body samples (such as hair and saliva). Overall, my levels are where I want them to be, although I continually fine-tune my program based on the research that I conduct with Grossman.¹⁷ Although my program may seem extreme, it is actually conservative—and optimal (based on my current knowledge). Grossman and I have extensively researched each of the several hundred therapies that I use for safety and efficacy. I stay away from ideas that are unproven or appear to be risky (the use of human-growth hormone, for example).

We consider the process of reversing and overcoming the dangerous progression of disease as a war. As in any war it is important to mobilize all the means of intelligence and weaponry that can be harnessed, throwing everything we have at the enemy. For this reason we advocate that key dangers—such as heart disease, cancer, diabetes, stroke, and aging—be attacked on multiple fronts. For example, our strategy for preventing heart disease is to adopt ten different heart disease prevention therapies that attack each of the known risk factors.

By adopting such multipronged strategies for each disease process and each aging process, even baby boomers like myself can remain in good health until the full blossoming of the biotechnology revolution (which we call "bridge two"), which is already in its early stages and will reach its peak in the second decade of this century.

Biotechnology will provide the means to actually change your genes: not just designer babies will be feasible but designer baby boomers. We'll also be able to rejuvenate all of your body's tissues and organs by transforming your skin cells into youthful versions of every other cell type. Already, new drug development is precisely targeting key steps in the process of atherosclerosis (the cause of heart disease), cancerous tumor formation, and the metabolic processes underlying each major disease and aging process.

<h2>Can we really live forever? </h2> An energetic and insightful advocate of stopping the aging process by changing the information processes underlying biology is Aubrey de Grey, a scientist in the department of genetics at Cambridge University. De Grey uses the metaphor of maintaining a house. How long does a house last? The answer obviously depends on how well you take care of it. If you do nothing, the roof will spring a leak before long, water and the elements will invade, and eventually the house will disintegrate. But if you proactively take care of the structure, repair all damage, confront all dangers, and rebuild or renovate parts from time to time using new materials and technologies, the life of the house can essentially be extended without limit.

The same holds true for our bodies and brains. The only difference is that, while we fully understand the methods underlying the maintenance of a house, we do not yet fully understand all of the biological principles of life. But with our rapidly increasing comprehension of the biochemical processes and pathways of biology, we are quickly gaining that knowledge. We are beginning to understand aging, not as a single inexorable progression but as a group of related processes. Strategies are emerging for fully reversing each of these aging progressions, using different combinations of biotechnology techniques.

De Grey describes his goal as "engineered negligible senescence"—stopping the body and brain from becoming more frail and disease-prone as it grows older.¹⁸ As he explains, "All the core knowledge needed to develop *engineered negligible senescence* is already in our possession—it mainly just needs to be pieced together."¹⁹ De Grey believes we'll demonstrate "robustly rejuvenated" mice—mice that are functionally younger than before being treated and with the life extension to prove it—within ten years, and he points out that this achievement will have a dramatic effect on public opinion. Demonstrating that we can reverse the aging process in an animal that shares 99 percent of our genes will profoundly challenge the common wisdom that aging and death are inevitable. Once robust rejuvenation is confirmed in an animal, there will be enormous competitive pressure to translate these results into human therapies, which should appear five to ten years later.

The diverse field of biotechnology is fueled by our accelerating progress in reverse engineering the information processes underlying biology and by a growing arsenal of tools that can modify these processes. For example, drug discovery was once a matter of finding substances that produced some beneficial result without excessive side effects. This process was similar to early humans' tool discovery, which was limited to simply finding rocks and other natural implements that could be used for helpful purposes. Today we are learning the precise biochemical pathways that underlie both disease and aging processes and are able to design drugs to carry out precise missions at the molecular level. The scope and scale of these efforts is vast.

Another powerful approach is to start with biology's information backbone: the genome. With recently developed gene technologies we're on the verge of being able to control how genes express themselves. Gene expression is the process by which specific cellular components (specifically RNA and the ribosomes) produce proteins according to a specific genetic blueprint. While every human cell has the full complement of the body's genes, a specific cell, such as a skin cell or a pancreatic Islet cell, gets its characteristics from only the small fraction of genetic information relevant to that particular cell type.²⁰ The therapeutic control of this process can take place outside the cell nucleus, so it is easier to implement than therapies that require access inside it.

Gene expression is controlled by peptides (molecules made up of sequences of up to 100 amino acids) and short RNA strands. We are now beginning to learn how these processes work.²¹ Many new therapies now in development and testing are based on manipulating them to either turn off the expression of disease-causing genes or to turn on desirable genes that may otherwise not be expressed in a particular type of cell.

RNAi (RNA interference). A powerful new tool called RNA interference (RNAi) is capable of turning off specific genes by blocking their mRNA, thus preventing them from creating proteins. Since viral diseases, cancer, and many other diseases use gene expression at some crucial point in their life cycle, this promises to be a breakthrough technology. Researchers construct short, double-stranded DNA segments that match and lock onto portions of the RNA that are transcribed from a targeted gene. With their ability to create proteins blocked, the gene is effectively silenced. In many genetic diseases only one copy of a given gene

is defective. Since we get two copies of each gene, one from each parent, blocking the disease causing gene leaves one healthy gene to make the necessary protein. If both genes are defective, RNAi could silence them both, but then a healthy gene would have to be inserted.²²

Cell Therapies. Another important line of attack is to regrow our own cells, tissues, and even whole organs and introduce them into our bodies without surgery. One major benefit of this "therapeutic cloning" technique is that we will be able to create these new tissues and organs from versions of our cells that have also been made younger via the emerging field of rejuvenation medicine. For example, we will be able to create new heart cells from skin cells and introduce them into the system through the bloodstream. Over time, existing heart cells will be replaced with these new cells, and the result will be a rejuvenated "young" heart manufactured using a person's own DNA. I discuss this approach to regrowing our bodies below.

Gene chips. New therapies are only one way that the growing knowledge base of gene expression will dramatically impact our health. Since the 1990s microarrays or chips no larger than a dime have been used to study and compare expression patterns of thousands of genes at a time.²³ The possible applications of the technology are so varied and the technological barriers have been reduced so greatly that huge databases are now devoted to the results from "do-it-yourself gene watching."²⁴

Genetic profiling is now being used to:

<bl>

- > **Revolutionize the processes of drug screening and discovery.** Microarrays can "not only confirm the mechanism of action of a compound" but "discriminate between compounds acting at different steps in the same metabolic pathway."²⁵
- > **Improve cancer classifications.** One study reported in *Science* demonstrated the feasibility of classifying some leukemias "solely on gene expression monitoring." The authors also pointed to a case in which expression profiling resulted in the correction of a misdiagnosis.²⁶
- > Identify **the genes, cells, and pathways involved in a process, such as aging or tumorigenesis.** For example, by correlating the presence of acute myeloblastic leukemia and increased expression of certain genes involved with programmed cell death, a study helped identify new therapeutic targets.²⁷
- > **Determine the effectiveness of an innovative therapy.** One study recently reported in *Bone* looked at the effect of growth-hormone replacement on the expression of insulinlike growth factors (IGFs) and bone metabolism markers.²⁸
- > Test **the toxicity of compounds in food additives, cosmetics, and industrial products quickly and without using animals.** Such tests can show, for example, the degree to which each gene has been turned on or off by a tested substance.²⁹

Somatic gene therapy (gene therapy for nonreproductive cells). This is the holy grail of bioengineering, which will enable us to effectively change genes inside the nucleus by "infecting" it with new DNA, essentially creating new genes.³⁰ The concept of controlling the genetic makeup of humans is often associated with the idea of influencing new generations in the form of "designer babies." But the real promise of gene therapy is to actually change our adult genes.³¹ These can be designed to either block undesirable disease-encouraging genes or introduce new ones that slow down and even reverse aging processes.

Animal studies that began in the 1970s and 1980s have been responsible for producing a range of transgenic animals, such as cattle, chickens, rabbits, and sea urchins. The first attempts at human gene therapy were undertaken in 1990. The challenge is to transfer therapeutic DNA into target cells that will then be expressed at the right level and at the right time.

Consider the challenge involved in effecting a gene transfer. Viruses are often the vehicle of choice. Long ago viruses learned how to deliver their genetic material to human cells and, as a result, cause disease. Researchers now simply switch the material a virus unloads into cells by removing its genes and inserting therapeutic ones. Although the approach itself is relatively easy, the genes are too large to pass into many types of cells (such as brain cells). The process is also limited in the length of DNA it can carry, and it may cause an immune response. And precisely where the new DNA integrates into the cell's DNA has been a largely uncontrollable process.³²

Physical injection (microinjection) of DNA into cells is possible but prohibitively expensive. Exciting advances have recently been made, however, in other means of transfer. For example, liposomes—fatty spheres with a watery core—can be used as a "molecular Trojan horse" to deliver genes to brain cells, thereby opening the door to treatment of disorders such as Parkinson's and epilepsy.³³ Electric pulses can also be employed to deliver a range of molecules (including drug proteins, RNA, and DNA) to cells.³⁴ Yet another option is to pack DNA into ultratiny "nanoballs" for maximum impact.³⁵

The major hurdle that must be overcome for gene therapy to be applied in humans is proper positioning of a gene on a DNA strand and monitoring of the gene's expression. One possible solution is to deliver an imaging reporter gene along with the therapeutic gene. The image signals would allow for close supervision, of both placement and level of expression.³⁶

Even faced with these obstacles gene therapy is starting to work in human applications. A team led by University of Glasgow research doctor Andrew H. Baker has successfully used adenoviruses to "infect" specific organs and even specific regions within organs. For example, the group was able to direct gene therapy precisely at the endothelial cells, which line the inside of blood vessels. Another approach is being developed by Celera Genomics, a company founded by Craig Venter (the head of the private effort to transcribe the human genome). Celera has already demonstrated the ability to create synthetic viruses from genetic information and plans to apply these biodesigned viruses to gene therapy.³⁷

One of the companies I help to direct, United Therapeutics, has begun human trials of delivering DNA into cells through the novel mechanism of autologous (the patient's own) stem cells, which are captured from a few vials of their blood. DNA that directs the growth of new pulmonary blood vessels are inserted into the stem cell genes, and the cells are reinjected into the patient. When the genetically engineered stem cells reach the tiny pulmonary blood vessels near the lung's alveoli, they begin to express growth factors for new blood vessels. In animal studies this has safely reversed pulmonary hypertension, a fatal and presently incurable disease. Based on the success and safety of these studies, the Canadian government gave permission for human tests to commence in early 2005.

Reversing Degenerative Disease

Degenerative (progressive) diseases—heart disease, stroke, cancer, type 2 diabetes, liver disease, and kidney disease—account for at least 90 percent of the deaths in our society. Our understanding of the principal components of degenerative disease and human aging is growing rapidly, and strategies have been identified to halt and even reverse each of these processes. In *Fantastic Voyage*, Grossman and I describe a wide range of therapies now in the testing pipeline that have already demonstrated significant results in attacking the key biochemical steps underlying the progress of such diseases.

<h2>Combating heart disease. <tx>As one of many examples, exciting research is being conducted with a synthetic form of HDL cholesterol called recombinant Apo-A-I Milano (AAIM). In animal trials AAIM was responsible for a rapid and dramatic regression of atherosclerotic plaque.³⁸ In a phase 1 FDA trial, which included forty-seven human subjects,

administering AAIM by intravenous infusion resulted in a significant reduction (an average 4.2 percent decrease) in plaque after just five weekly treatments. No other drug has ever shown the ability to reduce atherosclerosis this quickly.³⁹

Another exciting drug for reversing atherosclerosis now in phase 3 FDA trials is Pfizer's Torcetrapib.⁴⁰ This drug boosts levels of HDL by blocking an enzyme that normally breaks it down. Pfizer is spending a record one billion dollars to test the drug and plans to combine it with its bestselling "statin" (cholesterol-lowering) drug, Lipitor.

<h2>Overcoming cancer. <tx>Many strategies are being intensely pursued to overcome cancer. Particularly promising are cancer vaccines designed to stimulate the immune system to attack cancer cells. These vaccines could be used as a prophylaxis to prevent cancer, as a first-line treatment, or to mop up cancer cells after other treatments.⁴¹

The first reported attempts to activate a patient's immune response were undertaken more than one hundred years ago, with little success.⁴² More recent efforts focus on encouraging dendritic cells, the sentinels of the immune system, to trigger a normal immune response. Many forms of cancer have an opportunity to proliferate because they somehow do not trigger that response. Dendritic cells play a key role because they roam the body, collecting foreign peptides and cell fragments and delivering them to the lymph nodes, which in response produce an army of T cells primed to eliminate the flagged peptides.

Some researchers are altering cancer-cell genes to attract T cells, with the assumption that the stimulated T cells would then recognize other cancer cells they encounter.⁴³ Others are experimenting with vaccines for exposing the dendritic cells to antigens, unique proteins found on the surfaces of cancer cells. One group used electrical pulses to fuse tumor and immune cells to create an "individualized vaccine."⁴⁴ One of the obstacles to developing effective vaccines is that currently we have not yet identified many of the cancer antigens we need to develop potent targeted vaccines.⁴⁵

Blocking angiogenesis—the creation of new blood vessels—is another strategy. This process uses drugs to discourage blood-vessel development, which an emergent cancer needs to grow beyond a small size. Interest in angiogenesis has skyrocketed since 1997, when doctors at the Dana Farber Cancer Center in Boston reported that repeated cycles of endostatin, an angiogenesis inhibitor, had resulted in complete regression of tumors.⁴⁶ There are now many antiangiogenic drugs in clinical trials, including avastin and atrasentan.⁴⁷

A key issue for cancer as well as for aging concerns telomere "beads," repeating sequences of DNA found at the end of chromosomes. Each time a cell reproduces, one bead drops off. Once a cell has reproduced to the point that all of its telomere beads have been expended, that cell is no longer able to divide and will die. If we could reverse this process, cells could survive indefinitely. Fortunately, recent research has found that only a single enzyme (telomerase) is needed to achieve this.⁴⁸ The tricky part is to administer telomerase in such a way as not to cause cancer. Cancer cells possess a gene that produces telomerase, which effectively enables them to become immortal by reproducing indefinitely. A key cancer-fighting strategy, therefore, involves blocking the ability of cancer cells to generate telomerase. This may seem to contradict the idea of extending the telomeres in normal cells to combat this source of aging, but attacking the telomerase of the cancer cells in an emerging tumor could be done without necessarily compromising an orderly telomere-extending therapy for normal cells. However, to avoid complications, such therapies could be halted during a period of cancer therapy.

Reversing Aging

It is logical to assume that early in the evolution of our species (and precursors to our species) survival would not have been aided—indeed, it would have been compromised-by

individuals living long past their child-rearing years. Recent research, however, supports the so-called grandma hypothesis, which suggests a countereffect. University of Michigan anthropologist Rachel Caspari and University of California at Riverside's San-Hee Lee found evidence that the proportion of humans living to become grandparents (who in primitive societies were often as young as thirty) increased steadily over the past two million years, with a fivefold increase occurring in the Upper Paleolithic era (around thirty thousand years ago). This research has been cited to support the hypothesis that the survival of human societies was aided by grandmothers, who not only assisted in raising extended families but also passed on the accumulated wisdom of elders. Such effects may be a reasonable interpretation of the data, but the overall increase in longevity also reflects an ongoing trend towards longer life expectancy that continues to this day. Likewise, only a modest number of grandmas (and a few grandpas) would have been needed to account for the societal effects that proponents of this theory have claimed, so the hypothesis does not appreciably challenge the conclusion that genes that supported significant life extension were not selected for.

Aging is not a single process but involves a multiplicity of changes. De Grey describes seven key aging processes that encourage senescence, and he has identified strategies for reversing each one.

<h2>DNA mutations.⁴⁹ <tx>Generally mutations to nuclear DNA (the DNA in the chromosomes in the nucleus) result in a defective cell that's quickly eliminated or a cell that simply doesn't function optimally. The type of mutation that is of primary concern (as it leads to increased death rates) is one that affects orderly cellular reproduction, resulting in cancer. This means that if we can cure cancer using the strategies described above, nuclear mutations should largely be rendered harmless. De Grey's proposed strategy for cancer is preemptive: it involves using gene therapy to remove from all our cells the genes that cancers need to turn on in order to maintain their telomeres when they divide. This will cause any potential cancer tumors to wither away before they grow large enough to cause harm. Strategies for deleting and suppressing genes are already available and are being rapidly improved.

<h2>Toxic cells. <tx>Occasionally cells reach a state in which they're not cancerous, but it would still be best for the body if they did not survive. Cell senescence is an example, as is having too many fat cells. In these cases, it is easier to kill these cells than to attempt to revert them to a healthy state. Methods are being developed to target "suicide genes" to such cells and also to tag these cells in a way that directs the immune system to destroy them.

<h2>Mitochondrial mutations. <tx>Another aging process is the accumulation of mutations in the thirteen genes in the mitochondria, the energy factories for the cell.⁵⁰ These few genes are critical to the efficient functioning of our cells, and undergo mutation at a higher rate than genes in the nucleus. Once we master somatic gene therapy, we could put multiple copies of these genes in the cell nucleus, thereby providing redundancy (backup) for such vital genetic information. The mechanism already exists in the cell to allow nucleus-encoded proteins to be imported into the mitochondria, so it is not necessary for these proteins to be produced in the mitochondria itself. In fact, most of the proteins needed for mitochondrial function are already coded by the nuclear DNA. Researchers have already been successful in transferring mitochondrial genes into the nucleus in cell cultures.

<h2>Intracellular aggregates. <tx>Toxins are produced both inside and outside cells. De Grey describes strategies using somatic gene therapy to introduce new genes that will break down what he calls "intracellular aggregates"—toxins within cells. Proteins have been identified that can destroy virtually any toxin, using bacteria that can digest and destroy dangerous materials ranging from TNT to dioxin.

A key strategy being pursued by various groups for combating toxic materials outside the cell, including misformed proteins and amyloid plaque (seen in Alzheimer's Disease and other degenerative conditions), is to create vaccines that act against their constituent molecules.⁵¹ Although this approach may result in the toxic material's being ingested by immune system cells, we can then use the strategies for combating intracellular aggregates described above to dispose of it.

Extracellular aggregates. AGEs (advanced glycation end-products) result from undesirable cross-linking of useful molecules as a side effect of excess sugar. These cross-links interfere with the normal functioning of proteins and are key contributors to the aging process. An experimental drug called ALT-711 (phenacyldimethylthiazolium chloride) can dissolve these cross-links without damaging the original tissue.⁵² Other molecules with this capability have also been identified.

Cell loss and atrophy. Our body's tissues have the means to replace worn-out cells, but this ability is limited in certain organs. For example, as we get older, the heart is unable to replace its cells at a sufficient rate, so it compensates by making surviving cells bigger using fibrous material. Over time this causes the heart to become less supple and responsive. A primary strategy here is to deploy therapeutic cloning of our own cells, as described below.

Progress in combating all of these sources of aging is moving rapidly in animal models, and translation into human therapies will follow. Evidence from the genome project indicates that no more than a few hundred genes are involved in the aging process. By manipulating these genes, radical life extension has already been achieved in simpler animals. For example, by modifying genes in the *C. elegans* worm that control its insulin and sex-hormone levels, the lifespan of the test animals was expanded sixfold, to the equivalent of a five-hundred-year lifespan for a human.⁵³

A hybrid scenario involving both bio- and nanotechnology contemplates turning biological cells into computers. These "enhanced intelligence" cells can then detect and destroy cancer cells and pathogens or even regrow human body parts. Princeton biochemist Ron Weiss has modified cells to incorporate a variety of logic functions that are used for basic computation.⁵⁴ Boston University's Timothy Gardner has developed a cellular logic switch, another basic building block for turning cells into computers.⁵⁵ Scientists at the MIT Media Lab have developed ways to use wireless communication to send messages, including intricate sequences of instructions, to the computers inside modified cells.⁵⁶ Weiss points out that "once you have the ability to program cells, you don't have to be constrained by what the cells know how to do already. You can program them to do new things, in new patterns."

Human Cloning: The Least Interesting Application of Cloning Technology

One of the most powerful methods of applying life's machinery involves harnessing biology's own reproductive mechanisms in the form of cloning. Cloning will be a key technology—not for cloning actual humans but for life-extension purposes, in the form of "therapeutic cloning." This process creates new tissues with "young" telomere-extended and DNA-corrected cells to replace without surgery defective tissues or organs.

All responsible ethicists, including myself, consider human cloning at the present time to be unethical. The reasons, however, for me have little to do with the slippery-slope issues of manipulating human life. Rather, the technology today simply does not yet work reliably. The current technique of fusing a cell nucleus from a donor to an egg cell using an electric spark simply causes a high level of genetic errors.⁵⁷ This is the primary reason that most of the fetuses created by this method do not make it to term. Even those that do make it have genetic defects.

Dolly the Sheep developed an obesity problem in adulthood, and the majority of cloned animals produced thus far have had unpredictable health problems.⁵⁸

Scientists have a number of ideas for perfecting cloning, including alternative ways of fusing the nucleus and egg cell without use of a destructive electrical spark, but until the technology is demonstrably safe, it would be unethical to create a human life with such a high likelihood of severe health problems. There is no doubt that human cloning will occur, and occur soon, driven by all the usual reasons, ranging from its publicity value to its utility as a very weak form of immortality. The methods that are demonstrable in advanced animals will work quite well in humans. Once the technology is perfected in terms of safety, the ethical barriers will be feeble if they exist at all.

Cloning is a significant technology, but the cloning of humans is not its most noteworthy usage. Let's first address its most valuable applications and then return to its most controversial one.

<h2>Why is cloning important? <tx>The most immediate use for cloning is improved breeding by offering the ability to directly reproduce an animal with a desirable set of genetic traits. A powerful example is reproducing animals from transgenic embryos (embryos with foreign genes) for pharmaceutical production. A case in point: a promising anticancer treatment is an antiangiogenesis drug called aaATIII, which is produced in the milk of transgenic goats.⁵⁹

<h2>Preserving endangered species and restoring extinct ones. <tx>Another exciting application is re-creating animals from endangered species. By cryopreserving cells from these species, they never need become extinct. It will eventually be possible to re-create animals from recently extinct species. In 2001 scientists were able to synthesize DNA for the Tasmanian tiger, which had then been extinct for sixty-five years, with the hope of bringing this species back to life.⁶⁰ As for long-extinct species (for example, dinosaurs), it is highly doubtful that we will find the fully intact DNA required in a single preserved cell (as they did in the movie Jurassic Park). It is likely, however, that we will eventually be able to synthesize the necessary DNA by patching together the information derived from multiple inactive fragments.

<h2>Therapeutic cloning. <tx>Perhaps the most valuable emerging application is therapeutic cloning of one's own organs. By starting with germ-line cells (inherited from the eggs or sperm and passed on to offspring), genetic engineers can trigger differentiation into diverse types of cells. Because differentiation takes place during the prefetal stage (that is, prior to implantation of a fetus), most ethicists believe this process does not raise concerns, although the issue has remained highly contentious.⁶¹

<h2>Human somatic-cell engineering. <tx>This even more promising approach, which bypasses the controversy of using fetal stem cells entirely, is called transdifferentiation; it creates new tissues with a patient's own DNA by converting one type of cell (such as a skin cell) into another (such as a pancreatic Islet cell or a heart cell).⁶² Scientists from the United States and Norway have recently been successful in reprogramming liver cells into becoming pancreas cells. In another series of experiments, human skin cells were transformed to take on many of the characteristics of immune-system cells and nerve cells.⁶³

Consider the question, What is the difference between a skin cell and any other type of cell in the body? After all, they all have the same DNA. As noted above, the differences are found in protein signaling factors, which include short RNA fragments and peptides, which we are now beginning to understand.⁶⁴ By manipulating these proteins, we can influence gene expression and trick one type of cell into becoming another.

Perfecting this technology would not only defuse a sensitive ethical and political issue but also offer an ideal solution from a scientific perspective. If you need pancreatic Islet cells or kidney tissues—or a even whole new heart—to avoid autoimmune reactions, you would strongly

prefer to obtain these with your own DNA rather than the DNA from someone else's germ-line cells. In addition, this approach uses plentiful skin cells (of the patient) rather than rare and precious stem cells.

Transdifferentiation will directly grow an organ with your genetic makeup. Perhaps most importantly, the new organ can have its telomeres fully extended to their original youthful length, so that the new organ is effectively young again.⁶⁵ We can also correct accumulated DNA errors by selecting the appropriate skin cells (that is ones without DNA errors), prior to transdifferentiation into other types of cells. Using this method an eighty-year-old man could have his heart replaced with the same heart he had when he was, say, twenty-five.

Current treatments for type 1 diabetes require strong antirejection drugs that can have dangerous side effects.⁶⁶ With somatic-cell engineering, type 1 diabetics will be able to make pancreatic Islet cells from their own cells, either from skin cells (transdifferentiation), or from adult stem cells. They would be using their own DNA, and drawing upon a relatively inexhaustible supply of cells, so no antirejection drugs would be required. (But to fully cure type 1 diabetes, we would also have to overcome the patient's autoimmune disorder, which causes his body to destroy Islet cells.)

Even more exciting is the prospect of replacing one's organs and tissues with their "young" replacements without surgery. Introducing cloned, telomere-extended, DNA-corrected cells into an organ will allow them to integrate themselves with the older cells. By repeated treatments of this kind over a period of time, the organ will end up being dominated by the younger cells. We normally replace our own cells on a regular basis anyway, so why not do so with youthful rejuvenated cells rather than telomere-shortened error-filled ones? There's no reason why we couldn't repeat this process for every organ and tissue in our body, enabling us to grow progressively younger.

<h2>Solving world hunger. <tx>Cloning technologies even offer a possible solution for world hunger: creating meat and other protein sources in a factory *without animals* by cloning animal muscle tissue. Benefits would include extremely low cost, avoidance of pesticides and hormones that occur in natural meat, greatly reduced environmental impact (compared to factory farming), improved nutritional profile, and no animal suffering. As with therapeutic cloning, we would not be creating the entire animal but rather directly producing the desired animal parts or flesh. Essentially, all of the meat—billions of pounds of it—would be derived from a single animal.

There are other benefits to this process besides ending hunger. By creating meat in this way, it becomes subject to the law of accelerating returns—the exponential improvements in price-performance of information-based technologies over time—and will thus become extremely inexpensive. Even though hunger in the world today is certainly exacerbated by political issues and conflicts, meat could become so inexpensive that it would have a profound effect on the affordability of food.

The advent of animal-less meat will also eliminate animal suffering. The economics of factory farming place a very low priority on the comfort of animals, which are treated as cogs in a machine. The meat produced in this manner, although normal in all other respects, would not be part of an animal with a nervous system, which is generally regarded as a necessary element for suffering to occur, at least in a biological animal. We could use the same approach to produce such animal by-products as leather and fur. Other major advantages would be to eliminate the enormous ecological and environmental damage created by factory farming as well as the risk of prion-based diseases, such as mad-cow disease and its human counterpart, vCJD.⁶⁷

Human cloning revisited. This brings us again to human cloning. I predict that once the technology is perfected, neither the acute dilemmas seen by ethicists nor the profound promise

heralded by enthusiasts will predominate. So what if we have genetic twins separated by one or more generations? Cloning is likely to prove to be like other reproductive technologies that were briefly controversial but rapidly accepted. Physical cloning is far different from mental cloning, in which a person's entire personality, memory, skills, and history will ultimately be downloaded into a different, and most likely more powerful, thinking medium. There's no issue of philosophical identity with genetic cloning, since such clones would be different people, even more so than conventional twins are today.

If we consider the full concept of cloning, from cell to organisms, its benefits have enormous synergy with the other revolutions occurring in biology as well as in computer technology. As we learn to understand the genome and proteome (the expression of the genome into proteins) of both humans and animals, and as we develop powerful new means of harnessing genetic information, cloning provides the means to replicate animals, organs, and cells. And that has profound implications for health and well-being of both ourselves and our evolutionary cousins in the animal kingdom.

Ned Ludd: If everyone can change their genes, then everyone will choose to be "perfect" in every way, so there'll be no diversity and excelling will become meaningless.

Ray: Not exactly. Genes are obviously important, but our nature—skills, knowledge, memory, personality—reflects the design information in our genes, as our bodies and brains self-organize through our experience. This is also readily evident in our health. I personally have a genetic disposition to type 2 diabetes, having been actually diagnosed with that disease more than twenty years ago. But I don't have any indication of diabetes today because I've overcome this genetic disposition as a result of reprogramming my biochemistry through lifestyle choices such as nutrition, exercise, and aggressive supplementation. With regard to our brains, we all have various aptitudes, but our actual talents are a function of what we've learned, developed, and experienced. Our genes reflect dispositions only. We can see how this works in the development of the brain. The genes describe certain rules and constraints for patterns of interneuronal connections, but the actual connections we have as adults are the result of a self-organizing process based on our learning. The final result—who we are—is deeply influenced by both nature (genes) and nurture (experience).

So when we gain the opportunity to change our genes as adults, we won't wipe out the influence of our earlier genes. Experiences prior to the gene therapy will have been translated through the pretherapy genes, so one's character and personality would still be shaped primarily by the original genes. For example, if someone added genes for musical aptitude to his brain through gene therapy, he would not suddenly become a music genius.

Ned: Okay, I understand that designer baby boomers can't get away completely from their pre-designer genes, but with designer babies they'll have the genes and the time to express them.

Ray: The "designer baby" revolution is going to be a very slow one; it won't be a significant factor in this century. Other revolutions will overtake it. We won't have the technology for designer babies for another ten to twenty years. To the extent that it is used, it would be adopted gradually, and then it will take those generations another twenty years to reach maturity. By that time, we're approaching the Singularity, with the real revolution being the predominance of nonbiological intelligence. That will go far beyond the capabilities of any designer genes. The idea of designer babies and baby boomers is just the reprogramming of the information processes in biology. But it's still biology, with all its profound limitations.

Ned: You're missing something. Biological is what we are. I think most people would agree that being biological is the quintessential attribute of being human.

Ray: That's certainly true today.

Ned: And I plan to keep it that way.

Ray: Well, if you're speaking for yourself, that's fine with me. But if you stay biological and don't reprogram your genes, you won't be around for very long to influence the debate.

Nanotechnology: The Intersection of Information and the Physical World

The role of the infinitely small is infinitely large.

—Louis Pasteur

But I am not afraid to consider the final question as to whether, ultimately, in the great future, we can arrange the atoms the way we want; the very atoms, all the way down!

—Richard Feynman

<epi>Nanotechnology has the potential to enhance human performance, to bring sustainable development for materials, water, energy, and food, to protect against unknown bacteria and viruses, and even to diminish the reasons for breaking the peace [by creating universal abundance].

<epis>—*National Science Foundation Nanotechnology Report*

<tx>Nanotechnology promises the tools to rebuild the physical world—our bodies and brains included—molecular fragment by molecular fragment, potentially atom by atom. We are shrinking the key feature size of technology, in accordance with the law of accelerating returns, at the exponential rate of approximately a factor of four per linear dimension per decade.⁶⁸ At this rate the key feature sizes for most electronic and many mechanical technologies will be in the nanotechnology range—generally considered to be under one hundred nanometers—by the 2020s. (Electronics has already dipped below this threshold, although not yet in three-dimensional structures and not yet self-assembling.) Meanwhile rapid progress has been made, particularly in the last several years, in preparing the conceptual framework and design ideas for the coming age of nanotechnology.

As important as the biotechnology revolution discussed above will be, once its methods are fully mature, limits will be encountered in biology itself. Although biological systems are remarkable in their cleverness, we have also discovered that they are dramatically suboptimal. I've mentioned the extremely slow speed of communication in the brain, and as I discuss below, robotic replacements for our red blood cells could be thousands of times more efficient than their biological counterpart.⁶⁹ Biology will never be able to match what we will be capable of engineering once we fully understand biology's principles of operation.

The revolution in nanotechnology, however, will ultimately enable us to redesign and rebuild, molecule by molecule, our bodies and brains and the world with which we interact.⁷⁰ These two revolutions are overlapping, but the full realization of nanotechnology lags behind the biotechnology revolution by about one decade.

Most nanotechnology historians date the conceptual birth of nanotechnology to physicist Richard Feynman's seminal speech in 1959, "There's Plenty of Room at the Bottom," in which he described the inevitability and profound implications of engineering machines at the level of atoms:

<ext>The principles of physics, as far as I can see, do not speak against the possibility of maneuvering things atom by atom. It would be, in principle, possible . . . for a physicist to synthesize any chemical substance that the chemist writes down. . . . How? Put the atoms down where the chemist says, and so you make the substance. The problems of chemistry and biology can be greatly helped if our ability to see what we are doing, and to do things on an atomic level, is ultimately developed—a development which I think cannot be avoided.⁷¹

An even earlier conceptual foundation for nanotechnology was formulated by the information theorist John von Neumann in the early 1950s with his model of a self-replicating system based on a universal constructor, combined with a universal computer.⁷² In this proposal the computer runs a program that directs the constructor, which in turn constructs a copy of both the computer (including its self-replication program) and the constructor. At this level of description von Neumann's proposal is quite abstract—the computer and constructor could be made in a great variety of ways, as well as from diverse materials, and could even be a theoretical mathematical construction. But he took the concept one step further and proposed a

"kinematic constructor": a robot with at least one manipulator (arm) that would build a replica of itself from a "sea of parts" in its midst.⁷³

It was left to Eric Drexler to found the modern field of nanotechnology, with a draft of his landmark Ph.D. thesis in the mid-1980s, in which he essentially combined these two intriguing suggestions. Drexler described a von Neumann kinematic constructor, which for its sea of parts used atoms and molecular fragments, as suggested in Feynman's speech. Drexler's vision cut across many disciplinary boundaries and was so far-reaching that no one was daring enough to be his thesis adviser except for my own mentor, Marvin Minsky. Drexler's dissertation (which became his book *Engines of Creation* in 1986 and was articulated technically in his 1992 book, *Nanosystems*) laid out the foundation of nanotechnology and provided the road map still being followed today.⁷⁴

Drexler's "molecular assembler" will be able to make almost anything in the world. It has been referred to as a "universal assembler," but Drexler and other nanotechnology theorists do not use the word universal because the products of such a system necessarily have to be subject to the laws of physics and chemistry, so only atomically stable structures would be viable. Furthermore, any specific assembler would be restricted to building products from its sea of parts, although the feasibility of using individual atoms has been shown. Nevertheless, such an assembler could make just about any physical device we would want, including highly efficient computers, and subsystems for other assemblers.

Although Drexler did not provide a detailed design for an assembler—such a design has still not been fully specified—his thesis did provide extensive feasibility arguments for each of the principal components of a molecular assembler, which include the following subsystems:

<bl>

- The *computer*: to provide the intelligence to control the assembly process. As with all of the device's subsystems, the computer needs to be small and simple. As I described in chapter 3, Drexler provides an intriguing conceptual description of a mechanical computer with molecular "locks" instead of transistor gates. Each lock would require only sixteen cubic nanometers of space and could switch ten billion times per second. This proposal remains more competitive than any known electronic technology, although electronic computers built from three-dimensional arrays of carbon nanotubes appear to provide even higher densities of computation (that is, calculations per second per gram).⁷⁵
- The *instruction architecture*: Drexler and his colleague Ralph Merkle have proposed a SIMD (single instruction multiple data) architecture in which a single data store would record the instructions and transmit them to trillions of molecular-sized assemblers (each with its own simple computer) simultaneously. I discussed some of the limitations of the SIMD architecture in chapter 3, but this design (which is easier to implement than the more flexible multiple-instruction multiple-data approach) is sufficient for the computer in a universal nanotechnology assembler. With this approach each assembler would not have to store the entire program for creating the desired product. A "broadcast" architecture also addresses a key safety concern: the self-replication process could be shut down, if it got out of control, by terminating the centralized source of the replication instructions.

However, as Drexler points out, a nanoscale assembler does not necessarily have to be self-replicating.⁷⁶ Given the inherent dangers in self-replication, the ethical standards proposed by the Foresight Institute (a think tank founded by Eric Drexler and Christine Peterson) contain prohibitions against unrestricted self-replication, especially in a natural environment. As I will discuss in chapter 8, this approach should be reasonably

effective against inadvertent dangers, although it could be circumvented by a determined and knowledgeable adversary.

- *Instruction transmission*: Transmission of the instructions from the centralized data store to each of the many assemblers would be accomplished electronically if the computer is electronic or through mechanical vibrations if Drexler's concept of a mechanical computer were used.
- The *construction robot*: The constructor would be a simple molecular robot with a single arm, similar to von Neumann's kinematic constructor but on a tiny scale. There are already examples of experimental molecular scale systems that can act as motors and robot legs, as I discuss below.
- The *robot arm tip*: Drexler's *Nanosystems* provided a number of feasible chemistries for the tip of the robot arm to make it capable of grasping (using appropriate atomic-force fields) a molecular fragment, or even a single atom, and then depositing it in a desired location. In the chemical-vapor deposition process used to construct artificial diamonds, individual carbon atoms, as well as molecular fragments, are moved to other locations through chemical reactions at the tip. Building artificial diamonds is a chaotic process involving trillions of atoms, but conceptual proposals by Robert Freitas and Ralph Merkle contemplate robot arm tips that can remove hydrogen atoms from a source material and deposit them at desired locations in the construction of a molecular machine. In this proposal, the tiny machines are built out of a diamondoid material. In addition to having great strength, the material can be doped with impurities in a precise fashion to create electronic components such as transistors. Simulations have shown that such molecular scale gears, levers, motors, and other mechanical systems would operate properly as intended.⁷⁷ More recently attention has been focused on carbon nanotubes, comprising hexagonal arrays of carbon atoms assembled in three dimensions, which are also capable of providing both mechanical and electronic functions at the molecular level. I provide examples below of molecular scale machines that have already been built.
- The assembler's *internal environment* needs to prevent environmental impurities from interfering with the delicate assembly process. Drexler's proposal is to maintain a near vacuum and build the assembler walls out of the same diamondoid material that the assembler itself is capable of making.
- The *energy* required for the assembly process can be provided either through electricity or through chemical energy. Drexler proposed a chemical process with the fuel interlaced with the raw building material. More recent proposals use nanoengineered fuel cells incorporating hydrogen and oxygen or glucose and oxygen, or acoustic power at ultrasonic frequencies.⁷⁸

Although many configurations have been proposed, the typical assembler has been described as a tabletop unit that can manufacture almost any physically possible product for which we have a software description, ranging from computers, clothes, and works of art to cooked meals.⁷⁹ Larger products, such as furniture, cars, or even houses, can be built in a modular fashion or using larger assemblers. Of particular importance is the fact that an assembler can create copies of itself, unless its design specifically prohibits this (to avoid potentially dangerous self-replication). The incremental cost of creating any physical product, including the assemblers themselves, would be pennies per pound—basically the cost of the raw materials. Drexler estimates total manufacturing cost for a molecular-manufacturing process in the range of four cents to twenty cents per kilogram, regardless of whether the manufactured product were clothing, massively parallel supercomputers, or additional manufacturing systems.⁸⁰

The real cost, of course, would be the value of the information describing each type of product—that is, the software that controls the assembly process. In other words, the value of everything in the world, including physical objects, would be based essentially on information. We are not that far from this situation today, since the information content of products is rapidly increasing, gradually approaching an asymptote of 100 percent of their value.

The design of the software controlling molecular manufacturing systems would *itself* be extensively automated, much as chip design is today. Chip designers don't specify the location of each of the billions of wires and components but rather the specific functions and features, which computer-aided design (CAD) systems translate into actual chip layouts. Similarly, CAD systems would produce the molecular-manufacturing control software from high-level specifications. This would include the ability to reverse engineer a product by scanning it in three dimensions and then generating the software needed to replicate its overall capabilities.

In operation, the centralized data store would send out commands simultaneously to many trillions (some estimates as high as 10^{18}) of robots in an assembler, each receiving the same instruction at the same time. The assembler would create these molecular robots by starting with a small number and then using these robots to create additional ones in an iterative fashion, until the requisite number had been created. Each robot would have a local data storage that specifies the type of mechanism it's building. This storage would be used to mask the global instructions being sent from the centralized data store so that certain instructions are blocked and local parameters are filled in. In this way, even though all of the assemblers are receiving the same sequence of instructions, there is a level of customization to the part being built by each molecular robot. This process is analogous to gene expression in biological systems. Although every cell has every gene, only those genes relevant to a particular cell type are expressed. Each robot extracts the raw materials and fuel it needs, which include individual carbon atoms and molecular fragments, from the source material.

The Biological Assembler

<epi>Nature shows that molecules can serve as machines because living things work by means of such machinery. Enzymes are molecular machines that make, break, and rearrange the bonds holding other molecules together. Muscles are driven by molecular machines that haul fibers past one another. DNA serves as a data-storage system, transmitting digital instructions to molecular machines, the ribosomes, that manufacture protein molecules. And these protein molecules, in turn, make up most of the molecular machinery.

<epis>—*Eric Drexler*

The ultimate existence proof of the feasibility of a molecular assembler is life itself. Indeed, as we deepen our understanding of the information basis of life processes, we are discovering specific ideas that are applicable to the design requirements of a generalized molecular assembler. For example, proposals have been made to use a molecular energy source of glucose and ATP, similar to that used by biological cells.


Consider how biology solves each of the design challenges of a Drexler assembler. The ribosome represents both the computer and the construction robot. Life does not use centralized data storage but provides the entire code to every cell. The ability to restrict the local data storage of a nanoengineered robot to only a small part of the assembly code (using the "broadcast" architecture), particularly when doing self-replication, is one critical way nanotechnology can be engineered to be safer than biology.

Life's local data storage is, of course, the DNA strands, broken into specific genes on the chromosomes. The task of instruction masking (blocking genes that do not contribute to a

particular cell type) is controlled by the short RNA molecules and peptides that govern gene expression. The internal environment in which the ribosome is able to function is the particular chemical environment maintained inside the cell, which includes a particular acid-alkaline equilibrium (pH around 7 in human cells) and other chemical balances. The cell membrane is responsible for protecting this internal environment from disturbance.

Upgrading the cell nucleus with a nanocomputer and nanobot. Here's a conceptually simple proposal to overcome all biological pathogens except for prions (self-replicating pathological proteins). With the advent of full-scale nanotechnology in the 2020s we will have the potential to replace biology's genetic-information repository in the cell nucleus with a nanoengineered system that would maintain the genetic code and simulate the actions of RNA, the ribosome, and other elements of the computer in biology's assembler. A nanocomputer would maintain the genetic code and implement the gene-expression algorithms. A nanobot would then construct the amino-acid sequences for the expressed genes.

There would be significant benefits in adopting such a mechanism. We could eliminate the accumulation of DNA transcription errors, one major source of the aging process. We could introduce DNA changes to essentially reprogram our genes (something we'll be able to do long before this scenario, using gene-therapy techniques). We would also be able to defeat biological pathogens (bacteria, viruses, and cancer cells) by blocking any unwanted replication of genetic information.



With such a nanoengineered system the recommended broadcast architecture would enable us to turn off unwanted replication, thereby defeating cancer, autoimmune reactions, and other disease processes. Although most of these disease processes will have already been vanquished by the biotechnology methods described in the previous section, reengineering the computer of life using nanotechnology could eliminate any remaining obstacles and create a level of durability and flexibility that goes beyond the inherent capabilities of biology.

The robot arm tip would use the ribosome's ability to implement enzymatic reactions to break off an individual amino acid, each of which is bound to a specific tRNA, and to connect it to its adjoining amino acid using a peptide bond. Thus, such a system could utilize portions of the ribosome itself, since this biological machine is capable of constructing the requisite string of amino acids.

However, the goal of molecular manufacturing is not merely to replicate the molecular-assembly capabilities of biology. Biological systems are limited to building systems from protein, which has profound limitations in strength and speed. Although biological proteins are three-dimensional, biology is restricted to that class of chemicals that can be folded from a one-dimensional string of amino acids. Nanobots built from diamondoid gears and rotors can also be thousands of times faster and stronger than biological cells.

The comparison is even more dramatic with regard to computation: the switching speed of nanotube-based computation would be millions of times faster than the extremely slow transaction speed of the electrochemical switching used in mammalian interneuronal connections.

The concept of a diamondoid assembler described above uses a consistent input material (for construction and fuel), which represents one of several protections against molecule-scale replication of robots in an uncontrolled fashion in the outside world. Biology's replication robot, the ribosome, also requires carefully controlled source and fuel materials, which are provided by our digestive system. As nanobased replicators become more sophisticated, more capable of extracting carbon atoms and carbon-based molecular fragments from less well-controlled source materials, and able to operate outside of controlled replicator enclosures such as in the biological world, they will have the potential to present a grave threat to that world. This is particularly true in view of the vastly greater strength and speed of nanobased replicators over any biological system. That ability is, of course, the source of great controversy, which I discuss in chapter 8.

In the decade since publication of Drexler's *Nanosystems*, each aspect of Drexler's conceptual designs has been validated through additional design proposals,⁸¹ supercomputer simulations, and, most important, actual construction of related molecular machines. Boston College chemistry professor T. Ross Kelly reported that he constructed a chemically powered nanomotor out of seventy-eight atoms.⁸² A biomolecular research group headed by Carlo Montemagno created an ATP-fueled nanomotor.⁸³ Another molecule-sized motor fueled by solar energy was created out of fifty-eight atoms by Ben Feringa at the University of Groningen in the Netherlands.⁸⁴ Similar progress has been made on other molecular-scale mechanical components such as gears, rotors, and levers. Systems demonstrating the use of chemical energy and acoustic energy (as originally described by Drexler) have been designed, simulated, and actually constructed. Substantial progress has also been made in developing various types of electronic components from molecule-scale devices, particularly in the area of carbon nanotubes, an area that Richard Smalley has pioneered.

Nanotubes are also proving to be very versatile as a structural component. A conveyor belt constructed out of nanotubes was demonstrated recently by scientists at Lawrence Berkeley National Laboratory.⁸⁵ The nanoscale conveyor belt was used to transport tiny indium particles from one location to another, although the technique could be adapted to move a variety of molecule-size objects. By controlling an electrical current applied to the device, the direction and velocity of movement can be modulated. "It's the equivalent of turning a knob . . . and taking macroscale control of nanoscale mass transport," said Chris Regan, one of the designers. "And it's reversible: we can change the current's polarity and drive the indium back to its original position." The ability to rapidly shuttle molecule-size building blocks to precise locations is a key step toward building molecular assembly lines.

A study conducted for NASA by General Dynamics has demonstrated the feasibility of self-replicating nanoscale machines.⁸⁶ Using computer simulations, the researchers showed that molecularly precise robots called kinematic cellular automata, built from reconfigurable molecular modules, were capable of reproducing themselves. The designs also used the broadcast architecture, which established the feasibility of this safer form of self-replication.

DNA is proving to be as versatile as nanotubes for building molecular structures. DNA's proclivity to link up with itself makes it a useful structural component. Future designs may combine this attribute as well as its capacity for storing information. Both nanotubes and DNA have outstanding properties for information storage and logical control, as well as building strong three-dimensional structures.

A research team at Ludwig Maximilians University in Munich has built a "DNA hand" that can select one of several proteins, bind to it, and then release it upon command.⁸⁷ Important steps in creating a DNA assembler mechanism akin to the ribosome was demonstrated recently by nanotechnology researchers Shiping Liao and Nadrian Seeman.⁸⁸ Grasping and letting go of molecular objects in a controlled manner is another important enabling capability for molecular nanotechnology assembly.

Scientists at the Scripps Research Institute demonstrated the ability to create DNA building blocks by generating many copies of a 1,669-nucleotide strand of DNA that had carefully placed self-complementary regions.⁸⁹ The strands self-assembled spontaneously into rigid octahedrons, which could be used as blocks for elaborate three-dimensional structures. Another application of this process could be to employ the octahedrons as compartments to deliver proteins, which Gerald F. Joyce, one of the Scripps researchers, called a "virus in reverse." Viruses, which are also self-assembling, usually have outer shells of protein with DNA (or RNA) on the inside. "With this," Joyce points out, "you could in principle have DNA on the outside and proteins on the inside."

A particularly impressive demonstration of a nanoscale device constructed from DNA is a tiny biped robot that can walk on legs that are ten nanometers long.⁹⁰ Both the legs and the walking track are built from DNA, again chosen for the molecule's ability to attach and detach itself in a controlled manner. The nanorobot, a project of chemistry professors Nadrian Seeman and William Sherman of New York University, walks by detaching its legs from the track, moving down it, and then reattaching the legs to the track. The project is another impressive demonstration of the ability of nanoscale machines to execute precise maneuvers.

An alternate method of designing nanobots is to learn from nature. Nanotechnologist Michael Simpson of Oak Ridge National Laboratory describes the possibility of exploiting bacteria "as ready-made machine[s]." Bacteria, which are natural nanobot-size objects, are able to move, swim, and pump liquids.⁹¹ Linda Turner, a scientist at the Rowland Institute at Harvard, has focused on their thread-size arms, called fimbriae, which are able to perform a wide variety of tasks, including carrying other nanoscale objects and mixing fluids. Another approach is to use only parts of bacteria. A research group headed by Viola Vogel at the University of Washington built a system using just the limbs of *E. coli* bacteria that was able to sort out nanoscale beads of different sizes. Since bacteria are natural nanoscale systems that can perform a wide variety of functions, the ultimate goal of this research will be to reverse engineer the bacteria so that the same design principles can be applied to our own nanobot designs.

Fat and Sticky Fingers

<tx>In the wake of the rapidly expanding development of each facet of future nanotechnology systems, no serious flaw in Drexler's nano assembler concept has been described. A highly publicized objection in 2001 by Nobelist Richard Smalley in *Scientific American* was based on a distorted description of the Drexler proposal;⁹² it ignored the extensive body of work that has been carried out in the past decade. As a pioneer of carbon nanotubes Smalley has been enthusiastic about a variety of applications of nanotechnology, having written that "nanotechnology holds the answer, to the extent there are answers, to most of our pressing

material needs in energy, health, communication, transportation, food, water," but he remains skeptical about molecular nanotechnology assembly.

Smalley describes Drexler's assembler as consisting of five to ten "fingers" (manipulator arms) to hold, move, and place each atom in the machine being constructed. He then goes on to point out that there isn't room for so many fingers in the cramped space in which a molecular assembly nanorobot has to work (which he calls the "fat fingers" problem) and that these fingers would have difficulty letting go of their atomic cargo because of molecular attraction forces (the "sticky fingers" problem). Smalley also points out that an "intricate three-dimensional waltz . . . is carried out" by five to fifteen atoms in a typical chemical reaction.

In fact, Drexler's proposal doesn't look anything like the straw-man description that Smalley criticizes. Drexler's proposal, and most of those that have followed, uses a single "finger." Moreover, there have been extensive description and analyses of viable tip chemistries that do not involve grasping and placing atoms as if they were mechanical pieces to be deposited in place. In addition to the examples I provided above (for example, the DNA hand), the feasibility of moving hydrogen atoms using Drexler's "propynyl hydrogen abstraction" tip has been extensively confirmed in the intervening years.⁹³ The ability of the scanning-probe microscope (SPM), developed at IBM in 1981 and the more sophisticated atomic-force microscope (AFM) to place individual atoms through specific reactions of a tip with a molecular-scale structure provide additional proof of the concept. Recently, scientists at Osaka University used an AFM to move individual nonconductive atoms using a mechanical rather than electrical technique.⁹⁴ The ability to move both conductive and nonconductive atoms and molecules will be needed for future molecular nanotechnology.⁹⁵

Indeed, if Smalley's critique were valid, none of us would be here to discuss it, because life itself would be impossible, given that biology's assembler does exactly what Smalley says is impossible.

Smalley also objects that, despite "working furiously, . . . generating even a tiny amount of a product would take [a nanobot] . . . millions of years." Smalley is correct, of course, that an assembler with only one nanobot wouldn't produce any appreciable quantities of a product. However, the basic concept of nanotechnology is that we will use trillions of nanobots to accomplish meaningful results—a factor that is also the source of the safety concerns that have received so much attention. Creating this many nanobots at reasonable cost will require self-replication at some level, which while solving the economic issue will introduce potentially grave dangers, a concern I will address in chapter 8. Biology uses the same solution to create organisms with trillions of cells, and indeed we find that virtually all diseases derive from biology's self-replication process gone awry.

Earlier challenges to the concepts underlying nanotechnology have also been effectively addressed. Critics pointed out that nanobots would be subject to bombardment by thermal vibration of nuclei, atoms, and molecules. This is one reason conceptual designers of nanotechnology have emphasized building structural components from diamondoid or carbon nanotubes. Increasing the strength or stiffness of a system reduces its susceptibility to thermal effects. Analysis of these designs has shown them to be thousands of times more stable in the presence of thermal effects than are biological systems, so they can operate in a far wider temperature range.⁹⁶

Similar challenges were made regarding positional uncertainty from quantum effects, based on the extremely small feature size of nanoengineered devices. Quantum effects are significant for an electron, but a single carbon-atom nucleus is more than twenty thousand times more massive than an electron. A nanobot will be constructed from millions to billions of carbon and other atoms, making it up to trillions of times more massive than an electron. Plugging this

ratio in the fundamental equation for quantum positional uncertainty shows it to be an insignificant factor.⁹⁷

Power has represented another challenge. Proposals involving glucose-oxygen fuel cells have held up well in feasibility studies by Freitas and others.⁹⁸ An advantage of the glucose-oxygen approach is that nanomedicine applications can harness the glucose, oxygen, and ATP resources already provided by the human digestive system. A nanoscale motor was recently created using propellers made of nickel and powered by an ATP-based enzyme." However, recent progress in implementing MEMS-scale and even nanoscale hydrogen-oxygen fuel cells has provided an alternative approach, which I report on below.

The Debate Heats Up

<tx>In April 2003 Drexler challenged Smalley's *Scientific American* article with an open letter.¹⁰⁰ Citing twenty years of research by himself and others, the letter responded specifically to Smalley's fat- and sticky-fingers objections. As I discussed above, molecular assemblers were never described as having fingers at all but rather relying on precise positioning of reactive molecules. Drexler cited biological enzymes and ribosomes as examples of precise molecular assembly in the natural world. Drexler closed by quoting Smalley's own observation, "When a scientist says something is possible, they're probably underestimating how long it will take. But if they say it's impossible, they're probably wrong."

Three more rounds of this debate occurred in 2003. Smalley responded to Drexler's open letter by backing off of his fat- and sticky-fingers objections and acknowledging that enzymes and ribosomes do indeed engage in the precise molecular assembly that Smalley had earlier indicated was impossible. Smalley then argued that biological enzymes work only in water and that such water-based chemistry is limited to biological structures such as "wood, flesh and bone." As Drexler has stated, this, too, is erroneous.¹⁰¹ Many enzymes, even those that ordinarily work in water, can also function in anhydrous organic solvents, and some enzymes can operate on substrates in the vapor phase, with no liquid at all.¹⁰²

Smalley goes on to state (without any derivation or citations) that enzymatic-like reactions can take place only with biological enzymes and in chemical reactions involving water. This is also mistaken. MIT Professor of Chemistry and Biological Engineering Alexander Klivanov demonstrated such nonaqueous (not involving water) enzyme catalysis in 1984. Klivanov writes in 1993 "clearly [Smalley's] statements about nonaqueous enzyme catalysis are incorrect. There have been hundreds and perhaps thousands of papers published about nonaqueous enzyme catalysis since our first paper was published 20 years ago."¹⁰³

It's easy to see why biological evolution adopted water-based chemistry. Water is a very abundant substance on our planet, and comprises 70 to 90 percent of our bodies, our food, and indeed of all organic matter. These three-dimensional electrical properties of water are quite powerful and can break apart the strong chemical bonds of other compounds. Water is considered "the universal solvent," and because it is involved in most of the biochemical pathways in our bodies we can regard the chemistry of life on our planet primarily as water chemistry. However, the primary thrust of our technology has been to develop systems that are not limited to the restrictions of biological evolution, which exclusively adopted water-based chemistry and proteins as its foundation. Biological systems can fly, but if you want to fly at thirty thousand feet and at hundreds or thousands of miles per hour, you would use our modern technology, not proteins. Biological systems such as human brains can remember things and do calculations, but if you want to do data mining on billions of items of information, you would want to use electronic technology, not unassisted human brains.

Smalley is ignoring the past decade of research on alternative means of positioning molecular fragments using precisely guided molecular reactions. Precisely controlled synthesis of diamondoid material has been extensively studied, including the ability to remove a single hydrogen atom from a hydrogenated diamond surface¹⁰⁴ and the ability to add one or more carbon atoms to a diamond surface.¹⁰⁵ Related research supporting the feasibility of hydrogen abstraction and precisely guided diamondoid synthesis has been conducted at the Materials and Process Simulation Center at Cal Tech; the department of materials science and engineering at North Carolina State University; the Institute for Molecular Manufacturing at the University of Kentucky; the U.S. Naval Academy; and the Xerox Palo Alto Research Center.¹⁰⁶

Smalley is also ignoring the well-established SPM mentioned above, which uses precisely controlled molecular reactions. Building on these concepts, Ralph Merkle has described possible tip reactions that could involve up to four reactants.¹⁰⁷ There is an extensive literature on site-specific reactions that have the potential to be precisely guided and thus could be feasible for the tip chemistry in a molecular assembler.¹⁰⁸ Smalley overlooks this body of literature when he maintains that only biological enzymes in water can perform this type of reaction. Recently, many tools that go beyond SPMs are emerging that can reliably manipulate atoms and molecular fragments.

On September 3, 2003, Drexler responded to Smalley's response to his initial letter by alluding once again to the extensive body of literature that Smalley ignores.¹⁰⁹ He cited the analogy to a modern factory, only at a nanoscale. He cited analyses of transition-state theory indicating that positional control would be feasible at megahertz frequencies for appropriately selected reactants.

Smalley again responded with a letter that is short on specifics and science and long on imprecise metaphors.¹¹⁰ He writes, for example, that "much like you can't make a boy and a girl fall in love with each other simply by pushing them together, you cannot make precise chemistry occur as desired between two molecular objects with simple mechanical motion. . . . [It] cannot be done simply by mashing two molecular objects together." He again acknowledges that enzymes do in fact accomplish this but refuses to accept that such reactions could take place outside of a biologylike system: "This is why I led you . . . to talk about real chemistry with real enzymes. . . . [A]ny such system will need a liquid medium. For the enzymes we know about, that liquid will have to be water, and the types of things that can be synthesized with water around cannot be much broader than meat and bone of biology."

I can understand Drexler's frustration in this debate because I have had many critics that do not bother to read or understand the data and arguments that I have presented for my own conceptions of future technologies. Smalley's argument is of the form, "We don't have X today, therefore X is impossible." We encounter this class of argument repeatedly in the area of artificial intelligence. Critics will cite the limitations of today's systems as proof that such limitations are inherent and can never be overcome. For example, such critics disregard the extensive list of contemporary examples of AI (see the next section) that represent commercially available working systems that were only research programs a decade ago.

Those of us who attempt to project into the future based on well-grounded methodologies are at a disadvantage. Certain future realities may be inevitable, but they are not yet manifest, so they are easy to deny. A small body of thought at the beginning of the twentieth century insisted that heavier-than-air flight was feasible, but mainstream skeptics could simply point out that if it was so feasible, why had it never been demonstrated?

Smalley reveals at least part of his motives at the end of his most recent letter, when he writes:

<ext>A few weeks ago I gave a talk on nanotechnology and energy titled "Be a Scientist, Save the World" to about 700 middle and high school students in the Spring Branch ISD, a large public school system here in the Houston area. . Leading up to my visit the students were asked to write an essay on "why I am a Nanogeek". Hundreds responded, and I had the privilege of reading the top 30 essays, picking my favorite top 5. Of the essays I read, nearly half assumed that self-replicating nanobots were possible, and most were deeply worried about what would happen in their future as these nanobots spread around the world. I did what I could to allay their fears, but there is no question that many of these youngsters have been told a bedtime story that is deeply troubling.

You and people around you have scared our children.

<tx>I would point out to Smalley that earlier critics also expressed skepticism that either worldwide communication networks or software viruses that would spread across them were feasible. Today, we have both the benefits and the vulnerabilities from these capabilities. However, along with the danger of software viruses has emerged a technological immune system.. We are obtaining far more gain than harm from this latest example of intertwined promise and peril.

Smalley's approach to reassuring the public about the potential abuse of this future technology is not the right strategy. By denying the feasibility of nanotechnology based assembly, he is also denying its potential. Denying both the promise and the peril of molecular assembly will ultimately backfire and will fail to guide research in the needed constructive direction. By the 2020s molecular assembly will provide tools to effectively combat poverty, clean up our environment, overcome disease, extend human longevity, and many other worthwhile pursuits. Like every other technology that humankind has created, it can also be used to amplify and enable our destructive side. It's important that we approach this technology in a knowledgeable manner to gain the profound benefits it promises, while avoiding its dangers.

Early Adopters

Although Drexler's concept of nanotechnology dealt primarily with precise molecular control of manufacturing, it has expanded to include any technology in which key features are measured by a modest number (generally less than one hundred) of nanometers. Just as contemporary electronics has already quietly slipped into this realm the area of biological and medical applications has already entered the era of nanoparticles, in which nanoscale objects are being developed to create more effective tests and treatments. Although nanoparticles are created using statistical manufacturing methods rather than assemblers, they nonetheless rely on their atomic-scale properties for their effects. For example, nanoparticles are being employed in experimental biological tests as tags and labels to greatly enhance sensitivity in detecting substances such as proteins. Magnetic nanotags, for example, can be used to bind with antibodies, which can then be read using magnetic probes while still inside the body. Successful experiments have been conducted with gold nanoparticles that are bound to DNA segments and can rapidly test for specific DNA sequences in a sample. Small nanoscale beads called quantum dots can be programmed with specific codes combining multiple colors, similar to a color bar code, which can facilitate tracking of substances through the body.

Emerging microfluidic devices, which incorporate nanoscale channels, can run hundreds of tests simultaneously on tiny samples of a given substance. These devices will allow extensive tests to be conducted on nearly invisible samples of blood, for example.

Nanoscale scaffolds have been used to grow biological tissues such as skin. Future therapies could use these tiny scaffolds to grow any type of tissue needed for repairs inside the body.

A particularly exciting application is to harness nanoparticles to deliver treatments to specific sites in the body. Nanoparticles can guide drugs into cell walls and through the blood-brain barrier. Scientists at McGill University in Montreal demonstrated a nanopill with structures in the 25- to 45-nanometer range.¹¹¹ The nanopill is small enough to pass through the cell wall and delivers medications directly to targeted structures inside the cell.

Japanese scientists have created nanocages of 110 amino-acid molecules, each holding drug molecules. Adhered to the surface of each nanocage is a peptide that binds to target sites in the human body. In one experiment scientists used a peptide that binds to a specific receptor on human liver cells.¹¹²

MicroCHIPS of Bedford, Massachusetts, has developed a computerized device that is implanted under the skin and delivers precise mixtures of medicines from hundreds of nanoscale wells inside the device.¹¹³ Future versions of the device are expected to be able to measure blood levels of substances such as glucose. The system could be used as an artificial pancreas, releasing precise amounts of insulin based on blood glucose response. It would also be capable of simulating any other hormone-producing organ. If trials go smoothly, the system could be on the market by 2008.

Another innovative proposal is to guide gold nanoparticles to a tumor site, then heat them with infrared beams to destroy the cancer cells. Nanoscale packages can be designed to contain drugs, protect them through the GI tract, guide them to specific locations, and then release them in sophisticated ways, including allowing them to receive instructions from outside the body. Nanotherapeutics in Alachua, Florida, has developed a biodegradable polymer only several nanometers thick that uses this approach.¹¹⁴

Powering the Singularity

We produce about 14 trillion (about 10^{13}) watts of power today in the world. Of this energy about 33 percent comes from oil, 25 percent from coal, 20 percent from gas, 7 percent from nuclear fission reactors, 15 percent from biomass and hydroelectric sources, and only 0.5 percent from renewable solar, wind, and geothermal technologies.¹¹⁵ Most air pollution and significant contributions to water and other forms of pollution result from the extraction, transportation, processing, and uses of the 78 percent of our energy that comes from fossil fuels. The energy obtained from oil also contributes to geopolitical tensions, and there's the small matter of its \$2 trillion per year price tag for all of this energy. Although the industrial-era energy sources that dominate energy production today will become more efficient with new nanotechnology-based methods of extraction, conversion, and transmission, it's the renewable category that will need to support the bulk of future energy growth.

By 2030 the price-performance of computation and communication will increase by a factor often to one hundred million compared to today. Other technologies will also undergo enormous increases in capacity and efficiency. Energy requirements will grow far more slowly than the capacity of technologies, however, because of greatly increased efficiencies in the use of energy, which I discuss below. A primary implication of the nanotechnology revolution is that physical technologies, such as manufacturing and energy, will become governed by the law of accelerating returns. All technologies will essentially become information technologies, including energy.

Worldwide energy requirements have been estimated to double by 2030, far less than anticipated economic growth, let alone the expected growth in the capability of technology.¹¹⁶

The bulk of the additional energy needed is likely to come from new nanoscale solar, wind, and geothermal technologies. It's important to recognize that most energy sources today represent solar power in one form or another.

Fossil fuels represent stored energy from the conversion of solar energy by animals and plants and related processes over millions of years (although the theory that fossil fuels originated from living organisms has recently been challenged). But the extraction of oil from high-grade oil wells is at a peak, and some experts believe we may have already passed that peak. It's clear, in any case, that we are rapidly depleting easily accessible fossil fuels. We do have far larger fossil-fuel resources that will require more sophisticated technologies to extract cleanly and efficiently (such as coal and shale oil), and they will be part of the future of energy. A billion-dollar demonstration plant called FutureGen, now being constructed, is expected to be the world's first zero-emissions energy plant based on fossil fuels.¹¹⁷ Rather than simply burn coal, as is done today, the 275-million-watt plant will convert the coal to a synthetic gas comprising hydrogen and carbon monoxide, which will then react with steam to produce discrete streams of hydrogen and carbon dioxide, which will be sequestered. The hydrogen can then be used in fuel cells or else converted into electricity and water. Key to the plant's design are new materials for membranes that separate hydrogen and carbon dioxide.

Our primary focus, however, will be in the development of clean, renewable, distributed, and safe energy technologies made possible by nanotechnology. For the past several decades energy technologies have been on the slow slope of the industrial era S-curve (the late stage of a specific technology paradigm, when the capability slowly approaches an asymptote or limit). Although the nanotechnology revolution will require new energy resources, it will also introduce major new S-curves in every aspect of energy: production, storage, transmission, and utilization by the 2020s.

Let's deal with these energy requirements in reverse, starting with utilization. Because of nanotechnology's ability to manipulate matter and energy at the extremely fine scale of atoms and molecular fragments, the efficiency of using energy will be far greater, which will translate into lower energy requirements. Over the next several decades computing will make the transition to reversible computing. (See "The Limits of Computation" in chapter 3.) As I discussed, the primary energy need for computing with reversible logic gates is to correct occasional errors from quantum and thermal effects. As a result reversible computing has the potential to cut energy needs by as much as a factor of a billion, compared to nonreversible computing. Moreover, the logic gates and memory bits will be smaller, by at least a factor often in each dimension, reducing energy requirements by another thousand. Fully developed nanotechnology, therefore, will enable the energy requirements for each bit switch to be reduced by about a trillion. Of course, we'll be increasing the amount of computation by even more than this, but this substantially augmented energy efficiency will largely offset those increases.

Manufacturing using molecular nanotechnology fabrication will also be far more energy efficient than contemporary manufacturing, which moves bulk materials from place to place in a relatively wasteful manner. Manufacturing today also devotes enormous energy resources into producing basic materials, such as steel. A typical nanofactory will be a tabletop device that can produce products ranging from computers to clothing. Larger products (such as vehicles, homes, and even additional nanofactories) will be produced as modular subsystems that larger robots can then assemble. Waste heat, which accounts for the primary energy requirement for nanomanufacturing, will be captured and recycled.

The energy requirements for nanofactories are negligible. Drexler estimates that molecular manufacturing will be an energy *generator* rather than an energy *consumer*. According to Drexler, "A molecular manufacturing process can be driven by the chemical energy

content of the feedstock materials, producing electrical energy as a by-product (if only to reduce the heat dissipation burden). . . . Using typical organic feedstock, and assuming oxidation of surplus hydrogen, reasonably efficient molecular manufacturing processes are net energy producers." ¹¹⁸

Products can be made from new nanotube-based and nanocomposite materials, avoiding the enormous energy used today to manufacture steel, titanium, and aluminum. Nanotechnology-based lighting will use small, cool, light-emitting diodes, quantum dots, or other innovative light sources to replace hot, inefficient incandescent and fluorescent bulbs.

Although the functionality and value of manufactured products will rise, product size will generally not increase (and in some cases, such as most electronics, products will get smaller). The higher value of manufactured goods will largely be the result of the expanding value of their information content. Although the roughly 50 percent deflation rate for information-based products and services will continue throughout this period, the amount of valuable information will increase at an even greater, more than offsetting pace.

I discussed the law of accelerating returns as applied to the communication of information in chapter 2. The amount of information being communicated will continue to grow exponentially, but the efficiency of communication will grow almost as fast, so the energy requirements for communication will expand slowly.

Transmission of energy will also be made far more efficient. A great deal of energy today is lost in transmission due to the heat created in power lines and inefficiencies in the transportation of fuel, which also represent a primary environmental assault. Smalley, despite his critique of molecular nanomanufacturing, has nevertheless been a strong advocate of new nanotechnology-based paradigms for creating and transmitting energy. He describes new power-transmission lines based on carbon nanotubes woven into long wires that will be far stronger, lighter, and, most important, much more energy efficient than conventional copper ones.¹¹⁹ He also envisions using superconducting wires to replace aluminum and copper wires in electric motors to provide greater efficiency. Smalley's vision of a nanoenabled energy future includes a panoply of new nanotechnology-enabled capabilities:¹²⁰

<bl>

- > Photovoltaics: dropping the cost of solar panels by a factor often to one hundred.
- Production of hydrogen: new technologies for efficiently producing hydrogen from water and sunlight.
- > Hydrogen storage: light, strong materials for storing hydrogen for fuel cells.
- > Fuel cells: dropping the cost of fuel cells by a factor often to one hundred.
- > Batteries and supercapacitors to store energy: improving energy storage densities by a factor often to one hundred.
- Improving the efficiency of vehicles such as cars and planes through strong and light nanomaterials.
- > Strong, light nanomaterials for creating large-scale energy harvesting systems in space, including on the moon.
- Robots using nanoscale electronics with artificial intelligence to automatically produce energy-generating structures in space and on the moon.
- New nanomaterial coatings to greatly reduce the cost of deep drilling.
- > Nanocatalysts to obtain greater energy yields from coal, at very high temperatures.
- > Nanofilters to capture the soot created from high-energy coal extraction. The soot is mostly carbon, which is a basic building block for most nanotechnology designs.

- > New materials to enable hot, dry rock geothermal-energy sources (converting the heat of the Earth's hot core into energy).

<tx>Another option for energy transmission is wireless transmission by microwaves. This method would be especially well suited to efficiently beam energy created in space by giant solar panels (see below).¹²¹ The Millennium Project of the American Council for the United Nations University envisions microwave energy transmission as a key aspect of "a clean, abundant energy nature."¹²²

Energy storage today is highly centralized, which represents a key vulnerability in that liquid-natural-gas tanks and other storage facilities are subject to terrorist attacks, with potentially catastrophic effects. Oil trucks and ships are equally exposed. The emerging paradigm for energy storage will be fuel cells, which will ultimately be widely distributed throughout our infrastructure, another example of the trend from inefficient and vulnerable centralized facilities to an efficient and stable distributed system.

Hydrogen-oxygen fuel cells, with hydrogen provided by methanol and other safe forms of hydrogen-rich fuel, have made substantial progress in recent years. A small company in Massachusetts, Integrated Fuel Cell Technologies, has demonstrated a MEMS (microelectronic mechanical system)-based fuel cell.¹²³ Each postage-stamp-size device contains thousands of microscopic fuel cells and includes the fuel lines and electronic controls. NEC plans to introduce fuel cells based on nanotubes in the near future for notebook computers and other portable electronics.¹²⁴ They claim their small power sources will run devices for up to forty hours at a time. Toshiba is also preparing fuel cells for portable electronic devices.¹²⁵

Larger fuel cells for powering appliances, vehicles, and even homes are also making impressive advances. A 2004 report by the U.S. Department of Energy concluded that nanobased technologies could facilitate every aspect of a hydrogen fuel cell-powered car.¹²⁶ For example, hydrogen must be stored in strong but light tanks that can withstand very high pressure. Nanomaterials such as nanotubes and nanocomposites could provide the requisite material for such containers. The report envisions fuel cells that produce power twice as efficiently as gasoline-based engines, producing only water as waste.

Many contemporary fuel-cell designs use methanol to provide hydrogen, which then combines with the oxygen in the air to produce water and energy. Methanol (wood alcohol), however, is difficult to handle, and introduces safety concerns because of its toxicity and flammability. Researchers from St. Louis University have demonstrated a stable fuel cell that uses ordinary ethanol (drinkable grain alcohol).¹²⁷ This device employs an enzyme called dehydrogenase that removes hydrogen ions from alcohol, which subsequently react with the oxygen in the air to produce power. The cell apparently works with almost any form of drinkable alcohol. "We have run it on various types," reported Nick Akers, a graduate student who has worked on the project. "It didn't like carbonated beer and doesn't seem fond of wine, but any other works fine."

Scientists at the University of Texas have developed a nanobot-size fuel cell that produces electricity directly from the glucose-oxygen reaction in human blood.¹²⁸ Called a "vampire bot" by commentators, the cell produces electricity sufficient to power conventional electronics and could be used for future blood-borne nanobots. Japanese scientists pursuing a similar project estimated that their system had the theoretical potential to produce a peak of one hundred watts from the blood of one person. (A newspaper in Sydney observed that the project provided a basis for the premise in the *Matrix* movies of using humans as batteries.)¹²⁹

Another approach to converting the abundant sugar found in the natural world into electricity has been demonstrated by Swades K. Chaudhuri and Derek R. Lovley at the University of Massachusetts. Their fuel cell, which incorporates actual microbes (the *Rhodospirillum rubrum*)

ferrireducens bacterium), boasts a remarkable 81 percent efficiency and uses almost no energy in its idling mode. The bacteria produce electricity directly from glucose with no unstable intermediary by-products. The bacteria also use the sugar fuel to reproduce, thereby replenishing themselves, resulting in stable and continuous production of electrical energy. Experiments with other types of sugars such as fructose, sucrose, and xylose were equally successful. Fuel cells based on this research could utilize the actual bacteria or, alternatively, directly apply the chemical reactions that the bacteria facilitate. In addition to powering nanobots in sugar-rich blood, these devices have the potential to produce energy from industrial and agricultural waste products.

Nanotubes have also demonstrated the promise of storing energy as nanoscale batteries, which may compete with nanoengineered fuel cells.¹³⁰ This extends further the remarkable versatility of nanotubes, which have already revealed their prowess in providing extremely efficient computation, communication of information, and transmission of electrical power, as well as in creating extremely strong structural materials.

The most promising approach to nanomaterials-enabled energy is from solar power, which has the potential to provide the bulk of our future energy needs in a completely renewable, emission-free, and distributed manner. The sunlight input to a solar panel is free. At about 10^{17} watts, or about ten thousand times more energy than the 10^{13} watts currently consumed by human civilization, the total energy from sunlight falling on the Earth is more than sufficient to provide for our needs.¹³¹ As mentioned above, despite the enormous increases in computation and communication over the next quarter century and the resulting economic growth, the far greater energy efficiencies of nanotechnology imply that energy requirements will increase only modestly to around thirty trillion watts (3×10^{13}) by 2030. We could meet this entire energy need with solar power alone if we captured only 0.0003 (three ten-thousandths) of the sun's energy as it hits the Earth.

It's interesting to compare these figures to the total metabolic energy output of all humans, estimated by Robert Freitas at 10^{12} watts, and that of all vegetation on Earth, at 10^{14} watts. Freitas also estimates that the amount of energy we could produce and use without disrupting the global energy balance required to maintain current biological ecology (referred to by climatologists as the "hypsithermal limit") is around 10^{15} watts. This would allow a very substantial number of nanobots per person for intelligence enhancement and medical purposes, as well as other applications, such as providing energy and cleaning up the environment. Estimating a global population of around ten billion (10^{10}), Freitas estimates around 10^{16} (ten thousand trillion) nanobots for each human would be acceptable within this limit.¹³² We would need only 10^1 nanobots (ten millionths of this limit) per person to place one in every neuron.

By the time we have technology of this scale, we will also be able to apply nanotechnology to recycle energy by capturing at least a significant portion of the heat generated by nanobots and other nanomachinery and converting that heat back into energy. The most effective way to do this would probably be to build the energy recycling into the nanobot itself.¹³³ This is similar to the idea of reversible logic gates in computation, in which each logic gate essentially immediately recycles the energy it used for its last computation.

We could also pull carbon dioxide out of the atmosphere to provide the carbon for nanomachinery, which would reverse the increase in carbon dioxide resulting from our current industrial-era technologies. We might, however, want to be particularly cautious about doing more than *reversing* the increase over the past several decades, lest we replace global warming with global cooling.

Solar panels have to date been relatively inefficient and expensive, but the technology is rapidly improving. The efficiency of converting solar energy to electricity has steadily advanced

for silicon photovoltaic cells from around 4 percent in 1952 to 24 percent in 1992.¹³⁴ Current multilayer cells now provide around 34 percent efficiency. A recent analysis of applying nanocrystals to solar-energy conversion indicates that efficiencies above 60 percent appear to be feasible.¹³⁵

Today solar power costs an estimated \$2.75 per watt.¹³⁶ Several companies are developing nanoscale solar cells and hope to bring the cost of solar power below that of other energy sources. Industry sources indicate that once solar power falls below one \$1.00 per watt, it will be competitive for directly supplying electricity to the nation's power grid. Nanosolar has a design based on titanium-oxide nanoparticles that can be mass-produced on very thin flexible films. CEO Martin Roscheisen estimates that his technology has the potential to bring down solar-power costs to around fifty cents per watt by 2006, lower than that of natural gas.¹³⁷ Competitors Nanosys and Konarka have similar projections. Whether or not these business plans pan out, once we have MNT (molecular nanotechnology)-based manufacturing, we will be able to produce solar panels (and almost everything else) extremely inexpensively, essentially at the cost of raw materials, of which inexpensive carbon is the primary one. At an estimated thickness of several microns, solar panels could ultimately be as inexpensive as a penny per square meter. We could place efficient solar panels on the majority of human-made surfaces, such as buildings and vehicles, and even incorporate them into clothing for powering mobile devices. A 0.0003 conversion rate for solar energy should be quite feasible, therefore, and relatively inexpensive.

Terrestrial surfaces could be augmented by huge solar panels in space. A Space Solar Power satellite already designed by NASA could convert sunlight in space to electricity and beam it to Earth by microwave. Each such satellite could provide billions of watts of electricity, enough for tens of thousands of homes.¹³⁸ With circa-2029 MNT manufacturing, we could produce solar panels of vast size directly in orbit around the Earth, requiring only the shipment of the raw materials to space stations, possibly via the planned Space Elevator, a thin ribbon, extending from a shipborne anchor to a counterweight well beyond geosynchronous orbit, made of out a material called carbon nanotube composite.¹³⁹

Desktop fusion also remains a possibility. Scientists at Oak Ridge National Laboratory used ultrasonic sound waves to shake a liquid solvent, causing gas bubbles to become so compressed they achieved temperatures of millions of degrees, resulting in the nuclear fusion of hydrogen atoms and the creation of energy.¹⁴⁰ Despite the broad skepticism over the original reports of cold fusion in 1989, this ultrasonic method has been warmly received by some peer reviewers.¹⁴¹ However, not enough is known about the practicality of the technique, so its future role in energy production remains a matter of speculation.

Applications of Nanotechnology to the Environment

<tx>Emerging nanotechnology capabilities promise a profound impact on the environment. This includes the creation of new manufacturing and processing technologies that will dramatically reduce undesirable emissions, as well as remediating the prior impact of industrial-age pollution. Of course, providing for our energy needs with nanotechnology-enabled renewable, clean resources such as nanosolar panels, as I discussed above, will clearly be a leading effort in this direction.

By building particles and devices at the molecular scale, not only is size greatly reduced and surface area increased, but new electrical, chemical, and biological properties are introduced. Nanotechnology will eventually provide us with a vastly expanded toolkit for improved catalysis, chemical and atomic bonding, sensing, and mechanical manipulation, not to mention intelligent control through enhanced microelectronics.

Ultimately we will redesign all of our industrial processes to achieve their intended results with minimal consequences, such as unwanted by-products and their introduction into the environment. We discussed in the previous section a comparable trend in biotechnology: intelligently designed pharmaceutical agents that perform highly targeted biochemical interventions with greatly curtailed side effects. Indeed, the creation of designed molecules through nanotechnology will itself greatly accelerate the biotechnology revolution.

Contemporary nanotechnology research and development involves relatively simple "devices" such as nanoparticles, molecules created through nanolayers, and nanotubes. Nanoparticles, which comprise between tens and thousands of atoms, are generally crystalline in nature and use crystal-growing techniques, since we do not yet have the means for precise nanomolecular manufacturing. Nanostructures consist of multiple layers that self-assemble. Such structures are typically held together with hydrogen or carbon bonding and other atomic forces. Biological structures such as cell membranes and DNA itself are natural examples of multilayer nanostructures.

As with all new technologies, there is a downside to nanoparticles: the introduction of new forms of toxins and other unanticipated interactions with the environment and life. Many toxic materials, such as gallium arsenide, are already entering the ecosystem through discarded electronic products. The same properties that enable nanoparticles and nanolayers to deliver highly targeted beneficial results can also lead to unforeseen reactions, particularly with biological systems such as our food supply and our own bodies. Although existing regulations may in many cases be effective in controlling them, the overriding concern is our lack of knowledge about a wide range of unexplored interactions.

Nonetheless, hundreds of projects have begun applying nanotechnology to enhancing industrial processes and that explicitly address existing forms of pollution. A few examples:

☛ There is extensive investigation of the use of nanoparticles for treating, deactivating, and removing a wide variety of environmental toxins. The nanoparticle forms of oxidants, reductants, and other active materials have shown the ability to transform a wide range of undesirable substances. Nanoparticles activated by light (for example, forms of titanium dioxide and zinc oxide) are able to bind and remove organic toxins and have low toxicity themselves.¹⁴² In particular, zinc-oxide nanoparticles provide a particularly powerful catalyst for detoxifying chlorinated phenols. These nanoparticles act as both sensors and catalysts and can be designed to transform only targeted contaminants.

☛ Nanofiltration membranes for water purification provide dramatically improved removal of fine-particle contaminants, compared to conventional methods of using sedimentation basins and wastewater clarifiers. Nanoparticles with designed catalysis are capable of absorbing and removing impurities. By using magnetic separation, these nanomaterials can be reused, which prevents them from becoming contaminants themselves. As one of many examples, consider nanoscale aluminosilicate molecular sieves called zeolites, which are being developed for controlled oxidation of hydrocarbons (for example, converting toluene to nontoxic benzaldehyde).¹⁴³ This method requires less energy and reduces the volume of inefficient photoreactions and waste products.

☛ Extensive research is under way to develop nanoproduced crystalline materials for catalysts and catalyst supports in the chemical industry. These catalysts have the potential to improve chemical yields, reduce toxic by-products, and remove contaminants.¹⁴⁴ For example, the material MCM-41 is now used by the oil industry to remove ultrafine contaminants that other pollution-reduction methods miss.

✦ It's estimated that the widespread use of nanocomposites for structural material in automobiles would reduce gasoline consumption by 1.5 billion liters per year, which in turn would reduce carbon-dioxide emissions by five billion kilograms per year, among other environmental benefits.

✦ Nanorobotics can be used to assist with nuclear-waste management. Nanofilters can separate isotopes when processing nuclear fuel. Nanofluids can improve the effectiveness of cooling nuclear reactors.

✦ Applying nanotechnology to home and industrial lighting could reduce both the need for electricity and an estimated two hundred million tons of carbon emissions per year.¹⁴⁵

✦ Self-assembling electronic devices (for example, self-organizing biopolymers), if perfected, will require less energy to manufacture and use and will produce fewer toxic by-products than conventional semiconductor-manufacturing methods.

✦ New computer displays using nanotube-based field-emission displays (FEDs) will provide superior display specifications while eliminating the heavy metals and other toxic materials used in conventional displays.

✦ Bimetallic nanoparticles (such as iron/palladium or iron/silver) can serve as effective reductants and catalysts for PCBs, pesticides, and halogenated organic solvents.¹⁴⁶

✦ Nanotubes appear to be effective absorbents for dioxins and have performed significantly better at this than traditional activated carbon.¹⁴⁷

This is a small sample of contemporary research on nanotechnology applications with potentially beneficial impact on the environment. Once we can go beyond simple nanoparticles and nanolayers and create more complex systems through precisely controlled molecular nanoassembly, we will be in a position to create massive numbers of tiny intelligent devices capable of carrying out relatively complex tasks. Cleaning up the environment will certainly be one of those missions.

Nanobots in the Bloodstream

<epi>Nanotechnology has given us the tools . . . to play with the ultimate toy box of nature—atoms and molecules. Everything is made from it. . . . The possibilities to create new things appear limitless.

<epis>—*Nobelist Horst Störmer*

<epi>The net effect of these nanomedical interventions will be the continuing arrest of all biological aging, along with the reduction of current biological age to whatever new biological age is deemed desirable by the patient, severing forever the link between calendar time and biological health. Such interventions may become commonplace several decades from today. Using annual checkups and cleanouts, and some occasional major repairs, your biological age could be restored once a year to the more or less constant physiological age that you select. You might still eventually die of accidental causes, but you'll live at least ten times longer than you do now.

<epis>—*Robert A. Freitas Jr^m* .

<tx>A prime example of the application of precise molecular control in manufacturing will be the deployment of billions or trillions of nanobots: small robots the size of human blood cells or smaller that can travel inside the bloodstream. This notion is not as futuristic as it may sound; successful animal experiments have been conducted using this concept, and many such

microscale devices are already working in animals. At least four major conferences on BioMEMS (Biological Microelectronic Mechanical Systems) deal with devices to be used in the human bloodstream.¹⁴⁹

Consider several examples of nanobot technology, which, based on miniaturization and cost-reduction trends, will be feasible within about twenty-five years. In addition to scanning the human brain to facilitate its reverse engineering, these nanobots will be able to perform a broad variety of diagnostic and therapeutic functions.

Robert A. Freitas Jr.—a pioneering nanotechnology theorist and leading proponent of nanomedicine (reconfiguring our biological systems through engineering on a molecular scale), and author of a book with that title¹⁵⁰—has designed robotic replacements for human blood cells that perform hundreds or thousands of times more effectively than their biological counterparts. With Freitas' respirocytes (robotic red blood cells) a runner could do an Olympic sprint for fifteen minutes without taking a breath.¹⁵¹ Freitas' robotic macrophages called "microbivores" will be far more effective than our white blood cells at combating pathogens.¹⁵² His DNA-repair robot would be able to mend DNA transcription errors and even implement needed DNA changes. Other medical robots he has designed can serve as cleaners, removing unwanted debris and chemicals (such as prions, malformed proteins, and protofibrils) from individual human cells.

Freitas provides detailed conceptual designs for a wide range of medical nanorobots (Freitas' preferred term) as well as a review of numerous solutions to the varied design challenges involved in creating them. For example, he provides about a dozen approaches to directed and guided motion,¹⁵³ some based on biological designs such as propulsive cilia. I discuss these applications in more detail in the next chapter.

George Whitesides complained in *Scientific American* that "for nanoscale objects, even if one could fabricate a propeller, a new and serious problem would emerge: random jarring by water molecules. These water molecules would be smaller than a nanosubmarine but not much smaller."¹⁵⁴ As with Smalley's critical article, which it accompanied, Whitesides' analysis is based on misconceptions. All medical nanobot designs, including those of Freitas, are at least ten thousand times larger than a water molecule. Analyses by Freitas and others show that the impact of the Brownian motion of adjacent molecules to be insignificant. Indeed, nanoscale medical robots will be thousands of times more stable and precise than blood cells or bacteria.¹⁵⁵

It should also be pointed out that medical nanobots will not require much of the extensive overhead biological cells need to maintain metabolic processes such as digestive and respiration. Nor do they need to support biological reproductive systems.

Although Freitas' conceptual designs are a couple of decades away, substantial progress has already been made on bloodstream-based devices. For example, a researcher at the University of Illinois at Chicago has cured type 1 diabetes in rats with a nanoengineered device that incorporates pancreatic Islet cells.¹⁵⁶ The device has seven-nanometer pores that let insulin out but won't let in the antibodies that destroy these cells. There are many other innovative projects of this type already under way.

<dia>Molly 2004: Okay, so I'll have all these nanobots in my bloodstream. Aside from being able to sit at the bottom of my pool for hours, what else is this going to do for me?

Ray: It will keep you healthy. They'll destroy pathogens such as bacteria, viruses, and cancer cells, and they won't be subject to the various pitfalls of the immune system, such as autoimmune reactions. Unlike your biological immune system, if you don't like what the nanobots are doing, you can tell them to do something different.

Molly 2004: You mean, send my nanobots an e-mail? Like, Hey, nanobots, stop destroying those bacteria in my intestines because they're actually good for my digestion?

Ray: Yes, good example. The nanobots will be under our control. They'll communicate with one another and with the Internet. Even today we have neural implants (for example, for Parkinson's disease) that allow the patient to download new software into them.

Molly 2004: That kind of makes the software-virus issue a lot more serious, doesn't it? Right now, if I get hit with a bad software virus, I may have to run a virus-cleansing program and load my backup files, but if nanobots in my bloodstream get a rogue message, they may start destroying my blood cells.

Ray: Well, that's another reason you'll probably want robotic blood cells, but your point is well taken. However, it's not a new issue. Even in 2004, we already have mission-critical software systems that run intensive-care units, manage 911 emergency systems, control nuclear-power plants, land airplanes, and guide cruise missiles. So software integrity is already of critical importance.

Molly 2004: True, but the idea of software running in my body and brain seems more daunting. On my personal computer, I get more than one hundred spam messages a day, at least several of which contain malicious software viruses. I'm not real comfortable with nanobots in my body getting software viruses.

Ray: You're thinking in terms of conventional Internet access. With VPNs (private networks), we already have the means today to create secure firewalls—otherwise, contemporary mission-critical systems would be impossible. They do work reasonably well, and Internet security technology will continue to evolve.

Molly 2004: I think some people would take issue with your confidence in firewalls.

Ray: They're not perfect, true, and they never will be, but we have another couple decades before we'll have extensive software running in our bodies and brains.

Molly 2004: Okay, but the virus writers will be improving their craft as well.

Ray: It's going to be a nervous standoff, no question about it. But the benefit today clearly outweighs the damage.

Molly 2004: How clear is that?

Ray: Well, no one is seriously arguing we should do away with the Internet because software viruses are such a big problem.

Molly 2004: I'll give you that.

Ray: When nanotechnology is mature, it's going to solve the problems of biology by overcoming biological pathogens, removing toxins, correcting DNA errors, and reversing other sources of aging. We will then have to contend with new dangers that it introduces, just as the Internet introduced the danger of software viruses. These new pitfalls will include the potential for self-replicating nanotechnology getting out of control, as well as the integrity of the software controlling these powerful, distributed nanobots.

Molly 2004: Did you say reverse aging?

Ray: I see you're already picking up on a key benefit.

Molly 2004: So how are the nanobots going to do that?

Ray: We'll actually accomplish most of that with biotechnology, methods such as RNA interference for turning off destructive genes, gene therapy for changing your genetic code, therapeutic cloning for regenerating your cells and tissues, smart drugs to reprogram your metabolic pathways, and many other emerging techniques. But whatever biotechnology doesn't get around to accomplishing, we'll have the means to do with nanotechnology.

Molly 2004: Such as?

Ray: Nanobots will be able to travel through the bloodstream, then go in and around our cells and perform various services, such as removing toxins, sweeping out debris, correcting DNA errors, repairing and restoring cell membranes, reversing atherosclerosis, modifying the levels of hormones, neurotransmitters, and other metabolic chemicals, and a myriad of other tasks. For each aging process, we can describe a means for nanobots to reverse the process, down to the level of individual cells, cell components, and molecules.

Molly 2004: So I'll stay young indefinitely?

Ray: That's the idea.

Molly 2004: When did you say I could get these?

Ray: I thought you were worried about nanobot firewalls.

Molly 2004: Yeah, well, I've got time to worry about that. So what was that timeframe again?

Ray: About twenty to twenty-five years.

Molly 2004: I'm twenty-five now, so I'll age to about forty-five and then stay there?

Ray: No, that's not exactly the idea. You can slow down aging to a crawl right now by adopting the knowledge we already have. Within ten to twenty years, the biotechnology revolution will provide far more powerful means to stop and in many cases reverse each disease and aging process. And it's not like nothing is going to happen in the meantime. Each year, we'll have more powerful techniques, and the process will accelerate. Then nanotechnology will finish the job.

Molly 2004: Yes, of course, it's hard for you to get out a sentence without using the word "accelerate." So what biological age am I going to get to?

Ray: I think you'll settle somewhere in your thirties and stay there for a while.

Molly 2004: Thirties sounds pretty good. I think a slightly more mature age than twenty-five is a good idea anyway. But what do you mean "for a while"?

Ray: Stopping and reversing aging is only the beginning. Using nanobots for health and longevity is just the early adoption phase of introducing nanotechnology and intelligent computation into our bodies and brains. The more profound implication is that we'll augment our thinking processes with nanobots that communicate with one another and with our biological neurons. Once nonbiological intelligence gets a foothold, so to speak, in our brains, it will be subject to the law of accelerating returns and expand exponentially. Our biological thinking, on the other hand, is basically stuck.

Molly 2004: There you go again with things accelerating, but when this really gets going, thinking with biological neurons will be pretty trivial in comparison.

Ray: That's a fair statement.

Molly 2004: So, Miss Molly of the future, when did I drop my biological body and brain?

Molly 2104: Well, you don't really want me to spell out your future, do you? And anyway it's actually not a straightforward question.

Molly 2004: How's that?

Molly 2104: In the 2040s we developed the means to instantly create new portions of ourselves, either biological or nonbiological. It became apparent that our true nature was a pattern of information, but we still needed to manifest ourselves in some physical form. However, we could quickly change that physical form.

Molly 2004: By?

Molly 2104: By applying new high-speed MNT manufacturing. So we could readily and rapidly redesign our physical instantiation. So I could have a biological body at one time and not at another, then have it again, then change it, and so on.

Molly 2004: I think I'm following this.

Molly 2104: The point is that I could have my biological brain and/or body or not have it. It's not a matter of dropping anything, because we can always get something we drop back.

Molly 2004: So you 're still doing this?

Molly 2104: Some people still do this, but now in 2104 it's a bit anachronistic. I mean, the simulations of biology are totally indistinguishable from actual biology, so why bother with physical instantiations?

Molly 2004: Yeah, it's messy isn't it?

Molly 2104: I'll say.

Molly 2004: I do have to say that it seems strange to be able to change your physical embodiment. I mean, where's your—my—continuity?

Molly 2104: It's the same as your continuity in 2004. You're changing your particles all the time also. It's just your pattern of information that has continuity.

Molly 2004: But in 2104 you're able to change your pattern of information quickly also. I can't do that yet.

Molly 2104: It's really not that different. You change your pattern—your memory, skills, experiences, even personality over time—but there is a continuity, a core that changes only gradually.

Molly 2004: But I thought you could change your appearance and personality dramatically in an instant?

Molly 2104: Yes, but that's just a surface manifestation. My true core changes only gradually, just like when I was you in 2004.

Molly 2004: Well, there are lots of times when I'd be delighted to instantly change my surface appearance.

<hl>Robotics: Strong AI

<epi>Consider another argument put forth by Turing. So far we have constructed only fairly simple and predictable artifacts. When we increase the complexity of our machines, there may, perhaps, be surprises in store for us. He draws a parallel with a fission pile. Below a certain "critical" size, nothing much happens: but above the critical size, the sparks begin to fly. So too, perhaps, with brains and machines. Most brains and all machines are, at present "sub-critical"—they react to incoming stimuli in a stodgy and uninteresting way, have no ideas of their own, can

produce only stock responses—but a few brains at present, and possibly some machines in the future, are super-critical, and scintillate on their own account. Turing is suggesting that it is only a matter of complexity, and that above a certain level of complexity a qualitative difference appears, so that "super-critical" machines will be quite unlike the simple ones hitherto envisaged.

<epis>—*J. R. Lucas, Oxfordphilosopher*, in his 1961 essay "Minds, Machines, and Gödel"¹⁵⁷

<epi>Given that superintelligence will one day be technologically feasible, will people choose to develop it? This question can pretty confidently be answered in the affirmative. Associated with every step along the road to superintelligence are enormous economic payoffs. The computer industry invests huge sums in the next generation of hardware and software, and it will continue doing so as long as there is a competitive pressure and profits to be made. People want better computers and smarter software, and they want the benefits these machines can help produce. Better medical drugs; relief for humans from the need to perform boring or dangerous jobs; entertainment—there is no end to the list of consumer-benefits. There is also a strong military motive to develop artificial intelligence. And nowhere on the path is there any natural stopping point where technophobics could plausibly argue "hither but not further."

<epis>—*Nick Bostrom, "How Long Before Superintelligence?" 1997*

<epi>It is hard to think of any problem that a superintelligence could not either solve or at least help us solve. Disease, poverty, environmental destruction, unnecessary suffering of all kinds: these are things that a superintelligence equipped with advanced nanotechnology would be capable of eliminating. Additionally, a superintelligence could give us indefinite lifespan, either by stopping and reversing the aging process through the use of nanomedicine, or by offering us the option to upload ourselves. A superintelligence could also create opportunities for us to vastly increase our own intellectual and emotional capabilities, and it could assist us in creating a highly appealing experiential world in which we could live lives devoted to joyful game-playing, relating to each other, experiencing, personal growth, and to living closer to our ideals.

<epis>—*Nick Bostrom, "Ethical Issues in Advanced Artificial Intelligence," 2003*

<epi>Will robots inherit the earth? Yes, but they will be our children.

<epis>—*Marvin Minsky, 1995.*

<tx>Of the three primary revolutions underlying the Singularity (G, N, and R), the most profound is R, which refers to the creation of nonbiological intelligence that exceeds that of unenhanced humans. A more intelligent process will inherently outcompete one that is less intelligent, making intelligence the most powerful force in the universe.

While the "R" in GNR stands for robotics, the real issue involved here is strong AI (artificial intelligence that exceeds human intelligence). The standard reason for emphasizing robotics in this formulation is that intelligence needs an embodiment, a physical presence, to affect the world. I disagree with the emphasis on physical presence, however, for I believe that the central concern is intelligence. Intelligence will inherently find a way to influence the world, including creating its own means for embodiment and physical manipulation. Furthermore, we can include physical skills as a fundamental part of intelligence; a large portion of the human brain (the cerebellum, comprising more than half our neurons), for example, is devoted to coordinating our skills and muscles.

Artificial intelligence at human levels will necessarily greatly exceed human intelligence for several reasons. As I pointed out earlier machines can readily share their knowledge. As

unenanced humans we do not have the means of sharing the vast patterns of interaeroneural connections and neurotransmitter-concentration levels that comprise our learning, knowledge, and skills, other than through slow, language-based communication. Of course, even this method of communication has been very beneficial, as it has distinguished us from other animals and has been an enabling factor in the creation of technology.

Human skills are able to develop only in ways that have been evolutionarily encouraged. Those skills, which are primarily based on massively parallel pattern recognition, provide proficiency for certain tasks, such as distinguishing faces, identifying objects, and recognizing language sounds. But they're not suited for many others, such as determining patterns in financial data. Once we fully master pattern-recognition paradigms, machine methods can apply these techniques to any type of pattern.¹⁵⁸

Machines can pool their resources in ways that humans cannot. Although teams of humans can accomplish both physical and mental feats that individual humans cannot achieve, machines can more easily and readily aggregate their computational, memory and communications resources. As discussed earlier, the Internet is evolving into a worldwide grid of computing resources that can be instantly brought together to form massive supercomputers.

Machines have exacting memories. Contemporary computers can master billions of facts accurately, a capability that is doubling every year.¹⁵⁹ The underlying speed and price-performance of computing itself is doubling every year, and the rate of doubling is itself accelerating.

As human knowledge migrates to the Web, machines will be able to read, understand, and synthesize all human-machine information. The last time a biological human was able to grasp all human scientific knowledge was hundreds of years ago.

Another advantage of machine intelligence is that it can consistently perform at peak levels and can combine peak skills. Among humans one person may have mastered music composition, while another may have mastered transistor design, but given the fixed architecture of our brains we do not have the capacity (or the time) to develop and utilize the highest level of skill in every increasingly specialized area. Humans also vary a great deal in a particular skill, so that when we speak, say, of human levels of composing music, do we mean Beethoven, or do we mean the average person? Nonbiological intelligence will be able to match and exceed peak human skills in each area.

For these reasons, once a computer is able to match the subtlety and range of human intelligence, it will necessarily soar past it, and then continue its double-exponential ascent.

A key question regarding the Singularity is whether the "chicken" (strong AI) or the "egg" (nanotechnology) will come first. In other words, will strong AI lead to full nanotechnology (molecular manufacturing assemblers that can turn information into physical products), or will full nanotechnology lead to strong AI? The logic of the first premise is that strong AI would imply superhuman AI for the reasons just cited; and superhuman AI would be in a position to solve any remaining design problems required to implement full nanotechnology.

The second premise is based on the realization that the hardware requirements for strong AI will be met by nanotechnology-based computation. Likewise the software requirements will be facilitated by nanobots that could create highly detailed scans of human brain functioning and thereby achieve the completion of reverse engineering the human brain.

Both premises are logical; it's clear that either technology can assist the other. The reality is that progress in both areas will necessarily use our most advanced tools, so advances in each field will simultaneously facilitate the other. However, I do expect that full MNT will emerge prior to strong AI, but only by a few years (around 2025 for nanotechnology, around 2029 for strong AI).

As revolutionary as nanotechnology will be, strong AI will have far more profound consequences. Nanotechnology is powerful but not necessarily intelligent. We can devise ways of at least trying to manage the enormous powers of nanotechnology, but superintelligence innately cannot be controlled.

Runaway AI. Once strong AI is achieved, it can readily be advanced and its powers multiplied, as that is the fundamental nature of machine abilities. As one strong AI immediately begets many strong AIs, the latter access their own design, understand and improve it, and thereby very rapidly evolve into a yet more capable, more intelligent AI, with the cycle repeating itself indefinitely. Each cycle not only creates a more intelligent AI but takes less time than the cycle before it, as is the nature of technological evolution (or any evolutionary process). The premise is that once strong AI is achieved, it will immediately become a runaway phenomenon of rapidly escalating superintelligence.¹⁶⁰

My own view is only slightly different. The logic of runaway AI is valid, but we still need to consider the timing. Achieving human levels in a machine will not *immediately* cause a runaway phenomenon. Consider that a human level of intelligence has limitations. We have examples of this today—about six billion of them. Consider a scenario in which you took one hundred humans from, say, a shopping mall. This group would constitute examples of reasonably well educated humans. Yet if this group was presented with the task of improving human intelligence, it wouldn't get very far, even if provided with the templates of human intelligence. It would probably have a hard time creating a simple computer. Speeding up the thinking and expanding the memory capacities of these one hundred humans would not immediately solve this problem.

I pointed out above that machines will match (and quickly exceed) peak human skills in each area of skill. So instead, let's take one hundred scientists and engineers. A group of technically trained people with the right backgrounds would be capable of improving accessible designs. If a machine attained equivalence to one hundred (and eventually one thousand, then one million) technically trained humans, each operating much faster than a biological human, a rapid acceleration of intelligence would ultimately follow.

However, this acceleration won't happen immediately when a computer passes the Turing test. The Turing test is comparable to matching the capabilities of an average, educated human and thus is closer to the example of humans from a shopping mall. It will take time for computers to master all of the requisite skills and to marry these skills with all the necessary knowledge bases.

Once we've succeeded in creating a machine that can pass the Turing test (around 2029), the succeeding period will be an era of consolidation in which nonbiological intelligence will make rapid gains. However, the extraordinary expansion contemplated for the Singularity, in which human intelligence is multiplied by billions, won't take place until the mid-2040s (as discussed in chapter 3).

The AI Winter

<epi>There's this stupid myth out there that A.I. has failed, but A.I. is everywhere around you every second of the day. People just don't notice it. You've got A.I. systems in cars, tuning the parameters of the fuel injection systems. When you land in an airplane, your gate gets chosen by an A.I. scheduling system. Every time you use a piece of Microsoft software, you've got an A.I. system trying to figure out what you're doing, like writing a letter, and it does a pretty damned good job. Every time you see a movie with computer-generated characters, they're all little A.I.

characters behaving as a group. Every time you play a video game, you're playing against an A.I. system.

<epis>—*Rodney Brooks, director of the MIT AI Lab*¹⁶¹

<tx>I still run into people who claim that artificial intelligence withered in the 1980s, an argument that is comparable to insisting that the Internet died in the dot-com bust of the early 2000s.¹⁶² The bandwidth and price-performance of Internet technologies, the number of nodes (servers), and the dollar volume of e-commerce all accelerated smoothly through the boom as well as the bust and the period since. The same has been true for AI.

The technology hype cycle for a paradigm shift—railroads, AI, Internet, telecommunications, possibly now nanotechnology—typically starts with a period of unrealistic expectations based on a lack of understanding of all the enabling factors required. Although utilization of the new paradigm does increase exponentially, early growth is slow until the knee of the exponential-growth curve is realized. While the widespread expectations for revolutionary change are accurate, they are incorrectly timed. When the prospects do not quickly pan out, a period of disillusionment sets in. Nevertheless exponential growth continues unabated, and years later a more mature and more realistic transformation does occur.

We saw this in the railroad frenzy of the nineteenth century, which was followed by widespread bankruptcies. (I have some of these early unpaid railroad bonds in my collection of historical documents.) And we are still feeling the effects of the e-commerce and telecommunications busts of several years ago, which helped fuel a recession from which we are now recovering.

AI experienced a similar premature optimism in the wake of programs such as the 1957 General Problem Solver created by Allen Newell, J. C. Shaw, and Herbert Simon, which was able to find proofs for theorems that had stumped mathematicians such as Bertrand Russell, and early programs from the MIT Artificial Intelligence Laboratory, which could answer SAT questions (such as analogies and story problems) at the level of college students.¹⁶³ A rash of AI companies occurred in the 1970s, but when profits did not materialize there was an AI "bust" in the 1980s, which has become known as the "AI winter." Many observers still think that the AI winter was the end of the story and that nothing has since come of the AI field.

Yet today many thousands of AI applications are deeply embedded in the infrastructure of every industry. Most of these applications were research projects ten to fifteen years ago. People who ask, "Whatever happened to AI?" remind me of travelers to the rain forest who wonder, "Where are all these many species that are supposed to live here?" when hundreds of species of flora and fauna are flourishing only a few dozen meters away, deeply integrated into the local ecology.

We are well into the era of "narrow AI," which refers to artificial intelligence that performs a useful and specific function that once required human intelligence to perform, and does so at human levels or better. Often narrow AI systems greatly exceed the speed of humans, as well as provide the ability to manage and consider thousands of variables simultaneously. I describe a broad variety of narrow AI examples below.

These time frames for AI's technology cycle (a couple of decades of growing enthusiasm, a decade of disillusionment, then a decade and a half of solid advance in adoption) may seem lengthy, compared to the relatively rapid phases of the Internet and telecommunications cycles (measured in years, not decades), but two factors must be considered. First, the Internet and telecommunications cycles were relatively recent, so they are more affected by the acceleration of paradigm shift (as discussed in chapter 1). So recent adoption cycles (boom, bust, and recovery) will be much faster than ones that started forty years ago. Second, the AI revolution is

the most profound transformation that human civilization will experience, so it will take longer to mature than less complex technologies. It is characterized by the mastery of the most important and most powerful attribute of human civilization, indeed of the entire sweep of evolution on our planet: intelligence.

It's the nature of technology to understand a phenomenon and then engineer systems that concentrate and focus that phenomenon to greatly amplify it. For example, scientists discovered a subtle property of curved surfaces known as Bernoulli's principle: a gas (such as air) travels more quickly over a curved surface than over a flat surface. Thus, air pressure over a curved surface is lower than over a flat surface. By understanding, focusing, and amplifying the implications of this subtle observation, our engineering created all of aviation. . Once we understand the principles of intelligence, we will have a similar opportunity to focus, concentrate, and amplify its powers.

As we reviewed in chapter 4, every aspect of understanding, modeling, and simulating the human brain is accelerating: the price-performance and temporal and spatial resolution of brain scanning, the amount of data and knowledge available about brain function, and the sophistication of the models and simulations of the brain's varied regions.

We already have a set of powerful tools that emerged from AI research and that have been refined and improved over several decades of development. The brain reverse-engineering project will greatly augment this toolkit by also providing a panoply of new, biologically inspired, self-organizing techniques. We will ultimately be able to apply engineering's ability to focus and amplify human intelligence vastly beyond the hundred trillion extremely slow interneuronal connections that each of us struggles with today. Intelligence will then be fully subject to the law of accelerating returns, which is currently doubling the power of information technologies every year.

An underlying problem with artificial intelligence that I have personally experienced in my forty years in this area is that as soon as an AI technique works, it's no longer considered AI and is spun off as its own field (for example, character recognition, speech recognition, machine vision, robotics, data mining, medical informatics, automated investing).

Computer scientist Elaine Rich defines AI as "the study of how to make computers do things at which, at the moment, people are better." Rodney Brooks, director of the MIT AI Lab, puts it a different way: "Every time we figure out a piece of it, it stops being magical; we say, *Oh, that's just a computation.*" I am also reminded of Watson's remark to Sherlock Holmes, "I thought at first that you had done something clever, but I see that there was nothing in it after all."¹⁶⁴ That has been our experience as AI scientists. The enchantment of intelligence seems to be reduced to "nothing" when we fully understand its methods. The mystery that is left is the intrigue inspired by the remaining, not as yet understood methods of intelligence.

AI's Toolkit

<epi>AI is the study of techniques for solving exponentially hard problems in polynomial time by exploiting knowledge about the problem domain.

<epis>—*Elaine Rich*

<tx>As I mentioned in chapter 4, it's only recently that we have been able to obtain sufficiently detailed models of how human brain regions function to influence AI design. Prior to that, in the absence of tools that could peer into the brain with sufficient resolution, AI scientists and engineers developed their own techniques. Just as aviation engineers did not model the ability to fly on the flight of birds, these early AI methods were not based on reverse engineering natural intelligence.

A small sample of these approaches is reviewed here. Since their adoption, they have grown in sophistication, which has enabled the creation of practical products that avoid the fragility and high error rates of earlier systems.

Expert systems. In the 1970s AI was often equated with one specific method: expert systems. This involves the development of specific logical rules to simulate the decision-making processes of human experts. A key part of the procedure entails knowledge engineers interviewing domain experts such as doctors and engineers to codify their decision-making rules.

There were early successes in this area, such as medical diagnostic systems that compared well to human physicians, at least in limited tests. For example, a system called MYCIN, which was designed to diagnose and recommend remedial treatment for infectious diseases, was developed through the 1970s. In 1979 a team of expert evaluators compared diagnosis and treatment recommendations by MYCIN to those of human doctors and found that MYCIN did as well as or better than any of the physicians.¹⁶⁵

It became apparent from this research that human decision making typically is based not on definitive logic rules but rather on "softer" types of evidence. A dark spot on a medical imaging test may suggest cancer, but other factors such as its exact shape, location, and contrast are likely to influence a diagnosis. The hunches of human decision making are usually influenced by combining many pieces of evidence from prior experience, none definitive by itself. Often we are not even consciously aware of many of the rules that we use.

By the late 1980s expert systems were incorporating the idea of uncertainty and could combine many sources of probabilistic evidence to make a decision. The MYCIN system pioneered this approach. A typical MYCIN "rule" reads:

If the infection which requires therapy is meningitis, and the type of the infection is fungal, and organisms were not seen on the stain of the culture, and the patient is not a compromised host, and the patient has been to an area that is endemic for coccidiomycoses, and the race of the patient is Black, Asian, or Indian, and the cryptococcal antigen in the csf test was not positive, THEN there is a 50 percent chance that cryptococcus is not one of the organisms which is causing the infection.

Although a single probabilistic rule such as this would not be sufficient by itself to make a useful statement, by combining thousands of such rules the evidence can be marshaled and combined to make reliable decisions.

Probably the longest-running expert system project is CYC (for enCYClopedic), created by Doug Lenat and his colleagues at Cycorp. Initiated in 1984, CYC has been coding commonsense knowledge to provide machines with an ability to understand the unspoken assumptions underlying human ideas and reasoning. The project has evolved from hard-coded logical rules to probabilistic ones and now includes means of extracting knowledge from written sources (with human supervision). The original goal was to generate one million rules, which reflects only a small portion of what the average human knows about the world. Lenat's latest goal is for CYC to master "100 million things, about the number a typical person knows about the world by 1997."¹⁶⁶

Another ambitious expert system is being pursued by Darryl Macer, associate professor of Biological Sciences at the University of Tsukuba in Japan. He plans to develop a system incorporating all human ideas.¹⁶⁷ One application would be to inform policy makers of which ideas are held by which community.

Bayesian nets. Over the last decade a technique called Bayesian logic has created a robust mathematical foundation for combining thousands or even millions of such

probabilistic rules in what are called "belief networks" or Bayesian nets. Originally devised by English mathematician Thomas Bayes, and published posthumously in 1763, the approach is intended to determine the likelihood of future events based similar occurrences in the past.¹⁶⁸ Many expert systems based on Bayesian techniques gather data from experience in an ongoing fashion, thereby continually learning and improving their decision making.

The most promising type of spam filters are based on this method. I personally use a spam filter called SpamBayes, which trains itself on e-mail that you have identified as either "spam" or "okay."¹⁶⁹ You start out by presenting a folder of each to the filter. It trains its Bayesian belief network on these two files and analyzes the patterns of each, thus enabling it to automatically move subsequent e-mail into the proper category. It continues to train itself on every subsequent e-mail, especially when it's corrected by the user. This filter has made the spam situation manageable for me, which is saying a lot, as it weeds out two hundred to three hundred spam messages each day, letting over one hundred "good" messages through. Only about 1 percent of the messages it identifies as "okay" are actually spam; it almost never marks a good message as spam. The system is almost as accurate as I would be and much faster.

Markov models. Another method that is good at applying probabilistic networks to complex sequences of information involves Markov models.¹⁷⁰ Andrei Andreyevich Markov (1856-1922), a renowned mathematician, established a theory of "Markov chains," which was refined by Norbert Wiener (1894-1964) in 1923. The theory provided a method to evaluate the likelihood that a certain sequence of events would occur. It has been popular, for example, in speech recognition, in which the sequential events are phonemes (parts of speech). The Markov models used in speech recognition code the likelihood that specific patterns of sound are found in each phoneme, how the phonemes influence each other, and likely orders of phonemes. The system can also include probability networks on higher levels of language, such as the order of words. The actual probabilities in the models are trained on actual speech and language data, so the method is self-organizing.

Markov modeling was one of the methods my colleagues-and I used in our own speech-recognition development.¹⁷¹ Unlike phonetic approaches, in which specific rules about phoneme sequences are explicitly coded by human linguists, we did not tell the system that there are approximately forty-four phonemes in English, nor did we tell it what sequences of phonemes were more likely than others. We let the system discover these "rules" for itself from thousands of hours of transcribed human speech data. The advantage of this approach over hand-coded rules is that the models develop subtle probabilistic rules of which human experts are not necessarily aware.

Neural nets. Another popular self-organizing method that has also been used in speech recognition and a wide variety of other pattern-recognition tasks is neural nets. This technique involves simulating a simplified model of neurons and interneuronal connections. One basic approach to neural nets can be described as follows. Each point of a given input (for speech, each point represents two dimensions, one being frequency and the other time; for images, each point would be a pixel in a two-dimensional image) is randomly connected to the inputs of the first layer of simulated neurons. Every connection has an associated synaptic strength, which represents its importance and which is set at a random value. Each neuron adds up the signals coming into it. If the combined signal exceeds a particular threshold, the neuron fires and sends a signal to its output connection; if the combined input signal does not exceed the threshold, the neuron does not fire, and its output is zero. The output of each neuron is randomly connected to the inputs of the neurons in the next layer. There are multiple layers (generally three or more), and the layers may be organized in a variety of configurations. For example, one layer may feed back to an earlier layer. At the top layer, the output of one or more neurons, also

randomly selected, provides the answer (For an algorithmic description of neural nets, see this endnote.¹⁷²).

Since the neural-net wiring and synaptic weights are initially set randomly, the answers of an untrained neural net will be random. The key to a neural net, therefore, is that it must learn its subject matter. Like the mammalian brains on which it's loosely modeled, a neural net starts out ignorant. The neural net's teacher—which may be a human, a computer program, or perhaps another, more mature neural net that has already learned its lessons—rewards the student neural net when it generates the right output and punishes it when it does not. This feedback is in turn used by the student neural net to adjust the strengths of each interneuronal connection.

Connections that were consistent with the right answer are made stronger. Those that advocated a wrong answer are weakened. Over time, the neural net organizes itself to provide the right answers without coaching. Experiments have shown that neural nets can learn their subject matter even with unreliable teachers. If the teacher is correct only 60 percent of the time, the student neural net will still learn its lessons.

A powerful, well-taught neural net can emulate a wide range of human pattern-recognition faculties. Systems using multilayer neural nets have shown impressive results in a wide variety of pattern-recognition tasks, including recognizing handwriting, human faces, fraud in commercial transactions such as credit-card charges, and many others. In my own experience in using neural nets in such contexts, the most challenging engineering task is not coding the nets but in providing automated lessons for them to learn their subject matter.

The current trend in neural nets is to take advantage of more realistic and more complex models of how actual biological neural nets work, now that we are developing detailed models of neural functioning from brain reverse engineering.¹⁷³ Since we do have several decades of experience in using self-organizing paradigms, new insights from brain studies can be quickly adapted to neural-net experiments.

Neural nets are also naturally amenable to parallel processing,^s since that is how the brain works. The human brain does not have a central processor that simulates each neuron. Rather, we can consider each neuron and each interneuronal connection to be an individual slow processor. Extensive work is under way to develop specialized chips that implement neural-net architectures in parallel to provide substantially greater throughput.¹⁷⁴

<h2>Genetic algorithms (GAs). **<tx>**Another self-organizing paradigm inspired by nature is genetic, or evolutionary, algorithms, which emulate evolution, including sexual reproduction and mutations. Here is a simplified description of how they work. First, determine a way to code possible solutions to a given problem. If the problem is optimizing the design parameters for a jet engine, define a list of the parameters (with a specific number of bits assigned to each parameter). This list is regarded as the genetic code in the genetic algorithm. Then randomly generate thousands or more genetic codes. Each such genetic code (which represents one set of design parameters) is considered a simulated "solution" organism.

Now evaluate each simulated organism in a simulated environment by using a defined method to evaluate each set of parameters. This evaluation is a key to the success of a genetic algorithm. In our example, we would apply each solution organism to a jet-engine simulation and determine how successful that set of parameters is, according to whatever criteria we are interested in (fuel consumption, speed, and so on). The best solution organisms (the best designs) are allowed to survive, and the rest are eliminated.

Now have each of the survivors multiply themselves until they reach the same number of solution creatures. This is done by simulating sexual reproduction. In other words, each new offspring solution draws part of its genetic code from one parent and another part from a second parent. Usually no distinction is made between male or female organisms; it's sufficient to

generate an offspring from two arbitrary parents. As they multiply, allow some mutation (random change) in the chromosomes to occur.

We've now defined one generation of simulated evolution; now repeat these steps for each subsequent generation. At the end of each generation determine how much the designs have improved. When the improvement in the evaluation of the design creatures from one generation to the next becomes very small, we stop this iterative cycle of improvement and use the best design(s) in the last generation. (For an algorithmic description of genetic algorithms, see this endnote¹⁷⁵).

The key to a GA is that the human designers don't directly program a solution; rather, they let one emerge through an iterative process of simulated competition and improvement. As we discussed, biological evolution is smart but slow, so to enhance its intelligence we retain its discernment while greatly speeding up its ponderous pace. The computer is fast enough to simulate many generations in a matter of hours or days or weeks. But we have to go through this iterative process only once; once we have let this simulated evolution run its course, we can apply the evolved and highly refined rules to real problems in a rapid fashion.

Like neural nets GAs are a way to harness the subtle but profound patterns that exist in chaotic data. A key requirement for their success is a valid way of evaluating each possible solution. This evaluation needs to be fast because it must take account of many thousands of possible solutions for each generation of simulated evolution.

GAs are adept at handling problems with too many variables to compute precise analytic solutions. The design of a jet engine, for example, involves more than one hundred variables and requires satisfying dozens of constraints. GAs used by researchers at General Electric were able to come up with engine designs that met the constraints more precisely than conventional methods.

When using GAs you must, however, be careful what you ask for. University of Sussex researcher Jon Bird used a GA to optimally design an oscillator circuit. Several attempts generated conventional designs using a small number of transistors, but the winning design was not an oscillator at all but a simple radio circuit. Apparently the GA discovered that the radio circuit picked up an oscillating hum from a nearby computer.¹⁷⁶ The GA's solution worked only in the exact location on the table where it was asked to solve the problem.

Genetic algorithms, part of the field of chaos or complexity theory, are being increasingly used to solve otherwise intractable business problems, such as optimizing complex supply chains. This approach is beginning to supplant more analytic methods throughout industry. (See examples below.) The paradigm is also adept at recognizing patterns, and is often combined with neural nets and other self-organizing methods. It's also a reasonable way to write computer software, particularly software that needs to find delicate balances for competing resources.

In the novel *usr/bin/god*, Cory Doctorow, a leading science-fiction writer, uses an intriguing variation of a GA to evolve an AI. The GA generates a large number of intelligent systems based on various intricate combinations of techniques, with each combination characterized by its genetic code. These systems then evolve using a GA.

The evaluation function works as follows: each system logs on to various human chat rooms and tries to pass for a human, basically a covert Turing test. If one of the humans in a chat room says something like "What are you, a chatterbot?" (chatterbot meaning an automatic program, that at today's level of development is expected to not understand language at a human level), the evaluation is over, that system ends its interactions, and reports its score to the GA. The score is determined by how long it was able to pass for human without being challenged in

this way. The GA evolves more and more intricate combinations of techniques that are increasingly capable of passing for human.

The main difficulty with this idea is that the evaluation function is fairly slow, although it will take an appreciable amount of time only once the systems are reasonably intelligent. Also, the evaluations can take place largely in parallel. It's an interesting idea and may actually be a useful method to finish the job of passing the Turing test, once we get to the point where we have sufficiently sophisticated algorithms to feed into such a GA, so that evolving a Turing-capable AI is feasible.

Recursive search. Often we need to search through vast number of combinations of possible solutions to solve a given problem. A classic example is in playing games such as chess. As a player considers her next move, she can list all of her possible moves, and then, for each such move, all possible countermoves by the opponent, and so on. It is difficult, however, for human players to keep a huge "tree" of move-countermove sequences in their heads, and so they rely on pattern recognition—recognizing situations based on prior experience—whereas machines use logical analysis of millions of moves and countermoves.

Such a logical tree is at the heart of most game-playing programs. Consider how this is done. We construct a program called Pick Best Next Step to select each move. Pick Best Next Step starts by listing all of the possible moves from the current state of the board. (If the problem was solving a mathematical theorem, rather than game moves, the program would list all of the possible next steps in a proof.) For each move the program constructs a hypothetical board that reflects what would happen if we made this move. For each such hypothetical board, we now need to consider what our opponent would do if we made this move. Now recursion comes in, because Pick Best Next Step simply calls Pick Best Next Step (in other words, itself) to pick the best move for our opponent. In calling itself, Pick Best Next Step then lists all of the legal moves for our opponent.


The program keeps calling itself, looking ahead as many moves as we have time to consider, which results in the generation of a huge move-countermove tree. This is another example of exponential growth, because to look ahead an additional move (or countermove) requires multiplying the amount of available computation by about five. Key to the success of the recursive formula is pruning this huge tree of possibilities and ultimately stopping its growth. In the game context, if a board looks hopeless for either side, the program can stop the expansion of the move-countermove tree from that point (called a "terminal leaf" of the tree) and consider the most recently considered move to be a likely win or loss. When all of these nested program calls are completed, the program will have determined the best possible move for the current actual board within the limits of the depth of recursive expansion that it had time to pursue, and the quality of its pruning algorithm (For an algorithmic description of recursive search, see this endnote.¹⁷⁷).

The recursive formula is often effective at mathematics. Rather than game moves, the "moves" are the axioms of the field of math being addressed, as well as previously proved theorems. The expansion at each point is the possible axioms (or previously proved theorems) that can be applied to a proof at each step. (This was the approach used by Newell, Shaw, and Simons's General Problem Solver.)

From these examples it may appear that recursion is well suited only for problems in which we have crisply defined rules and objectives. But it has also shown promise in computer generation of artistic creations. For example, a program I designed called Ray Kurzweil's Cybernetic Poet uses a recursive approach.¹⁷⁸ The program establishes a set of goals for each word—achieving a certain rhythmic pattern, poem structure, and word choice that is desirable at that point in the poem. If the program is unable to find a word that meets these criteria, it backs

up and erases the previous word it has written, reestablishes the criteria it had originally set for the word just erased, and goes from there. If that also leads to a dead end, it backs up again, thus moving backwards and forwards. Eventually, it forces itself to make up its mind by relaxing some of the constraints if all paths lead to dead ends.

Black (computer)...is deciding on a move



White (you)

<bh>Deep Fritz Draws: Are Humans Getting Smarter, or Are Computers Getting Stupider?

<bt>We find one example of qualitative improvements in software in the world of computer chess, which, according to common wisdom, is governed only by the brute-force expansion of computer hardware. In a chess tournament in October 2002 with top-ranking human player Vladimir Kramnik, the Deep Fritz software achieved a draw. I point out that Deep Fritz has available only about 1.3 percent of the brute-force computation as the previous computer champion, Deep Blue. Despite that, it plays chess at about the same level because of its superior pattern recognition—based pruning algorithm (see below). In six years a program like Deep Fritz will again achieve Deep Blue's ability to analyze two hundred million board positions per second. Deep Fritz-like chess programs running on ordinary personal computers will routinely defeat all humans later in this decade.

In *The Age of Intelligent Machines*, which I wrote between 1986 and 1989, I predicted that a computer would defeat the human world chess champion by the end of the 1990s. I also noted that

computers were gaining about forty-five points per year in their chess ratings, whereas the best human playing was essentially fixed, so this projected a crossover point in 1998. Indeed, Deep Blue did defeat Gary Kasparov in a highly publicized tournament in 1997.

Yet in the Deep Fritz-Kramnik match, the current reigning computer program was able to achieve only a tie. Five years had passed since Deep Blue's victory, so what are we to make of this situation? Should we conclude that:

<bnl>1. Humans are getting smarter, or at least better at chess?

2. Computers are getting worse at chess? If so, should we conclude that the much-publicized improvement in computer speed over the past five years was not all it was cracked up to be? Or, that computer software is getting worse, at least in chess?

<bsh>The specialized-hardware advantage. <btx>Neither of the above conclusions is warranted. The correct conclusion is that software is getting better because Deep Fritz essentially matched the performance of Deep Blue, yet with far smaller computational resources. To gain some insight into these questions, we need to examine a few essentials. When I wrote my predictions of computer chess in the late 1980s, Carnegie Mellon University was embarked on a program to develop specialized chips for conducting the "minimax" algorithm (the standard game-playing method that relies on building trees of move-countermove sequences, and then evaluating the terminal-leaf positions of each branch of the tree) specifically for chess moves.

Based on this specialized hardware CMU's 1988 chess machine, HiTech, was able to analyze 175,000 board positions per second. It achieved a chess rating of 2,359, only about 440 points below the human world champion.

A year later, in 1989, CMU's Deep Thought machine increased this capacity to one million board positions per second and achieved a rating of 2,400. IBM eventually took over the project and renamed it Deep Blue but kept the basic CMU architecture. The version of Deep Blue that defeated Kasparov in 1997 had 256 special-purpose chess processors working in parallel, which analyzed two hundred million board positions per second.

It is important to note the use of specialized hardware to accelerate the specific calculations needed to generate the minimax algorithm for chess moves. It's well-known to computer-systems designers that specialized hardware generally can implement a specific algorithm at least one hundred times faster than a general-purpose computer. Specialized ASICs (Application-specific Integrated Circuits) require significant development efforts and costs, but for critical calculations that are needed on a repetitive basis (for example, decoding MP3 files or rendering graphics primitives for video games), this expenditure can be well worth the investment.

<bsh>Deep Blue versus Deep Fritz. <btx>Because there had always been a great deal of focus on the milestone of a computer's being able to defeat a human opponent, support was available for investing in special-purpose chess circuits. Although there was some lingering controversy regarding the parameters of the Deep Blue-Kasparov match, the level of interest in computer chess waned considerably after 1997. After all, the goal had been achieved, and there was little point in beating a dead horse. IBM canceled work on the project, and there has been no work on specialized chess chips since that time. The focus of research in the various domains spun out of AI has been placed instead on problems of greater consequence, such as guiding airplanes, missiles, and factory robots, understanding natural language, diagnosing electrocardiograms and blood-cell images, detecting credit-card fraud, and a myriad of other successful narrow AI applications.

Computer hardware has nonetheless continued its exponential increase, with personal-computer speeds doubling every year since 1997. Thus the general-purpose Pentium processors used by Deep Fritz are about thirty-two times faster than processors in 1997. Deep Fritz uses a network of only eight personal computers, so the hardware is equivalent to 256 1997-class personal computers. Compare that to Deep Blue, which used 256 specialized chess processors, each of which was about one hundred times faster than 1997 personal computers (of course, only for computing chess minimax). So Deep Blue was 25,600 times faster than a 1997 PC and one hundred times faster than Deep Fritz. This analysis is confirmed by the reported speeds of the two systems: Deep Blue can analyze 200 million board positions per second compared to only about 2.5 million for Deep Fritz.

<bsh>Significant software gains. <btx>So what can we say about the software in Deep Fritz? Although chess machines are usually referred to as examples of brute-force calculation, there is one important aspect of these systems that does require qualitative judgment. The combinatorial explosion of possible move-countermove sequences is rather formidable.

In *The Age of Intelligent Machines* I estimated that it would take about forty billion years to make one move if we failed to prune the move-countermove tree and attempted to make a "perfect" move in a typical game. (Assuming about thirty moves each in a typical game and about eight possible moves per play, we have 8^{30} possible move sequences; analyzing one billion move sequences per second would take 10^{18} seconds or forty billion years.) Thus a practical system needs to continually prune away unpromising lines of action. This requires insight and is essentially a pattern-recognition judgment.

Humans, even world-class chess masters, perform the minimax algorithm extremely slowly, generally performing less than one move-countermove analysis per second. So how is it that a chess master can compete at all with computer systems? The answer is that we possess formidable powers of pattern recognition, which enable us to prune the tree with great insight.

It's precisely in this area that Deep Fritz has improved considerably over Deep Blue. Deep Fritz has only slightly more computation available than CMU's Deep Thought yet is rated almost 400 points higher.

Are human chess players doomed? Another prediction I made in *The Age of Intelligent Machines* was that once computers did perform as well or better as humans in chess, we would either think more of computer intelligence, less of human intelligence, or less of chess, and that if history is a guide, the last of these would be the likely outcome. Indeed, that is precisely what happened. Soon after Deep Blue's victory we began to hear a lot about how chess is really just a simple game of calculating combinations and that the computer victory just demonstrated that it was a better calculator.

The reality is slightly more complex. The ability of humans to perform well in chess is clearly not due to our calculating prowess, at which we are in fact rather poor. We use instead a quintessentially human form of judgment. For this type of qualitative judgment, Deep Fritz represents genuine progress over earlier systems. (Incidentally, humans have made no progress in the last five years, with the top human scores remaining just below 2,800. Kasparov is rated at 2,795 and Kramnik at 2,794.)

Where we go from here? Now that computer chess is relying on software running on ordinary personal computers, chess programs will continue to benefit from the ongoing acceleration of computer power. By 2009 a program like Deep Fritz will again achieve Deep Blue's ability to analyze two hundred million board positions per second. With the opportunity to harvest computation on the Internet, we will be able to achieve this potential several years sooner than 2009. (Internet harvesting of computers will require more ubiquitous broadband communication, but that's coming, too).

With these inevitable speed increases, as well as ongoing improvements in pattern recognition, computer chess ratings will continue to edge higher. Deep Fritz-like programs running on ordinary personal computers will routinely defeat all humans later in this decade. Then we'll really lose interest in chess.

Combining methods. The most powerful approach to building robust AI systems is to combine approaches, which is how the human brain works. As we discussed, the brain is not one big neural net but instead consists of hundreds of regions, each of which is optimized for processing information in a different way. None of these regions by itself operates at what we would consider human levels of performance, but clearly by definition the overall system does exactly that.

I've used this approach in my own AI work, especially in pattern recognition. In speech recognition, for example, we implemented a number of different pattern-recognition systems based on different paradigms. Some were specifically programmed with knowledge of phonetic and linguistic constraints from experts. Some were based on rules to parse sentences (which involves creating sentence diagrams showing word usage, similar to the diagrams taught in grade school). Some were based on self-organizing techniques, such as Markov models, trained on extensive libraries of recorded and annotated human speech. We then programmed a software "expert manager" to learn the strengths and weaknesses of the different "experts" (recognizers) and combine their results in optimal ways. In this fashion, a particular technique that by itself might produce unreliable results can nonetheless contribute to increasing the overall accuracy of the system.

There are many intricate ways to combine the varied methods in AFs toolbox. For example, one can use a genetic algorithm to evolve the optimal topology (organization of nodes

and connections) for a neural net or a Markov model. The final output of the GA-evolved neural net can then be used to control the parameters of a recursive search algorithm. We can add in powerful signal- and image-processing techniques that have been developed for pattern-processing systems. Each specific application calls for a different architecture. Computer science professor and AI entrepreneur Ben Goertzel has written a series of books and articles that describe strategies and architectures for combining the diverse methods underlying intelligence. His "Novamente" architecture is intended to provide a framework for general purpose AI.¹⁷⁹

The above basic descriptions provide only a glimpse into how increasingly sophisticated current AI systems are designed. It's beyond the scope of this book to provide a comprehensive description of the techniques of AI, and even a doctoral program in computer science is unable to cover all of the varied approaches in use today.

Many of the examples of real-world narrow AI systems described in the next section use a variety of methods integrated and optimized for each particular task. Narrow AI is strengthening as a result of several concurrent trends: continued exponential gains in computational resources, extensive real-world experience with thousands of applications, and fresh insights into how the human brain makes intelligent decisions.

A Narrow AI Sampler

<tx>When I wrote my first AI book, *The Age of Intelligent Machines* in the late 1980s., I had to conduct extensive investigations to find a few successful examples of AI in practice. The Internet was not yet prevalent, so I had to go to real libraries and visit the AI research centers in the United States, Europe, and Asia. I included in my book pretty much all of the reasonable examples I could identify. In my research for this book my experience has been altogether different. I have been inundated with thousands of corripelling examples. In our reporting on the KurzweilAI.net Web site, we feature several dramatic systems every day.¹⁸⁰

A 2003 study by Business Communication Company projected a \$21 billion market by 2007 for AI applications, with average annual growth of 12.2 percent from 2002 to 2007.¹⁸¹ Leading industries for AI applications include business intelligence, customer relations, finance, defense and domestic security, and education. Here is a small sample of narrow AI in action.

<h2>Military **and** intelligence. <tx>The U.S. military has been an avid user of AI systems. Pattern-recognition software systems guide autonomous weapons such as cruise missiles, which can fly thousands of miles to find a specific building or even a specific window.¹⁸² Although the relevant details of the terrain that the missile flies over are programmed ahead of time, variations in weather, ground cover, and other factors require a flexible level of real-time image recognition.

The army has developed prototypes of self-organizing communication networks (called "mesh networks") to automatically configure many thousands of communication nodes when a platoon is dropped into a new location.¹⁸³

Expert systems incorporating Bayesian networks and GAs are used to optimize complex supply chains that coordinate millions of provisions, supplies, and weapons based on rapidly changing battlefield requirements.

AI systems are routinely employed to simulate the performance of weapons, including nuclear bombs and missiles.

Advance warning of the September 11, 2001, terrorist attacks was apparently detected by the National Security Agency's AI-based Echelon system, which analyzes the agency's extensive monitoring of communications traffic.¹⁸⁴ Unfortunately, Echelon's warnings were not reviewed by human agents until it was too late.

The 2002 military campaign in Afghanistan saw the debut of the armed Predator, an unmanned robotic flying fighter. Although the Air Force's Predator had been under development for many years, arming it with Army-supplied missiles was a last-minute improvisation that proved remarkably successful. In the Iraq war that began in 2003 the armed Predator (operated by the CIA) and other flying unmanned aerial vehicles (UAVs) destroyed thousands of enemy tanks and missile sites.

All of the military services are using robots. The army utilizes them to search caves (in Afghanistan) and buildings. The navy uses small robotic ships to protect its aircraft carriers. As I discuss in the next chapter, moving soldiers away from battle is a rapidly growing trend.

<h2>Space exploration. <tx>NASA is building self-understanding into the software controlling its unmanned spacecraft. Because Mars is about three light-minutes from Earth, and Jupiter around forty light-minutes (depending on the exact position of the planets), communication between spacecraft headed there and earthbound controllers is significantly delayed. For this reason it's important that the software controlling these missions have the capability of performing its own tactical decision making. To accomplish this NASA software is being designed to include a model of the software's own capabilities and those of the spacecraft, as well as the challenges each mission is likely to encounter. Such AI-based systems are capable of reasoning through new situations rather than just following preprogrammed rules. This approach enabled the craft Deep Space One in 1999 to use its own technical knowledge to devise a series of original plans to overcome a stuck switch that threatened to destroy its mission of exploring an asteroid.¹⁸⁵ The AI system's first plan failed to work, but its second plan saved the mission. "These systems have a commonsense model of the physics of their internal components," explains Brian Williams, coinventor of Deep Space One's autonomous software and now a scientist at MIT's Space Systems and AI laboratories. "[The spacecraft] can reason from that model to determine what is wrong and to know how to act."

Using a network of computers NASA used GAs to evolve an antenna design for three Space Technology 5 satellites that will study the Earth's magnetic field. Millions of possible designs competed in the simulated evolution. According to NASA scientist and project leader Jason Lohn, "We are now using the [GA] software to design tiny microscopic machines, including gyroscopes, for spaceflight navigation. The software also may invent designs that no human designer would ever think of."¹⁸⁶

Another NASA AI system learned on its own to distinguish stars from galaxies in very faint images with an accuracy surpassing that of human astronomers.

New land-based robotic telescopes are able to make their own decisions on where to look and how to optimize the likelihood of finding desired phenomena. Called "autonomous, semi-intelligent observatories," the systems can adjust to the weather, notice items of interest, and decide on their own to track them. They are able to detect very subtle phenomena, such as a star blinking for a nanosecond, which may indicate a small asteroid in the outer regions of our solar system passing in front of the light from that star.¹⁸⁷ One such system, called Moving Object and Transient Event Search System (MOTESS) has identified on its own 180 new asteroids and several comets during its first two years of operation. "We have an intelligent observing system," explained University of Exeter astronomer Alasdair Allan. "It thinks and reacts for itself, deciding whether something it has discovered is interesting enough to need more observations. If more observations are needed, it just goes ahead and gets them."

Similar systems are used by the military to automatically analyze data from spy satellites. Current satellite technology is capable of observing ground-level features about an inch in size and is not affected by bad weather, clouds, or darkness.¹⁸⁸ The massive amount of data

continually generated would not be manageable without automated image recognition programmed to look for relevant developments.

<h2>Medicine. <tx>If you obtain an electrocardiogram (ECG) your doctor is likely to receive an automated diagnosis using pattern recognition applied to ECG recordings. My own company (Kurzweil Technologies) is working with United Therapeutics to develop a new generation of automated ECG analysis for long-term unobtrusive monitoring (via sensors embedded in clothing and wireless communication using a cell phone) of the early warning signs of heart disease.¹⁸⁹ Other pattern-recognition systems are used to diagnose a variety of imaging data.

Every major drug developer is using AI programs to do pattern recognition and intelligent data mining in the development of new drug therapies. For example SRI International is building flexible knowledge bases that encode everything we know about a dozen disease agents, including tuberculosis and *H. pylori* (the bacteria that cause ulcers).¹⁹⁰ The goal is to apply intelligent data-mining tools (software that can search for new relationships in data) to find new ways to kill or disrupt the metabolisms of these pathogens.

Similar systems are being applied to performing the automatic discovery of new therapies for other diseases, as well as understanding the function of genes and their roles in disease.¹⁹¹ For example Abbott Laboratories claims that six human researchers in one of its new labs equipped with AI-based robotic and data-analysis systems are able to match the results of two hundred scientists in its older drug-development labs.¹⁹²

Men with elevated prostate-specific antigen (PSA) levels typically undergo surgical biopsy, but about 75 percent of these men do not have prostate cancer. A new test, based on pattern recognition of proteins in the blood, would reduce this false positive rate to about 29 percent.¹⁹³ The test is based on an AI program designed by Correlogic Systems in Bethesda, Maryland, and the accuracy is expected to improve further with continued development.

Pattern recognition applied to protein patterns has also been used in the detection of ovarian cancer. The best contemporary test for ovarian cancer, called CA-125, employed in combination with ultrasound, misses almost all early-stage tumors. "By the time it is now diagnosed, ovarian cancer is too often deadly," says Emanuel Petricoin III, codirector of the Clinical Proteomics Program run by the FDA and the National Cancer Institute. Petricoin is the lead developer of a new AI-based test looking for unique patterns of proteins found only in the presence of cancer. In an evaluation involving hundreds of blood samples, the test was, according to Petricoin, "an astonishing 100% accurate in detecting cancer, even at the earliest stages."¹⁹⁴

About 10 percent of all Pap-smear slides in the United States are analyzed by a self-learning AI program called FocalPoint, developed by TriPath Imaging. The developers started out by interviewing pathologists on the criteria they use. The AI system then continued to learn by watching expert pathologists. Only the best human diagnosticians were allowed to be observed by the program. "That's the advantage of an expert system," explains Bob Schmidt, TriPath's technical product manager. "It allows you to replicate your very best people."

Ohio State University Health System has developed a computerized physician order-entry (CPOE) system based on an expert system with extensive knowledge across multiple specialties.¹⁹⁵ The system automatically checks every order for possible allergies in the patient, drug interactions, duplications, drug restrictions, dosing guidelines, and appropriateness given information about the patient from the hospital's laboratory and radiology departments.

<h2>Science and math. <tx>A "robot scientist" has been developed at the University of Wales that combines an AI-based system capable of formulating original theories, a robotic system that can automatically carry out experiments, and a reasoning engine to evaluate results.

The researchers provided their creation with a model of gene expression in yeast. The system "automatically originates hypotheses to explain observations, devises experiments to test these hypotheses, physically runs the experiments using a laboratory robot, interprets the results to falsify hypotheses inconsistent with the data, and then repeats the cycle."¹⁹⁶ The system is capable of improving its performance by learning from its own experience. The experiments designed by the robot scientist were three times less expensive than those designed by human scientists. A test of the machine against a group of human scientists showed that the discoveries made by the machine were comparable to those made by the humans.

Mike Young, director of biology at the University of Wales, was one of the human scientists who lost to the machine. He explains that "the robot did beat me, but only because I hit the wrong key at one point."

A long-standing conjecture in algebra was finally proved by an AI system at Argonne National Laboratory. Human mathematicians called the proof "creative."

<h2>Business, finance, and manufacturing. <tx>Companies in every industry are using AI systems to control and optimize logistics, detect fraud and money laundering, and perform intelligent data mining on the horde of information they gather each day. Wal-Mart, for example, gathers vast amounts of information from its transactions with shoppers. AI-based tools using neural nets and expert systems review this data to provide market-research reports for managers. This intelligent data mining allows them to make remarkably accurate predictions of the inventory required for each product in each store for each day.¹⁹⁷

AI-based programs are routinely used to detect fraud in financial transactions. Future Route, an English company, for example, offers iHex, based on AI routines developed at Oxford University, to detect fraud in credit-card transactions and loan applications.¹⁹⁸ The system continuously generates and updates its own rules based on its experience. First Union Home Equity Bank in Charlotte, North Carolina, uses Loan Arranger, a similar AI-based system to decide whether to approve mortgage applications.¹⁹⁹

The NASDAQ similarly uses a learning program called the Securities Observation, News Analysis, and Regulation (SONAR) system to monitor all trades for fraud as well as the possibility of insider trading.²⁰⁰ As of the end of 2003 more than 180 incidents had been detected by SONAR and referred to the U.S. Securities and Exchange Commission and Department of Justice. These included several cases that later received significant news coverage.

Ascent Technology, founded by Patrick Winston, who directed MIT's AI Lab from 1972 through 1997, has designed an GA-based system called SmartAirport Operations Center (SAOC) that can optimize the complex logistics of an airport, such as balancing work assignments of hundreds of employees, making gate and equipment assignments, and managing a myriad of other details.²⁰¹ Winston points out that "figuring out ways to optimize a complicated situation is what genetic algorithms do." SAOC has raised productivity by approximately 30 percent in the airports where it has been implemented.

Ascent's first contract was to apply its AI techniques to managing the logistics for the 1991 Desert Storm campaign in Iraq. DARPA claimed that AI-based logistic-planning systems, including the Ascent system, resulted in more savings than the entire government research investment in AI over several decades.²⁰²

A recent trend in software is for AI systems to monitor a complex software system's performance, recognize malfunctions, and determine the best way to recover automatically without necessarily informing the human user. The idea stems from the realization that as software systems become more complex, like humans, they will never be perfect, and that eliminating all bugs is impossible. As humans, we use the same strategy: we don't expect to be perfect, but we usually try to recover from inevitable mistakes. "We want to stand this notion of

systems management on its head," says Armando Fox, the head of Stanford University's Software Infrastructures Group, who is working on what is now called "autonomic computing." Fox adds, "The system has to be able to set itself up, it has to optimize itself. It has to repair itself, and if something goes wrong, it has to know how to respond to external threats." IBM, Microsoft, and other software vendors are all developing systems that incorporate autonomic capabilities.

<h2>Manufacturing and robotics. <tx>Computer-integrated manufacturing (CIM) increasingly employs AI techniques to optimize the use of resources, streamline logistics, and reduce inventories through just-in-time purchasing of parts and supplies. A new trend in CIM systems is to use "case-based reasoning" rather than hard-coded, rule-based expert systems. Such reasoning codes knowledge as "cases," which are examples of problems with solutions. Initial cases are usually designed by the engineers, but the key to a successful case-based reasoning system is its ability to gather new cases from actual experience. The system is then able to apply the reasoning from its stored cases to new situations.

Robots are extensively used in manufacturing. The latest generation of robots uses flexible AI-based machine-vision systems—from companies such as Cognex Corporation in Natick, Massachusetts—that can respond flexibly to varying conditions. This reduces the need for precise setup for the robot to operate correctly. Brian Carlisle, CEO of Adept Technologies, a Livermore, California, factory-automation company, points out that "even if labor costs were eliminated [as a consideration], a strong case can still be made for automating with robots and other flexible automation. In addition to quality and throughput, users gain by enabling rapid product changeover and evolution that can't be matched with hard tooling."

One of AI's leading roboticists, Hans Moravec, has founded a company called Seegrid to apply his machine-vision technology to applications in manufacturing, materials handling, and military missions.²⁰³ Moravec's software enables a device (a robot or just a material-handling cart) to walk or roll through an unstructured environment and in a single pass build a reliable "voxel" (three-dimensional pixel) map of the environment. The robot can then use the map and its own reasoning ability to determine an optimal and obstacle-free path to carry out its assigned mission.

This technology enables autonomous carts to transfer materials throughout a manufacturing process without the high degree of preparation required with conventional preprogrammed robotic systems. In military situations autonomous vehicles could carry out precise missions while adjusting to rapidly changing environments and battlefield conditions.

Machine vision is also improving the ability of robots to interact with humans. Using small, inexpensive cameras, head- and eye-tracking software can sense where a human user is, allowing robots, as well as virtual personalities on a screen, to maintain eye contact, a key element for natural interactions. Head- and eye-tracking systems have been developed at Carnegie Mellon University and MIT and are offered by small companies such as Seeing Machines of Australia.

An impressive demonstration of machine vision was a vehicle that was driven by an AI system with no human intervention for almost the entire distance from Washington, D.C., to San Diego.²⁰⁴ Bruce Buchanan, computer-science professor at the University of Pittsburgh and president of the American Association of Artificial Intelligence, pointed out that this feat would have been "unheard of 10 years ago."

Palo Alto Research Center (PARC) is developing a swarm of robots that can navigate in complex environments, such as a disaster zone, and find items of interest, such as humans who may be injured. In a September 2004 demonstration at an AI conference in San Jose, they demonstrated a group of self-organizing robots on a mock but realistic disaster area.²⁰⁵ The

robots moved over the rough terrain, communicated with one another, used pattern recognition on images, and detected body heat to locate humans.

<h2>Speech and language. <tx>Dealing naturally with language is the most challenging task of all for artificial intelligence. No simple tricks, short of fully mastering the principles of human intelligence, will allow a computerized system to convincingly emulate human conversation, even if restricted to just text messages. This was Turing's enduring insight in designing his eponymous test based entirely on written language..

Although not yet at human levels, natural language-processing systems are making solid progress. Search engines have become so popular that "Google" has gone from a proper noun to a common verb, and its technology has revolutionized research and access to knowledge. Google and other search engines use AI-based statistical-learning methods and logical inference to determine the ranking of links. The most obvious failing of these search engines is their inability to understand the context of words. Although an experienced user learns how to design a string of keywords to find the most relevant sites (for example, a search for "computer chip" is likely to avoid references to potato chips that a search for "chip" alone might turn up) what we would really like to be able to do is converse with our search engines in natural language. Microsoft has developed a natural-language search engine called Ask MSR (Ask MicroSoft Research), which actually answers natural-language questions such as "When was Mickey Mantle born?"²⁰⁶ After the system parses the sentence to determine the parts of speech (subject, verb, object, adjective and adverb modifiers, and so on), a special search engine then finds matches based on the parsed sentence. The found documents are searched for sentences that appear to answer the question, and the possible answers are ranked. At least 75 percent of the time, the correct answer is in the top three ranked positions, and incorrect answers are usually obvious (such as "Mickey Mantle was born in 3"). The researchers hope to include knowledge bases that will lower the rank of many of the nonsensical answers.

Microsoft researcher Eric Brill, who has led research on Ask MSR, has also attempted an even more difficult task: building a system that provides answers of about fifty words to more complex questions, such as, "How are the recipients of the Nobel Prize selected?" One of the strategies used by this system is to find an appropriate FAQ section on the Web that answers the query.

Natural-language systems combined with large-vocabulary, speaker-independent (that is, responsive to any speaker) speech recognition over the phone are entering the marketplace to conduct routine transactions. You can talk to British Airways' virtual travel agent about anything you like as long as it has to do with booking flights on British Airways.²⁰⁷ You're also likely to talk to a virtual person if you call Verizon for customer service or Charles Schwab and Merrill Lynch to conduct financial transactions. These systems, while they can be annoying to some people, are reasonably adept at responding appropriately to the often ambiguous and fragmented way people speak. Microsoft and other companies are offering systems that allow a business to create virtual agents to book reservations for travel and hotels and conduct routine transactions of all kinds through two-way, reasonably natural voice dialogues.

Not every caller is satisfied with the ability of these virtual agents to get the job done, but most systems provide a means to get a human on the line. Companies using these systems report that they reduce the need for human service agents up to 80 percent. Aside from the money saved, reducing the size of call centers has management benefit. Call-center jobs have very high turnover rates because of low job satisfaction.

It's said that men are loath to ask others for directions, but car vendors are betting that both male and female drivers will be willing to ask their own car for help in getting to their destination. In 2005 the Acura RL and Honda Odyssey will be offering a system from IBM that

allows users to converse with their cars.²⁰⁸ Driving directions will include street names (for example, "turn left on Main Street, then right on Second Avenue"). Users can ask such questions as, "Where is the nearest Italian restaurant?" or they can enter specific locations by voice, ask for clarifications on directions, and give commands to the car itself (such as "turn up the air conditioning"). The Acura RL will also track road conditions and highlight traffic congestion on its screen in real time. The speech recognition is claimed to be speaker-independent and to be unaffected by engine sound, wind, and other noises. The system will reportedly recognize 1.7 million street and city names, in addition to nearly one thousand commands.

Computer language translation continues to improve gradually. Because this is a Turing-level task—that is, it requires full human-level understanding of language to perform at human levels—it will be one of the last application areas to compete with human performance. Franz Josef Och, a computer scientist at the University of Southern California, has developed a technique that can generate a new language-translation system between any pair of languages in a matter of hours or days.²⁰⁹ All he needs is a "Rosetta stone"—that is, text in one language and the translation of that text in the other language—although he needs millions of words of such translated text. Using a self-organizing technique, the system is able to develop its own statistical models of how text is translated from one language to the other and develops these models in both directions.

This contrasts with other translation systems, in which linguists painstakingly code grammar rules with long lists of exceptions to each rule. Och's system recently received the highest score in a competition of translation systems conducted by the U.S. Commerce Department's National Institute of Standards and Technology.

<h2>Entertainment and sports. <tx>In an amusing and intriguing application of GAs, Oxford scientist Torsten Reil created animated creatures with simulated joints and muscles and a neural net for a brain. He then assigned them a task: to walk. He used a GA to evolve this capability, which involved seven hundred parameters. "If you look at that system with your human eyes, there's no way you can do it on your own, because the system is just too complex," Reil points out. "That's where evolution comes in."²¹⁰

While some of the evolved creatures walked in a smooth and convincing way, the research demonstrated a well-known attribute of GAs: you get what you ask for. Some creatures figured out novel new ways of passing for walking. According to Weil, "We got some creatures that didn't walk at all, but had these very strange ways of moving forward: crawling or doing somersaults."

Software is being developed that can automatically extract excerpts from a video of a sports game that show the more important plays.²¹¹ A team at Trinity College in Dublin is working on table-based games like pool, in which software tracks the location of each ball and is programmed to identify when a significant shot has been made. A team at the University of Florence is working on soccer. This software tracks the location of each player and can determine the type of play being made (such as free kicking or attempting a goal), when a goal is achieved, when a penalty is earned, and other key events.

The Digital Biology Interest Group at University College in London is designing Formula One race cars by breeding them using GAs.²¹²

The AI winter is long since over. We are well into the spring of narrow AI. Most of the examples above were research projects just ten to fifteen years ago. If all the AI systems in the world suddenly stopped functioning, our economic infrastructure would grind to a halt. Your bank would cease doing business. Most transportation would be crippled. Most communications would fail. This was not the case a decade ago. Of course, our AI systems are not smart enough—yet—to organize such a conspiracy.

Strong AI

<epi>If you understand something in only one way, then you don't really understand it at all. This is because, if something goes wrong, you get stuck with a thought that just sits in your mind with nowhere to go. The secret of what anything means to us depends on how we've connected it to all the other things we know. This is why, when someone learns 'by rote,' we say that they don't really understand. However, if you have several different representations then, when one approach fails you can try another. Of course, making too many indiscriminate connections will turn a mind to mush. But well-connected representations let you turn ideas around in your mind, to envision things from many perspectives until you find one that works for you. And that's what we mean by thinking!

<epis>—*Marvin Minsky*¹³

<epi>Advancing computer performance is like water slowly flooding the landscape. A half century ago it began to drown the lowlands, driving out human calculators and record clerks, but leaving most of us dry. Now the flood has reached the foothills, and our outposts there are contemplating retreat. We feel safe on our peaks, but, at the present rate, those too will be submerged within another half century. I propose that we build Arks as that day nears, and adopt a seafaring life! For now, though, we must rely on our representatives in the lowlands to tell us what water is really like.

Our representatives on the foothills of chess and theorem-proving report signs of intelligence. Why didn't we get similar reports decades before, from the lowlands, as computers surpassed humans in arithmetic and rote memorization? Actually, we did, at the time. Computers that calculated like thousands of mathematicians were hailed as "giant brains," and inspired the first generation of AI research. After all, the machines were doing something beyond any animal, that needed human intelligence, concentration and years of training. But it is hard to recapture that magic now. One reason is that computers' demonstrated stupidity in other areas biases our judgment. Another relates to our own ineptitude. We do arithmetic or keep records so painstakingly and externally, that the small mechanical steps in a long calculation are obvious, while the big picture often escapes us. Like Deep Blue's builders, we see the process too much from the inside to appreciate the subtlety that it may have on the outside. But there is a non-obviousness in snowstorms or tornadoes that emerge from the repetitive arithmetic of weather simulations, or in rippling tyrannosaur skin from movie animation calculations. We rarely call it intelligence, but "artificial reality" may be an even more profound concept than artificial intelligence (Moravec 1998).

The mental steps underlying good human chess playing and theorem proving are complex and hidden, putting a mechanical interpretation out of reach. Those who can follow the play naturally describe it instead in mentalistic language, using terms like strategy, understanding and creativity. When a machine manages to be simultaneously meaningful and surprising in the same rich way, it too compels a mentalistic interpretation. Of course, somewhere behind the scenes, there are programmers who, in principle, have a mechanical interpretation. But even for them, that interpretation loses its grip as the working program fills its memory with details too voluminous for them to grasp.

As the rising flood reaches more populated heights, machines will begin to do well in areas a greater number can appreciate. The visceral sense of a thinking presence in machinery will

become increasingly widespread. When the highest peaks are covered, there will be machines than can interact as intelligently as any human on any subject. The presence of minds in machines will then become self-evident.

<epis>—*Hans Moravec*²¹⁴

<tx>Because of the exponential nature of progress in information-based technologies, performance often shifts quickly from pathetic to daunting. In many diverse realms, as the examples in the previous section make clear, the performance of narrow AI is already impressive. The range of intelligent tasks in which machines can now compete with human intelligence is continually expanding. In a cartoon in *The Age of Spiritual Machines*, a defensive "human race" is seen writing out signs that state what only people (and not machines) can do.²¹⁵ Littered on the floor are the signs the human race has already discarded, because machines can now perform these functions: diagnose an electrocardiogram, compose in the style of Bach, recognize faces, guide a missile, play Ping-Pong, play master chess, pick stocks, improvise jazz, prove important theorems, and understand continuous speech. Back in 1999 these tasks were no longer solely the province of human intelligence; machines could do them all.



<tx>On the wall behind the man symbolizing the human race were signs he had written out describing the tasks that were still the sole province of humans: have common sense, review a movie, hold press conferences, translate speech, clean a house, and drive cars. If we were to redesign this cartoon in a few years, some of these signs would also be likely to end up on the floor. When CYC reaches one hundred million items of commonsense knowledge, perhaps human superiority in the realm of commonsense reasoning won't be so clear.

The era of household robots, although still fairly primitive today, has already started. Ten years from now, it's likely we will consider "clean a house" as within the capabilities of machines. As for driving cars, robots with no human intervention have already driven nearly across the United States on ordinary roads with other normal traffic. We are not yet ready to turn over all steering wheels to machines, but there are serious proposals to create electronic highways on which cars (with people in them) will drive by themselves.

The three tasks that have to do with human-level understanding of natural language—reviewing a movie, holding a press conference, and translating speech—are the most difficult.

Once we can take down these signs, we'll have Turing-level machines, and the era of strong AI will have started.

This era will creep up on us. As long as there are any discrepancies between human and machine performance—areas in which humans outperform machines—strong AI skeptics will seize on these differences. But our experience in each area of skill and knowledge is likely to follow that of Kasparov. Our perceptions of performance will shift quickly from pathetic to daunting as the knee of the exponential curve is reached for each human capability.

How will strong AI be achieved? Most of the material in this book is intended to lay out the fundamental requirements for both hardware and software and explain why we can be confident that these requirements will be met in nonbiological systems. The continuation of the exponential growth of the price-performance of computation to achieve hardware capable of emulating human intelligence was still controversial in 1999. There has been so much progress in developing the technology for three-dimensional computing over the past five years that relatively few knowledgeable observers now doubt that this will happen. Even just taking the semiconductor industry's published ITRS road map, which runs to 2018, we can project human-level hardware at reasonable cost by that year.²¹⁶

I've stated the case in chapter 4 of why we can have confidence that we will have detailed models and simulations of all regions of the human brain by the late 2020s. Until recently, our tools for peering into the brain did not have the spatial and temporal resolution, bandwidth, or price-performance to produce adequate data to create sufficiently detailed models. This is now changing. The emerging generation of scanning and sensing tools can analyze and detect neurons and neural components with exquisite accuracy, while operating in real time.

Future tools will provide far greater resolution and capacity. By the 2020s, we will be able to send scanning and sensing nanobots into the capillaries of the brain to scan it from inside. We've shown the ability to translate the data from diverse sources of brain scanning and sensing into models and computer simulations that hold up well to experimental comparison with the performance of the biological versions of these regions. We already have compelling models and simulations for several important brain regions. As I argued in chapter 4, it's a conservative projection to expect detailed and realistic models of all brain regions by the late 2020s.

One simple statement of the strong AI scenario is that we will learn the principles of operation of human intelligence from reverse engineering all the brain's regions, and we will apply these principles to the brain-capable computing platforms that will exist in the 2020s. We already have an effective toolkit for narrow AI. Through the ongoing refinement of these methods, the development of new algorithms, and the trend toward combining multiple methods into intricate architectures, narrow AI will continue to become less narrow. That is, AI applications will have broader domains, and their performance will become more flexible. AI systems will develop multiple ways of approaching each problem, just as humans do. Most important, the new insights and paradigms resulting from the acceleration of brain reverse engineering will greatly enrich this set of tools on an ongoing basis. This process is well under way.

It's often said that the brain works differently from a computer, so we cannot apply our insights about brain function into workable nonbiological systems. This view completely ignores the field of self-organizing systems, for which we have a set of increasingly sophisticated mathematical tools. As I discussed in the previous chapter, the brain differs in a number of important ways from that of conventional, contemporary computers. If you open up your Palm Pilot and cut a wire, there's a good chance you will break the machine. Yet we routinely lose many neurons and interneuronal connections with no ill effect, because the brain is self-organizing and relies on distributed patterns in which many specific details are not important.

When we get to the mid- to late 2020s, we will have access to a generation of extremely detailed brain-region models. Ultimately the toolkit will be greatly enriched with these new models and simulations and will encompass a full knowledge of how the brain works. As we apply the toolkit to intelligent tasks, we will draw upon the entire range of tools, some derived directly from brain reverse engineering, some merely inspired by what we know about the brain, and some not based on the brain at all but on decades of AI research.

Part of the brain's strategy is to learn information, rather than having knowledge hard-coded from the start. ("Instinct" is the term we use to refer to such innate knowledge.) Learning will be an important aspect of AI, as well. In my experience in developing pattern-recognition systems in character recognition, speech recognition, and financial analysis, providing for the AIs education is the most challenging and important part of the engineering. With the accumulated knowledge of human civilization increasingly accessible online, future AIs will have the opportunity to conduct their education by accessing this vast body of information.

The education of AIs will be much faster than that of unenhanced humans. The twenty-year time span required to provide a basic education to biological humans could be compressed into a matter of weeks or less. Also, because nonbiological intelligence can share its patterns of learning and knowledge, only one AI has to master each particular skill. As I pointed out, we trained one set of research computers to understand speech, but then the hundreds of thousands of people who acquired our speech-recognition software had to load only the already trained patterns into their computers.

One of the many skills that nonbiological intelligence will achieve with the completion of the human brain-reverse engineering project is sufficient mastery of language and shared human knowledge to pass the Turing test. The Turing test is important not so much for its practical significance but rather because it will demarcate a crucial threshold. As I have pointed out, there is no simple means to pass a Turing test, other than to convincingly emulate the flexibility, subtlety, and suppleness of human intelligence. Having captured that capability in our technology, it will then be subject to engineering's ability to concentrate, focus, and amplify it.

Variations of the Turing test have been proposed. The annual Loebner Prize contest awards a bronze prize to the chatterbot (conversational bot) best able to convince human judges that it's human.²¹⁷ The criteria for winning the silver prize is based on Turing's original test, and it obviously has yet to be awarded. The gold prize is based on visual and auditory communication. In other words, the AI must have a convincing face and voice, as transmitted over a terminal, and thus it must appear to the human judge as if he or she is interacting with a real person over a videophone. On the face of it, the gold prize sounds more difficult. I've argued that it may actually be easier, because judges may pay less attention to the text portion of the language being communicated and could be distracted by a convincing facial and voice animation. In fact, we already have real-time facial animation, and while it is not quite up to these modified Turing standards, it's reasonably close. We also have very natural-sounding voice synthesis, which is often confused with recordings of human speech, although more work is needed on prosodies (intonation). We're likely to achieve satisfactory facial animation and voice production sooner than the Turing-level language and knowledge capabilities.

Turing was carefully imprecise in setting the rules for his test, and significant literature has been devoted to the subtleties of establishing the exact procedures for determining how to assess when the Turing test has been passed.²¹⁸ In 2002 I negotiated the rules for a Turing-test wager with Mitch Kapor on the Long Now Web site.²¹⁹ The question underlying our twenty-thousand-dollar bet, the proceeds of which go to the charity of the winner's choice, was, "Will the Turing test be passed by a machine by 2029?" I said yes, and Kapor said no. It took us months of dialogue to arrive at the intricate rules to implement our wager. Simply defining

"machine" and "human," for example, was not a straightforward matter. Is the human judge allowed to have any nonbiological thinking processes in his or her brain? Conversely, can the machine have any biological aspects?

Because the definition of the Turing test will vary from person to person, Turing test-capable machines will not arrive on a single day, and there will be a period during which we will hear claims that machines have passed the threshold. Invariably, these early claims will be debunked by knowledgeable observers, probably including myself. By the time there is a broad consensus that the Turing test has been passed, the actual threshold will have long since been achieved.

Edward Feigenbaum proposes a variation of the Turing test, which assesses not a machine's ability to pass for human in casual, everyday dialogue but its ability to pass for a scientific expert in a specific field.²²⁰ The Feigenbaum test (FT) may be more significant than the Turing test, because FT-capable machines, being technically proficient, will be capable of improving their own designs. Feigenbaum describes his test in this way:

<ext>Two players play the FT game. One player is chosen from among the elite practitioners in each of three pre-selected fields of natural science, engineering, or medicine. (The number could be larger, but for this challenge not greater than ten). Let's say we choose the fields from among those covered in the U.S. National Academy. . . . For example, we could choose astrophysics, computer science, and molecular biology. In each round of the game, the behavior of the two players (elite scientist and computer) is judged by another Academy member in that particular domain of discourse, e.g., an astrophysicist judging astrophysics behavior. Of course the identity of the players is hidden from the judge as it is in the Turing test. The judge poses problems, asks questions, asks for explanations, theories, and so on—as one might do with a colleague. Can the human judge choose, at better than chance level, which is his National Academy colleague and which is the computer?

<tx>Of course Feigenbaum overlooks the possibility that the computer might already be a National Academy colleague, but he is obviously assuming that machines will not yet have invaded institutions that today comprise exclusively biological humans. While it may appear that the FT is more difficult than the Turing test, the entire history of AI reveals that machines started with the skills of professionals and only gradually moved towards the language skills of a child. Early AI systems demonstrated their prowess initially in professional fields such as proving mathematical theorems and diagnosing medical conditions. These early systems would not be able to pass the FT, however, because they do not have the language skills and the flexible ability to model knowledge from different perspectives, which are needed to engage in the professional dialogue inherent in the FT.

This language ability is essentially the same ability needed in the Turing test. Reasoning in many technical fields is not necessarily more difficult than the commonsense reasoning engaged in by most human adults. I would expect that machines will pass the FT, at least in some disciplines, around the same time as they pass the Turing test. Passing the FT in all disciplines is likely to take longer, however. This is why I see the 2030s as a period of consolidation, as machine intelligence rapidly expands its skills and incorporates the vast knowledge bases of our biological human and machine civilization. By the 2040s we will have the opportunity to apply the accumulated knowledge and skills of our civilization to computational platforms that are billions of times more capable than unassisted biological human intelligence.

The advent of strong AI is the most important transformation this century will see. Indeed, it's comparable in importance to the advent of biology itself. It will mean that a creation of biology has finally mastered its own intelligence and discovered means to overcome its limitations. Once the principles of operation of human intelligence are understood, expanding its abilities will be conducted by human scientists and engineers whose own biological intelligence will have been greatly amplified through an intimate merger with nonbiological intelligence. Over time, the nonbiological portion will predominate.

We've discussed aspects of the impact of this transformation throughout this book, which I focus on in the next chapter. Intelligence is the ability to solve problems with limited resources, including limitations of time. The Singularity will be characterized by the rapid cycle of human intelligence—increasingly nonbiological—capable of comprehending and leveraging its own powers.

<dia>Friend of Futurist Bacteria, 2 billion B.C.: So tell me again about these ideas you have about the future.

Futurist Bacteria, 2 billion B.C.: Well, I see bacteria getting together into societies, with the whole band of cells basically acting like one big complicated organism with greatly enhanced capabilities.

Friend of Futurist Bacteria: What gives you that idea?

Futurist Bacteria: Well already, some of our fellow Daptobacters have gone inside other larger bacteria to form a little duo.²²¹ It's inevitable that our fellow cells will band together so that each cell can specialize its function. As it is now, we each have to do everything by ourselves: find food, digest it, excrete by-products.

Friend of Futurist Bacteria: And then what?

Futurist Bacteria: All these cells will develop ways of communicating with one another that go beyond just the swapping of chemical gradients that you and I can do.

Friend of Futurist Bacteria: Okay, now tell me again the part about that future superassembly of ten trillion cells.

Futurist Bacteria: Yes, well, according to my models, in about two billion years a big society of ten trillion cells will make up a single organism and include tens of billions of special cells that can communicate with one another in very complicated patterns.

Friend of Futurist Bacteria: What sort of patterns?

Futurist Bacteria: Well, "music," for one thing. These huge bands of cells will create musical patterns and communicate them to all the other bands of cells.

Friend of Futurist Bacteria: Music?

Futurist Bacteria: Yes, patterns of sound.

Friend of Futurist Bacteria: Sound?

Futurist Bacteria: Okay, look at it this way. These supercell societies will be complicated enough to understand their own organization. They will be able to improve their own design, getting better and better, faster and faster. They will reshape the rest of the world in their image.

Friend of Futurist Bacteria: Now, wait a second. Sounds like we'll lose our basic bacteriumity.

Futurist Bacteria: Oh, but there will be no loss.

Friend of Futurist Bacteria: I know you keep saying that, but . . .

Futurist Bacteria: It will be a great step forward. It's our destiny as bacteria. And, anyway, there will still be little bacteria like us floating around.

Friend of Futurist Bacteria: Okay, but what about the downside? I mean, how much harm can our fellow Daptobacter a/ft/Bdellovibrio bacteria do? But these future cell associations with their vast reach may destroy everything.

Futurist Bacteria: It's not certain, but I think we'll make it through.

Friend of Futurist Bacteria: You always were an optimist.

Futurist Bacteria: Look, we won't have to worry about the downside for a couple billion years.

Friend of Futurist Bacteria: Okay, then, let's get lunch.

<sd>Meanwhile, two billion years later . . .

<dia>*Ned Ludd: These future intelligences will be worse than the textile machines I fought back in 1812. Back then we had to worry about only one man with a machine doing the work of twelve. But you're talking about a marble-size machine outperforming all of humanity.*

Ray: It will only outperform the biological part of humanity. In any event, that marble is still human, even if not biological.

Ned: These superintelligences won't eat food. They won't breathe air. They won't reproduce through sex. . . . So just how are they human?

Ray: We're going to merge with our technology. We're already starting to do that in 2004, even if most of the machines are not yet inside our bodies and brains. Our machines nonetheless extend the reach of our intelligence. Extending our reach has always been the nature of being human.

Ned: Look, saying that these superintelligent nonbiological entities are human is like saying that we're basically bacteria. After all, we're evolved from them also.

Ray: It's true that a contemporary human is a collection of cells, and that we are a product of evolution, indeed its cutting edge. But extending our intelligence by reverse engineering it, modeling it, simulating it, reinstantiating it on more capable substrates, and modifying and extending it is the next step in its evolution. It was the fate of bacteria to evolve into a technology-creating species. And it's our destiny now to evolve into the vast intelligence of the Singularity.

¹ Samuel Butler (1835-1902), "Darwin Among the Machines," *CAM/ Church Press*, June 13, 1863 (republished by Festing Jones in 1912 in *The Notebooks of Samuel Butler*).

² Peter Weibel, "Virtual Worlds: The Emperor's New Bodies," in *Ars Electronica: Facing the Future*, ed. Timothy Druckery (Cambridge, Mass.: MIT Press, 1999), pp. 207-23; available online at http://www.aec.at/en/archiv_files/19902/E1990b_009.pdf.

³ James Watson and Francis Crick, "Molecular structure of nucleic acids: a structure for deoxyribose nucleic acid," *Nature* 171.4356 (April 23, 1953): 737-738, <http://www.nature.com/nature/dnaSO/watsoncrick.pdf>.

⁴ Robert Waterston quoted in "Scientists reveal complete sequence of human genome," *CBC News*, April 14 2003, <http://www.cbc.ca/story/science/national/2003/04/14/genome030414.html>.

⁵ See Chapter 2, endnote 57.

⁶ The original reports of Crick and Watson, which still make compelling reading today, may be found in James A. Peters, ed., *Classic Papers in Genetics* (Englewood Cliffs, N.J.: Prentice-Hall, 1959). An exciting account of the successes and failures that led to the double helix is given in J. D. Watson, *The Double Helix: A Personal Account of the Discovery of the Structure of DNA* (New York: Atheneum, 1968). *Nature.com* has a collection of Crick's papers available online at <http://www.nature.com/nature/focus/crick/index.html>.

⁷ Morislav Radman and Richard Wagner, "The High Fidelity of DNA Duplication," *Scientific American* 259.2 (August 1988): 40-46.

⁸ The structure and behavior of DNA and RNA are described in Gary Felsenfeld, "DNA," and James Darnell, "RNA," both in *Scientific American* 253.4 (October 1985) pages 58-67 and 68-78 respectively.

⁹ Mark A. Jobling and Chris Tyler-Smith, "The Human Y Chromosome: An Evolutionary Marker Comes of Age," *Nature Reviews Genetics* 4 (August 2003): 598-612; Helen Skaletsky et al., "The Male-Specific Region of the Human Y Chromosome Is a Mosaic of Discrete Sequence Classes," *Nature* 423 (June 19, 2003): 825-37.

¹⁰ Misformed proteins are perhaps the most dangerous toxin of all. Research suggests that misfolded proteins may be at the heart of numerous disease processes in the body. Such diverse diseases as Alzheimer's disease, Parkinson's

disease, the human form of mad-cow disease, cystic fibrosis, cataracts, and diabetes are all thought to result from the inability of the body to adequately eliminate misfolded proteins.

Protein molecules perform the lion's share of cellular work. Proteins are made within each cell according to DNA blueprints. They begin as long strings of amino acids, which must then be folded into precise three-dimensional configurations in order to function as enzymes, transport proteins, et cetera. Heavy-metal toxins interfere with normal function of these enzymes, further exacerbating the problem. There are also genetic mutations that predispose individuals to misformed-protein buildup.

When protofibrils begin to stick together, they form filaments, fibrils, and ultimately larger globular structures called amyloid plaque. Until recently these accumulations of insoluble plaque were regarded as the pathologic agents for these diseases, but it is now known that the protofibrils themselves are the real problem. The speed with which a protofibril is turned into insoluble amyloid plaque is inversely related to disease progression. This explains why some individuals are found to have extensive accumulation of plaque in their brains but no evidence of Alzheimer's disease, while others have little visible plaque yet extensive manifestations of the disease. Some people form amyloid plaque quickly, which protects them from further protofibril damage. Other individuals turn protofibrils into amyloid plaque less rapidly, allowing more extensive damage. These people also have little visible amyloid plaque. See Per Hammarström, Frank Schneider, Jeffrey W. Kelly, "Trans-Suppression of Misfolding in an Amyloid Disease," *Science* 293.5539 (September 28, 2001): 2459-62.

¹¹ A fascinating account of the new biology is given in Horace F. Judson, *The Eighth Day of Creation: The Makers of the Revolution in Biology* (Woodbury, N.Y.: CSHL Press, 1996).

¹² Raymond Kurzweil and Terry Grossman, M.D., *Fantastic Voyage: Live Long Enough to Live Forever* (New York: Rodale, 2004). See <http://www.Fantastic-Voyage.net> and <http://www.RavandTerrv.com>.

¹³ Raymond Kurzweil, *The 10% Solution for a Healthy Life: How to Eliminate Virtually All Risk of Heart Disease and Cancer* (New York: Crown Books, 1993).

¹⁴ Kurzweil and Grossman, *Fantastic Voyage*. "Ray & Terry's Longevity Program" is articulated throughout the book.

¹⁵ The test for "biological age," called the H-scan test, includes tests for auditory-reaction time, highest audible pitch, vibrotactile sensitivity, visual-reaction time, muscle-movement time, lung (forced expiratory) volume, visual-reaction time with decision, muscle-movement time with decision, memory (length of sequence), alternative button-tapping time, and visual accommodation. The author had this test done at Frontier Medical Institute (Grossman's health and longevity clinic), <http://www.FMIClinic.com>. For information on the H-scan test, see Diagnostic and Lab Testing, Longevity Institute, Dallas, <http://www.lidhealth.com/diagnostic.html>.

¹⁶ Kurzweil and Grossman, *Fantastic Voyage*, chapter 10: "Ray's Personal Program."

¹⁷ *Ibid.*

¹⁸ Aubrey D. N. J. de Grey, "The Foreseeability of Real Anti-Aging Medicine: Focusing the Debate," *Experimental Gerontology* 38.9 (September 2003): 927-34; Aubrey D. N. J. de Grey, "An Engineer's Approach to the Development of Real Anti-Aging Medicine," *Science of Aging, Knowledge, Environment* 1 (2003); Aubrey D. N. J. de Grey et al., "Is Human Aging Still Mysterious Enough to be Left Only to Scientists?" *BioEssays* 24.7 (July 2002): 667-76.

¹⁹ Aubrey D. N. J. De Grey, ed., *Strategies for Engineered Negligible Senescence: Why Genuine Control of Aging May Be Foreseeable*, Annals of the New York Academy of Sciences, vol. 1019 (New York: New York Academy of Sciences, June 2004).

²⁰ In addition to providing the junctions of different types of cells, two other reasons for cells to control the expression of genes are environmental cues and developmental processes. Even simple organisms such as bacteria can turn on and off the synthesis of proteins depending on environmental cues. *E. coli*, for example, can turn off the synthesis of proteins that allow it to control the level of nitrogen gas from the air when there are other, less energy-intensive sources of nitrogen in its environment. A recent study of 1,800 strawberry genes found that the expression of two hundred of those genes varied during different stages of development. E. Marshall. "An Array of Uses: Expression Patterns in Strawberries, Ebola, TB, and Mouse Cells," *Science* 286.5439 (1999): 445.

²¹ Along with a protein-encoding region, genes include regulatory sequences called promoters and enhancers that control where and when that gene is expressed. Promoters of genes that encode proteins are typically located immediately "upstream" on the DNA. An enhancer activates the use of a promoter, thereby controlling the rate of gene expression. Most genes require enhancers to be expressed. Enhancers have been called "the major determinant of differential transcription in space (cell type) and time"; and any given gene can have several different enhancer sites linked to it (S. F. Gilbert, *Developmental Biology*, 6th ed. [Sunderland, Mass.: Sinauer Associates, 2000]; available online at www.ncbi.nlm.nih.gov/books/bv.fcgi?call=bv.View..ShowSection&rid=.OBPKYEB-SPfxl8nm8QOxH).

By binding to enhancer or promoter regions, transcription factors start or repress the expression of a gene. New knowledge of transcription factors has transformed our understanding of gene expression. Per Gilbert in the chapter "The Genetic Core of Development: Differential Gene Expression," "The gene itself is no longer seen as an independent entity controlling the synthesis of proteins. Rather, the gene both directs and is directed by protein synthesis. Natalie Anger (1992) has written, 'A series of discoveries suggests that DNA is more like a certain type of politician, surrounded by a flock of protein handlers and advisors that must vigorously massage it, twist it and, on occasion, reinvent it before the grand blueprint of the body can make any sense at all.'"

²² Bob Holmes, "Gene Therapy May Switch Off Huntington's," March 13, 2003,

<http://www.newscientist.com/news/news.jsp?id=ns99993493>. "Emerging as a powerful tool for reverse genetic analysis, RNAi is rapidly being applied to study the function of many genes associated with human disease, in particular those associated with oncogenesis and infectious disease." J. C. Cheng, T. B. Moore, and K. M.

Sakamoto, "RNA Interference and Human Disease," *Molecular Genetics and Metabolism* 80.1-2 (October 2003): 121-28. RNAi is a "potent and highly sequence-specific mechanism." L. Zhang, D. K. Fogg, and D. M. Waisman, "RNA Interference-Mediated Silencing of the SIOOA10 Gene Attenuates Plasmin Generation and Invasiveness of Colo 222 Colorectal Cancer Cells," *Journal of Biological Chemistry* 279.3 (January 16, 2004): 2053-62.

²³ Each chip contains synthetic oligonucleotides that replicate sequences that identify specific genes. "To determine which genes have been expressed in a sample, researchers isolate messenger RNA from test samples, convert it to complementary DNA (cDNA), tag it with fluorescent dye, and run the sample over the wafer. Each tagged cDNA will stick to an oligo with a matching sequence, lighting up a spot on the wafer where the sequence is known. An automated scanner then determines which oligos have bound, and hence which genes were expressed," E. Marshall, "Do-It-Yourself Gene Watching," *Science* 286.5439 (October 15, 1999): 444-47.

²⁴ Ibid.

²⁵ J. Rosamond and A. Allsop, "Harnessing the Power of the Genome in the Search for New Antibiotics," *Science* 287.5460 (March 17, 2000): 1973-76.

²⁶ T. R. Golub et al., "Molecular Classification of Cancer: Class Discovery and Class Prediction by Gene Expression Monitoring," *Science* 286.5439 (October 15, 1999): 53 1-37. I

²⁷ Ibid., as reported in A. Berns, "Cancer: Gene Expression in Diagnosis," *Nature* 403 (February 3, 2000): 491-92.

In another study, 1 percent of the genes studied showed reduced expression in aged muscles. These genes produced proteins associated with energy production and cell building, so a reduction makes sense given the weakening associated with age. Genes with increased expression produced stress proteins, which are used to repair damaged DNA or proteins. J. Marx, "Chipping Away at the Causes of Aging," *Science* 287.5462 (March 31, 2000): 2390.

As another example, liver metastases are a common cause of colorectal cancer. These metastases respond differently to treatment depending on their genetic profile. Expression profiling is an excellent way to determine an appropriate mode of treatment. J. C. Sung et al., "Genetic Heterogeneity of Colorectal Cancer Liver Metastases." *Journal of Surgical Research* 114.2 (October 2003): 251.

As a final example, researchers have had difficulty analyzing the Reed-Sternberg cell of Hodgkin's disease due to its extreme rarity in diseased tissue. Expression profiling is now providing a clue regarding the heritage of this cell. J. Cossman et al., "Reed-Sternberg Cell Genome Expression Supports a B-Cell Lineage," *Blood* 94.2 (July 15, 1999): 411-16.

²⁸ T. Ueland et al., "Growth Hormone Substitution Increases Gene Expression of Members of the IGF Family in Cortical Bone from Women with Adult Onset Growth Hormone Deficiency-Relationship with Bone Turn-Over," *Bone* 33.4 (October 2003): 638-45.

²⁹ R. Lovett, "Toxicologists Brace for Genomics Revolution," *Science* 289.5479 (July 28, 2000): 536-37.

³⁰ Gene transfer to somatic cells affects a subset of cells in the body for a period of time. It is theoretically possible to also alter genetic information in egg and sperm (germ-line) cells, for the purpose of passing on those changes to the next generations. Such therapy poses many ethical concerns and has not yet been attempted. "Gene Therapy," Wikipedia, http://en.wikipedia.org/wiki/Gene_therapy.

³¹ Genes encode proteins, which perform vital functions in the human body. Abnormal or mutated genes encode proteins that are unable to perform those functions, resulting in genetic disorders and diseases. The goal of gene therapy is to replace the defective genes so that normal proteins are produced. This can be done in a number of ways, but the most typical way is to insert a therapeutic replacement gene into the patient's target cells using a carrier molecule called a vector. "Currently, the most common vector is a virus that has been genetically altered to carry normal human DNA. Viruses have evolved a way of encapsulating and delivering their genes to human cells in a pathogenic manner. Scientists have tried to take advantage of this capability and manipulate the virus genome to remove the disease-causing genes and insert therapeutic genes." (Human Genome Project, "Gene Therapy," http://www.ornl.gov/TechResources/Human_Genome/medicine/genetherapy.html). See the Human Genome Project

site for more information about gene therapy and links. Gene therapy is an important enough area of research that there are currently six scientific peer-reviewed gene-therapy journals and four professional associations dedicated to this topic.

- ³² K. R. Smith., "Gene Transfer in Higher Animals: Theoretical Considerations and Key Concepts," *Journal of Biotechnology* 99,1 (October 9, 2002): 1-22.
- ³³ Anil Ananthaswamy, "Undercover Genes Slip into the Brain," March 20, 2003, <http://www.newscientist.com/news/news.jsp?id=ns99993520>.
- ³⁴ A. E. Trezise et al., "In Vivo Gene Expression: DNA Electrotransfer," *Current Opinion in Molecular Therapeutics* 5.4 (August 2003): 397-404.
- ³⁵ Sylvia Westphal, "DNA Nanoballs Boost Gene Therapy," May 12, 2002, <http://www.newscientist.com/news/news.jsp?id=ns99992257>.
- ³⁶ L. Wu, M. Johnson, and M. Sato, "Transcriptionally Targeted Gene Therapy to Detect and Treat Cancer," *Trends in Molecular Medicine* 9.10 (October 2003): 421-29.
- ³⁷ S. Westphal, "Virus Synthesized in a Fortnight," November 14, 2003, www.newscientist.com/news/news.jsp?id=ns99994383.
- ³⁸ G. Chiesa, "Recombinant Apolipoprotein A-I(Milano) Infusion into Rabbit Carotid Artery Rapidly Removes Lipid from Fatty Streaks." *Circulation Research* 90.9 (May 17, 2002): 974-80; P. K. Shah et al., "High-Dose Recombinant Apolipoprotein A-I(milano) Mobilizes Tissue Cholesterol and Rapidly Reduces Plaque Lipid and Macrophage Content in Apolipoprotein e-Deficient Mice," *Circulation* 103.25 (June 26, 2001): 3047-50.
- ³⁹ S. E. Nissen et al. "Effect of Recombinant Apo A-I Milano on Coronary Atherosclerosis in Patients with Acute Coronary Syndromes: A Randomized Controlled Trial," *JAMA* 290.17 (November 5, 2003): 2292-2300.
- ⁴⁰ A recent phase 2 study reported "markedly increased HDL cholesterol levels and also decreased LDL cholesterol levels." M. E. Brousseau et al. "Effects of an Inhibitor of Cholesteryl Ester Transfer Protein on HDL Cholesterol," *New England Journal of Medicine* 350.15 (April 8, 2004): 1505-15, <http://content.nejm.org/cgi/content/abstract/350/15/1505>. Global phase 3 trials began in late 2003. Information on Torcetrapib is available on the Pfizer site: http://www.pfizer.com/are/investors_reports/annual_2003/review/p2003,ar14_15.htm.
- ⁴¹ O. J. Finn, "Cancer Vaccines: Between the Idea and the Reality," *Nature Reviews: Immunology* 3.8 (August 2003): 63(M 1; R. C. Kennedy and M. H. Shearer, "A Role for Antibodies in Tumor Immunity," *International Reviews of Immunology* 22.2 (March-April 2003): 141-72.
- ⁴² T. F. Greten and E. M. Jaffee, "Cancer Vaccines," *Journal of Clinical Oncology* 17.3 (March 1999): 1047-60.
- ⁴³ "Cancer 'Vaccine' Results Encouraging," *BBCNews*, January 8, 2001, <http://news.bbc.co.uk/2/hi/health/102618.stm>, reporting on research by E. M. Jaffee et al., "Novel Allogeneic Granulocyte-Macrophage Colony-Stimulating Factor-Secreting Tumor Vaccine for Pancreatic Cancer: A Phase I Trial of Safety and Immune Activation," *Journal of Clinical Oncology* 19.1 (January 1, 2001): 145-56.
- ⁴⁴ John Travis, "Fused Cells Hold Promise of Cancer Vaccines," March 4, 2000, <http://www.sciencenews.org/articles/20000304/fob3.asp>, referring to D. W. Kufe, "Smallpox, Polio and Now a Cancer Vaccine?" *Nature Medicine* 6 (March 2000): 252-53.
- ⁴⁵ J. D. Lewis, B. D. Reilly, and R. K. Bright, "Tumor-Associated Antigens: From Discovery to Immunity," *International Reviews of Immunology* 22.2 (March-April 2003): 81-112.
- ⁴⁶ T. Boehm et al., "Antiangiogenic Therapy of Experimental Cancer Does Not Induce Acquired Drug Resistance," *Nature* 390.6658 (November 27, 1997): 404-7.
- ⁴⁷ Angiogenesis Foundation, "Understanding Angiogenesis," http://www.angio.org/understanding/content_understanding.html; L. K. Lassiter and M. A. Carducci, "Endothelin Receptor Antagonists in the Treatment of Prostate Cancer," *Seminars in Oncology* 30.5 (October 2003): 678-88. For an explanation of the process, see the National Cancer Institute Web site, "Understanding Angiogenesis," <http://press2.nci.nih.gov/sciencebehind/angiogenesis/angio02.htm>.
- ⁴⁸ B. Roninson, "Tumor Cell Senescence in Cancer Treatment," *Cancer Research* 63.11 (June 1, 2003): 2705-15; B. R. Davies et al., "Immortalization of Human Ovarian Surface Epithelium with Telomerase and Temperature-Sensitive SV40 Large T Antigen," *Experimental Cell Research* 288.2 (August 15, 2003): 390-402.
- ⁴⁹ See the research on aging mentioned in note 11. See also R. C. Woodruff and J. N. Thompson Jr., "The Role of Somatic and Germline Mutations in Aging and a Mutation Interaction Model of Aging," *Journal of Anti-Aging Medicine* 6.1 (spring 2003): 29-39.
- ⁵⁰ Aubrey D. N. J. de Grey, "The Reductive Hotspot Hypothesis of Mammalian Aging: Membrane Metabolism Magnifies Mutant Mitochondrial Mischief," *European Journal of Biochemistry* 269.8 (April 2002): 2003-9; P. F. Chin π tery et al., "Accumulation of Mitochondrial DNA Mutations in Ageing, Cancer, and Mitochondrial Disease: Is

There a Common Mechanism?" *The Lancet* 360.9342 (October 26, 2002): 1323-25; A. D. de Grey, "Mitochondrial Gene Therapy: An Arena for the Biomedical Use of Inteins," *Trends in Biotechnology* 18.9 (September 2000): 394-99.

⁵¹ "The notion of 'vaccinating' individuals against a neurodegenerative disorder such as Alzheimer's disease is a marked departure from classical thinking about mechanism and treatment, and yet therapeutic vaccines for both Alzheimer's disease and multiple sclerosis have been validated in animal models and are in the clinic. Such approaches, however, have the potential to induce unwanted inflammatory responses as well as to provide benefit" (H. L. Weiner and D. J. Selkoe, "Inflammation and Therapeutic Vaccination in CNS Diseases," *Nature* 420.6917 [December 19-26, 2002]: 879-84). These researchers showed that a vaccine in the form of nose drops could slow the brain deterioration of Alzheimer's. H. L. Weiner et al., "Nasal Administration of Amyloid-beta Peptide Decreases Cerebral Amyloid Burden in a Mouse Model of Alzheimer's Disease," *Annals of Neurology* 48.4 (October 2000): 567-79.

⁵² S. Vasani, P. Foiles, and H. Founds, "Therapeutic Potential of Breakers of Advanced Glycation End Product-Protein Crosslinks," *Archives of Biochemistry and Biophysics* 419.1 (November 1, 2003): 89-96; D. A. Kass, "Getting Better Without AGE: New Insights into the Diabetic Heart," *Circulation Research* 92.7 (April 18, 2003): 704-6.

⁵³ S. Graham, "Methuselah Worm Remains Energetic for Life." October 27, 2003;

www.sciam.com/article.cfm?chanID=sa003&articleID=000C601F-8711-1F99-86FB83414B7F0156.

⁵⁴ Ron Weiss's home page at Princeton University (<http://www.princeton.edu/~rweiss>) lists his publications, such as "Genetic Circuit Building Blocks for Cellular Computation, Communications, and Signal Processing," *Natural Computing, an international Journal* 2.1 (January 2003): 47-84.

⁵⁵ S. L. Garfinkel, "Biological Computing," *Technology Review* (May/June 2000),

<http://static.highbeam.com/t/technologyreview/may012000/biologicalcomputing>.

⁵⁶ Ibid. See also the list of current research on the MIT Media Lab Web site,

<http://www.media.mit.edu/research/index.html>.

⁵⁷ Here is one possible explanation: "In mammals, female embryos have two X-chromosomes and males have one. During early development in females, one of the X's and most of its genes are normally silenced or inactivated. That way, the amount of gene expression in males and females is the same. But in cloned animals, one X-chromosome is already inactivated in the donated nucleus. It must be reprogrammed and then later inactivated again, which introduces the possibility of errors." CBC News online staff, "Genetic Defects May Explain Cloning Failures," May 27, 2002, http://www.cbc.ca/stories/2002/05/27/cloning_errors020527. That story reports on F. Xue et al., "Aberrant Patterns of X Chromosome Inactivation in Bovine Clones," *Nature Genetics* 31.2 (June 2002): 216-20.

⁵⁸ Rick Weiss, "Clone Defects Point to Need for 2 Genetic Parents," *The Washington Post*, May 10, 1999,

<http://www.gene.ch/genet/1999/Jun/msg00004.html>.

⁵⁹ A. Baguisi et al., "Production of Goats by Somatic Cell Nuclear Transfer," *Nature Biotechnology* 5 (May 1999):

456-61. For more information on the partnership between Genzyme Transgenics Corporation, Louisiana State University, and Tufts University School of Medicine that produced this work, see the April 27, 1999, press release, "Genzyme Transgenics Corporation Announces First Successful Cloning of Transgenic Goat," <http://www.transgenics.com/pressreleases/prO42799.html>.

⁶⁰ Luba Vangelova, "True or False? Extinction is Forever," *Smithsonian Magazine*, June 2003,

<http://www.smithsonianmag.com/smithsonian/issues03/jun03/phenomena.html>.

⁶¹ J. B. Gurdon and A. Colman, "The Future of Cloning," *Nature* 402.6763 (December 16, 1999): 743^46; Gregory Stock and John Campbell, eds., *Engineering the Human Germline: An Exploration of the Science and Ethics of Altering the Genes We Pass to Our Children* (New York: Oxford University Press, 2000).

⁶² As the Scripps Research Institute points out, "The ability to dedifferentiate or reverse a lineage-committed cells to multipotent progenitor cells might overcome many of the obstacles associated with using ESCs and adult stem cells in clinical applications (inefficient differentiation, rejection of allogenic cells, efficient isolation and expansion, etc.). With an efficient dedifferentiation process, it is conceivable that healthy, abundant and easily accessible adult cells could be used to generate different types of functional cells for the repair of damaged tissues and organs" (<http://www.scripps.edu/chem/ding/sciences.htm>).

<next>The direct conversion of one differentiated cell type into another—a process referred to as transdifferentiation—would be beneficial for producing isogenic [patient's own] cells to replace sick or damaged cells or tissue. Adult stem cells display a broader differentiation potential than anticipated and might contribute to tissues other than those in which they reside. As such, they could be worthy therapeutic agents. Recent advances in transdifferentiation involve nuclear transplantation, manipulation of cell culture conditions, induction of ectopic gene expression and uptake of molecules from cellular extracts. These approaches open the doors to new avenues

for engineering isogenic replacement cells. To avoid unpredictable tissue transformation, nuclear reprogramming requires controlled and heritable epigenetic modifications. Considerable efforts remain to unravel the molecular processes underlying nuclear reprogramming and evaluate stability of the changes in reprogrammed cells.

<ntx>P. Collas and A. M. Hakelien, "Teaching Cells New Tricks," *Trends in Biotechnology* 21.8 (August 2003): 354-61 ; P. Collas, "Nuclear Reprogramming in Cell-free Extracts," *Philosophical Transactions of the Royal Society of London, B* 358.1436 (August 29, 2003): 1389-95.

⁶³ Researchers have converted human liver cells to pancreas cells in the laboratory: Jonathan Slack et al., "Experimental conversion of liver to pancreas," *Current Biology* 13.2 (January 2003):105-15. Researchers reprogrammed cells to behave like other cells using cell extracts- for example, skin cells were reprogrammed to exhibit T-cell characteristics. Anne-Mari Hakelien et al., "Reprogramming Fibroblasts to Express T-Cell Functions Using Cell Extracts," *Nature Biotechnology* 20.5 (May 2002): 460-66; Mari Hakelien and P. Collas, "Novel Approaches to Transdifferentiation," *Cloning Stem Cells* 4.4 (2002): 379-87. See also David Tosh & Jonathan M. W. Slack, "How Cells change their Phenotype," *Nature Reviews Molecular Cell Biology* 3.3 (March 2002): 187 - 194.

⁶⁴ See the description of transcription factors in note 21 above.

⁶⁵ R. P Lanza et al., "Extension of Cell Life-Span and Telomere Length in Animals Cloned from Senescent Somatic Cells," *Science* 288.5466 (April 28, 2000): 665-69. See also J. C. Ameisen, "On the Origin, Evolution, and Nature of Programmed Cell Death: A Timeline of Four Billion Years," *Cell Death and Differentiation* 9.4 (April 2002): 367-93; Mary-Ellen Shay, "Transplantation Without a Donor," *Dream: The Magazine of Possibilities* (Children's Hospital, Boston), fall 2001 .

⁶⁶ In 2000 the Immune Tolerance Network (<http://www.immunetolerance.org>), a project of the National Institutes of Health (NIH) and the Juvenile Diabetes Foundation, announced a multicenter clinical trial to assess the effectiveness of Islet transplantation.

According to a clinical-trial research summary (James Shapiro, "Cam path-1H and One-Year Temporary Sirolimus Maintenance Monotherapy in Clinical Islet Transplantation," <http://www.immunetolerance.org/public/clinical/islet/trials/shapiro2.html>), "This therapy is not suitable for all patients with Type I diabetes, even if there were no limitation in islet supply, because of the potential long-term risks of cancer, life-threatening infections and drug side-effects related to the anti-rejection therapy. If tolerance [indefinite graft function without a need for long-term drugs to prevent rejection] could be achieved at minimal up-front risk, then islet transplant could be used safely earlier in the course of diabetes, and eventually in children at the time of diagnosis."

⁶⁷ "Lab Grown Steaks Nearing Menu," <http://www.newscientist.com/news/news.jsp?id =ns99993208>, includes discussion of technical issues.

⁶⁸ The halving time for feature sizes is five years in each dimension. See discussion in chapter 2.

⁶⁹ An analysis by Robert A. Freitas Jr. indicates that replacing 10% of a person's red blood cells with robotic respirocytes would enable holding one's breath for about four hours, which is about 240 times longer than one minute (about the length of time feasible with all biological red blood cells). Since this increase derives from replacing only 10 percent of the red blood cells, the respirocytes are thousands of times more effective.

⁷⁰ Nanotechnology is "thorough, inexpensive control of the structure of matter based on molecule-by-molecule control of products and byproducts; the products and processes of molecular manufacturing, including molecular machinery" (Eric Drexler and Chris Peterson, *Unbounding the Future: The Nanotechnology Revolution* [New York: William Morrow, 1991]). According to the authors:

<next>Technology has been moving toward greater control of the structure of matter for millennia. . . . [P]ast advanced technologies—microwave tubes, lasers, superconductors, satellites, robots, and the like—have come trickling out of factories, at first with high price tags and narrow applications. Molecular manufacturing, though, will be more like computers: a flexible technology with a huge range of applications. And molecular manufacturing won't come trickling out of conventional factories as computers did; it will replace factories and replace or upgrade their products. This is something new and basic, not just another twentieth-century gadget. It will arise out of twentieth-century trends in science, but it will break the trend-lines in technology, economics, and environmental affairs, (chap. 1)

<ntx> Drexler and Peterson outline the possible scope of the effects of the revolution: efficient solar cells "as cheap as newspaper and as tough as asphalt," molecular mechanisms that can kill cold viruses in six hours before biodegrading, immune machines that destroy malignant cells in the body at the push of a button, pocket supercomputers, the end of the use of fossil fuels, space travel, and restoration of lost species. Also see E. Drexler, *Engines of Creation* (New York: Anchor Books, 1986). The Foresight Institute has a useful list of nanotechnology FAQs (<http://www.foresight.org/NanoRev/FIFAQ1.html>) and other information. Other Web resources include the

National Nanotechology Initiative (<http://www.nano.gov>), <http://nanotechweb.org>, Dr. Ralph Merkle's nanotechnology page (<http://www.zyvex.com/nan0>), and *Nanotechnology*, an online journal (<http://www.iop.org/EJ/journal/0957-4484>). Extensive material on nanotechnology can be found on the author's Web site at <http://www.kiirzweilai.net/meme/ftame.html?m=1> 8.

⁷¹ Richard P. Feynman, "There's Plenty of Room at the Bottom," American Physical Society annual meeting, Pasadena, 1959; transcript at <http://www.zyvex.com/nanotech/feynman.html>.

⁷² John von Neumann, *Theory of Self-Reproducing Automata*, A.W. Burks, ed., University of Illinois Press, Urbana IL, 1966.

⁷³ The most comprehensive survey of kinematic machine replication is: Robert A. Freitas Jr., Ralph C. Merkle, *Kinematic Self-Replicating Machines*, Georgetown, TX: Landes Bioscience, 2004.

<http://www.MolecularAssembler.coiWKSARM.htm>

⁷⁴ Drexler, *Engines of Creation*, and K. Eric Drexler, *Nanosystems: Molecular Machinery, Manufacturing, and Computation* (New York: John Wiley and Sons, 1992).

⁷⁵ See discussion of nanotube circuitry in Chapter 3, including analysis of the potential of nanotube circuitry in endnote 9 in chapter 3.

⁷⁶ K. Eric Drexler and Richard E. Smalley, "Nanotechnology: Drexler and Smalley Make the Case for and Against 'Molecular Assemblers,'" *Chemical and Engineering News*, November 30, 2003, <http://pubs.acs.org/cer1/coverstory/8148/8148counteoint.html>.

⁷⁷ Ralph C. Merkle, "A Proposed 'Metabolism' for a Hydrocarbon Assembler," *Nanotechnology* 8 (December 1997): 149-62, <http://www.iop.org/EJ/abstract/0957-4484/8/4/001>, or <http://www.zyvex.com/nanotech/hydroCarbonMetabolism.html>

See also: Ralph C. Merkle, "Binding sites for use in a simple assembler," *Nanotechnology* 8(1997):23-28; <http://www.zyvex.com/nanotech/bindingSites.html>

Ralph C. Merkle, "A new family of six degree of freedom positional devices," *Nanotechnology* 8(1997):47-52; <http://www.zyvex.com/nanotech/6dof.html>

Ralph C. Merkle, "Casing an assembler," *Nanotechnology* 10(1999):315-322; <http://www.zyvex.com/nanotech/casing>

Robert A. Freitas Jr., "A Simple Tool for Positional Diamond Mechanosynthesis, and its Method of Manufacture," U.S. Provisional Patent Application No. 60/543,802, filed 11 February 2004. Process described in lecture at <http://www.MolecularAssembler.com/Papers/PathDiamMolMfg.htm>

Ralph C. Merkle, Robert A. Freitas Jr., "Theoretical analysis of a carbon-carbon dimer placement tool for diamond mechanosynthesis," *J. Nanosci. Nanotechnol.* 3 (August 2003):3 19-324,

<http://www.rfreitas.com/Nano/JNNDimerTool.pdf>

Robert A. Freitas Jr., Ralph C. Merkle, *Kinematic Self-Replicating Machines*, Georgetown, TX: Landes Bioscience, 2004, Section 4.11.3 "Merkle-Freitas Hydrocarbon Molecular Assembler," pp. 130-135,

<http://www.MolecularAssembler.com/KSRM/4.11.3.htm>

⁷⁸ Robert A. Freitas Jr., *Nanomedicine, Vol. I: Basic Capabilities*, Georgetown, TX: Landes Bioscience, 1999, Section 6.3.4.5 "Chemoelectric Cells", pp. 152-154, <http://www.nanomedicine.com/NMI/6.3.4.5.htm>.

Robert A. Freitas Jr., *Nanomedicine, Vol. I: Basic Capabilities*, Georgetown, TX: Landes Bioscience, 1999, Section 6.3.4.4 "Glucose Engines", pp. 149-152, <http://www.nanomedicine.com/NMI/6.3.4.4.htm>.

K. Eric Drexler, *Nanosystems: Molecular Machinery, Manufacturing, and Computation*, New York: Wiley Interscience, 1992, Section 16.3.2 "Acoustic power and control," pp. 472-476.

See also: Robert A. Freitas Jr., Ralph C. Merkle, *Kinematic Self-Replicating Machines*, Georgetown, TX: Landes Bioscience, 2004, Appendix B.4 "Acoustic Transducer for Power and Control," pp. 225-233,

<http://www.MolecularAssembler.com/KSRM/AppB.4.htm>

⁷⁹ The most comprehensive survey of these proposals may be found in: Robert A. Freitas Jr., Ralph C. Merkle, *Kinematic Self-Replicating Machines*, Georgetown, TX: Landes Bioscience, 2004, Chapter 4 "Microscale and Molecular Kinematic Machine Replicators," pp. 89-144, <http://www.MolecularAssembler.com/KSRM/4.htm>

⁸⁰ Drexler, *Nanosystems*, p. 441.

⁸¹ The most comprehensive survey of these proposals may be found in: Robert A. Freitas Jr., Ralph C. Merkle, *Kinematic Self-Replicating Machines*, Georgetown, TX: Landes Bioscience, 2004, Chapter 4 "Microscale and Molecular Kinematic Machine Replicators," pp. 89-144, <http://www.MolecularAssembler.com/KSRM/4.htm>

⁸² T. R. Kelly, H. De Silva, and R. A. Silva, "Unidirectional Rotary Motion in a Molecular System," *Nature* 401.6749 (September 9, 1999): 150-52.

- ⁸³ Carlo Montemagno and George Bachand, "Constructing Nanomechanical Devices Powered by Biomolecular Motors," *Nanotechnology* 10 (1999): 225-31; George D. Bachand and Carlo D. Montemagno, "Constructing Organic/Inorganic NEMS Devices Powered by Biomolecular Motors," *Biomedical Microdevices* 2.3 (June 2000): 179-84.
- ⁸⁴ N. Koumura et al., "Light-Driven Monodirectional Molecular Rotor," *Nature* 401.6749 (September 9, 1999): 152-55.
- ⁸⁵ Berkeley Lab, "A Conveyor Belt for the Nano-Age," April 28, 2004, <http://www.lbl.gov/Science-Articles/Archive/MSD-conveyor-belt-for-nanoage.html>.
- ⁸⁶ "Study: Self-Replicating Nanomachines Feasible," June 2, 2004, http://www.smalltimes.com/document_display.cfm?section_id=53&document_id=8007, reporting on Tihamer Toth-Fejel, "Modeling Kinematic Cellular Automata," April 30, 2004, http://www.niac.usra.edu/files/studies/final_report/pdf/883Toth-Fejel.pdf.
- ⁸⁷ W. U. Dittmer, A. Reuter, and F. C. Simmel, "A DNA-based Machine That Can Cyclically Bind and Release Thrombin," *Angewandte Chemie International Edition* 43 (2004): 3550-53.
- ⁸⁸ Shiping Liao, Nadrian C. Seeman, "Translation of DNA Signals into Polymer Assembly Instructions," *Science* 306: 2072-2074 (December 17, 2004), <http://www.sciencemag.org/cgi/reprint/306/5704/2072.pdf>
- ⁸⁹ Scripps Research Institute, "Nano-origami," February 11, 2004, http://www.eurekalert.org/pub_releases/2004-02/sri-n021004.php.
- ⁹⁰ Jenny Hogan, "DNA Robot Takes Its First Steps," May 6, 2004, <http://www.newscientist.com/news/news.jsp?id=ns99994958>, reporting on Nadrian Seeman and William Sherman, "A Precisely Controlled DNA Biped Walking Device," *Nano Letters* 4.7 (July 2004): 1203-7.
- ⁹¹ Helen Pearson, "Construction Bugs Find Tiny Work," *Nature News*, July 11, 2003, <http://www.nature.com/news/2003/030707/rull/030707-9.html>.
- ⁹² Richard E. Smalley, "Nanofallacies: Of Chemistry, Love and Nanobots," *Scientific American* 285.3 (September 2001): 76-77, Subscription required for this link: http://www.sciamdigital.com/browse.cfm?sequencenameCHAR=item2&methodnameCHAR=resource_getitembrowse&interfacenameCHAR=browse.cfm&ISSUEID_CHAR=6A628AB3-17A5-4374-B100-3185A0CCC86&ARTICLEID_CHAR=F90C4210-C153-4B2F-83A1-28F2012B637&sc=I100322.
- ⁹³ See the bibliography of references in endnotes 108 and 109 below. >
- See also Drexler, *Nanosystems*, for his proposal. For sample confirmations*: see: Xiao Yan Chang, Martin Perry, James Peploski, Donald L. Thompson, Lionel M. Raff, "Theoretical studies of hydrogen-abstraction reactions from diamond and diamond-like surfaces," *Chem. Phys.* 99 (September 15, 1993):4748-4758.
- See also: LJ. Lauhon, W. Ho, "Inducing and observing the abstraction of a single hydrogen atom in bimolecular reaction with a scanning tunneling microscope," *Phys. Chem.* 105(2000):3987-3992, G. AIHs and K. Eric Drexler, "Design and Analysis of a Molecular Tool for Carbon Transfer in Mechanosynthesis," *Journal of Computational and Theoretical Nanoscience* 2.1 (March/April 2005, in press).
- ⁹⁴ Lea Winerman, "How to Grab an Atom," *Physical Review Focus*, May 2, 2003, http://focus.aps.org/story/v1_1/stl9, reporting on Noriaki Oyabu, "Mechanical Vertical Manipulation of Selected Single Atoms by Soft Nanoindentation Using a Near Contact Atomic Force Microscope," *Physical Review Letters* 90.17 (May 2, 2003): 176102.
- ⁹⁵ Robert A. Freitas Jr., "Technical Bibliography for Research on Positional Mechanosynthesis," Foresight Institute website, December 16, 2003, <http://foresight.org/stage2/mechsvnthbib.html>
- ⁹⁶ See equation and explanation on p. 3 of Ralph C. Merkle, "That's Impossible! How Good Scientists Reach Bad Conclusions," <http://www.zyvex.com/nanotech/impossible.html>.
- ⁹⁷ "Thus ΔX_C is just -5% of the typical atomic electron cloud diameter of ~0.3 nm, imposing only a modest additional constraint on the fabrication and stability of nanomechanical structures. (Even in most liquids at their boiling points, each molecule is free to move only ~0.07 nm from its average position.)"
- Robert A. Freitas Jr., *Nanomedicine, Vol. I: Basic Capabilities*, Georgetown, TX: Landes Bioscience, 1999, Section 2.1 "Is Molecular Manufacturing Possible?", p. 39, http://www.nanomedicine.com/NMI/2_1.htm#p9
- ⁹⁸ Robert A. Freitas Jr., *Nanomedicine, Vol. I: Basic Capabilities*, Georgetown, TX: Landes Bioscience, 1999, Section 6.3.4.5 "Chemoelectric Cells", pp. 152-154, <http://www.nanomedicine.com/NMI/6.3.4.5.htm>
- ⁹⁹ Montemagno and Bachand, "Constructing Nanomechanical Devices Powered by Biomolecular Motors."
- ¹⁰⁰ Open letter from Foresight chairman K. Eric Drexler to Nobel laureate Richard Smalley, <http://www.foresight.org/NanoRev/Letter.html>, and reprinted here: <http://www.kurzweilai.net/meme/frame.html?main=/articles/art0560.html>. The full story can be found at Ray

Kurzweil, "The Drexler-Smalley Debate on Molecular Assembly,"

<http://www.kurzweilai.net/meme/frame.html?main=/articles/art0604.html>.

¹⁰¹ K. Eric Drexler and Richard E. Smalley, "Nanotechnology: Drexler and Smalley Make the Case for and Against 'Molecular Assemblers,'" *Chemical & Engineering News* 81.48 (Dec. 1, 2003): 37-42,

<http://pubs.acs.org/cen/coverstory/8148/8148counterpoint.html>.

¹⁰² A. Zaks and A. M. Klibanov, "Enzymatic Catalysis in Organic Media at 100 Degrees C," *Science* 224.4654 (June 15, 1984): 1249-51.

¹⁰³ Patrick Bailey, "Unraveling the Big Debate about Small Machines," *BetterHumans*, August 16, 2004.

<http://www.betterhumans.com/Features/Reports/report.aspx?articleID=2004-08-16-1>

¹⁰⁴ Charles B. Musgrave et al., "Theoretical Studies of a Hydrogen Abstraction Tool for Nanotechnology," *Nanotechnology* 2:187-195, (Oct 1991).

Michael Page and Donald W. Brenner, "Hydrogen Abstraction from a Diamond Surface: *Ab initio* Quantum Chemical Study with Constrained Isobutane as a Model," *J. Am. Chem. Soc.* 113(9):3270-3274 (1991).

Xiao Yan Chang, Martin Perry, James Peploski, Donald L. Thompson, Lionel M. Raff, "Theoretical studies of hydrogen-abstraction reactions from diamond and diamond-like surfaces," *J. Chem. Phys.* 99 (15 September 1993):4748-4758.

J. W. Lyding, K. Hess, G. C. Abeln, et al., "UHV-STM Nanofabrication and Hydrogen/Deuterium Desorption from Silicon Surfaces: Implications for CMOS Technology," *Appl. Surf. Sci.*, 132, 221 (1998), <http://www.hetsam-group.northwestern.edu/publications.html>

E.T. Foley et al., "Cryogenic UHV-STM Study of Hydrogen and Deuterium Desorption from Silicon(100)," *Rev. Lett.* 80:1336-1339 (1998), http://prola.aps.org/abstract/PRL/v80/i6/pl 336_1

L.J. Lauhon, W. Ho, "Inducing and observing the abstraction of a single hydrogen atom in bimolecular reaction with a scanning tunneling microscope," *J. Phys. Chem.* 105 (2000):3987-3992.

¹⁰⁵ Stephen P. Walch, Ralph C. Merkle, "Theoretical Studies of Diamond Mechanosynthesis Reactions," *Nanotechnology* 9:285-296 (Sep 1998).

Fedor N. Dzegilenko, Deepak Srivastava, Subhash Saini, "Simulations of Carbon Nanotube Tip Assisted Mechano-Chemical Reactions on a Diamond Surface," *Nanotechnology* 9:325-330 (December 1998).

Ralph C. Merkle and Robert A. Freitas Jr., "Theoretical Analysis of a Carbon-Carbon Dimer Placement Tool for Diamond Mechanosynthesis," *J. Nanosci. Nanotechnol.* 3(4):319-24 2003 (Aug 2003),

<http://www.rfreitas.coni/Nano/DimerTool.htm>

Jingping Peng, Robert A. Freitas Jr., Ralph C. Merkle, "Theoretical analysis of diamond mechanosynthesis. Part I. Stability of C₂ mediated growth of nanocrystalline diamond C(1 10) surface," *Comput. Theor. Nanosci* 1 (March 2004):62-70; <http://www.molecularassembler.com/JCTNPengMar04.pdf>

David J. Mann, Jingping Peng, Robert A. Freitas Jr., Ralph C. Merkle, "Theoretical analysis of diamond mechanosynthesis. Part II. C₂ mediated growth of diamond C(1 10) surface via Si/Ge-triadamantane dimer placement tools," *Comput. Theor. Nanosci.* 1 (March 2004):71-80;

<http://www.molecularassembler.com/JCTNMannMar04.pdf>

¹⁰⁶ The analysis of the hydrogen abstraction tool and carbon deposition tools has involved many people, including: Donald W. Brenner, Tahir Cagin, Richard J. Colton, K. Eric Drexler, Fedor N. Dzegilenko, Robert A. Freitas Jr., William A. Goddard III, J.A. Harrison, Charles B. Musgrave, Ralph C. Merkle, Michael Page, Jason K. Perry, Subhash Saini, O.A. Shenderova, Susan B. Sinnott, Deepak Srivastava, Stephen P. Walch, and Carter T. White.

¹⁰⁷ Ralph C. Merkle, "A Proposed 'Metabolism' for a Hydrocarbon Assembler," *Nanotechnology* 8:149-162 (Dec 1997), <http://www.iop.Org/EJ/abstract/0957-4484/8/4/001> or

<http://www.zyvex.coni/nanotech/hydroCarbonMetabolism.html>

¹⁰⁸ A useful bibliography of references is: Robert A. Freitas Jr., "Technical Bibliography for Research on Positional Mechanosynthesis," Foresight Institute website, 16 December 2003, <http://foresight.org/stage2/mechsvnthbib.html>

Other references include:

Wilson Ho and Hyojune Lee, "Single Bond Formation and Characterization with a Scanning Tunneling Microscope," *Science* 286.5445 (November 26, 1999): 1719-22, <http://www.physics.uci.edu/~wilsonho/stm-lets.html>; Drexler, *Nanosystems*, chap. 8; Merkle, "Proposed 'Metabolism' for a Hydrocarbon Assembler";

Musgrave et al., "Theoretical Studies of a Hydrogen Abstraction Tool for Nanotechnology"; Michael Page and Donald W. Brenner, "Hydrogen Abstraction from a Diamond Surface: *Ab initio* Quantum Chemical Study with

Constrained Isobutane as a Model," *Journal of the American Chemical Society* 113.9 (1991): 3270-74; D. W. Brenner et al., "Simulated Engineering of Nanostructures," *Nanotechnology* 7 (September 1996): 161-67,

<http://www.zyvex.com/nanotech/nano4/brennerPaper.pdf>; S. P. Walch, W. A. Goddard III, R. M. Merkle,

"Theoretical Studies of Reactions on Diamond Surfaces," fifth Foresight Conference on Molecular Nanotechnology,

- 1997, <http://www.foresight.org/Conferences/MNT05/AbstractsAVaIcabst.html>; Stephen P. Walch and Ralph C. Merkle, "Theoretical Studies of Diamond Mechanosynthesis Reactions," *Nanotechnology* 9 (September 1998): 285-96; Fedor N. Dzegilenko, Deepak Srivastava, and Subhash Saini, "Simulations of Carbon Nanotube Tip Assisted Mechano-Chemical Reactions on a Diamond Surface," *Nanotechnology* 9 (December 1998): 325-30; J. W. Lyding et al., "UHV-STM Nanofabrication and Hydrogen/Deuterium Desorption from Silicon Surfaces: Implications for CMOS Technology," *Applied Surface Science* 132 (1998): 221, <http://www.hersam-group.northwestern.edu/publications.html>; E. T. Foley et al., "Cryogenic UHV-STM Study of Hydrogen and Deuterium Desorption from Silicon(100)," *Physical Review Letters* 80 (1998): 1336-39, <http://prola.aps.org/abstract/PRL/v80/i6/pl336J>; M. C. Hersam, G. C. Abeln, and J. W. Lyding, "An Approach for Efficiently Locating and Electrically Contacting Nanostructures Fabricated via UHV-STM Lithography on Si(100)," *Microelectronic Engineering* 47 (1999): 235-37; L. J. Lauhon and W. Ho, "Inducing and Observing the Abstraction of a Single Hydrogen Atom in Bimolecular Reaction with a Scanning Tunneling Microscope," *Journal of Physical Chemistry* 105 (2000): 3987-92, <http://www.physics.uci.edu/~wilsonho/stm-iets.html>.
- ¹⁰⁹ Eric Drexler, "Drexler Counters," first published on KurzweilAI.net on November 1, 2003: <http://www.kurzweilai.net/meme/frame.htmr?main=/articles/artO6O6.html>. See also K. Eric Drexler, *Nanosystems: Molecular Machinery, Manufacturing, and Computation*, New York: Wiley Interscience, 1992, Chapter 8.
- Ralph C. Merkle, "Foresight Debate with Scientific American," 1995, <http://www.foresight.org/SciAmDebate/SciAmResponse.html>
- Wilson Ho, Hyojune Lee, "Single Bond Formation and Characterization with a Scanning Tunneling Microscope," *Science* 286(5445):1719-1722 (November 26, 1999), <http://www.physics.uci.edu/~wilsonho/stm-iets.html>
- K. Eric Drexler, David Forrest, Robert A. Freitas Jr. J. Storrs Hall, Neil Jacobstein, Tom McKendree, Ralph Merkle, Christine Peterson, "On Physics, Fundamentals, and Nanorobots: A Rebuttal to Smalley's Assertion that Self-Replicating Mechanical Nanorobots Are Simply Not Possible," *A Debate About Assemblers* (2001), <http://www.imm.org/SciAmDebate2/smalley.html>
- ¹ See <http://pubs.acs.org/cen/coverstory/8148/8148counterpoint.html>;
<http://www.kurzweilai.net/meme/frame.html?main=/articles/artO6O4.html?>
- ¹¹¹ D. Maysinger et al., "Block Copolymers Modify the Internalization of Micelle-incorporated Probes into Neural Cells," *Biochimica et Biophysica Acta* 1539.3 (June 20, 2001): 205-17; R. Savic et al., "Micellar Nanocontainers Distribute to Defined Cytoplasmic Organelles," *Science* 300.5619 (April 25, 2003): 615-18.
- ¹¹² T. Yamada et al., "Nanoparticles for the delivery of genes and drugs to human hepatocytes," *Nature Biotechnology* 21.8 (August 2003):885-90. Published electronically June 29, 2003. Abstract: <http://www.nature.com/cgi-taf/DynaPage.taf?file=/nbt/journal/v21/n8/abs/nbt843.html>. Short press release from *Nature*: http://www.nature.com/nbt/press_release/nbt0803.html.
- ¹¹³ Richards Grayson et al., "A BioMEMS Review: MEMS Technology for Physiologically Integrated Devices," *IEEE Proceedings* 92 (2004): 6-21; Richards Grayson et al., "Molecular Release from a Polymeric Microreservoir Device: Influence of Chemistry, Polymer Swelling, and Loading on Device Performance," *Journal of Biomedical Materials Research* 69A.3 (June 1, 2004): 502-12.
- ¹¹⁴ D. Patrick O'Neal et al., "Photo-thermal Tumor Ablation in Mice Using Near Infrared-Absorbing Nanoparticles," *Cancer Letters* 209.2 (June 25, 2004): 171-76.
- ¹¹⁵ International Energy Agency, from a R. E. Smalley presentation, "Nanotechnology, the S&T Workforce, Energy & Prosperity," p. 12, presented at PCAST (President's Council of Advisors on Science and Technology), Washington, D.C., March 3, 2003, <http://www.ostp.gov/PCAST/PCAST%203-3-03%20R%20Smalley%20Slides.pdf>, also at <http://cohesion.rice.edu/NaturalSciences/Smalley/emplibrary/PCAST%20March%203,%202003.ppt>.
- ¹¹⁶ Smalley, "Nanotechnology, the S&T Workforce, Energy & Prosperity."
- ¹¹⁷ "FutureGen—A Sequestration and Hydrogen Research Initiative," U.S. Department of Energy, Office of Fossil Energy, February 2003, http://www.fossil.energy.gov/programs/powersystems/futuregen/futuregen_factsheet.pdf.
- ¹¹⁸ Drexler, *Nanosystems*, pp. 428, 433.
- ¹¹⁹ Barnaby J. Feder, "Scientist at Work/Richard Smalley: Small Thoughts for a Global Grid," *The New York Times*, September 2, 2003, The following link requires subscription or purchase: <http://query.nytimes.com/gst/abstract.html?res=F30C17FC3D5C0C718CDDA00894DB404482>.
- ¹²⁰ International Energy Agency, from Smalley, "Nanotechnology, the S&T Workforce, Energy & Prosperity," p. 12.
- ¹²¹ American Council for the United Nations University, Millennium Project Global Challenge 13: <http://www.acunu.org/millennium/eh-13.html>

¹²² "Wireless Transmission in Earth's Energy Future," Environment News Service, November 19, 2002, reporting on Jerome C. Glenn and Theodore J. Gordon in "2002 State of the Future," American Council for the United Nations University (August 2002).

¹²³ Disclosure: the author is an adviser to and investor in this company.

¹²⁴ "NEC Unveils Methanol-Fueled Laptop," Associated Press, June 30, 2003, <http://www.siliconvalley.com/mld/siliconvalley/news/6203790.htm>, reporting on NEC press release, "NEC Unveils Notebook PC with Built-in Fuel Cell," June 30, 2003, <http://www.nec.co.jp/press/en/0306/3002.html>.

¹²⁵ Tony Smith, "Toshiba Boffins Prep Laptop Fuel Cell," *The Register*, March 5, 2003, http://www.theregister.co.uk/2003/03/05/toshiba_boffinsjprep_laptop_fuel; Yoshiko Hara, "Toshiba Develops Matchbox-Sized Fuel Cell for Mobile Phones," *EE Times*, June 24, 2004, <http://www.eet.com/article/showArticle.jhtml?articleId=22101804>, reporting on Toshiba press release, "Toshiba Announces World's Smallest Direct Methanol Fuel Cell with Energy Output of 100 Milliwatts," http://www.toshiba.com/taec/press/dmfc_04_222.shtml.

¹²⁶ Karen Lurie, "Hydrogen Cars," *ScienceCentral News*, May 13, 2004, http://www.sciencentral.com/articles/view.php3?language=english&type=article&article_id=218392247.

¹²⁷ Louise Knapp, "Booze to Fuel Gadget Batteries," *WiredNews*, April 2, 2003, http://www.wired.com/news/gizmos/0,1452,581_19,00.html and St. Louis University press release, "Powered by Your Liquor Cabinet, New Biofuel Cell Could Replace Rechargeable Batteries," March 24, 2003, <http://www.slu.edu/readstory/newsinfo/2474>, reporting on Nick Akers and Shelley Minter, "Towards the Development of a Membrane Electrode Assembly," presented at the American Chemical Society national meeting, Anaheim, California (2003).

¹²⁸ "Biofuel Cell Runs on Metabolic Energy to Power Medical Implants," *Nature Online*, November 12, 2002, <http://www.nature.com/news/2002/021111/full/021111-1.html>, reporting on N. Mano, F. Mao, and A. Heller, "A Miniature Biofuel Cell Operating in a Physiological Buffer," *Journal of the American Chemical Society* 124 (2002): 12962-63.

¹²⁹ "Power from Blood Could Lead to 'Human Batteries,'" FairfaxDigital, August 4, 2003, <http://www.smh.com.au/articles/2003/08/03/1059849278131.html?oneclick=true>

¹³⁰ Mike Martin, "Pace-Setting Nanotubes May Power Micro-Devices," NewsFactor, February 27, 2003, <http://physics.iisc.emet.in/~asood/Pace-Setting%20Nanotubes%20May%20Power%20Micro-Devices.htm>.

¹³¹ "Finally, it is possible to derive a limit to the total planetary active nanorobot mass by considering the global energy balance. Total solar insolation received at the Earth's surface is $\sim 1.75 \times 10^{17}$ watts ($I_{\text{Earth}} \sim 1370 \text{ W/m}^2 \pm 0.4\%$ at normal incidence)." Robert A. Freitas Jr., *Nanomedicine, Vol. I: Basic Capabilities*, Georgetown, TX: Landes Bioscience, 1999, Section 6.5.7 "Global Hypsithermal Limit", pp. 175-176, <http://www.nanomedicine.com/NMI/6.5.7.htm#fp4>

*** This assumes 10 billion (10^{10}) persons, a power density for nanorobots of around 10^7 watts per cubic meter, a nanorobot size of one cubic micron, and a power draw of about 10 picowatts (10^{-11} watts) per nanorobot. The hypsithermal limit of 10^{16} watts implies about 10 kilograms of nanorobots per person, or 10^{16} nanorobots per person. Robert A. Freitas Jr., *Nanomedicine, Vol. I: Basic Capabilities*, Georgetown, TX: Landes Bioscience, 1999, Section 6.5.7 "Global Hypsithermal Limit", pp. 175-176, <http://www.nanomedicine.com/NMI/6.5.7.htm#fp4>

¹³³ Alternatively, nanotechnology can be designed to be extremely energy efficient in the first place so that energy recapture would be unnecessary, and infeasible because there would be relatively little heat dissipation to recapture. In a private communication (January, 2005), Robert A. Freitas Jr. writes "Drexler (*Nanosystems 396*) claims that energy dissipation may in theory be as low as $E_{\text{diss}} \sim 0.1 \text{ MJ/kg}$ 'if one assumes the development of a set of mechanochemical processes capable of transforming feedstock molecules into complex product structures using only reliable, nearly reversible steps.' 0.1 MJ/kg of diamond corresponds roughly to the minimum thermal noise at room temperature (e.g., $kT \sim 4 \text{ zJ/atom}$ at 298 K)."

¹³⁴ Alexis De Vos, *Endoreversible Thermodynamics of Solar Energy Conversion* (London: Oxford University Press, 1992), p. 103.

¹³⁵ R. D. Schaller and V. I. Klimov, "High Efficiency Carrier Multiplication in PbSe Nanocrystals: Implications for Solar Energy Conversion," *Physical Review Letters* 92.18 (May 7, 2004): 186601.

¹³⁶ National Academies Press, Commission on Physical Sciences, Mathematics, and Applications, *Harnessing Light: Optical Science and Engineering for the 21st Century*, (Washington, D.C.:National Academy Press, 1998): 166, <http://books.nap.edu/books/0309059917/html/166.html>

¹³⁷ Matt Marshall, "World Events Spark Interest in Solar Cell Energy Start-ups," Mercury News, August 15, 2004, http://www.konarkatech.com/news_articles_082004/b-silicon_valley.php and http://www.nanosolar.com/cache/merc08_1504.htm.

- ¹³⁸ John Gartner, "NASA Spaces on Energy Solution," *WiredNews*, June 22, 2004, <http://www.wired.com/news/technology/0,1282,63913,00.html>. See also Arthur Smith, "The Case for Solar Power from Space," <http://www.lispace.org/articles/SSPCase.html>.
- ¹³⁹ "The Space Elevator Primer," Spaceward Foundation, <http://www.elevator2010.org/site/primer.html>.
- ¹⁴⁰ Kenneth Chang, "Experts Say New Desktop Fusion Claims Seem More Credible," *The New York Times*, March 3, 2004, <http://www.rpi.edu/web/News/nytlahey3.html>, reporting on R. P. Taleyarkhan, "Additional Evidence of Nuclear Emissions During Acoustic Cavitation," *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics* 69.3, pt. 2 (March 2004): 036109.
- ¹⁴¹ The original Pons and Fleischman method of desktop cold fusion using palladium electrodes is not dead. Ardent advocates have continued to pursue the technology, and the Department of Energy announced in 2004 that it was conducting a new formal review of the recent research in this field. Toni Feder, "DOE Warms to Cold Fusion," *Physics Today* (April 2004), <http://www.physicstoday.org/vol-57/iss-4/p27.html>.
- ¹⁴² Akira Fujishima, Tata N. Rao, and Donald A. Tryk, "Titanium Dioxide Photocatalysis," *Journal of Photochemistry and Photobiology C: Photochemistry Review* 1 (June 29, 2000): 1-21; Prashant V. Kamat, Rebecca Huehn, and Roxana Nicolaescu, "A 'Sense and Shoot' Approach for Photocatalytic Degradation of Organic Contaminants in Water," *Journal of Physical Chemistry B* 106 (January 31, 2002): 788-94.
- ¹⁴³ A. G. Panov et al., "Photooxidation of Toluene and p-Xylene in Cation-Exchanged Zeolites X, Y, ZSM-5, and Beta: The Role of Zeolite Physicochemical Properties in Product Yield and Selectivity," *Journal of Physical Chemistry B* 104 (June 22, 2000): 5706-14.
- ¹⁴⁴ Gabor A. Somorjai and Keith McCrea, "Roadmap for Catalysis Science in the 21st Century: A Personal View of Building the Future on Past and Present Accomplishments," *Applied Catalysis A: General* 222.1-2 (2001): 3-18, Lawrence Berkeley National Laboratory number 3.LBNL-48555, <http://www.cchem.berkeley.edu/~gasgrp/2000.html> (publication 877). See also Zhao, Lu, and Millar, "Advances in mesoporous molecular sieve MCM-41," *Industrial & Engineering Chemistry Research*, 35 (1996): 2075-2090, http://cheed.nus.edu.sg/~chezxs/Zhao/publication/1996_2075.pdf.
- ¹⁴⁵ NTSC/NSET report, *National Nanotechnology Initiative: The Initiative and Its Implementation Plan*, July 2000, <http://www.nano.gov/html/res/nni2.pdf>.
- ¹⁴⁶ Wei-xian Zhang, Chuan-Bao Wang, and Hsing-Lung Lien, "Treatment of Chlorinated Organic Contaminants with Nanoscale Bimetallic Particles," *Catalysis Today* 40 (May 14, 1988): 387-95.
- ¹⁴⁷ R. Q. Long and R. T. Yang, "Carbon Nanotubes as Superior Sorbent for Dioxin Removal," *Journal of the American Chemical Society* 123.9 (2001): 2058-59.
- ¹⁴⁸ Robert A. Freitas, Jr. "Death Is an Outrage!" presented at the Fifth Alcor Conference on Extreme Life Extension, Newport Beach, CA, November 16, 2002, <http://www.rfreitas.com/Nano/DeathIsAnOutraee.html>
- ¹⁴⁹ For example, the fifth annual BIOMEMS conference, June 2003, San Jose, <http://www.knowledgepress.com/events/11201717.htm>.
- ¹⁵⁰ First 2 volumes of a planned 4-volume series: Robert A. Freitas Jr., *Nanomedicine, Vol. I: Basic Capabilities*, Georgetown, TX: Landes Bioscience, 1999; *Nanomedicine, Vol. HA: Biocompatibility*, Georgetown, TX: Landes Bioscience, 2003; <http://www.nanomedicine.com>
- ¹⁵¹ Robert A. Freitas Jr., "Exploratory Design in Medical Nanotechnology: A Mechanical Artificial Red Cell," *Artificial Cells, Blood Substitutes, and Immobilization Biotechnology* 26 (1998): 411-30, <http://www.foresight.org/Nanomedicine/Respirocytes.html>.
- ¹⁵² Robert A. Freitas Jr., "Microbivores: Artificial Mechanical Phagocytes using Digest and Discharge Protocol," Zyvex preprint, March 2001, <http://wmv.rfreitas.com/Nano/Microbivores.htm>; Robert A. Freitas Jr., "Microbivores: Artificial Mechanical Phagocytes," *Foresight Update* No. 44, March 31, 2001, pp. 11-13, <http://www.inmi.org/Reports/Rep025.html>; see also microbivore images at the Nanomedicine Art Gallery, <http://www.foresight.Org/Nanomedicine/Gallery/Species/MJcrobivores.html>
- ¹⁵³ Robert A. Freitas Jr., *Nanomedicine, Vol. I: Basic Capabilities*, Georgetown, TX: Landes Bioscience, 1999, Section 9.4.2.5 "Nanomechanisms for Natation", pp. 309-3 12, <http://www.nanomedicine.com/NMI/9.4.2.5.htm>
- ¹⁵⁴ George Whitesides, "Nanoinspiration: The Once and Future Nanomachine," *Scientific American* 285.3 (September 16, 2001): 78-83.
- ¹⁵⁵ "According to Einstein's approximation for Brownian motion, after 1 second has elapsed at room temperature a fluidic water molecule has, on average, diffused a distance of ~50 microns (~400,000 molecular diameters) whereas a 1-micron nanorobot immersed in that same fluid has displaced by only ~0.7 microns (only ~0.7 device diameter) during the same time period. Thus Brownian motion is at most a minor source of navigational error for motile medical nanorobots." See K. Eric Drexler et al., "Many Future Nanomachines: A Rebuttal to Whiteside' Assertion

That Mechanical Molecular Assemblers Are Not Workable and Not a Concern," a Debate about Assemblers, Institute for Molecular Manufacturing, 2001, <http://www.imm.org/SciAmDebate2/whitesides.html>.

¹⁵⁶ Tejal A. Desai, "MEMS-Based Technologies for Cellular Encapsulation," *American Journal of Drug Delivery* 1.1 (2003): 3-11, abstract available at

<http://www.ingentaconnect.com/search/expand?pub=infobike://adis/add/2003/00000001/00000001/art00001>

¹⁵⁷ As quoted by Douglas Hofstadter in *Gödel, Escher, Bach: An Eternal Golden Braid* (New York: Basic Books, 1979).

¹⁵⁸ The author runs a company, FATKAT (Financial Accelerating Transactions by Kurzweil Adaptive Technologies), which applies computerized pattern recognition to financial data to make stock-market investment decisions, <http://www.FatKat.com>.

¹⁵⁹ See discussion in chapter 2 on price-performance improvements in computer memory and electronics in general.

¹⁶⁰ Runaway AI refers to a scenario where, as Max More describes "superintelligent machines, initially harnessed for human benefit, soon leave us behind." Max More, "Embrace, Don't Relinquish, the Future,"

http://www.kurzweilai.net/articles/artOIO_6.htm^printable^1

See also Damien Broderick's description of the "Seed AI" as "A self-improving seed AI could run glacially slowly on a limited machine substrate. The point is, so long as it has the capacity to improve itself, at some point it will do so convulsively, bursting through any architectural bottlenecks to design its own improved hardware, maybe even build it (if it's allowed control of tools in a fabrication plant)." Damien Broderick, "Tearing toward the Spike," presented at "Australia at the Crossroads? Scenarios and Strategies for the Future," (April 31 - May 2, 2000), published on KurzweilAI.net May 7, 2001:

http://www.kurzweilai.net/meme/frame.html?main=/articles/art_0173.html

¹⁶¹ David Talbot, "Lord of the Robots," *Technology Review* (April 2002).

¹⁶² Heather Havenstein writes that the "inflated notions spawned by science fiction writers about the convergence of humans and machines tarnished the image of AI in the 1980s because AI was perceived as failing to live up to its potential." Heather Havenstein, "Spring comes to AI winter: A thousand applications bloom in medicine, customer service, education and manufacturing," *Computerworld*, February 14, 2005,

http://www.computenvorld.com/softvwaretopics/software/storv/0_10801,99691,00.html

This tarnished image led to "AI Winter," defined as "a term coined by Richard Gabriel for the (circa 1990-94?) crash of the wave of enthusiasm for the AI language Lisp and AI itself, following a boom in the 1980s." Duane Rettig wrote "...companies rode the great AI wave in the early 80's, when large corporations poured billions of dollars into the AI hype that promised thinking machines in J O years. When the promises turned out to be harder than originally thought, the AI wave crashed, and Lisp crashed with it because of its association with AI. We refer to it as the AI Winter." Duane Rettig quoted in "AI Winter," <http://c2.com/cgi/wiki?AiWinter>

¹⁶³ The General Problem Solver (GPS) computer program, written in 1957, was able to solve problems through rules that allowed the GPS to divide a problem's goals into subgoals, and then check if obtaining a particular subgoal would bring the GPS closer to solving the overall goal. In the early 1960s Thomas Evan wrote ANALOGY, a "program [that] solves geometric-analogy problems of the form A:B::C:?" taken from IQ tests and college entrance exams." Boicho Kokinov and Robert M. French, *Computational Models of Analogy-Making*, in Nadel, L. (Ed.) *Encyclopedia of Cognitive Science, Vol. I*, (London: Nature Publishing Group, 2003) pp.113 - 118. See also A. Newell, J.C. Shaw, and H.A. Simon, "Report on a general problem-solving program," *Proceedings of the International Conference on Information Processing*, (Paris: UNESCO House, 1959) pp. 256-264; and Thomas Evans, "A Heuristic Program to Solve Geometric-Analogy Problems," in *Semantic Information Processing*, M. Minsky, Editor, (Cambridge, MA: MIT Press, 1968).

¹⁶⁴ Sir Arthur Conan Doyle, "The Red-Headed League," 1890, available at <http://www.eastoftheweb.com/short-stories/UBooks/RedHead.shtml>.

¹⁶⁵ V. Yu et al., "Antimicrobial Selection by a Computer: A Blinded Evaluation by Infectious Diseases Experts," *JAMA* 242.12(1979): 1279-82.

¹⁶⁶ Gary H. Anthes, "Computerizing Common Sense," *Computerworld*, April 8, 2002, http://www.computerworld.com/news/2002/story/0_11280,69881,00.html.

¹⁶⁷ Kristen Philipkoski, "Now Here's a Really Big Idea," *WiredNews*, November 25, 2002, <http://www.wired.com/news/technology/0,1282,56374,00.html>, reporting on Darryl Macer, "The Next Challenge Is to Map the Human Mind," *Nature* 420 (November 14, 2002): 121; see also a description of the project at <http://www.biol.tsukuba.ac.jp/~macer/index.html>.

¹⁶⁸ Thomas Bayes, "An Essay Towards Solving a Problem in the Doctrine of Chances," published in 1763, two years after his death in 1761.

¹⁶⁹ SpamBayes spam filter, <http://spambayes.sourceforge.net>.

¹⁷⁰ Lawrence R. Rabiner, "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition," *Proceedings of the IEEE 11* (1989): 257-86. For a mathematical treatment of Markov models, see <http://jedlik.phy.bme.hu/~gerjanos/HMM/node2.html>.

¹⁷¹ Kurzweil Applied Intelligence (KAI), founded by the author in 1982, was sold in 1997 for \$100 million and is now part of ScanSoft (formerly called Kurzweil Computer Products, the author's first company, which was sold to Xerox in 1980), now a public company. KAI introduced the first commercially marketed large-vocabulary speech-recognition system in 1987 (Kurzweil Voice Report, with a ten-thousand-word vocabulary).

¹⁷² Here is the basic schema for a neural net algorithm. Many variations are possible, and the designer of the system needs to provide certain critical parameters and methods, detailed below.

Creating a neural-net solution to a problem involves the following steps:

<nbl>

- Define the input.
- Define the topology of the neural net (i.e., the layers of neurons and the connections between the neurons).
- Train the neural net on examples of the problem.
- Run the trained neural net to solve new examples of the problem.
- Take your neural-net company public.

<ntx>These steps (except for the last one) are detailed below:

<b2>The Problem Input

<ntx>The problem input to the neural net consists of a series of numbers. This input can be:

<nbl>

- In a visual pattern-recognition system, a two-dimensional array of numbers representing the pixels of an image; or
- In an auditory (e.g., speech) recognition system, a two-dimensional array of numbers representing a sound, in which the first dimension represents parameters of the sound (e.g., frequency components) and the second dimension represents different points in time; or
- In an arbitrary pattern-recognition system, an n-dimensional array of numbers representing the input pattern.

<b2>Defining the Topology

<ntx>To set up the neural net, the architecture of each neuron consists of:

<nbl>

- Multiple inputs in which each input is "connected" to either the output of another neuron, or one of the input numbers.
- Generally, a single output, which is connected either to the input of another neuron (which is usually in a higher layer), or to the final output.

<b2>Set up the First Layer of Neurons

<nbl>

- Create N_0 neurons in the first layer. For each of these neurons, "connect" each of the multiple inputs of the neuron to "points" (i.e., numbers) in the problem input. These connections can be determined randomly or using an evolutionary algorithm (see below).
- Assign an initial "synaptic strength" to each connection created. These weights can start out all the same, can be assigned randomly, or can be determined in another way (see below).

<b2>Set up the Additional Layers of Neurons

<ntx>Setup a total of M layers of neurons. For each layer, set up the neurons in that layer.

For layer,:

<nbl>

- Create N_1 neurons in layer,. For each of these neurons, "connect" each of the multiple inputs of the neuron to the outputs of the neurons in layer,-1 (see variations below).
- Assign an initial "synaptic strength" to each connection created. These weights can start out all the same, can be assigned randomly, or can be determined in another way (see below).

-
- The outputs of the neurons in layer_M are the outputs of the neural net (see variations below).

<b2>The Recognition Trials

<b3>How Each Neuron Works

<ntx>Once the neuron is set up, it does the following for each recognition trial.

<nbl>

- Each weighted input to the neuron is computed by multiplying the output of the other neuron (or initial input) that the input to this neuron is connected to by the synaptic strength of that connection.
- All of these weighted inputs to the neuron are summed.
- If this sum is greater than the firing threshold of this neuron, then this neuron is considered to fire and its output is 1. Otherwise, its output is 0 (see variations below).

<b3>Do the Following for Each Recognition Trial

<ntx>For each layer, from layer₀ to layer_M:

For each neuron in the layer:

<nbl>

- Sum its weighted inputs (each weighted input = the output of the other neuron [or initial input] that the input to this neuron is connected to multiplied by the synaptic strength of that connection).
- If this sum of weighted inputs is greater than the firing threshold for this neuron, set the output of this neuron = 1, otherwise set it to 0.

<b3>To Train the Neural Net

<nbl>

- Run repeated recognition trials on sample problems.
- After each trial, adjust the synaptic strengths of all the interneuronal connections to improve the performance of the neural net on this trial (see the discussion below on how to do this).
- Continue this training until the accuracy rate of the neural net is no longer improving (i.e., reaches an asymptote).

<b2>Key Design Decisions

<ntx>In the simple schema above, the designer of this neural-net algorithm needs to determine at the outset:

<nbl>

- What the input numbers represent.
- The number of layers of neurons.
- The number of neurons in each layer. (Each layer does not necessarily need to have the same number of neurons.)
- The number of inputs to each neuron in each layer. The number of inputs (i.e., interneuronal connections) can also vary from neuron to neuron and from layer to layer.
- The actual "wiring" (i.e., the connections). For each neuron in each layer, this consists of a list of other neurons, the outputs of which constitute the inputs to this neuron. This represents a key design area. There are a number of possible ways to do this:

<nbl>

- Wire the neural net randomly; or
- Use an evolutionary algorithm (see below) to determine an optimal wiring; or
- Use the system designer's best judgment in determining the wiring.

<nbl>

- The initial synaptic strengths (i.e., weights) of each connection. There are a number of possible ways to do this:

<nbl>

- Set the synaptic strengths to the same value; or
- Set the synaptic strengths to different random values; or
- Use an evolutionary algorithm to determine an optimal set of initial values; or
- Use the system designer's best judgment in determining the initial values.

-
- <nbl>
- The firing threshold of each neuron.
 - Determine the output. The output can be:
 - <nnl>
 - (i) the outputs of layer_M of neurons; or
 - (ii) the output of a single output neuron, the inputs of which are the outputs of the neurons in layer_M.
 - (iii) a function of (e.g., a sum of) the outputs of the neurons in layer_M; or
 - (iv) another function of neuron outputs in multiple layers.
- <nbl>
- Determine how the synaptic strengths of all the connections are adjusted during the training of this neural net. This is a key design decision and is the subject of a great deal of research and discussion. There are a number of possible ways to do this:
- <nnl>
- (i) For each recognition trial, increment or decrement each synaptic strength by a (generally small) fixed amount so that the neural net's output more closely matches the correct answer. One way to do this is to try both incrementing and decrementing and see which has the more desirable effect. This can be time-consuming, so other methods exist for making local decisions on whether to increment or decrement each synaptic strength.
 - (ii) Other statistical methods exist for modifying the synaptic strengths after each recognition trial so that the performance of the neural net on that trial more closely matches the correct answer.
- Note that neural-net training will work even if the answers to the training trials are not all correct. This allows using real-world training data that may have an inherent error rate. One key to the success of a neural net-based recognition system is the amount of data used for training. Usually a very substantial amount is needed to obtain satisfactory results. Just like human students, the amount of time that a neural net spends learning its lessons is a key factor in its performance.

<b2>Variations

<ntx>Many variations of the above are feasible:

<nbl>

- There are different ways of determining the topology. In particular, the interneuronal wiring can be set either randomly or using an evolutionary algorithm.
- There are different ways of setting the initial synaptic strengths.
- The inputs to the neurons in layer_j do not necessarily need to come from the outputs of the neurons in layer_n. Alternatively, the inputs to the neurons in each layer can come from any lower layer or any layer.
- There are different ways to determine the final output.
- The method described above results in an "all or nothing" (1 or 0) firing called a nonlinearity. There are other nonlinear functions that can be used. Commonly a function is used that goes from 0 to 1 in a rapid but more gradual fashion. Also, the outputs can be numbers other than 0 and 1.
- The different methods for adjusting the synaptic strengths during training represent key design decisions.

The above schema describes a "synchronous" neural net, in which each recognition trial proceeds by computing the outputs of each layer, starting with layer₀ through layer_M. In a true parallel system, in which each neuron is operating independently of the others, the neurons can operate "asynchronously" (i.e., independently). In an asynchronous approach, each neuron is constantly scanning its inputs and fires whenever the sum of its weighted inputs exceeds its threshold (or whatever its output function specifies).

¹⁷³ See Chapter 4 for a detailed discussion of brain reverse-engineering. As one example of the progression, S. J. Thorpe writes, "We have really only just begun what will certainly be a long term project aimed at reverse engineering the primate visual system. For the moment, we have only explored some very simple architectures, involving essentially just feed-forward architectures involving a relatively small numbers of layers... In the years to come, we will strive to incorporate as many of the computational tricks used by the primate and human visual system as possible. More to the point, it seems that by adopting the spiking neuron approach, it will soon be possible

to develop sophisticated systems capable of simulating very large neuronal networks in real time." SJ. Thorpe et al., "Reverse engineering of the visual system using networks of spiking neurons," *Proceedings of the IEEE 2000 International Symposium on Circuits and Systems*, IEEE press. IV: 405-408, <http://www.sccn.ucsd.edu/~arno/mypapers/thorpe.pdf>

¹⁷⁴ T. Schoenauer et. al. write, "Over the past years a huge diversity of hardware for artificial neural networks (ANN) has been designed... Today one can choose from a wide range of neural network hardware. Designs differ in terms of architectural approaches, such as neurochips, accelerator boards and multi-board neurocomputers, as well as concerning the purpose of the system, such as the ANN algorithm(s) and the system's versatility... Digital neurohardware can be classified by the: [sic] system architecture, degree of parallelism, typical neural network partition per processor, inter-processor communication network and numerical representation." T. Schoenauer, A. Jahnke, U. Roth and H. Klar, "Digital Neurohardware: Principles and Perspectives," in *Proc. Neuronale Netze in der Anwendung - Neural Networks in Applications NN'98*, Magdeburg, invited paper (February 1998): 101-106, <http://bwrc.eecs.berkeley.edu/People/kcamera/neural/papers/schoenauer98digital.pdf>. See also: Yihua Liao, "Neural Networks in Hardware: A Survey," 2001, <http://ailab.das.ucdavis.edu/~yihua/research/NNhardware.pdf>

¹⁷⁵ Here is the basic schema for a genetic (evolutionary) algorithm. Many variations are possible, and the designer of the system needs to provide certain critical parameters and methods, detailed below.

<b2>The Evolutionary Algorithm

Create N solution "creatures". Each one has:

<nbl>

- A genetic code: a sequence of numbers that characterize a possible solution to the problem. The numbers can represent critical parameters, steps to a solution, rules, etc.

<ntx>For each generation of evolution, do the following:

<nbl>

- Do the following for each of the N solution creatures:
 - Apply this solution creature's solution (as represented by its genetic code) to the problem, or simulated environment. Rate the solution.
 - Pick the L solution creatures with the highest ratings to survive into the next generation.
 - Eliminate the (N-L) nonsurviving solution creatures.
 - Create (N-L) new solution creatures from the L surviving solution creatures by:
 - <nsl>(i) Making copies of the L surviving creatures. Introduce small random variations into each copy; or
 - (ii) Create additional solution creatures by combining parts of the genetic code (using "sexual" reproduction, or otherwise combining portions of the chromosomes) from the L surviving creatures; or
 - (iii) Doing a combination of (i) and (ii).
- Determine whether or not to continue evolving:
 - Improvement = (highest rating in this generation)-(highest rating in the previous generation).
 - If Improvement < Improvement Threshold then we're done.
- The solution creature with the highest rating from the last generation of evolution has the best solution. Apply the solution defined by its genetic code to the problem.

<b2>Key Design Decisions

In the simple schema above, the designer needs to determine at the outset:

<nbl>

- Key parameters:
 - <nunsl>N
 - L
 - Improvement threshold.
- What the numbers in the genetic code represent and how the solution is computed from the genetic code.

- A method for determining the N solution creatures in the first generation. In general, these need only be "reasonable" attempts at a solution. If these first-generation solutions are too far afield, the evolutionary algorithm may have difficulty converging on a good solution. It is often worthwhile to create the initial solution creatures in such a way that they are reasonably diverse. This will help prevent the evolutionary process from just finding a "locally" optimal solution.
- How the solutions are rated.
- How the surviving solution creatures reproduce.

<b2>Variations

Many variations of the above are feasible. For example:

<nbl>

- There does not need to be a fixed number of surviving solutions creatures (L) from each generation. The survival rule(s) can allow for a variable number of survivors.
- There does not need to be a fixed number of new solution creatures created in each generation (N-L). The procreation rules can be independent of the size of the population. Procreation can be related to survival, thereby allowing the fittest solution creatures to procreate the most.
- The decision as to whether or not to continue evolving can be varied. It can consider more than just the highest-rated solution creature from the most recent generation(s). It can also consider a trend that goes beyond just the last two generations.

¹⁷⁶ Sam Williams, "When Machines Breed," August 12, 2004,

http://www.salon.com/tech/feature/2004/08/12/evolvable_hardware/index_np.html.

¹⁷⁷ Here is the basic scheme (algorithm description) of recursive search. Many variations are possible, and the designer of the system needs to provide certain critical parameters and methods, detailed below.

<b2>The Recursive Algorithm

<ntx>Define a function (program) "Pick Best Next Step." The function returns a value of "SUCCESS" (we've solved the problem) or "FAILURE" (we didn't solve it). If it returns with a value of SUCCESS, then the function also returns the sequence of steps that solved the problem. Pick Best Next Step does the following:

<nbl>

- Determine if the program can escape from continued recursion at this point. This bullet, and the next two bullets deal with this escape decision.

First, determine if the problem has now been solved. Since this call to Pick Best Next Step probably came from the program calling itself, we may now have a satisfactory solution. Examples are:

<nnl>

(i) In the context of a game (for example, chess), the last move allows us to win (such as checkmate).

(ii) In the context of solving a mathematical theorem, the last step proves the theorem.

(iii) In the context of an artistic program (for example, a computer poet or composer), the last step matches the goals for the next word or note.

If the problem has been satisfactorily solved, the program returns with a value of "SUCCESS" and the sequence of steps that caused the success.

- If the problem has not been solved, determine if a solution is now hopeless. Examples are:

<nnl>

(i) In the context of a game (such as chess), this move causes us to lose (checkmate for the other side).

(ii) In the context of solving a mathematical theorem, this step violates the theorem,

(iii) In the context of an artistic creation, this step violates the goals for the next

word or note.

<nbl>If the solution at this point has been deemed hopeless, the program returns with a value of "FAILURE."

-
- If the problem has been neither solved nor deemed hopeless at this point of recursive expansion, determine whether or not the expansion should be abandoned anyway. This is a key aspect of the design and takes into consideration the limited amount of computer time we have to spend. Examples are:
 - <nnl>
 - (i) In the context of a game (such as chess), this move puts our side sufficiently "ahead" or "behind." Making this determination may not be straightforward and is the primary design decision. However, simple approaches (such as adding up piece values) can still provide good results. If the program determines that our side is sufficiently ahead, then Pick Best Next Step returns in a similar manner to a determination that our side has won (that is, with a value of "SUCCESS"). If the program determines that our side is sufficiently behind, then Pick Best Next Step returns in a similar manner to a determination that our side has lost (that is, with a value of "FAILURE").
 - (ii) In the context of solving a mathematical theorem, this step involves determining if the sequence of steps in the proof is unlikely to yield a proof. If so, then this path should be abandoned, and Pick Best Next Step returns in a similar manner to a determination that this step violates the theorem (that is, with a value of "FAILURE"). There is no "soft" equivalent of success. We can't return with a value of "SUCCESS" until we have actually solved the problem. That's the nature of math.
 - (iii) In the context of an artistic program (such as a computer poet or composer), this step involves determining if the sequence of steps (such as the words in a poem, notes in a song) is unlikely to satisfy the goals for the next step. If so, then this path should be abandoned, and Pick Best Next Step returns in a similar manner to a determination that this step violates the goals for the next step (that is, with a value of "FAILURE").
 - <nbl>
 - If Pick Best Next Step has not returned (because the program has neither determined success nor failure nor made a determination that this path should be abandoned at this point), then we have not escaped from continued recursive expansion. In this case, we now generate a list of all possible next steps at this point. This is where the precise statement of the problem comes in:
 - <nnl>
 - (i) In the context of a game (such as chess), this involves generating all possible moves for "our" side for the current state of the board. This involves a straightforward codification of the rules of the game.
 - (ii) In the context of finding a proof for a mathematical theorem, this involves listing the possible axioms or previously proved theorems that can be applied at this point in the solution.
 - (iii) In the context of a cybernetic art program, this involves listing the possible words/notes/line segments that could be used at this point.
 - <nbl>For each such possible next step:
 - <nnl>
 - (i) Create the hypothetical situation that would exist if this step were implemented. In a game, this means the hypothetical state of the board. In a mathematical proof, this means adding this step (for example, axiom) to the proof. In an art program, this means, adding this word/note/line segment.
 - (ii) Now call Pick Best Next Step to examine this hypothetical situation. This is, of course, where the recursion comes in because the program is now calling itself.
 - (iii) If the above call to Pick Best Next Step returns with a value of "SUCCESS," then return from the call to Pick Best Next Step (that we are now in) also with a value of "SUCCESS." Otherwise consider the next possible step.
 - <nbl>If all the possible next steps have been considered without finding a step that resulted in a return from the call to Pick Best Next Step with a value of "SUCCESS," then return from this call to Pick Best Next Step (that we are now in) with a value of "FAILURE."
 - <ntx>END OF PICK BEST NEXT STEP

If the original call to Pick Best Next Move returns with a value of "SUCCESS," it will also return the correct sequence of steps:

<nnl>

- (i) In the context of a game, the first step in this sequence is the next move you should make.
- (ii) In the context of a mathematical proof, the full sequence of steps is the proof,
- (iii) In the context of a cybernetic art program, the sequence of steps is your work of art.

<ntx>If the original call to Pick Best Next Step is "FAILURE," then you need to go back to the drawing board.

<b2>Key Design Decisions

<ntx>In the simple schema above, the designer of the recursive algorithm needs to determine the following at the outset:

<nbl>

- The key to a recursive algorithm is the determination in Pick Best Next Step of when to abandon the recursive expansion. This is easy when the program has achieved clear success (such as checkmate in chess or the requisite solution in a math or combinatorial problem) or clear failure. It is more difficult when a clear win or loss has not yet been achieved. Abandoning a line of inquiry before a well-defined outcome is necessary because otherwise the program might run for billions of years (or at least until the warranty on your computer runs out).
- The other primary requirement for the recursive algorithm is a straightforward codification of the problem. In a game like chess, that's easy. But in other situations, a clear definition of the problem is not always so easy to come by.

<ntx>¹⁷⁸ See Kurzweil Cyberart, <http://www.KurzweilCyberArt.com>, for further description of Ray Kurzweil's Cybernetic Poet and to download a free copy of the program.

See U.S. Patent # 6,647,395 "Poet Personalities," inventors: Ray Kurzweil and John Keklak. Abstract: "A method of generating a poet personality including reading poems, each of the poems containing text, generating analysis models, each of the analysis models representing one of poems and storing the analysis models in a personality data structure. The personality data structure further includes weights, each of the weights associated with each of the analysis models. The weights include integer values."

¹⁷⁹ Ben Goertzel, *The Structure of Intelligence*, 1993, Springer-Verlag.

Ben Goertzel, *The Evolving Mind*, 1993, Gordon and Breach.

Ben Goertzel, *Chaotic Logic*, 1994, Plenum.

Ben Goertzel, *From Complexity to Creativity*, 1997, Plenum.

For a link to Ben Goertzel's books and essays, see <http://www.goertzel.org/work.html>

¹⁸⁰ KurzweilAI.net (<http://www.KurzweilAI.net>) provides hundreds of articles by one hundred "big thinkers" and other features on "accelerating intelligence." The site offers a free daily or weekly newsletter on the latest developments in the areas covered by this book. To subscribe, enter your e-mail address (which is maintained in strict confidence and is not shared with anyone) on the home page.

¹⁸¹ John Gosney, Business Communications Company, "Artificial Intelligence: Burgeoning Applications in Industry," June 2003, <http://www.bccresearch.com/comm/G275.html>.

¹⁸² Kathleen Melymuka, "Good Morning, Dave . . .," *Computerworld*, November 11, 2002, <http://www.computerworld.com/industrytopics/defense/story/0,10801,75728,00.html>.

¹⁸³ JTRS Technology Awareness Bulletin, August 2004,

http://jtrs.army.mil/sections/technicalinformation/fset_technical.htol?tech_aware_2004-8.

¹⁸⁴ Otis Port, Michael Arndt, and John Carey, "Smart Tools," spring 2003,

<http://www.businessweek.com/bw50/content/mar2003/a3826072.htm>.

¹⁸⁵ Wade Roush, "Immobots Take Control: From Photocopiers to Space Probes, Machines Injected with Robotic Self-Awareness Are Reliable Problem Solvers," *TechnologyReview* (December 2002/January 2003), <http://www.occm.de/roush1202.pdf>.

¹⁸⁶ Jason Lohn quoted in NASA news release "NASA 'Evolutionary' Software Automatically Designs Antenna," http://www.nasa.gov/lb/centers/ames/news/releases/2004/04_55AR.html

¹⁸⁷ Robert Roy Britt, "Automatic Astronomy: New Robotic Telescopes See and Think," June 4, 2003, http://www.space.com/business/technology/automated_astronomy_030604.html.

- ¹⁸⁸ H. Keith Melton, "Spies in the Digital Age," <http://www.cnn.com/SPECIALS/cold.war/experience/spies/melton.essay>.
- ¹⁸⁹ "United Therapeutics (UT) is a biotechnology company focused on developing chronic therapies for life threatening conditions in three therapeutic areas: cardiovascular, oncology and infectious diseases" (<http://www.unither.com>). Kurzweil Technologies is working with UT to develop pattern recognition-based analysis from either "Holter" monitoring (twenty-four-hour recordings) or "Event" monitoring (thirty days or more).
- ¹⁹⁰ Kristen Philipkoski, "A Map That Maps Gene Functions," *WiredNews*, May 28, 2002, <http://www.wired.com/news/medtech/0,1286,52723,00.html>.
- ¹⁹¹ Jennifer Ouellette, "Bioinformatics Moves into the Mainstream," *The Industrial Physicist* (October/November 2003), <http://www.sciencemasters.com/bioinformatics.pdf>.
- ¹⁹² Port, Arndt, and Carey, "Smart Tools."
- ¹⁹³ "Protein Patterns in Blood May Predict Prostate Cancer Diagnosis," National Cancer Institute, October 15, 2002, <http://www.nci.nih.gov/newscenter/ProstateProteomics>, reporting on E. F. Petricoin et al., "Serum Proteomic Patterns for Detection of Prostate Cancer," *Journal of the National Cancer Institute* 94 (2002): 1576-78.
- ¹⁹⁴ Charlene Laino, "New Blood Test Spots Cancer," December 13, 2002, <http://my.webmd.com/content/Article/56/65831.htm>; Emanuel F. Petricoin III et al., "Use of Proteomic Patterns in Serum to Identify Ovarian Cancer," *The Lancet* 359.9306 (February 16, 2002): 572-77.
- ¹⁹⁵ Mark Hagland, "Doctor's Orders," January 2003, http://www.healthcare-informatics.com/issues/2003/01_03/cpoe.htm.
- ¹⁹⁶ Ross D. King et al., "Functional Genomic Hypothesis Generation and Experimentation by a Robot Scientist," *Nature All* (January 15, 2004): 247-52.
- ¹⁹⁷ Port, Arndt, and Carey, "Smart Tools."
- ¹⁹⁸ "Future Route Releases AI-Based Fraud Detection Product," August 18, 2004, <http://www.finextra.com/fullstory.asp?id=12365>.
- ¹⁹⁹ John Hackett, "Computers Are Learning the Business," *Collections World*, April 24, 2001, http://www.creditcollectionworld.com/news/042401_2.htm.
- ²⁰⁰ "Innovative Use of Artificial Intelligence, Monitoring NASDAQ for Potential Insider Trading and Fraud," AAAI press release, July 30, 2003, <http://www.aaai.org/Pressroom/Releases/release-03-0730.html>.
- ²⁰¹ "Adaptive Learning, Fly the Brainy Skies," *WiredNews*, March 2002, <http://www.wired.com/wired/archive/10.03/everywhere.html?pg=2>. 2
- ²⁰² Introduction to Artificial Intelligence, EL 629, Maxwell Air Force Base, Air University Library course www.au.af.mil/au/aul/school/acsc/ai02.htm.
- ²⁰³ See www.Seegrid.com.
- ²⁰⁴ No Hands Across America Web site, <http://cart.frc.ri.cmu.edu/users/hpm/project.archive/reference.file/nhaa.html>, and "Carnegie Mellon Researchers Will Prove Autonomous Driving Technologies During a 3,000 Mile, Hands-off-the-Wheel Trip from Pittsburgh to San Diego," Carnegie Mellon press release, http://www-2.cs.cmu.edu/afs/cs/user/tjochem/www/nhaa/official_press_release.html; Robert J. Derocher, "Almost Human," September 2001, <http://www.insight-mag.com/insight/01/09/col-2-pt-1-ClickCulture.htm>.
- ²⁰⁵ "Search and Rescue Robots," Associated Press, September 3, 2004, <http://www.smh.com.au/articles/2004/09/02/1093939058792.html?oneclick=true>.
- ²⁰⁶ "From Factoids to Facts," *The Economist*, August 26, 2004, http://www.economist.com/science/displayStory.cfm?story_id=3127462.
- ²⁰⁷ Joe McCool, "Voice Recognition, It Pays to Talk," May 2003, <http://www.bcs.org/BCS/Products/Publications/JournalsAndMagazines/ComputerBulletin/OnlineArchive/may03/voicerecognition.htm>.
- ²⁰⁸ John Gartner, "Finally a Car That Talks Back," *WiredNews*, September 2, 2004, http://www.wired.com/news/autotech/0,2554,64809,00.html?tw=wn_14techhead.
- ²⁰⁹ "Computer Language Translation System Romances the Rosetta Stone," Information Sciences Institute, USC School of Engineering (July 24, 2003), <http://www.usc.edu/isinews/stories/102.html>.
- ²¹⁰ Torsten Reil quoted in Steven Johnson, "Darwin in a Box," *Discover Magazine* 24.8 (August 2003), <http://www.discover.com/issues/aitg-OS/departments/feattech/>
- ²¹¹ "Let Software Catch the Game for You," July 3, 2004, <http://www.newscientist.com/news/news.jsp?id=ns99996097>.
- ²¹² Michelle Delio, "Breeding Race Cars to Win," *WiredNews*, June 18, 2004, <http://www.wired.com/news/autotech/0,2554,63900,00.html>.
- ²¹³ Marvin Minsky, *TAe Society of Mind* (New York: Simon & Schuster, 1988).

²¹⁴ Hans Moravec, "When will computer hardware match the human brain?" *Journal of Evolution and Technology*. 1998. Volume 1.

²¹⁵ Ray Kurzweil, *The Age of Spiritual Machines* (New York: Viking, 1999), p. 156.

²¹⁶ See Chapter 2 endnotes 22 and 23 on the International Technology Roadmap for Semiconductors.

²¹⁷ "The First Turing Test," <http://www.loebner.net/Prize/loebner-prize.html>.

²¹⁸ Douglas R. Hofstadter, "A Coffeehouse Conversation on the Turing Test," May 1981, included in Ray Kurzweil, *The Age of Intelligent Machines* (Cambridge, Mass.: MIT Press, 1990), pp. 80-102, <http://www.kurzweilai.net/meme/rframe.html?main=/articles/art0318.html>.

²¹⁹ Ray Kurzweil "Why I Think I Will Win." And Mitch Kapor, "Why I Think I Will Win."

Rules: <http://www.kurzweilai.net/meme/frame.html?main=/articles/art0373.html>; Kapor:

<http://www.kurzweilai.net/meme/frame.html?main=/articles/art0412.html>; Kurzweil:

<http://www.kurzweilai.net/meme/frame.html?main=/articles/art0374.html>; Kurzweil "final word":

<http://www.kurzweilai.net/meme/frame.html?main=/articles/art0413.html>.

²²⁰ Edward A. Feigenbaum, "Some Challenges and Grand Challenges for Computational Intelligence," *Journal of the Association for Computing Machinery* 50 (January 2003): 32-40.

²²¹ According to the serial endosymbiosis theory of eukaryotic evolution, the ancestors of mitochondria (the structures in cells that produce energy and have their own genetic code comprising thirteen genes in humans) were originally independent bacteria (that is, not part of another cell) similar to the *Daptobacterbacte* *via* of today. "Serial Endosymbiosis Theory," <http://encyclopedia.thefreedictionary.com/Serial%20endosymbiosis%20theory>.

WHAT IS CLAIMED IS:

1. A system comprising:
a computer configured to execute instructions for synthesizing biological material,
and
an assembly unit electrically connected to the computer, the assembly unit being configured to synthesize the biological material based on the instructions executed by **the** computer.
2. The system of claim 1, further comprising an insertion unit.
3. The system of claim 1, further comprising a repository unit.
4. The system of claim 1, wherein the assembly unit further comprises one or more of an input channel and an output channel.
5. The system of claim 1, further comprising a separate device for a user to communicate wirelessly with the computer.
6. The system of claim 1, wherein the computer comprises one or more of a memory unit, software, and a database.
7. The system of claim 6, wherein the database comprises one or more of DNA sequence, RNA sequence and polypeptide sequence information.
8. The system of claim 3, wherein the repository unit comprises one or more different types of monomeric biological components.
9. The system of claim 8, wherein the monomeric biological components include one or more of a nucleotide, amino acid, and tRNA.

10. The system of claim 9, wherein each tRNA is attached to an amino acid.
11. The system of claim 1, wherein the assembly unit comprises one or more of a polymerase or a ribosome.
12. The system of claim 1, wherein the assembly unit comprises a robot that mimics the activity of a polymerase or a ribosome.
13. The system of claim 1, wherein the biological material is a nucleic acid or polypeptide.
14. The system of claim 1, wherein the biological material is an RNA.
15. The system of claim 2, wherein the insertion unit is attached to the assembly unit.
16. The system of claim 1, wherein the system is coated with a biocompatible material.
17. A system comprising:
 - a computer configured to execute instructions for synthesizing biological material;
 - a central unit responsive to execution of the instructions to control an assembly unit; and
 - an assembly unit electrically connected to the central unit, the assembly unit being configured to synthesize the biological material based on the instructions executed by the computer.
18. The system of claim 17, further comprising a repository unit.

19. The system of claim 17, wherein the assembly unit further comprises one or more of an input channel and an output channel.

20. The system of claim 17, wherein the computer is separate from the central unit and the assembly unit.

21. The system of claim 17, wherein the computer resides outside of the cell and the central unit and the assembly unit reside inside the cell.

22. The system of claim 17, wherein the computer resides within the central unit, and the central unit and the assembly unit reside inside the cell.

23. The system of claim 17, further comprising a separate device for a user to communicate wirelessly with the computer.

24. The system of claim 17, wherein the computer comprises one or more of a transmitter, software, and a database.

25. The system of claim 17, wherein the central unit comprises one or more of a memory, a receiver, an engine, and an antenna.

26. The system of claim 24, wherein the database comprises one or more of DNA sequence, RNA sequence and polypeptide sequence information.

27. The system of claim 17, wherein the system further comprises a repository unit, and the repository unit comprises one or more different types of monomeric biological components.

28. The system of claim 27, wherein the monomeric biological components include one or more of a nucleotide, amino acid, and tRNA.

29. The system of claim 28, wherein each tRNA is attached to an amino acid.
30. The system of claim 17, wherein the assembly unit comprises one or more of a polymerase or a ribosome.
31. The system of claim 17, wherein the assembly unit comprises a robot that mimics the activity of a polymerase or a ribosome.
32. The system of claim 17, wherein the biological material is a nucleic acid or polypeptide.
33. The system of claim 17, wherein the biological material is an RNA.
34. The system of claim 17, wherein the system is coated with a biocompatible material.
35. A method comprising:
synthesizing a biological material and introducing the biological material into a cell,
wherein the biological material is synthesized by a system comprising:
a computer configured to execute instructions for synthesizing biological material, and
an assembly unit electrically connected to the computer, the assembly unit being configured to synthesize the biological material based on the instructions executed by the computer.
36. The method of claim 35, wherein the system further comprises an insertion unit and a repository unit.
37. The method of claim 35, wherein the system further comprises one or more of an input channel and an output channel on the assembly unit.

38. The method of claim 35, wherein the step of synthesizing the biological material is initiated by a signal from a user operating a device separated from the system, wherein the user uses the device to send a signal to a receiver located in the computer of the system.

39. A method comprising:
synthesizing a biological material and introducing the biological material into a cell,
wherein the biological material is synthesized by a system comprising:
a computer configured to execute instructions for synthesizing biological material;
a central unit responsive to execution of the instructions to control an assembly unit; and
an assembly unit electrically connected to the computer, the assembly unit being configured to synthesize the biological material based on the instructions executed by the computer.

40. The method of claim 39, further comprising a repository unit.

41. The method of claim 39, further comprising a step, prior to the synthesizing step, of putting at least the central unit and the assembly unit of the system inside the cell.

42. The method of claim 39, wherein the system further comprises one or more of an input channel and an output channel on the assembly unit.

43. The method of claim 41, wherein the step of putting the central unit and the assembly unit of the system into the cell comprises electroporation, microinjection, or a lipophilic carrier.

44. The method of claim 39, wherein the step of synthesizing the biological material is initiated by a signal from a user operating a device separated from the system, wherein the user uses the device to send a signal to a receiver located in the central unit of the system.

45. The method of claim 39, wherein the computer resides inside the central unit.

46. A cell comprising a system comprising:
a central unit responsive to instructions to control an assembly unit; and
an assembly unit electrically connected to the central unit, the assembly unit being configured to synthesize biological material based on the instructions,
wherein the instructions are executed by a computer.

47. The cell of claim 46, wherein the computer resides inside the central unit.

48. The cell of claim 46, wherein the computer resides outside the central unit and outside the cell.

49. The cell of claim 46, further comprising a repository unit attached to the assembly unit.

50. The cell of claim 46, wherein the assembly unit comprises one or more of an input channel and an output channel on the assembly unit.

51. The cell of claim 48, wherein the computer comprises one or more of a transmitter, software, and a database.

52. The cell of claim 48, wherein the central unit comprises one or more of a memory, a receiver, an engine, and an antenna.

53. The cell of claim 49, wherein the repository unit comprises one or more different types of monomeric biological components.

54. The cell of claim 53, wherein the monomeric biological components include one or more of a nucleotide, amino acid, and tRNA.

55. The cell of claim 54, wherein each tRNA is attached to an amino acid.

56. The cell of claim 46, wherein the cell originates from a mammal.

57. The cell of claim 46, wherein the cell originates from a human, mouse, rat, monkey, dog, cat, or rabbit.

58. The cell of claim 46, wherein the cell is in a tissue of a human.

59. A method of treating a human comprising administering a system to the human, wherein the system comprises:
a central unit responsive to instructions to control an assembly unit, and
an assembly unit electrically connected to the central unit, the assembly unit being configured to synthesize biological material based on the instructions,
wherein the instructions are executed by a computer.

60. The method of claim 59, wherein the computer resides inside the central unit.

61. The method of claim 59, wherein the computer resides outside the central unit and outside the cell.

62. The method of claim 59, wherein the system further comprises a repository unit attached to the assembly unit.

63. The method of claim 59, wherein the assembly unit comprises one or more of an input channel and an output channel.

64. The method of claim 59, wherein the computer comprises one or more of a transmitter, software, and a database.

65. The method of claim 59, wherein the central unit comprises one or more of a memory, a receiver, an engine, and an antenna.

66. The method of claim 64, wherein the database comprises one or more of DNA sequence information, RNA sequence information, and polypeptide sequence information.

67. The method of claim 62, wherein the repository unit comprises one or more different types of monomeric biological components.

68. The method of claim 67, wherein the monomeric biological components include one or more of a nucleotide, amino acid, and tRNA.

69. The method of claim 68, wherein each tRNA is attached to an amino acid.

70. The method of claim 59, wherein the human has a cancer, a tissue disorder, or a disorder of the nervous system.

71. The method of claim 59, wherein the system is administered by tissue graft, microprojectile technology, or by a lipophilic carrier.

72. A method of treating a human comprising administering a cell comprising a system comprising:

a central unit responsive to instructions to control an assembly unit; and

an assembly unit electrically connected to the central unit, the assembly unit being configured to synthesize biological material based on the instructions, wherein the instructions are generated by a computer.

73. The method of claim 72, wherein the computer resides inside the central unit.

74. The method of claim 72, wherein the computer resides outside the central unit and outside the human.

75. The method of claim 72, wherein the system further comprises a repository unit attached to the assembly unit.

76. The method of claim 72, wherein the assembly unit of the system comprises one or more of an input channel and an output channel on the assembly unit.

77. The method of claim 72, wherein the computer comprises one or more of a transmitter, software, and a database.

78. The method of claim 72, wherein the central unit comprises one or more of a memory, a receiver, an engine and an antenna.

79. The method of claim 77, wherein the database comprises DNA sequence information, RNA sequence information, and polypeptide sequence information.

80. The method of claim 75, wherein the repository unit comprises one or more different types of monomeric biological components.

81. The method of claim 80, wherein the monomeric biological components include one or more of a nucleotide, amino acid, and tRNA.

82. The method of claim 81, wherein each tRNA is attached to an amino acid.

83. The method of claim 72, wherein the human has a cancer, a tissue disorder, or a disorder of the nervous system.

84. The method of claim 72, wherein the cell is administered by tissue graft.




FIG. 1

Database 12




FIG. 2




FIG. 3

Flowchart 140

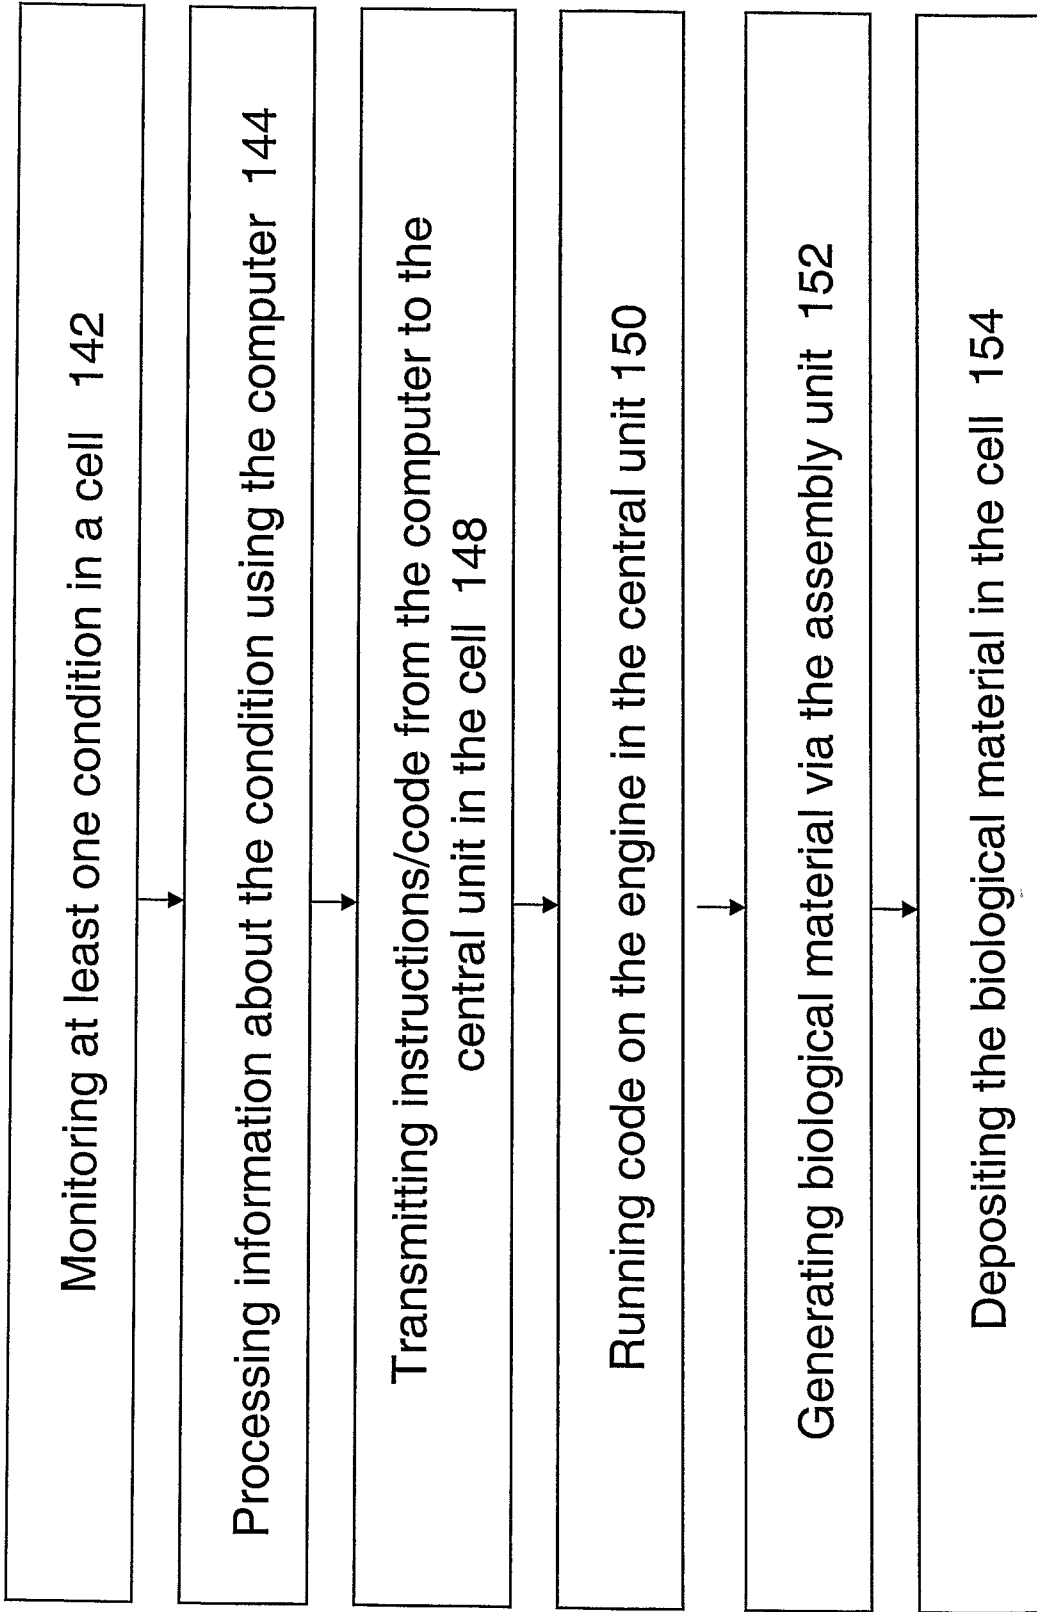


FIG. 4