HOW TO GROW A ROBOT

Developing Human-Friendly, Social AI

Mark Lee

HOW TO GROW A ROBOT

HOW TO GROW A ROBOT

DEVELOPING HUMAN-FRIENDLY, SOCIAL AI

MARK H. LEE

THE MIT PRESS   CAMBRIDGE, MASSACHUSETTS   LONDON, ENGLAND

This book was set in Stone Serif by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Lee, Mark H., author.
Title: How to grow a robot : developing human-friendly, social AI / Mark Lee.
Description: Cambridge : The MIT Press, 2020. | Includes bibliographical references and index.
Identifiers: LCCN 2019025740 | ISBN 9780262043731 (hardcover) | ISBN 9780262357852 (ebook)
Subjects: LCSH: Robots—Control systems. | Robots—Social aspects. | Machine learning. | Artificial intelligence—Forecasting. | Human-computer interaction.
Classification: LCC TJ211.35 .L43 2020 | DDC 006.301—dc23
LC record available at https://lccn.loc.gov/2019025740

CONTENTS

To Elizabeth, with love and gratitude for living with this unwelcome interloper for so long.

PREFACE

I have always been fascinated by the relationship between machines and humans. In my youth, I was very interested in electrical machines and engineering generally, which, with hindsight, I now see as an incredibly creative discipline. But during my training as an engineer, I was particularly attracted to systems that appeared to mimic some aspects of human behavior. Even simple feedback systems can be captivating (e.g., watching a ship constantly correct its course while being buffeted about by winds and tides). This leads to questions like: What kind of mental machinery do humans actually use? I somehow managed to work on projects that combined engineering with human-centered issues: speech encoding, color vision processing, and autonomous control. I found that psychology was the vital missing element, and my PhD, in modeling sensorimotor control and coordination, was a precursor of the work described here. The relationships among computers, brains, and machines are many and fascinating.

When I tried to pursue this line of research, in the 1980s, it was completely out of favor. The new commercialism of artificial intelligence (AI) had just begun, and AI software was becoming highly marketable. Before then, the early pioneers of AI saw no large distinctions between human and computer intelligence and thought both could be studied together. The drive for software products and applications caused AI to become largely separate from the study of human cognition. I remember discussing my work with Andy Barto, Richard Sutton, and others at the University of Massachusetts–Amherst when they were in the middle of developing reinforcement learning. In the UK, it was very difficult to get any funding for basic, curiosity-driven AI research, so I started work on AI in industrial robotics, tactile sensing, and error recovery and diagnostic systems, still exploring how humans behave during such tasks and how AI could be used in robotics. Fortunately, the twenty-first century has seen a swing back to topics like learning, autonomous agents, and renewed attention to human performance, and there is now a strong global research community in the new field of developmental robotics.

I have written this book partly in response to some common misconceptions about robots and AI. A great deal of misinformation about robots and AI has been bandied about in the media, some of which is clearly unreasoned nonsense. While there will be some amazing technology produced in the near future, there are limitations inherent in AI, and there are ethical issues that involve us all.

However, this is not a negative story because, after explaining the difficulties and nature of the problem, I then go on to describe an alternative approach that is currently delivering results and showing real potential for developing general-purpose robots. I give details from my own and related research to show how psychology and developmental ideas can get nearer to humanlike behavior—which is fundamentally general purpose but can adapt in order to learn to accomplish specialized tasks.

I try to avoid the minutiae of the latest technology or subtle changes of fashion; instead, I concentrate on the trends, the way technology develops, and the role of technology in our lives. I feel that it is important that everyone should be much more engaged in appreciating the technological developments of the day in order to get a better perspective on what is feasible, what is reasonable, and what is wild hyperbole. By applying basic common sense to many of the claims and predictions, we can evaluate, and therefore have a much larger influence on, the role of technology in our lives (and future lives). This book is a response to these concerns and contains some of my findings and insights into these fascinating problems. I hope that it will give a perspective that offers a better way of approaching some of the wider issues raised by intelligent robots.

ACKNOWLEDGMENTS

There are many colleagues, friends, and students that have helped me in my career, and it is impracticable to list them all. This includes colleagues in the Computer Science Department at Aberystwyth University, as well as many others in universities and other organizations in the UK, Europe, Japan, and the US. Research is a truly global community. I am grateful to all these friends and colleagues who generously supported, influenced, and encouraged the approach described here. I should particularly thank those who were directly involved in the work reported in this book: Raphael Braud, Fei Chao, Kevin Earland, Tao Geng, Alexandros Giagkos, Richard Gunstone, Martin Hülse, Ian and Tom Izzett, Suresh Kumar, James Law, Daniel Lewkowicz, Sebastian McBride, Qinggang Meng, Marcos Rodrigues, Patricia Shaw, Michael Sheldon, Qiang Shen, James Wilson, and Xing Zhang. Of course, the opinions expressed here are solely my responsibility, as are the errors and omissions.

Others who have helped in more mysterious ways include Anna Borghi, Merideth Gattis, Frank Guerin, Kevin Gurney, David Llewelyn, Giorgio Metta, David Middleton, Kevin O'Regan, Peter Redgrave, Peter Tallack, Raymond Tallis, and my late but dear friends Brendan McGonigle and Ulrich Nehmzow. Forgive me for not listing all the others; you are not forgotten, and I am grateful to you all.

I much appreciate the expertise and kind support of Marie L. Lee at MIT Press, and I thank all the copy-editing and other staff who have helped in this project.

The research described here was mainly carried out during four projects: two funded by the British Engineering and Physical Science Research Council, and two funded by the FP7 program by the European Commission (EC). The iCub humanoid robot originated in an EC-funded project, and the EC had the foresight to support the supply of these robots to new research projects, with the result that iCubs are now being used in more than twenty-five robotics research laboratories worldwide.

I am grateful to the Royal Society of Arts for permission to reuse and modify material taken from my article "A Frame of Mind," *Royal Society of the Arts Journal*, no. 2, 2018, 46–47. This appears in chapter 9, in the section entitled "The Brain Is a Machine—So What?"

I am grateful for the use of many libraries, particularly six local libraries. This includes the Bronglais Postgraduate Center and the wonderful resources of Llyfrgell Genedlaethol Cymru, the National Library of Wales. Finally, I am most grateful for having such a positive, helpful, and brilliant family, and particularly for my dear wife, Elizabeth, who has encouraged and supported this project in so many ways.

I

WHAT'S WRONG WITH ARTIFICIAL INTELLIGENCE?

1

THE NATURE OF THE PROBLEM

Common sense—Ordinary sensible understanding; one's basic intelligence which allows for plain understanding and without which good decisions or judgments cannot be made.
—*Wiktionary*, July 16, 2019

Keepon is a tiny robot, only 12 centimeters high. It's made of two balls of yellow silicone rubber: a head and a body. All Keepon can do is roll its head from side to side, nod backward and forward, and rock and bounce its torso. But humans find Keepon very interesting and get involved in "conversations," as Keepon nods its head and jigs about in excited response to their vocalizations.

Keepon has no arms or legs. It cannot move around or handle objects. It has no facial muscles; its face is simply two tiny eyes and a circular nose (actually, the eyes contain small cameras, and the nose is a microphone). Any expression or communication has to be through bodily movement or gesture.[1] Yet this extremely minimal robot is quite engaging for humans. Why?

This phenomenon of engagement is also seen in animated films, where moving line drawings, minimal representations, somehow convince us that the depicted characters actually exist. The answer, of course, is that they are animate; they move in exactly the right way—the way that living

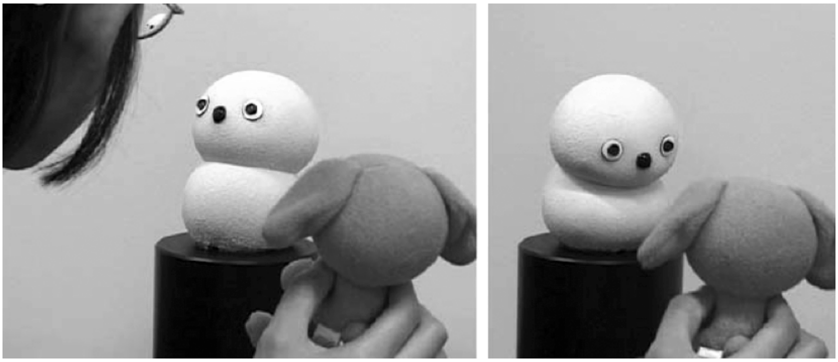

1.1 Keepon displaying eye contact and joint attention (looking at an object that the human has just looked at). Used with permission from Professor Hideki Kozima (Tohoku University).

creatures move. Walt Disney created an industry, and made a fortune, from this effect.

Keepon is captivating because it *looks at humans when they talk* and also performs what's known as *joint attention*—it looks at the same objects that the human looks at, giving the impression of shared interest (see figure 1.1). It's not necessary for a social robot to pretend to be human, or even resemble us. We are happy talking to all kinds of animals: Just look at the fuss we make over a puppy or a kitten! But we do need that animal-like behavior and response.

This raises a big question: If Keepon and some other robots can give a passable imitation of animal movements, why do the vast majority of robots seem so clunky, so "mechanical"? If we want future robots to become useful assistants, companions, or supernumeraries, then they will have to be much more humanlike, friendly, and engaging.

ACTING AND THINKING

Robotic devices have fascinated both scientists and the public at large for a very long time. Some of the earliest ones were treated as religious marvels, such as the automaton monk of 1560, a small figurine that prayed while "walking" across a table top.[2] From the seventeenth century onward, automata became very popular as entertainments, and the eighteenth century was a golden age for autonomous mechanical figures and animals. Thanks to the skills of watchmakers, really impressive models were constructed, such as the life-size Silver Swan of 1773, which moves its head and neck realistically, preens itself, and catches a wriggling fish from the simulated moving water on which it sits.[3] The automatic human figures of those times often played musical instruments or danced; sometimes they wrote or drew on paper; but they always mimicked human behavior. The key to their attraction was that their actions were humanlike.

From the clockwork toys and automata that so amused our ancestors to the amazing science fiction creations seen in modern film and television, we have always had a deep empathy for machines that can mimic human actions. This fascination extends across a wide range of movement styles, from total mimicry to imitations of bodily actions. We enjoy watching the grace and elegance of the human form in athletics or ballet, but we are also intrigued by the motions of factory machines that assemble components or paint cars ("poetry in motion," as I saw emblazoned on the side of a heavy lifting crane).

It's not just movement: We are also fascinated by displays of purely mental activity. Champion chess players, quiz masters, and expert players at word and number games all draw large audiences; and thinkers like Isaac Newton, Charles Darwin, and Albert Einstein are revered as scientific heroes. Future robots will have to be smart too—very smart. If they are going to keep up with humans, then they will have to understand us and what we want, and this requires a lot of brain power. They will have to perceive us, as humans, and continually learn about our changing wants and needs. And all this must happen in real time; that is, *in time with us*—not too fast or too slow. We soon ditch technology when it gets boring, inflexible, or unresponsive.

These two requirements, *appropriate behavior* and *mental ability*, are the subject of this book. Some robots are good at moving; some computers are good at thinking; but both skills are necessary for really intelligent robots. We need robots that humans can engage with, robots that are perceptive, animated, and responsive to us. This book explores if and how this can be achieved; and what is necessary to support such an advance. In short, the question is: *How can we build intelligent robots that think and behave like humans?*

THE SOCIAL ROBOT

Among the robots that already exist, very few even attempt to behave like humans. Keepon is in a very small minority. Most robots can be called *autonomous agents*, but very few are *social agents*. This is the issue that this book examines: Robots that think and behave like humans must be social agents, freely interacting with humans and accepted by humans.

Just to be clear, a social robot is more than just a fancy interface, avatar, or conversational system. Social robots act like a companion or friend, with frequent interactions over an extended period. They won't necessarily look like, or pretend to be human, but they will understand human gestures, behaviors, and communications and be able to respond in ways that we find socially satisfactory. They will have limited "emotional" behavior but they will learn and adapt through their social life, recognizing different people and remembering their personalities from previous encounters. Of course, they will also do useful work, but rather than being programmed for specific tasks, they will be trained by their users and learn through demonstration and human guidance, as we do. If we want to create long-term, useful robot assistants, then social interaction is the key.

We will see in chapter 8 how competence in certain kinds of dialogue has been demonstrated. But this is a narrow application: The big question is how far we can progress toward sustained and meaningful social intercourse.

This book is not concerned with any specific application of social robotics; the social robot is proposed here as a milestone for the highest level of robotic achievement: thinking and behaving like humans. This is a research goal, not an application target. For example, I don't advocate humanoid robots for the care of elderly people; this important area still requires the human touch.

But once we know how to create social robots, once we have developed the technology, there will be plenty of scope for applications, practically anywhere that human-robot interactions regularly occur.

THE ROLE OF ARTIFICIAL INTELLIGENCE

So, how can we build social robots? What techniques are required, and what knowledge do we need? The growth of artificial intelligence (AI) is widely seen as the answer. AI is viewed as the bedrock on which future robots will be created. (I use a simple definition of the term here: *AI* refers to computer systems that achieve results that are assumed to require intelligence.)

It's true that AI is making enormous strides, producing impressive advances and products. For example, IBM's Watson system can answer questions posed in everyday language. Watson dissects the question and assembles answers using huge data sources: over 200 million pages of data, including the full text of Wikipedia! This quiz machine, which won $1 million on the *Jeopardy!* game show playing against two highly expert contestants, is now being adapted for use in health care, finance, telecommunications, and biotechnology.

In other areas, like speech and image analysis, new deep learning techniques have advanced AI. Computer recognition of images now demonstrates error rates of less than 1 percent. Image processing, like analyzing X-ray images or labeling the content of pictures, is producing impressive applications and products. This new boom in AI is generating a great deal of investment. Small AI companies are frequently bought for hundreds of millions of dollars. The heads of big companies believe that they need an AI department; and often go out and buy one—or several! Robotics has always involved AI, and many future robot systems will be powered by these new AI methods.

INTELLIGENCE IN GENERAL

There are three broad strands to human intelligence: the intrinsic abilities and adaptive potential of the individual brain, the joint results of combining several brains in social interactions, and the collective cultural knowledge produced by and available to humanity (see box 1.1). AI has mostly concentrated on building versions of mental skills and competences found in the brain: an *individual* brain. This ignores the interactions *between* brains that form a crucial context for human behavior

Box 1.1
Sources of intelligence

Brains The origins and expression of intelligence are found in animal brains. This obvious fact has led to most AI systems being constructed as a single agent, thus reinforcing the idea of the individual as the sole source of intelligence.

Social interactions Any activity that requires a group of individuals to complete, such as problem-solving in emergency situations, will involve social cooperation and organization. Social conversation involves communication, shared understanding of events and experiences, and learning through language.

Culture The term *culture*, in this context, refers to the cumulative collection of all written text, data, diagrams, reference materials, and other works that humanity has compiled. Through reuse and training, such knowledge enables individuals to become more intelligent and also, through study, to discover further knowledge.

(box 1.1). In our early years, our comprehension, worldly understanding, knowledge, skills, and other cognitive abilities all depended on the interactions of at least one attendant caregiver. Without parental care, we do not develop and thrive and often do not even survive. In adult life, we form groups, societies, organizations, and networks, which are important for our continued development. Experts consult their peers and predecessors, thereby increasing their intellectual skill and knowledge. The third aspect of intelligence is culture. By *culture*, I do not mean trends, fashions, or popular activities, but rather the totality of recorded human experience across the global population. Language gives humans these extensions of intelligence into social and cultural dimensions; this is perhaps our only real difference from other animals, as will be discussed in later chapters.

Many AI systems are individual packages of intellectual skills, such as game-playing programs. On the other hand, speech and textual processing systems (e.g., translation) rely on enormous data banks of externalized knowledge. Cumulative cultural knowledge is recorded and disseminated in all kinds of ways; recently, this has become important for AI and is often known as *Big Data*—large data repositories that capture vast quantities of text, images, and other information.

So brains and culture (or intellect and knowledge) are well represented in AI, but social interaction between individuals is less in evidence as a research area. Social systems *are* studied in AI, two examples being the pioneering work on the robot Kismet (Breazeal, 2002) and the research program of Dautenhahn (e.g., Dautenhahn, 2007) (see Fong et al., 2003 for a general review). Social interaction is a feature of human interfaces, interagent communications, and robot-human interactions, but these nearly all focus on specific tasks or applications; they often do not consider the social sciences and the way that humans actually behave and think. AI generally has to face the fact that human intelligence is not entirely contained within the skull; it is also diffuse and distributed, and it often requires social interaction and close cooperation.

BRAINS NEED BODIES

Many scientists and philosophers have argued that bodies are essential for intelligent agents, and modern robotics has contributed evidence to this concept. Embodiment, together with constructivism and primacy of action, are key principles in understanding how the life process, the life cycle of individuals, affects their cognitive growth. We will examine these ideas in Part II.

Social agents, both humans and robots, need real-time, real-life experience: They have to have a body! The body is not just an input-output device; rather, it is the source of the experiences that shape cognitive processes. You have to *possess* a body *and* be prepared to talk about it. Social robots need a model of the "self"; they have to know that they have a body and be aware of its properties and behavior. Then they can take a subjective stance; they can view their experience from their own perspective, not that of their designer. This may appear to be a radical viewpoint. It requires us to examine what AI offers, look at alternative approaches, and consider and compare the possible outcomes.

THE STRUCTURE AND THEME OF THIS BOOK

In this book, I have avoided the two extremes of (1) the detailed and arcane complexities of modern technology and (2) the heavy philosophizing that AI attracts. Instead, I have adopted an *engineering* approach that focuses on the purpose of the products and how they serve our interests. I believe most of these issues can be explained in practical terms. I aim to ask simple but probing questions to cut through the politics, bias, and exaggeration and expose the most fundamental issues upon which judgments can be made. The results are given in a section entitled "Observations" at the end of each chapter.

It is necessary to understand the basic ideas before we can form an opinion about them, and unfortunately, the complexities of computing technology often seem to obscure the concepts and principles at the heart of the matter. A general understanding, with sufficient relevant information, allows common sense to draw sensible and sound conclusions. For example, car drivers need to know (and actually are mandated by law to know) quite a lot about how a vehicle behaves in response to its controls. There are various things to learn about how and when to apply the brakes on bends, at speed, and in wet conditions. Drivers need to master complex relationships between the controls and the braking effects desired. These are often subtle and take time to learn. But drivers do not need to know what kind of brakes they are using, how they operate, or what engineering design considerations they embody. That information isn't necessary to be a good driver or to make informed opinions about what features are desirable for our transport needs.

The same should be true with computing technology. We may know that there is some kind of digital processor inside the latest device, but we don't need the fine details; we need a good idea of what it does, a rough idea of how it does it, and an understanding of how we can use it—what use it is to us.

Thinking and behaving are characteristics that don't often occur together in machines. At the "behaving but no thinking" end of the spectrum, we find robotic devices that can perform all manner of physical movements but are controlled, by either fixed recordings (e.g., paint-spraying robots) or remote human operators (e.g., nuclear reactor maintenance robots). In both these cases, the robot is not considered to be thinking. At the other extreme, we find computers that just answer questions (like HAL in *2001: A Space Odyssey*, or IBM's Watson, described in detail in chapter 8). These are technically not robots because they don't have bodies and don't move about.

We must consider both robotic hardware and pure thinking machines because, as we will see later, behaving and thinking are intimately interrelated in humans and animals, and probably need to be so in robots too. You have to be able to think *and* move the pieces to play a game like chess;[4] but who would have thought that physically moving the pieces would turn out to be the harder task to automate?

PART I: THE LIMITATIONS OF AI FOR ROBOTICS

The first part of this book looks into the state of current robotics and AI. We begin in the next chapter, with a brief overview of a range of commercial robot offerings that is presented in order for readers to appreciate the scope of modern robotic devices and what they do. This gives a picture of their current capabilities, achievements, and limitations. This is followed in chapter 3 by an overview of research robotics, including the way that innovation is supported and funded. A summary table shows two key factors that determine the difficulty of implementing successful robots: the nature of the tasks that they are set to do, and the complexity and chaos in the environments in which they have to work.

Because AI is key to future robotics, part I contains several chapters looking at how AI actually works, what it can deliver, and what it has trouble doing. Each chapter introduces a key topic that readers can use to build an understanding of the ideas and issues in modern AI and their relevance for robotics.

It turns out that there are some serious obstacles when it comes to humanlike behavior. Despite media claims, AI is nothing like human intelligence and has not made much progress in moving in this direction. AI is very powerful for specific tasks and excellent for solving individual problems, but it still struggles to reproduce the general-purpose intelligence characteristic of the human species. No program has yet been written that can pretend to be a human in any convincing way for more than a few minutes. The AI chapters lead to a striking conclusion: *Engineering methods and AI are not sufficient (or even appropriate) for creating humanlike robots*.

Humans are *generalists*, unlike many animals. We readily adapt to fit different niches; we even adapt niches to suit us. This requires a wide-ranging ability to adapt to new situations. Humans are autonomous, motivated by internal drives as much as external forces, and the key to our flexibility is our open-ended learning ability. This allows us to face new challenges, learn about them, and find ways of dealing with them. Some animals are prewired by their genes; ducklings will follow the first "parent" they see after hatching. But humans are able to learn how to master completely new tasks from scratch. Even our failures teach us important lessons. So a humanlike robot really requires general-purpose intelligence that has the potential to deal with any task.

General-purpose intelligence is the "holy grail" of AI, but thus far it has proved elusive. Despite sixty years of research and theorizing, general-purpose AI has not made any progress worth mentioning, and this could even turn out to be an impossible problem (at least, impossible for software-based AI). However, even if general-purpose AI could somehow be realized in a computer (which I doubt), this would not be enough to solve the problem of thinking and behaving like humans. This is because human intelligence is *embodied*: All our thoughts, sensations, intentions, relationships, concepts, and meanings only make sense in terms of a body. Our brains do not run abstract codes that control our behavior—that's the wrong metaphor. Brains (i.e., our thoughts) are deeply dependent on, and shaped by, our bodily experience, including social and cultural experience. We are so fully integrated into our living environments that our internal thinking structures mirror and match these in an almost symbiotic way.

The AI chapters build up from basic concepts toward recent advances in Big Data and Deep Learning (chapter 7). These developments have made some remarkable progress and are being heralded as the solution to all remaining AI problems. The idea is that AI learning methods are now (or soon will be) powerful enough to learn any task required, and thus evolve into general-purpose intelligence. This would be a serious and important event in the history of AI, and so it has profound implications. Toward the end of part I, chapters 9 and 10, on brain models and AI integration, respectively, explain the issues with general AI and the nature of agency and embodiment. These two chapters contain the core of the argument against AI in social robots.

It is possible to skip some of the AI details, but do be sure to read the list of short points ("Observations") at the end of each chapter. These state the bottom lines that summarize all the important things to remember, and they cumulatively build up the book's arguments.

PART II: DEVELOPMENTAL ROBOTICS

The obvious question then is: If AI faces obstacles in building humanlike robots, are there other ways to go? Are there any other paradigms we can explore? The second part of this book presents a very different approach. This new viewpoint says: If you want a humanlike robot, why not start from what we know about humans? The idea here is not to "build your own robot," but rather to "grow your own robot." Don't load up your robot with piles of preprepared human knowledge; allow it to explore for itself and build its own knowledge about the things that come to have meaning for it, as we do. This is an embodied, enactive approach that allows cognitive functions to grow in step with experience. The idea is to shift the source of knowledge and meaning from the robot designer, or from Big Data, to the robot itself. Inspiration and data for this approach come from studies of early infant behavior, where many new and striking ideas on learning and interactive behavior can be found.

Infants have very limited perceptual and motor skills at birth but rapidly become physically and socially adept. They learn to communicate, they build an understanding of themselves as a "self," and they learn to recognize other social agents. By modeling some of these mental growth processes, robots will be able to *develop* and learn social interaction abilities, rather than being given social algorithms by their designers or extracting behavioral patterns from human data files.

This part of the book is supported by my own research into development as a robot learning mechanism. Experiments on the iCub humanoid robot have refined the approach and produced interesting results. The robot has developed from very poor levels of sensory-motor competence (just a few reflex actions and weak sensory resolution) to gaining mastery of the muscles in its limbs, eyes, head, and torso, perceiving and reaching for interesting stimuli, and playing with objects. These longitudinal experiments have taken the robot from equivalent infant stages of newborn to about nine months. At present, the iCub can reach and play with objects, associate single word sounds with familiar objects or colors, make meaningful gestures, and produce new actions for new situations.

Fortunately, the way that babies play and interact is well documented in the large body of theoretical and experimental work found in developmental psychology. This provides both inspiration and fiducial milestones for the new research field of developmental robotics. The origins of this kind of robotics are explained in chapter 11, as is embodiment, which is fundamental and different from AI. Then, the way that the iCub works is described and the principles made clear. It becomes obvious that robots can have a kind of subjective experience because everything learned is based on active experience *as seen by the robot*. By developing learning, they generalize and refine their most basic experiences so that meaningful patterns and correlations are stored, while unreliable phenomena are discarded. Note that these robots develop individually; they may begin identical, but they will produce somewhat different mental models and behaviors. This reminds us of the uniqueness of humans: Each person has her or his own history of lifetime experiences, and these contextualize particular social interactions.

This "growing" methodology, involving repeated refinements through experimentation, is synthetic rather than analytic and, I believe, gives us a much better route toward human-friendly robots with more humanlike characteristics.

PART III: FURTHER AHEAD, LIKELY DEVELOPMENTS, AND FUTUROLOGY

The implications for the future of both AI robots and humanlike robots are outlined in the final part of this book. The likely near-term advances in behavior and performance are estimated and compared. For instance, AI robots are given goals by their designer; they have an objective stance because they are structured from without. So the contrast with developmental robots is sharp and revealing. For example, future AI robots will have vast data resources available to them. They will be able to answer all kinds of questions, they will solve problems, and they will learn. So entering an unknown room, recognizing objects, and carrying out tasks using the objects in the room will become quite routine. Their abilities with objects and environments will initially be much superior to developmental robots. But their behavior will not really be humanlike. They will use Big Data on human behavior to simulate common human patterns, often convincingly. But they will solve problems differently, reason differently, and construct different concepts and percepts. And these differences will be detectable, and actually *noticeable*, particularly in social situations. AI robots will know humans only as objects—active agents, but objects nonetheless. They will fail to empathize, to appreciate another social agent as the same as themselves. This is because models of self cannot be designed, but rather are dynamic and must be created from experience.

In contrast, developmental robots will have an internal model of "self" (their body space, their actions and effects, and their environment), and they will be able to match their experience of another similar agent or similar behavior with their own self-model and recognize a human as a separate agent: another "self." This allows much better interaction and social relationships because the robot can anticipate and appreciate why humans behave as they do. Learning will improve the model of the interactive partner, and exchanges will become more grounded in the context and understanding of both individual and joint intentions and goals.

After exploring these differences, I then take a somewhat wider look at the impacts that we might expect from this kind of technology. The inevitable spread of AI will affect all realms of human life. Such so-called disruptive technology offers both threats and benefits. I comment on some of the claims and alarms generated by the current research climate. This includes the threat to jobs, the role of trust, and the positive benefits of keeping humans involved in combined human/computer systems. It becomes clear that the threats from robotics and AI are part of the context of wider threats from general digital technology. In this part of the book, the observations spread out to apply to our total computing infrastructure.

The final chapters consider a number of projections for future progress and examine their feasibility. All sorts of nonsense is written in the media about robots, and it can be hard to know what to believe. Journalists love robotics stories because they allow huge scope for speculation, outrageous predictions, and scaremongering. Headlines like "Robots Could Be Deemed as an 'Invasive Species' and Threaten the Future of the Human Race Because They Are Evolving So Quickly," which appeared about a BBC program on July 25, 2017, are not uncommon. The news media continually produce stories about the dangers of intelligent robots taking over practically everything, and this is often backed up by futurists and academics who predict superpowerful breakthroughs at regular intervals. While amazing advances both in the power and scope of applications offered by new technology, do occur regularly, there is no reason to blindly project this success into areas that have repeatedly failed to live up to predictions. The fact that a great many technological problems have been solved in a very short time is not sufficient evidence that we are on the verge of solving *all* outstanding problems.

I try to examine the current situation from a less hysterical perspective. From an analysis of the documented predictions of over 200 experts, I categorize and average their claims to produce a metasummary. This reveals a more balanced view of future developments from practicing engineers who are working in the field and actually solving difficult AI and robotics problems. This includes a small group of the leading experts at the major high-tech companies—those with great experience, those who have made major breakthroughs, or both. They recognize the limitations of technological progress and offer a mature, positive, and more stable image of the future. Nevertheless, the threats are real, if somewhat different from the media hype, and we must all be involved in the safe and beneficial management of future digital technology.

COPING WITH THE PACE OF CHANGE

Three factors are responsible for the recent surge in AI applications: (1) the doubling of raw computer power every two years has been enhanced by special hardware for graphics and neural networks, as well as new hardware for running new data-learning methods; (2) massive data storage facilities are becoming common, providing huge, readily available data banks on all kinds of topics; and (3) programming and software tools are much improved, and it is getting easier to rapidly implement complex software for new applications.

Box 1.2
Terminology for programs

Algorithms are plans for a computation. An algorithm gives a simplified description of what a program does or should do to achieve a defined job. Algorithms are the language of software, used among designers, programmers, and developers. Algorithms are to programs as recipes are to cooking.[5]

Programs are the written instructions for a computer that are intended to perform a specific task. Programs usually implement an algorithm. A wide range of programming languages exist for different uses (e.g., business, scientific, gaming, embedded, interactive, web-based).

Code is a set of instructions running inside a computer. Programs are converted into code during the implementation of software.

Software is a program or collection of programs that make up an application or system. *Software* is a mass noun, like *money*, so it frequently refers to many programs.

App is short for *application software*. Apps perform a specific task and can be easily installed into a system. They are widely used for smartphone applications, but there are also web apps and general computer apps.

Bots or **softbots** are programs that operate autonomously to find or record particular data on the internet or within a large data store.

Chatbots are bots that interact with humans, often on the internet, in a conversational mode. They are used to assist customers on websites and commercial services. Personal assistant devices are another example of chatbots.

These advances mean that new ideas can be tried out more quickly, and previously impractical ideas become feasible. The range of tools and ideas available for creating new applications has increased enormously.

For these reasons, it is not sensible to write about only the latest developments; basing analysis on the very latest products is a futile effort. By the time this book is published, the very latest technology will be quite different. But what we can do is examine the growth trends and general principles behind what's going on. By looking back at past milestones, we can see how obstacles have been overcome and get general insight into how things develop. The history of a technology offers a perspective that is valuable for predicting and assessing the future. The more data that are available for consideration, the more confidence that we can have in any projections.[6]

A NOTE ON JARGON

I aim to avoid jargon, but I've already used the word *algorithm*! I will try to explain any overly technical terms that come up in the discussion. Actually, *algorithm* is a good example of a technical term that now has entered the vernacular. Once we talked about computer programs, then software, and now algorithms. At a colloquial level, all these can be taken as roughly equivalent, but to add a bit of discrimination, box 1.2 gives some details on the particulars of each term. Software usually consists of a collection of programs, and programs implement one or more algorithms; an algorithm is the essence of a computation. New terms have been coined for particular classes of application, as the last three illustrate.

I also try hard to avoid formulas and unexpanded acronyms—those dread, dead symbols that separate the technological laity from those "in the know."

For those who wish to follow up on some of the many fascinating topics, I include references to the relevant literature. I usually recommend either the classic papers or (for new concepts) introductory material, and other citations provide links to key technical sources.

BIBLIOGRAPHY

Ariely, Dan. *Predictably irrational*. HarperCollins, 2008.

Armstrong, Stuart, and Sotala, Kaj. "How we're predicting AI—or failing to." In J. Romportl, E. Zackova, and J. Kelemen, eds., *Beyond artificial intelligence: The disappearing human-machine divide* (11–29). Springer, 2015.

Ashby, W. Ross. *Design for a brain: The origin of adaptive behaviour*. 2nd ed. Chapman and Hall, 1976.

Banavar, Guruduth. *Learning to trust artificial intelligence systems: Accountability, compliance and ethics in the age of smart machines, 2016*. IBM technical report. Available from https://www.alain-bensoussan.com/wp-content/uploads/2017/06/34348524.pdf.

Barlow, Horace B. "Cerebral cortex as model builder." In D. Rose and V. G. Dobson eds., *Models of the visual cortex* (37–46). John Wiley, 1985.

Bastos, Andre M., Usrey, W. Martin, Adams, Rick A., Mangun, George R., Fries, Pascal, and Friston, Karl J. "Canonical microcircuits for predictive coding." *Neuron*, 76(4), 695–711, 2012.

Becker, J. D. "A model for the encoding of experiential information." In R. C. Schank and K. M. Colby, eds., *Computer models of thought and language* (396–434). W. H. Freeman and Company, 1973.

Bengio, Yoshua, Courville, Aaron, and Vincent, Pierre. "Representation learning: A review and new perspectives." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828, 2013.

Bernstein, Nikolai A. *The co-ordination and regulation of movements*. Pergamon Press Ltd., 1967.

Berthouze, L., and Lungarella, M. "Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom." *Adaptive Behavior*, 12(1), 47–64, 2004.

Bessen, James E. *How computer automation affects occupations: Technology, jobs, and skills*. Technical Report No. 15-49, Boston University, School of Law, Law and Economics Research Paper, 2016.

Bizony, Piers. *The making of Stanley Kubrick's 2001: A Space Odyssey*. Taschen, 2015.

Blum, Christian, Winfield, Alan F. T., and Hafner, Verena V. "Simulation-based internal models for safer robots." *Frontiers in Robotics and AI*, 4, 74, 2018. Retrieved from https://www.frontiersin.org/articles/10.3389/frobt.2017.00074/full.

Boden, Margaret, Bryson, Joanna, Caldwell, Darwin, et al. "Principles of robotics: regulating robots in the real world." *Connection Science*, 29(2), 124–129, 2017.

Boden, Margaret A. *AI: Its nature and future*. Oxford University Press, 2016.

Borghi, Anna M., Flumini, Andrea, Natraj, Nikhilesh, and Wheaton, Lewis A. "One hand, two objects: emergence of affordance in contexts." *Brain and Cognition*, 80(1), 64–73, 2012.

Bostrom, Nick. *Superintelligence: Paths, dangers, strategies*. Oxford University Press, 2014.

Bottou, Léon. "From machine learning to machine reasoning." *Machine Learning*, 94(2), 133–149, 2014.

Braitenberg, Valentino. "Brain size and number of neurons: An exercise in synthetic neuroanatomy." *Journal of Computational Neuroscience*, 10(1), 71–77, 2001.

Braud, Raphael, Giagkos, Alexandros, Shaw, Patricia, Lee, Mark, and Shen, Qiang. "Building representations of proto-objects with exploration of the effect on fixation times." In *7th International Conference on Development and Learning and on Epigenetic Robotics, ICDL-EpiRob 2017*, 296–303. 2017.

Breazeal, Cynthia L. *Designing sociable robots*. MIT Press, 2002.

Bridle, James. *New dark age: Technology and the end of the future*. Verso Books, 2018.

Brockman, John. *What to think about machines that think*. Harper Perennial, 2015.

Brooks, Rodney A. "Intelligence without representation." *Artificial Intelligence*, 47 (1–3), 139–159, 1991.

Bruinsma, Anne Hendrik. *Practical robot circuits: Electronic sensory organs and nerve systems*. Philips Technical Library, 1960.

Bruner, J. S., Jolly, A., and Sylva, K. *Play: Its role in development and evolution*. Basic Books, 1976.

Bruner, Jerome Seymour. *Toward a theory of instruction*. Harvard University Press, 1966.

Bruner, Jerome Seymour. *Acts of meaning*. Harvard University Press, 1990.

Bushnell, Emily W., and Boudreau, J. Paul. "Motor development and the mind: The potential role of motor abilities as a determinant of aspects of perceptual development." *Child Development*, 64(4), 1005–1021, 1993.

Busso, Carlos, Deng, Zhigang, Yildirim, Serdar, et al. "Analysis of emotion recognition using facial expressions, speech and multimodal information." In *Proceedings of the 6th International Conference on Multimodal Interfaces*, 205–211. ACM, 2004.

Campbell, Robert L., and Bickhard, Mark H. "Types of constraints on development: An interactivist approach." *Developmental Review*, 12(3), 311–338, 1992.

Campos, Joseph J., Anderson, David I., Barbu-Roth, Marianne A., Hubbard, Edward M., Hertenstein, Matthew J., and Witherington, David. "Travel broadens the mind." *Infancy*, 1(2), 149–219, 2000.

Cangelosi, A., Metta, G., Sagerer, G., et al. "Integration of action and language knowledge: A roadmap for developmental robotics." *IEEE Transactions on Autonomous Mental Development*, 2(3), 167–195, 2010.

Cangelosi, Angelo, and Schlesinger, Matthew. *Developmental robotics: From babies to robots*. MIT Press, 2015.

Caporale, Natalia, and Dan, Yang. "Spike timing-dependent plasticity: a Hebbian learning rule." *Annual Review of Neuroscience*, 31, 25–46, 2008.

Charness, Neil. "Search in chess: Age and skill differences." *Journal of Experimental Psychology: Human Perception and Performance*, 7(2), 467, 1981.

Chatila, Raja, Renaudo, Erwan, Andries, Mihai, et al. "Toward self-aware robots." *Frontiers in Robotics and AI*, 5, 88, 2017.

Clark, Andy. *Being there—Putting brain, body, and world together again*. MIT Press, 1998.

Clark, Andy. "Whatever next? Predictive brains, situated agents, and the future of cognitive science." *Behavioral and Brain Sciences*, 36(3), 181–204, 2013.

Cohen, Jack, and Stewart, Ian. *The collapse of chaos: Discovering simplicity in a complex world*. Penguin UK, 1995.

Craik, Kenneth J. W. *The nature of explanation*. Cambridge University Press, 1943. Reprinted 1967.

Crick, Francis. *The astonishing hypothesis*. Simon and Schuster, 1994.

Damasio, Antonio. *The feeling of what happens: Body and emotion in the making of consciousness*. Vintage, 2000.

Daugherty, Paul R., and Wilson, H. James. *Human + machine: Reimagining work in the age of AI*. Harvard Business Review Press, 2018.

Dautenhahn, Kerstin. "Socially intelligent robots: dimensions of human–robot interaction." *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 362(1480), 679–704, 2007.

Davis, Ernest, and Marcus, Gary. "The scope and limits of simulation in automated reasoning." *Artificial Intelligence*, 233, 60–72, 2016.

Davis, Josh P., Lander, Karen, Evans, Ray, and Jansari, Ashok. "Investigating predictors of superior face recognition ability in police super-recognisers." *Applied Cognitive Psychology*, 30(6), 827–840, 2016.

de Garis, Hugo, Shuo, Chen, Goertzel, Ben, and Ruiting, Lian. "A world survey of artificial brain projects, part 1: Large-scale brain simulations." *Neurocomputing*, 74(1–3), 3–29, 2010.

de Groot, Adriaan D. *Thought and choice in chess*, vol. 4. de Gruyter, 2014.

de Latil, Pierre. *Thinking by machine: A study of cybernetics*. Trans. Y. M. Golla. Sidgwick and Jackson, 1956.

Dennett, Daniel C. "Reflections on language and mind." In Peter Carruthers and Jill Boucher, eds., *Language and thought: Interdisciplinary themes*, 284–294, Cambridge University Press, 1998.

Denning, Peter J., ed. *Talking back to the machine: Computers and human aspiration*. Springer-Verlag, 1999.

Der, Ralf, and Martius, Georg. *The playful machine—Theoretical foundation and practical realization of self-organizing robots*. Cognitive Systems Monographs (COSMOS). Springer-Verlag, 2011.

Derdikman, D., and Moser, E. I. "A manifold of spatial maps in the brain." *Trends in Cognitive Sciences*, 14(12), 561–569, 2010.

de Vries, Johanna I. P., Visser, Gerard H. A., and Prechtl, Heinz F. R. "The emergence of fetal behaviour. I. Qualitative aspects." *Early Human Development*, 7(4), 301–322, 1982.

de Waal, Frans. *Are we smart enough to know how smart animals are?* W. W. Norton & Company, 2016.

Dickmanns, Ernst Dieter. *Dynamic vision for perception and control of motion*. Springer Science & Business Media, 2007.

Domingos, Pedro. *The master algorithm: How the quest for the ultimate learning machine will remake our world*. Basic Books, 2015.

Dorigo, Marco, and Colombetti, Marco. *Robot shaping: An experiment in behavior engineering*. MIT Press, 1998.

Drescher, Gary L. *Made up minds: A constructivist approach to artificial intelligence*. MIT Press, 1991.

Earland, Kevin, Lee, Mark, Shaw, Patricia, and Law, James. "Overlapping structures in sensory-motor mappings." *PloS* ONE, 9(1), e84240, 2014.

Edelman, Gerald M. *Bright air, brilliant fire: On the matter of the mind*. Basic Books, 1992.

Ehrlich, Paul R. *The population bomb*. Sierra Club/Ballantine Books, 1968.

Elman, Jeffrey L. "Learning and development in neural networks: The importance of starting small." *Cognition*, 48(1), 71–99, 1993.

Fauconnier, Gilles. *Mappings in thought and language*. Cambridge University Press, 1997.

Flavell, J. H. *The developmental psychology of Jean Piaget*. van Nostrand, 1963.

Fong, Terrence, Nourbakhsh, Illah, and Dautenhahn, Kerstin. "A survey of socially interactive robots." *Robotics and Autonomous Systems*, 42(3–4), 143–166, 2003.

Forbus, Kenneth D. *Qualitative representations: How people reason and learn about the continuous world*. MIT Press, 2018.

Forestier, Sébastien, and Oudeyer, Pierre-Yves. "Modular active curiosity-driven discovery of tool use." In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3965–3972. IEEE, 2016.

Frase, Peter. *Four futures: Life after capitalism*. Verso Books, 2016.

Fritzke, Bernd. "A growing neural gas network learns topologies." *Advances in Neural Information-Processing Systems*, 7, 625–632, 1995.

Gallagher, Shaun. "Philosophical conceptions of the self: implications for cognitive science." *Trends in Cognitive Sciences*, 4(1), 14–21, 2000.

Gallese, Vittorio. "Embodied simulation: From neurons to phenomenal experience." *Phenomenology and the Cognitive Sciences*, 4(1), 23–48, 2005.

Gates, Bill. "A robot in every home." *Scientific American*, 296(1), 58–65, 2007.

Gerson, Sarah A., and Woodward, Amanda L. "The joint role of trained, untrained, and observed actions at the origins of goal recognition." *Infant Behavior and Development*, 37(1), 94–104, 2014.

Giagkos, A., Lewkowicz, D., Shaw, P., Kumar, S., Lee, M., and Shen, Q. "Perception of localized features during robotic sensorimotor development." *IEEE Transactions on Cognitive and Developmental Systems*, 9(2), 127–140, 2017; doi:10.1109/TCDS.2017.2652129.

Ginsburg, K. R., and the Committee on Communications, and the Committee on Psychosocial Aspects of Child and Family Health. "The importance of play in promoting healthy child development and maintaining strong parent-child bonds." *Pediatrics*, 119(1), 182, 2007.

Gobet, Fernand. *Understanding expertise: A multi-disciplinary approach*. Macmillan International Higher Education, 2015.

Godfrey-Smith, Peter. *Other minds: The octopus and the evolution of intelligent life*. William Collins, 2017.

Goldin-Meadow, Susan. "The role of gesture in communication and thinking." *Trends in Cognitive Sciences*, 3(11), 419–429, 1999.

Gómez, Gabriel, Lungarella, Max, Eggenberger Hotz, Peter, Matsushita, Kojiro, and Pfeifer, Rolf. "Simulating development in a real robot: on the concurrent increase of sensory, motor, and neural complexity." In L. Berthouze, H. Kozima, C. G. Prince, et al. eds., *Proceedings of the Fourth International Workshop on Epigenetic Robotics*, Lund University Cognitive Studies, 117, 1–24 2004.

Good, Irving John. "Speculations concerning the first ultraintelligent machine." *Advances in Computers*, 6(31), 31–88, 1965.

Goodfellow, Ian J., Shlens, Jonathon, and Szegedy, Christian. "Explaining and harnessing adversarial examples." *arXiv preprint, arXiv:1412.6572*, 2014.

Gottlieb, Jacqueline, Oudeyer, Pierre-Yves, Lopes, Manuel, and Baranes, Adrien. "Information-seeking, curiosity, and attention: Computational and neural mechanisms." *Trends in Cognitive Sciences*, 17(11), 585–593, 2013.

Grand, Steve. *Growing up with Lucy: How to Build an Android in Twenty Easy Steps*. Phoenix, 2004.

Greenspan, Ralph J., and Van Swinderen, Bruno. "Cognitive consonance: complex brain functions in the fruit fly and its relatives." *Trends in Neurosciences*, 27(12), 707–711, 2004.

Hamann, Katharina, Warneken, Felix, Greenberg, Julia R., and Tomasello, Michael. "Collaboration encourages equal sharing in children but not in chimpanzees." *Nature*, 476(7360), 328–331, 2011.

Haque, Usman. "The architectural relevance of Gordon Pask." *Architectural Design*, 77(4), 54–61, 2007.

Harari, Yuval Noah. *Homo Deus: A brief history of tomorrow*. Random House, 2016.

Harnad, Stevan. "The symbol grounding problem." *Physica D: Nonlinear Phenomena*, 42(1), 335–346, 1990.

Hawkins, Jeff, and Blakeslee, Sandra. *On intelligence: How a new understanding of the brain will lead to the creation of truly intelligent machines*. Macmillan, 2007.

Hoffmann, Matej, Marques, Hugo, Arieta, Alejandro, Sumioka, Hidenobu, Lungarella, Max, and Pfeifer, Rolf. "Body schema in robotics: a review." *IEEE Transactions on Autonomous Mental Development*, 2(4), 304–324, 2010.

Holland, Owen. "Exploration and high adventure: the legacy of Grey Walter." *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical, and Engineering Sciences*, 361(1811), 2085–2121, 2003.

Hsu, Feng-Hsiung. *Behind Deep Blue: Building the computer that defeated the world chess champion*. Princeton University Press, 2004.

Hu, Huosheng. "Biologically inspired design of autonomous robotic fish at Essex." In *Proceedings of the IEEE SMC UK-RI Chapter Conference on Advances in Cybernetic Systems*,, 1–8. 2006.

Hu, Yue, Eljaik, Jorhabib, Stein, Kevin, Nori, Francesco, and Mombaur, Katja D. "Walking of the iCub humanoid robot in different scenarios: Implementation and performance analysis." *Computing Research Repository, abs/1607.08525*, 2016.

Hubel, David H., and Wiesel, Torsten N. "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex." *Journal of Physiology*, 160(1), 106, 1962.

Hughes, M., and Hutt, C. "Heart-rate correlates of childhood activities: play, exploration, problem-solving and day-dreaming." *Biological Psychology*, 8(4), 253–263, 1979.

Humby, Clive, Hunt, Terry, and Phillips, Tim. *Scoring points: How Tesco continues to win customer loyalty*. Kogan Page Publishers, 2008.

Humphries, Mark D., Gurney, Kevin, and Prescott, Tony J. "The brainstem reticular formation is a small-world, not scale-free, network." *Proceedings of the Royal Society of London B: Biological Sciences*, 273(1585), 503–511, 2006.

Hunnius, Sabine, and Bekkering, Harold. "What are you doing? How active and observational experience shape infants' action understanding." *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1644), 20130490, 2014.

Husbands, Phil, and Holland, Owen. "Warren McCulloch and the British cyberneticians." *Interdisciplinary Science Reviews*, 37(3), 237–253, 2012.

Husbands, Philip, and Holland, Owen. "The Ratio Club: A hub of British cybernetics." In Philip Husbands, Owen Holland, and Michael Wheeler, eds., *The mechanical mind in history*, 91–148. MIT Press, 2008.

Husbands, Philip, Holland, Owen, and Wheeler, Michael, eds. *The mechanical mind in history*. MIT Press, 2008.

Hutchins, Edwin. *Cognition in the wild*. MIT Press, 1995.

Jamone, Lorenzo, Bernardino, Alexandre, and Santos-Victor, José. "Benchmarking the grasping capabilities of the iCub hand with the YCB object and model set." *IEEE Robotics and Automation Letters*, 1(1), 288–294, 2016.

Johnson, M., Schuster, M., Le, Q. V., et al. "Google's multilingual neural machine translation system: Enabling zero-shot translation." *arXiv e-prints, arXiv:1611.04558*, 2016.

Johnson, Mark H., and de Haan, Michelle. *Developmental cognitive neuroscience: An introduction*. 4th ed. John Wiley & Sons, 2015.

Kaas, Jon H. "Topographic maps are fundamental to sensory processing." *Brain Research Bulletin*, 44(2), 107–112, 1997.

Kahneman, Daniel. *Thinking, fast and slow*. Penguin Books, 2012.

Karmiloff-Smith, Annette. *Beyond modularity: A developmental perspective on cognitive science*. MIT Press, 1995.

Kasparov, Garry. *Deep thinking: Where machine intelligence ends and human creativity begins*. John Murray, 2017.

Keil, Frank C. "Constraints on constraints: Surveying the epigenetic landscape." *Cognitive Science*, 14(1), 135–168, 1990.

Keysers, Christian, and Gazzola, Valeria. "Social neuroscience: Mirror neurons recorded in humans." *Current Biology*, 20(8), R353–R354, 2010.

King, Elizabeth. "Clockwork prayer: A sixteenth-century mechanical monk." *Blackbird: An Online Journal of Literature and the Arts*, 1(1), 1–29, 2002.

Kirby, Rachel, Forlizzi, Jodi, and Simmons, Reid. "Affective social robots." *Robotics and Autonomous Systems*, 58(3), 322–332, 2010.

Klatzky, Roberta L. "Allocentric and egocentric spatial representations: Definitions, distinctions, and interconnections." In C. Freksa, C. Habel, and K. F. Wender, eds., *Spatial cognition*. *Lecture Notes in Computer Science* 1404, 1–17. Springer, 1998.

Kober, Jens, Bagnell, J. Andrew, and Peters, Jan. "Reinforcement learning in robotics: A survey." *International Journal of Robotics Research*, 32(11), 1238–1274, 2013.

Kozima, Hideki, Michalowski, Marek P., and Nakagawa, Cocoro. "Keepon: A playful robot for research, therapy, and entertainment." *International Journal of Social Robotics*, 1(1), 3–18, 2009.

Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. "ImageNet classification with deep convolutional neural networks." In P. Bartlett, F. C. N. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, eds., *Advances in neural information processing systems*, 1097–1105. Curran Associates, 2012.

Kumar, Suresh, Shaw, Patricia, Giagkos, Alexandros, Braud, Raphael, Lee, Mark, and Shen, Qiang. "Developing hierarchical schemas and building schema chains through practice play behavior." *Frontiers in Neurorobotics*, 12(33), 2018; doi:10.3389/fnbot.2018.00033.

Kuniyoshi, Y., Yorozu, Y., Suzuki, S., et al. "Emergence and development of embodied cognition: A constructivist approach using robots." *Progress in Brain Research*, 164, 425–445, 2007.

Kuniyoshi, Yasuo, and Sangawa, Shinji. "Early motor development from partially ordered neural-body dynamics: experiments with a cortico-spinal-musculo-skeletal model." *Biological Cybernetics*, 95(6), 589–605, 2006.

Lake, Brenden M., Lawrence, Neil D., and Tenenbaum, Joshua B. "The emergence of organizing structure in conceptual representation." Special Issue: Memory, Learning, and Expertise. *Cognitive Science*, 42(S3), 809–832, 2018.

Lake, Brenden M., Ullman, Tomer D., Tenenbaum, Joshua B., and Gershman, Samuel J. "Building machines that learn and think like people." *Behavioral and Brain Sciences*, 40, e253, 2017; doi:10.1017/S0140525X16001837.

Lakoff, George. *Women, fire, and dangerous things: What categories reveal about the mind*. University of Chicago Press, 1987.

Law, James, Lee, Mark, Hülse, Martin, and Tomassetti, Alessandra. "The infant development timeline and its application to robot shaping." *Adaptive Behavior*, 19(5), 335–358, 2011.

Law, James A., Shaw, Patricia H., Earland, Kevin G., Sheldon, Michael T., and Lee, Mark H. "A psychology based approach for longitudinal development in cognitive robotics." *Frontiers in Neurorobotics*, 8(1), 1–19, 2014a; doi:10.3389/fnbot.2014.00001.

Law, James A., Shaw, Patricia H., and Lee, Mark H. "A biologically constrained architecture for developmental learning of eye-head gaze control on a humanoid robot." *Autonomous Robots*, 35(1), 77–92, 2013. doi:10.1007/s10514-013-9335-2.

Law, James A., Shaw, Patricia H., Lee, Mark H., and Sheldon, Michael T. "From saccades to grasping: A model of coordinated reaching through simulated development on a humanoid robot." *IEEE Transactions on Autonomous Mental Development*, 6(2), 93–109, 2014b; doi:10.1109/TAMD.2014.2301934.

LeCun, Yann, Bengio, Yoshua, and Hinton, Geoffrey. "Deep learning." *Nature*, volume 521(7553), 436–444, 2015.

LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. "Gradient-based learning applied to document recognition." *Proceedings of the IEEE*, 86(11), 2278–2324, 1998.

Lee, Mark. "Intrinsic activitity: from motor babbling to play." In *Prooceedings of the IEEE Joint International Conference on Development and Learning (ICDL) and Epigenetic Robotics (EpiRob)* 2, 1–6, 2011.

Lee, Mark, Law, James, and Hülse, Martin. "A developmental framework for cumulative learning robots." In *Computational and robotic models of the hierarchical organization of behavior*, 177–212. Springer, 2013.

Legg, Shane. "Is there an elegant universal theory of prediction?" In *Algorithmic Learning Theory, 17th International Conference*, 274–287. Springer, 2006.

Legg, Shane. *Machine super intelligence*. Ph.D thesis, University of Lugano, Switzerland, 2008.

Legg, Shane, and Hutter, Marcus. "A collection of definitions of intelligence." In B. Goertzel and P. Wang, eds., *Advances in artificial general intelligence*, vol. 157, 17–24. IOS Press, 2007.

Lepora, Nathan F., Martinez-Hernandez, Uriel, Evans, Mathew, Natale, Lorenzo, Metta, Giorgio, and Prescott, Tony J. "Tactile superresolution and biomimetic hyperacuity." *IEEE Transactions on Robotics*, 31(3), 605–618, 2015.

Levesque, Hector J., Davis, Ernest, and Morgenstern, Leora. "The Winograd Schema Challenge." In *13th International Conference on the Principles of Knowledge Representation and Reasoning, KR 2012*, 552–561. 2012.

Levine, Sergey, Pastor, Peter, Krizhevsky, Alex, and Quillen, Deirdre. "Learning hand-eye coordination for robotic grasping with Deep Learning and large-scale data collection." *Computing Research Repository (CoRR), arXiv:1603.02199*, 2016.

Lewis, Marc D. "The promise of dynamic systems approaches for an integrated account of human development." *Child Development*, 71(1), 36–43, 2000.

Lewis, Michael. *The Undoing Project: A friendship that changed our minds*. W. W. Norton, 2016.

Long, John. *Darwin's devices: What evolving robots can teach us about the history of life and the future of technology*. Basic Books, 2012.

Lungarella, M., and Berthouze, L. "Adaptivity through physical immaturity." In *Proceedings of the Second International Workshop on Epigenetic Robotics*, Lund University Cognitive Studies, 94, 79–86, 2002.

Lungarella, Max, Metta, Giorgio, Pfeifer, Rolf, and Sandini, Giulio. "Developmental robotics: a survey." *Connection Science*, 15(4), 151–190, 2003.

MacKay, Donald M. "Mind-like behaviour in artefacts." *Bulletin of the British Society for the History of Science*, 1(S5), 164–165, 1951.

MacKay, Donald M. "On comparing the brain with machines." *American Scientist*, 261–268, 1954.

MacKay, Donald M. *Information, mechanism, and meaning*. MIT Press, 1969.

Maddison, Chris J., Huang, Aja, Sutskever, Ilya, and Silver, David. "Move evaluation in go using deep convolutional neural networks." *arXiv preprint arXiv:1412.6564*, 2014.

Mahler, Jeffrey, Liang, Jacky, Niyaz, Sherdil, et al. "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics." *arXiv preprint arXiv:1703.09312*, 2017.

Mareschal, Denis, Johnson, Mark H., Sirois, Sylvain, Spratling, Michael, Thomas, Michael S. C., and Westermann, Gert. *Neuroconstructivism vol. I: How the brain constructs cognition*. Oxford University Press, 2007a.

Mareschal, Denis, Sirois, Sylvain, Westermann, Gert, and Johnson, Mark H. *Neuroconstructivism vol. II: Perspectives and prospects*. Oxford University Press, 2007b.

Marino, Lori. "Thinking chickens: A review of cognition, emotion, and behavior in the domestic chicken." *Animal Cognition*, 20, 127–147, 2017; doi:10.1007/s10071-016-1064-4.

Mason, Paul. *Postcapitalism: A guide to our future*. Macmillan, 2016.

Matin, Leonard, and Li, Wenxun. "Multimodal basis for egocentric spatial localization and orientation." *Journal of Vestibular Research: Equilibrium & Orientation*, 5(6), 499–518, 1995.

Maturana, Humberto R., and Varela, Francisco J. *Autopoiesis and cognition: The realization of the living*. Springer Science & Business Media, 1991.

McCarthy, J., Minsky, M. L., and Rochester, N. *A proposal for the Dartmouth Summer Research Project on Artificial Intelligence.* Technical report, Dartmouth College, Hanover, NH, available from http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf, 1955.

McColl, Derek, Hong, Alexander, Hatakeyama, Naoaki, Nejat, Goldie, and Benhabib, Beno. "A survey of autonomous human affect detection methods for social robots engaged in natural HRI." *Journal of Intelligent & Robotic Systems*, 82(1), 101–133, 2016.

McCorduck, Pamela. *Machines who think*. W. H. Freeman & Co., 1979, and 2nd edition 2004.

McCulloch, Warren S., and Pitts, Walter H. "A logical calculus of the ideas immanent in nervous activity." *Bulletin of Mathematical Biophysics* 5, 115–133, 1943.

Metta, G., Sandini, G., Vernon, D., Natale, L., and Nori, F. "The iCub humanoid robot: An open platform for research in embodied cognition." In *Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems*, ACM, 50–56. ACM, 2008.

Miikkulainen, Risto, Bednar, James A., Choe, Yoonsuck, and Sirosh, Joseph. *Computational maps in the visual cortex*. Springer Nature, 2006.

Minsky, Marvin L. "Logical versus analogical or symbolic versus connectionist or neat versus scruffy." *AI Magazine*, 12(2), 34–51, 1991.

Mnih, Volodymyr, Kavukcuoglu, Koray, Silver, David, et al. "Human-level control through deep reinforcement learning." *Nature*, 518(7540), 529–533, 2015.

Moorehead, Alan. *The Fatal impact: An account of the invasion of the South Pacific, 1767–1840*. Harper & Row, 1966.

Mori, Masahiro. "The uncanny valley." Translated by Karl F. MacDorman and Takashi Minato. *Energy*, 7(4), 33–35, 1970.

Morton, A. Jennifer, and Avanzo, Laura. "Executive decision-making in the domestic sheep." *PLoS ONE*, 6(1), e15752, 2011.

Muhammad, Wasif, and Spratling, Michael W. "A neural model of binocular saccade planning and vergence control." *Adaptive Behavior*, 23(5), 265–282, 2015.

Muhammad, Wasif, and Spratling, Michael W. "A neural model of coordinated head and eye movement control." *Journal of Intelligent & Robotic Systems*, 85(1), 107–126, 2017.

Müller, Vincent C., and Bostrom, Nick. "Future progress in artificial intelligence: A survey of expert opinion." In Vincent C Müller, ed., *Fundamental issues of artificial intelligence*, 376, 555–572. Springer, 2016.

Natale, Lorenzo, Nori, Francesco, Metta, Giorgio, et al. "The iCub platform: A tool for studying intrinsically motivated learning." In Gianluca Baldassarre and Marco Mirolli, eds., *Intrinsically motivated learning in natural and artificial systems*, 433–458. Springer Berlin Heidelberg, 2013; doi:10.1007/978-3-642-32375-1_17.

Nemes, T. N. *Cybernetic machines*. Iliffe Books Ltd., 1969. Translated from the 1962 Hungarian edition by I. Foldes.

Nguyen, Anh, Yosinski, Jason, and Clune, Jeff. "Deep neural networks are easily fooled: High confidence predictions for unrecognizable images." *arXiv preprint arXiv:1412.1897*, 2014.

O'Neil, Cathy. *Weapons of math destruction: How Big Data increases inequality and threatens democracy*. Crown Publishing Group (NY), 2016.

Oudeyer, Pierre-Yves, Baranes, Adrien, and Kaplan, Frédéric. "Intrinsically motivated learning of real-world sensorimotor skills with developmental constraints." In Gianluca Baldassarre and Marco Mirolli, eds., *Intrinsically motivated learning in natural and artificial systems*, 303–365. Springer, 2013.

Oudeyer, Pierre-Yves, Kaplan, F., and Hafner, V. V. "Intrinsic motivation systems for autonomous mental development." *IEEE Transactions on Evolutionary Computation*, 11(2), 265–286, 2007.

Pellis, Sergio, and Pellis, Vivien. *The playful brain: Venturing to the limits of neuroscience*. Oneworld, 2009.

Pfeifer, Rolf, and Bongard, Josh. *How the body shapes the way we think: A new view of intelligence*. MIT Press, 2006.

Piaget, J. *Play, dreams, and imitation in childhood*. Routledge, 1999.

Piek, J. P., and Carman, R. "Developmental profiles of spontaneous movements in infants." *Early Human Development*, 39(2), 109–126, 1994.

Pierson, Harry A., and Gashler, Michael S. "Deep learning in robotics: A review of recent research." *Computing Research Repository (CoRR), arXiv:1707.07217*, 2017.

Pinker, Steven. *The better angels of our nature: The decline of violence in history and its causes*. Penguin, 2011.

Pinker, Steven. *Enlightenment now: The case for reason, science, humanism, and progress*. Penguin, 2018.

Plummer, Robert Patrick. *A computer program which simulates sensorimotor learning*. Ph.D. thesis, the University of Texas at Austin, 1970.

Plummer, Robert Patrick. "A sensorimotor learning program." In *Proceedings of the 5th Australian Computer Conference*, 617–622. Australian Computer Society, 1972.

Pomerleau, Dean. "Neural network vision for robot driving." In M. H. Hebert, C. Thorpe, and A. Stentz, eds., *Intelligent unmanned ground vehicles*, 53–72. Springer, 1997.

Popper, Karl, Ryan, Alan, and Gombrich, E. H. *The open society and its enemies*. Princeton University Press, 2013.

Prince, C. G., Helder, N. A., and Hollich, G. J. "Ongoing emergence: A core concept in epigenetic robotics." In *Proceedings of the 5th International Workshop on Epigenetic Robotics*, Lund University Cognitive Studies, 63–70, 2005.

Purves, Dale, and Lichtman, Jeff W. "Elimination of synapses in the developing nervous system." *Science*, 210(4466), 153–157, 1980.

Quinlan, Philip T. *Connectionist models of development: Developmental processes in real and artificial neural networks*. Taylor & Francis, 2003.

Rizzolatti, Giacomo, Fadiga, Luciano, Gallese, Vittorio, and Fogassi, Leonardo. "Premotor cortex and the recognition of motor actions." *Cognitive Brain Research*, 3(2), 131–141, 1996.

Rochat, Philippe, and Striano, Tricia. "Perceived self in infancy." *Infant Behavior and Development*, 23(3–4), 513–530, 2000.

Rolf, Matthias, and Steil, Jochen J. "Goal babbling: A new concept for early sensorimotor exploration." In E. Ugur, Y. Nagai, E. Oztop, and M. Asada, eds., *Proceedings of IEEE International Conference on Humanoid Robots, Humanoids 2012 Workshop on Developmental Robotics*, 40–43, 2012.

Rosling, Hans, Rönnlund, Anna Rosling, and Rosling, Ola. *Factfulness: Ten reasons we're wrong about the world and why things are better than you think*. Flatiron Books, 2018.

Russell, Stuart, Dewey, Daniel, and Tegmark, Max. "Research priorities for robust and beneficial artificial intelligence." *AI Magazine*, 36(4), 105–114, 2015.

Russell, Stuart, and Norvig, Peter. *Artificial intelligence: A modern approach*. 3rd ed. Pearson, 2010.

Rutkowska, J. C. "Scaling up sensorimotor systems: Constraints from human infancy." *Adaptive Behaviour*, 2, 349–373, 1994.

Samuel, Arthur L. "Some studies in machine learning using the game of checkers." *IBM Journal of Research and Development*, 3(3), 210–229, 1959.

Samuel, Arthur L. "Some studies in machine learning using the game of checkers. II: Recent progress." *IBM Journal of Research and Development*, 11(6), 601–617, 1967.

Sandamirskaya, Yulia, and Storck, Tobias. "Neural-dynamic architecture for looking: Shift from visual to motor target representation for memory saccades." In *4th International Conference on Development and Learning and on Epigenetic Robotics*, 34–40. IEEE, 2014.

Schaeffer, Jonathan, Burch, Neil, Björnsson, Yngvi, et al. "Checkers is solved." *Science*, 317(5844), 1518–1522, 2007.

Schmidhuber, Jürgen. "Deep learning in neural networks: An overview." *Neural Networks*, 61, 85–117, 2015.

Shaw, Patricia, Law, James, and Lee, Mark. "Representations of body schemas for infant robot development." In *5th International Conference on Development and Learning and on Epigenetic Robotics*, 123–128. IEEE Press, 2015.

Shaw, Patricia H., Law, James A., and Lee, Mark H. "A comparison of learning strategies for biologically constrained development of gaze control on an iCub robot." *Autonomous Robots*, 37(1), 97–110, 2014; doi:10.1007/s10514-013-9378-4.

Sheldon, M. T. *Intrinsically motivated developmental learning of communication in robotic agents*. Ph.D thesis, Aberystwyth University, Wales, 2012.

Sheldon, M. T., and Lee, M. "PSchema: A developmental schema learning framework for embodied agents." In *Proceedings of the IEEE Joint International Conference on Development and Learning (ICDL) and Epigenetic Robotics*. IEEE, 2011.

Sherwood, Stephen L., ed. *The nature of psychology: A selection of papers, essays, and other writings by Kenneth J. W. Craik*. Cambridge University Press, 1966.

Shneiderman, Ben. *Leonardo's laptop: Human needs and the new computing technologies*. MIT Press, 2003.

Shultz, Thomas R. *Computational developmental psychology*. MIT Press, 2003.

Shultz, Thomas R. "A constructive neural-network approach to modeling psychological development." *Cognitive Development*, 27(4), 383–400, 2012.

Silver, David, Huang, Aja, Maddison, Chris J., et al. "Mastering the game of Go with deep neural networks and tree search." *Nature*, 529(7587), 484–489, 2016.

Silver, David, Hubert, Thomas, Schrittwieser, Julian, et al. "Mastering chess and shogi by self-play with a general reinforcement learning algorithm." *arXiv preprint arXiv:1712.01815*, 2017a.

Silver, David, Schrittwieser, Julian, Simonyan, Karen, et al. "Mastering the game of Go without human knowledge." *Nature*, 550(7676), 354, 2017b.

Smith, Linda, and Gasser, Michael. “The development of embodied cognition: Six lessons from babies.” *Artificial Life* 11(1–2), 13–29, 2005.

Smith, Linda B., and Thelen, Esther. “Development as a dynamic system.” *Trends in Cognitive Sciences*, 7(8), 343–348, 2003.

Smith, Peter K. *Children and play: Understanding children’s worlds*, vol. 12. John Wiley & Sons, 2009.

Smith, Roger. “James Cox’s silver swan: An eighteenth-century automaton in the Bowes Museum.” *Artefact: Techniques, histoire et sciences humaines*, 4(4), 361–365, 2016.

Snyder, Lawrence H., Grieve, Kenneth L., Brotchie, Peter, and Andersen, Richard A. “Separate body-and world-referenced representations of visual space in parietal cortex.” *Nature*, 394(6696), 887, 1998.

Spelke, Elizabeth S. “Principles of object perception.” *Cognitive Science*, 14(1), 29–56, 1990.

Sporns, O., and Edelman, G. M. “Solving Bernstein’s problem: A proposal for the development of coordinated movement by selection.” *Child Development*, 64(4), 960–981, 1993.

Standing, Guy. *The corruption of capitalism: Why rentiers thrive and work does not pay*. Biteback Publishing, 2016.

Steels, Luc. *The Talking Heads Experiment: Origins of words and meanings*. Language Science Press, 2015.

Stork, David G. *HAL’s legacy: 2001’s computer as dream and reality*. MIT Press, 1998.

Sutherland, Stuart. *Irrationality: The enemy within*. Constable and Company, 1992.

Szegedy, Christian, Zaremba, Wojciech, Sutskever, Ilya, et al. “Intriguing properties of neural networks.” *arXiv preprint arXiv:1312.6199*, 2013.

Tainter, Joseph. *The collapse of complex societies*. Cambridge University Press, 1988.

Tallis, Raymond. *The hand: A philosophical inquiry into human being*. Edinburgh University Press, 2003.

Tang, Yi-Yuan, Rothbart, Mary K., and Posner, Michael I. “Neural correlates of establishing, maintaining, and switching brain states.” *Trends in Cognitive Sciences*, 16(6), 330–337, 2012.

Terada, Yuuzi, and Yamamoto, Ikuo. “An animatronic system including lifelike robotic fish.” *Proceedings of the IEEE*, 92(11), 1814–1820, 2004.

Thaler, Richard H., and Sunstein, Cass R. *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press, 2008.

Thelen, E. “Grounded in the world: Developmental origins of the embodied mind.” *Infancy*, 1(1), 3–28, 2000.

Thelen, Esther. “Rhythmical stereotypies in normal human infants.” *Animal Behaviour*, 27, 699–715, 1979.

Tikhanoff, Vadim, Cangelosi, Angelo, and Metta, Giorgio. “Integration of speech and action in humanoid robots: iCub simulation experiments.” *IEEE Transactions on Autonomous Mental Development*, 3(1), 17–29, 2011.

Tomasello, M., and Herrmann, E. “Ape and human cognition: What’s the difference?” *Current Directions in Psychological Science*, 19(1), 3–8, 2010.

Topol, Eric J., and Hill, Dick. *The creative destruction of medicine: How the digital revolution will create better health care*. Basic Books, 2012.

Triantafyllou, Michael S., and Triantafyllou, George S. “An efficient swimming machine.” *Scientific American*, 272(3), 64–71, 1995.

Tronick, E. “Stimulus control and the growth of the infant’s effective visual field.” *Perception and Psychophysics*, 11(5), 373–376, 1972.

Turing, A. M. “Computing machinery and intelligence.” *Mind*, 59(236), 433–460, 1950.

Turing, Alan Mathison, and Copeland, B. Jack. *The essential Turing: Seminal writings in computing, logic, philosophy, artificial intelligence, and artificial life, plus the secrets of Enigma*. Oxford University Press, 2004.

Turkewitz, Gerald, and Kenny, Patricia A. “Limitations on input as a basis for neural organization and perceptual development: A preliminary theoretical statement.” *Developmental psychobiology*, 15(4), 357–368, 1982.

Ugur, Emre, Nagai, Yukie, Sahin, Erol, and Oztop, Erhan. “Staged development of robot skills: Behavior formation, affordance learning, and imitation with motionese.” *IEEE Transactions on Autonomous Mental Development*, 7(2), 119–139, 2015.

Vernon, David. “Enaction as a conceptual framework for developmental cognitive robotics.” *Paladyn*, 1(2), 89–98, 2010.

Vernon, David, Metta, Giorgio, and Sandini, Giulio. “The iCub cognitive architecture: Interactive development in a humanoid robot.” In *2007 IEEE 6th International Conference on Development and Learning*, 122–127. IEEE, 2007.

Vernon, David, von Hofsten, Claes, and Fadiga, Luciano. *A roadmap for cognitive development in humanoid robots*. Vol. 11 of *Cognitive Systems Monographs (COSMOS)*. Springer, 2010.

von Hofsten, C. “An action perspective on motor development.” *Trends in Cognitive Sciences*, 8(6), 266–272, 2004.

Wainer, Joshua, Dautenhahn, Kerstin, Robins, Ben, and Amirabdollahian, Farshid. "Collaborating with Kaspar: Using an autonomous humanoid robot to foster cooperative dyadic play among children with autism." In *Humanoid Robots–10th IEEE-RAS International Conference*, 631–638. IEEE Robotics and Automation Society, 2010.

Walter, W. Grey. *The living brain*. Gerald Duckworth & Co, 1953.

Walter, W. Grey. "The brain as a machine." *Proceedings of the Royal Society of Medicine*, 50(10), 799–808, 1957.

Webb, William. *Our digital future: Smart analysis of smart technology*. CreateSpace Independent Publishing Platform, 2017.

Westermann, G., and Mareschal, D. "From parts to wholes: Mechanisms of development in infant visual object processing." *Infancy*, 5(2), 131–151, 2004.

Westheimer, Gerald. "Optical superresolution and visual hyperacuity." *Progress in Retinal and Eye Research*, 31(5), 467–480, 2012.

White, B. L., Castle, P., and Held, R. "Observations on the development of visually-directed reaching." *Child Development*, 35(2), 349–364, 1964.

Wiener, Norbert. *Cybernetics: or control and communication in the animal and the machine*. 2nd ed. MIT Press, 1961.

Wilkins, David. "Using chess knowledge to reduce search." In Peter W. Frey, ed., *Chess skill in man and machine*, 211–242. Springer, 1983.

Wilmott, Paul. "The use, misuse and abuse of mathematics in finance." *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 358(1765), 63–73, 2000.

Wilmott, Paul. *Frequently asked questions in quantitative finance*. John Wiley & Sons, 2010.

Wolpert, Daniel M. "Probabilistic models in human sensorimotor control." *Human Movement Science*, 26(4), 511–524, 2007.

Woodward, Amanda L. "Infants' grasp of others' intentions." *Current Directions in Psychological Science*, 18(1), 53–57, 2009.

Yamada, Yasunori, Mori, Hiroki, and Kuniyoshi, Yasuo. "A Fetus and infant developmental scenario: Self-organization of goal-directed behaviors based on sensory constraints." In *Proceedings of 10th International Conference on Epigenetic Robotics. Lund University Cognitive Studies*, 149, 145–152. 2010.

Yosinski, Jason, Clune, Jeff, Nguyen, Anh, Fuchs, Thomas, and Lipson, Hod. "Understanding neural networks through deep visualization." *arXiv preprint arXiv: 1506.06579*, 2015.

Zeiler, Matthew D., and Fergus, Rob. "Visualizing and understanding convolutional networks." In *Computer Vision—European Conference on Computer Vision, 2014*, 818–833. Springer, 2014.

Zeki, Semir. *A vision of the brain*. Oxford University Press, 1993; now available at www.vislab.ucl.ac.uk.

Zuboff, Shoshana. *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile Books, 2019.