Artificial Intelligence and The Environmental Crisis

Can Technology Really Save the World?

Keith R Skene
Director, Biosphere Research Institute
Angus, Scotland, UK

Routledge
Taylor & Francis Group
A SCIENCE PUBLISHERS BOOK

Published by Routledge

2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

ISBN: 978-0-367-15200-0 (Hardback)

Contents

Preface

Section I: Artificial Intelligence and the Internet of Things

I.1. Nothing new under the Sun
I.2. Oh, for a nice cold soda: The birth of the internet of things
I.3. The two-month, ten-man project to transform the world
I.4. Getting to grips with the jargon: Symbolic and non-symbolic AI

Section II: Should I Stay or Should I Go? Ethics in AI

II.1. Choosing an ethical framework
II.2. The strange case of Asimov's laws
II.3. Free will and moral judgement
II.4. The confused owl of Minerva: Dangers of a moral vacuum
II.5. Who's in charge of the big bad wolf?
II.6. What should a declaration of AI rights look like?

Section III: Gender, Race, Culture and Fear

III.1. Gender issues in AI
III.2. Racial issues in AI
III.3. Cultural issues in AI
III.4. Fear and loathing in AI

Section IV: The Thinker: Human Intelligence

IV.1. Human intelligence: Carolus Linnaeus and his wise, wise men
IV.2. So what is human intelligence?
IV.3. Philosophy and intelligence: The framing of our thoughts

Section V: Other Modes of Intelligence: Thinking Outside the Human Box

V.1. Animal intelligence: Machiavellian sentience and the wisdom of the swarm
V.2. Plant intelligence: Headless, brainless, dispersed intelligence
V.3. Microbial intelligence: Gene-swapping revelry in the quorum

V.4. Ecosystem intelligence: Systems thinking in the cathedral of thought
V.5. Systems are non-linear
V.6. Systems are emergent
V.7. Systems are sub-optimal
V.8. Systems rely on real-time feedback

Section VI: Highway to Hell: The Existentialist Threat Facing Humankind

VI.1. A brief history of our path towards destruction
VI.2. The five clear road signs that point towards criticality
VI.3. Why ecological damage matters to us
VI.4. Adam Smith and his invisible hand
VI.5. Kuznets and his curve: How ninety five percent speculation led us badly astray

Section VII: Forget the Romans. What has AI Ever Done for Us?

VII.1. AI and economics: The best of things or the worst of things?
VII.2. AI and society
VII.3. AI and the environment
VII.4. Technology and sustainability: Bellicose bedfellows or Romeo and Juliet?

Section VIII: Imagining a New World

VIII.1. The swallow whose nest was stolen: A salutary tale
VIII.2. Blinded by the bling: Dashboard dogs and a disappearing sea
VIII.3. What needs changed and what change do we need?
VIII.4. The chains that bind: Taking responsibility for our footprints
VIII.5. The Ogiek people and the new, improved invisible hand
VIII.6. Lessons from the edge of the world: The St Kildan legacy
VIII.7. The Garden of Eden complex: How not to fix the world
VIII.8. The three cornerstones: Diversity, resilience and integration
VIII.9. The central role of AI in feedback: Shaping our new world

Section IX: Barriers to Change

IX.1. The five philosophical barriers
IX.2. Structural barriers to change
IX.3. The seven dragons: Psychological barriers
IX.4. How AI can help overcome these barriers

Section X: Transition

X.1. The nature of transition
X.2. Studies in transition
X.3. Why societal change is key

X.4. How to manage societal change
X.5. Requiem for the King of Phrygia

Glossary

References

Index

SECTION I

Artificial Intelligence and the Internet of Things

> *"The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it."*
>
> –Mark Weiser, *Scientific American* (Weiser, 1991)

Welcome! Diversity is a theme that will run through this entire book, as we explore the importance of this key concept in society, ecology and technology. Diversity lies at the heart of resilience and creativity, broadening the available solution space. Indeed, our ability to solve problems is more likely to be limited by a lack of diversity in thinking than by limitations imposed by the problems themselves. Perhaps the greatest challenge for computer technologists is how to embed and maintain diversity into artificial intelligence (AI), in a globalized, connected infosphere. The infosphere is the new technological world, where information, in digital form, creates another dimension of reality. This virtual space is increasingly affecting the 'real' world, and the two are beginning to merge.

All of us interact with the world around us in different ways as an outcome of our genes, our journey and our landscape. We also have different levels of knowledge on the subject matter of this book. Some will be very aware of artificial intelligence, sustainability, ecology and social science, others less so. This section is designed to bring everyone to a similar point at the outset. If you are already comfortable with areas such as symbolic and non-symbolic approaches, the internet of things and the issues surrounding whether intelligence can be artificial at all, then feel free to skip this section. Otherwise, read on.

In the next few pages, you will travel back in time to the earliest conception of artificial intelligence, some three thousand years ago, in China and Greece, before relocating to a soda machine in a Brutalist university building in Pittsburgh, the birthplace of the internet of things. Then we take a whistle top tour of the terminology surrounding artificial intelligence, while at the same time following its development. We'll meet some of the key characters along the way, from King Mu of Zhou to Marvin Minsky. We trace the euphoric rise, the barren AI winter and the resurgence of a field that most of us have only really viewed through the lens of a Hollywood blockbuster movie camera. So, let's make a start.

I.1. Nothing new under the Sun

We think of artificial intelligence as a really modern phenomenon, at the cutting edge of technology, representing the dawn of a brave new world, or an apocalyptic vision of robotic dominion and terror. Yet the concept of artificial intelligence actually reaches much further back in time, some three thousand years before present. Almost as soon as written records started appearing, humans have envisioned machines that could think, feel and act just like we can. Two separate strands of evidence show that, early in recorded history, people could imagine artificial life forms. The first of these strands comes from Ancient China, while the second emerged from Ancient Greece.

The *Liezi*, sometimes referred to as the *True Scripture of Excellence*, is a revered Daoist scripture and philosophical treatise, viewed as one of the three most important texts in Taoism, along with the Laozi and the Zhuangzi. Written around two thousand five hundred years ago, and recomposed one thousand years later, it consisted of around one hundred and forty parables. The Liezi is accredited to Li Yukou, the Daoist philosopher. Daoism emphasises the importance of achieving a resonance between our inner being and the outer, natural world, leading to a holistic integration.

While some question if Li Yukou ever actually existed, and others question if he wrote the text, or if it was merely a collection of plagiarised passages from other classic texts, the fact is that the book itself exists and is of great interest at many levels. However, in relation to our purposes, it is a story revolving around the wonderfully-named King Mu of Zhou that is of interest.

King Mu was a real king who ruled China around 950 BC. In Book Five of the Liezi, the story is told of the king meeting an engineer called Yen Shih, who had built a human-like machine that could walk, sing and dance. The king was really impressed, until the robot made advances at his concubines and winked and cavorted with them. King Mu ordered the execution of the engineer, and he was only saved when he destroyed the robot. We note in passing that this early account of an intelligent machine was laced with lust and jealousy. It is also of interest that King Mu did not hold the automaton responsible for its desire and depravity, but rather the engineer, who was sentenced to death, briefly. This will become of relevance when we consider ethical accountability, later in this book.

The second early conception of artificial intelligence occurs in the work of Homer. Questions exist over whether or not Homer even existed as well as whether or not he wrote the *Iliad* and *Odyssey*, the two great classic works of Ancient Greece. However, we do know that Troy existed, and that the Iliad was written in the 8th century BC. The dating of this work is as interesting as much of the content. In a similar way that genetic mutations through time have been used to elucidate when species separated from other species, so linguistic experts have been able to trace changes in language over time, i.e. linguistic mutations, and determine when particular works of literature were written. According to Eric Altschuler and colleagues, the Iliad was written sometime between 710–760 BC (Altschuler et al., 2013). It is considered to be the first work of the Western literary tradition.

The Iliad focuses on events surrounding the fall of Troy, an ancient city in what is now North West Turkey. The fall of Troy itself is generally thought to have occurred in the twelfth century BC. Of particular interest to us is an account of the banqueting hall of the gods, and a series of handmaids made out of gold. Mounted on tripods with gold wheels, they served the tables of the gods at a great feast. They were made by Hephaestus, the blacksmith god who also created Achilles' armour. These structures were more than mere golden statues. Not only could they move around and serve food independently, but they had the ability to speak. Homer states that they had "*understanding in their hearts*" (Iliad XVIII, 417-425). These were machines with understanding in their hearts. The concept of an advanced artificial intelligence had been conceived of in Ancient Greece as well as in China, almost three thousand years ago.

As Pablo Picasso, the Cubist painter, is reported to have exclaimed (though this is disputed), when viewing the ancient cave art in Lascaux in 1940, with its clever use of twisted perspective, for which Picasso was famous, "*We have invented nothing*" (Graff, 2006). As a noteworthy aside, the caves at Lascaux were discovered by some young local schoolchildren when their dog, Robot (indeed!) disappeared down a hole near the village they lived in. The children followed, in search of the dog, and found themselves in a suite of caverns, surrounded by vast, painted panoramas, unseen for thousands of years.

In the *Odyssey*, the second of Homer's epic poems, that describes the return of the conquering heroes from Troy, another example of artificial intelligence is found, but without the issues of lust or jealousy. Homer writes of the boats of the Phaeacians, a hedonistic people from the modern-day island of Corfu: "*Phaeacian ships have no helmsman or steering oar, for the ships themselves know our thoughts and wishes, and the cities of men, every fertile country, and hidden by mist and cloud they speed over the sea's wide gulf, and never fear damage or shipwreck.*" Not merely a satnav or autopilot here, but a boat that knows our thoughts and wishes.

And so, the dreams and conception of artificial intelligence have been with us for millennia and straddle the cultural globe. The Syrian writer, Lucien, was the author of what is considered the first sci fi novel, *True History* (Hickes, 1894). Lucian lived during the second century AD, and hailed from Samosata, a city whose remains now lie under the waters of the Atatürk Dam in southeast Turkey.

True History is a satirical classic, reminiscent of the eighteenth-century French tale, *Candide, oul'Optimisme*, whose author, François-Marie Arouet, was better known by his nom-de-plume, Voltaire. Both books ridicule supernatural religious practices. The work of Dante Alighieri, Jonathan Swift, Ludovico Ariosto, Cyrano de Bergerac, Edgar

Alan Poe, Jules Verne and Herbert George Wells all owe much to Lucian's groundbreaking approach.

In *True History*, at one point, Lucian and his crew arrive at Lychnopolis, or lamp city, apparently located in the constellation of Taurus (the bull) between two different star clusters, the Hyades (at the base of the horns of Taurus) and the Pleiades (at the tail of Taurus). Here they encounter a series of highly knowledgeable lamps. Finding one that represented Earth, Lucian writes that they "*spake unto it and questioned it of our affairs at home, and how all did there, which related everything unto us*".

And so, we see that the concept of artificial intelligence has been around for a very long time. But it was not until the twentieth century that these visions could be made in flesh and blood, or, rather, silicon and gold.

We now get to grips with some of the jargon associated with artificial intelligence, the context and development of the field and, more fundamentally, we define what is meant by artificial intelligence. If you are already comfortable with these areas, you can probably skip or skim this section. However, if you are only starting to study the whole subject, then it is probably essential reading. I'll try to make it as pain-free as possible.

First, some definitions. There is a glossary at the back of the book that you can access at any time, but let's deal with some of the main players here. When thinking about artificial intelligence, we are really dealing with two complementary strands, the *internet of things* (IoT) and *artificial intelligence* (AI). So, let's get started.

I.2. Oh, for a nice cold soda: The birth of the internet of things

Wean Hall is a big box of a building of rough-hewn concrete, an example of the Brutalist architectural tradition. It sits in a Beaux-Arts quadrangle, as if dropped from a height by modernist aliens just to offend and provoke. So ugly is it that the architect is unknown. I guess no-one wants to be associated with it (unless it was actually dropped by aliens). It houses the Department of Computer Science of Carnegie Mellon University (CMU) in Pittsburgh. Famous for being the alma mater of Andy Warhol, CMU has an even greater claim to fame: the home of the first 'thing' in the internet of things (or one of the first anyhow).

Computing science is thirsty, warm, long work. The hardware releases a lot of heat. It's also compulsive, and hours go by. Caffeine and cold beverages are essential. And in Wean Hall, there was a soda machine, promising both of these things in one bottle. I remember cramming for exams at the University of Illinois, Urbana-Champaign, on a diet of Mountain Dew and Pro Plus. But in Wean Hall, the soda machine was on the third floor, while the research labs were on the fourth floor. Stuck at your work station for hours, you become aware that your mouth is as dry as a frog in the Sonoran

Desert. You crawl, heavy-limbed and weary, downstairs to the machine of your dreams.

You know it's going to be OK. A light-headed sensation of hedonistic anticipation begins to build, elation growing as if you'd crawled across the aforementioned desert and were just one sand dune away from the oasis of heavenly delights with its pool of cold, effervescent liquid. Refreshment and awakening lay just along the corridor, down the stairs and a few short steps to the left. But imagine the overwhelming sorrow if, when you got there, armed with your nickel, the machine was empty. It's enough to fell a human being to the ground. An old line from an ancient hymn comes floating through the air.

"I tried the broken cisterns, Lord,
But, ah, the waters failed!
E'en as I stooped to drink they fled,
And mocked me as I wailed."

Or it could be even worse, if that is possible. Just before you descended the stairs to the corridor of bounteous beverages, one of the student volunteers had filled the machine with bottles from the hot store cupboard down the corridor. You arrive. It looks perfect, just how you imagined it would. Bottles all lined up bolt-upright, ready to satiate your cravings. A scene of true glamour and perfection, like a painting by the great former student, Warhol, himself. You deposit the nickel, reach down for the bottle and your world darkens, as thermoreceptors in the skin on your fingers transmit the devastating news to your brain by way of a series of electrochemical impulses, that it interprets and whispers to your conscious self: The.bottle.is.warm.

There is nothing worse than a warm soda. There's no way around it. It will take three hours to chill. But you've already removed it from the chiller, and that was your last nickel. You stand in the corridor, forlorn and friendless, and memories of all of the lowest, emptiest, most isolated feelings from your life and all the other crumpled lives that ever slumped against this wall through time immemorial come tumbling down around you.

You realize that you've slid to the floor, but you don't care. You groan that sound that comes from some deep cavern within your soul, but far more extensive than the known universe and the universes beyond. This sound is not the Vedic '*Om*', proclaiming the creative powers of the universe, but the polar opposite. This is not the big bang, but the thirsty, overheated, very, very sleepy big shrink. Dreams dashed, devastated and unsatiated, you feel exiled, brutalized and so very alone.

Yet, in our lowest moments, hope is the last to die, and with hope, springs imagination. In 1971, high in the Brutalist box that is Wean Hall, a computer scientist had a dream: to know if the soda machine had soda in it and to know how cold that soda was, without even leaving his work station. David Nichols, a graduate student, convinced two other students, Mike Kazar and Ivor Durham, as well as a research engineer, John Zsarnay, that they had the technology to make this dream a reality.

The soda machine, like all the other soda machines in the world, had indicator lights. When a drink was bought, a red light would flash on and off. If the last bottle on the rail was purchased, the light would stay on. Once the drinks had been replaced, the light would go off again. The students installed electronic sensors to monitor the activity of the lights in the machine, and a cable ran from the sensors to a mainframe computer. From here, data could be fed through the forerunner of the internet, called the Advanced Research Projects Agency Network (ARPANET).

The ARPANET linked around three hundred computers, mostly in the USA. As a result, anyone, from Santa Barbara, California to Cambridge, Massachusetts, could check on the status of the soda machine in Wean Hall. But, most importantly, the computer scientists upstairs could know if there was cold soda down below. The first really useful *smart device* was born. A smart device is an electronic device that can connect, share and interact with its user and other smart devices through the internet. A revolution had begun.

Today, there are over seven billion smart devices in the world, not including smartphones, and the numbers continue to soar. The British Government has a target to have smart electricity meters in every home by 2020. These devices are continuously feeding data into the internet. They are the eyes and ears of the infosphere, listening and watching, measuring and mapping our every move. Connected together, they form what is called the internet of things, or IoT, a term coined, in 1999, by the executive director of Auto-ID Laboratories, Kevin Auston.

Other names have been used, such as the web of things, internet of objects, embedded intelligence, cyber physical systems, pervasive computing, ubiquitous computing, calm technology, machine-to-machine (M2M) and human-computer interaction. But it is the internet of things that has stuck. Defined as "*an open and comprehensive network of intelligent objects that have the capacity to auto-organize, share information, data and resources, reacting and acting in face of situations and changes in the environment*" (Madakam, 2015), the IoT has become synonymous with the digital age. It is now embedded into many elements of critical infrastructure. Increasingly invisible, it cohabits the world with us, often without us even knowing. We live in smart houses, in smart cities, working in smart offices and watching smart televisions. Soon we'll all be commuting in smart cars and hanging out with smart robotic companions.

This pervasive connectivity, linking automation, integration and servitization, is steadily growing, to the extent that the US National Intelligence Council included the IoT as one of its six disruptive civil technologies, the other five being *clean coal technology* (posing a threat to the oil-dependent economy of the USA), *robotic technology* (threatening jobs and social fabric), *biofuels* (requiring dependence on other countries), *energy storage technology* (again threatening the oil-based economy) and *biogerontechnology* (increased life expectancy through technology, having huge cost implications and demographic consequences).

IoT has perhaps been the most sensitive area, particularly militarily, as it represents potential information security vulnerability. Information is the new oil, and represents power and control. This is why organizations such as Wikileaks are viewed as significant threats. Companies that control the connectivity of the internet are seen as particularly powerful, as they control the internet of things. The new 5G mobile internet connectivity is creating concerns among cyber-security experts. The potential problems arise because the companies that control the connectivity have the potential to tap into what is being communicated. The eyes and ears of the internet could become the eyes and ears of these companies, realizing information security risks. Even more concerning to national governments is the situation where such companies are under the control of potentially hostile governments. However, the cloak-and-dagger, allegation and counter-allegation world of industrial and military activity makes it difficult to know exactly where we sit.

It's a bit like the *Potemkin village*, as my precious, sadly departed friend, Klement Rejšek (to whom this book is dedicated), explained to me, while discussing Russian politics. Catherine the Great was Russia's longest reigning female monarch, ruling from 1762 to 1796. Her reign coincided with the Golden Age of Russia. Famous for her brilliant political strategy, she often had affairs with her most capable ministers.

One of the most important supporters of her reign was Grigory Potemkin, the military leader and nobleman. By 1787, he was governor of Crimea. Crimea had recently been taken from the Ottoman Empire after a bloody war. Potemkin wanted to attract Russians into the newly acquired territory in order to build stability and trade.

The story goes that Catherine the Great was due to visit in order to help her former lover with his mission. Rather than Catherine seeing the widespread starvation, poverty and ruin, Potemkin arranged for a mobile village to be assembled, with fake walls, fake villagers (his well-fed troops in disguise) and fake food on display. Catherine toured by boat, not landing at the village, and so could not see that it was fake. Each night, Potemkin had his troops disassemble the village and move it downstream, where, the following day, Catherine would again encounter what she thought was another perfect village.

The Potemkin village has been used as a metaphor in economics and politics, representing a subterfuge that appears to show things being better than they are. Of course, if it suits, the opposite trick can be as powerful, pretending things are bad or questionable in order to create distrust in a company, political party or country. And what an Achille's heel the connectivity of the internet of things represents: the very oxygen of the infosphere.

Yet spreading a rumour that oxygen is poisonous may not be the cleverest way to proceed. The few companies that can actually supply 5G are already heavily involved in many aspects of mobile technology service provision, to such an extent that it would be difficult to remove them, potentially disrupting this precious connectivity. Indeed, the internet of things is fundamental to any global business today, collecting data that allows predictability and foresight and to follow objects through the commodity chains in which they are situated. As we have already noted, information is power. The transfer of information is the life blood of our modern world.

And here lies the great irony. While companies want to gain information from competitors and consumers, they don't want their own information to be shared with anyone else. Unshared information is another form of power. But it is difficult to have it both ways. You can't have both off-grid privacy and interconnected sharing. Other issues with IoT include diverse standards and different technologies across the globe, leading to what are called *fragmented ecosystems*, wherein true functional integrity is not possible. Also, the sheer rate of expansion and modernization creates its own difficulties, as docontinuous, though necessary, regulatory changes.

Secure connectivity is crucial in order to build confidence and trust amongst increasingly connected humans, to avoid data theft, to protect health and safety, to prevent a loss of productivity and to avoid noncompliance to regulations targeting the aforementioned areas. Another challenge is that as the IoT expands, there is the increasing likelihood that devices will be connected to the system that have been set up by inexperienced coders, insufficiently schooled in the fine arts of cybersecurity. This can allow an entrance point for attackers, a weak link in the chainmail of the IoT knight. As the fearsome biblical warrior Goliath found, even a small shepherd boy armed with a sling and a pebble can land a killer blow if he hits a weak spot.

Many of the devices now connected are battery-powered and cheap, with insufficient power to host adequate security programmes. Also, it is increasingly recognized that software alone cannot provide the protection needed to ensure security. More often now, hardware security circuitry or parallel security processors are in use, but this again increases costs and power requirements.

As the IoT becomes a key player in essential and critical infrastructure, its value increases in terms of criminal interest. Vulnerability means leverage in terms of malevolent intentions, where the ability to disrupt technology becomes extremely lucrative. Technological hostage situations are already occurring, where individuals or companies are asked to pay a ransom, often in Bitcoins, in order to have services restored.

So, the IoT, born in the minds of four thirsty, weary computer scientists in Pittsburgh wanting a cold soda, has become a global Argos, the many- eyed giant of Ancient Greece, whose eyes never close (even without a caffeinated soda).

The concept is expanding. The internet of everything (IoE) refers to the interactions between the 'thing', people, data and process. Some argue that this is really just IoT, but others feel that it is more inclusive and allows sociological contexts to be discussed as part of the landscape (Evans, 2012). Smart cities are viewed as large scale expressions of the IoE (Mitchell et al., 2013). Interestingly though, there is no room for the biological environment here.

But it is a second great technological development of the infosphere that promises to truly transform our futures, for better or for worse. That development is *artificial intelligence*.

I.3. The two-month, ten-man project to transform the world

On first appearances, the meaning of the term *artificial intelligence* would seem to be self-evident: a form of intelligence that is artificial, or non-organic. The term was invented by John McCarthy and colleagues, in 1955, in a proposal for a conference (what would become the first Dartmouth Conference) where they wrote: "*We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it*" (McCarthy et al., 1955). They were granted the funding and the historic meeting went ahead. Two months and ten men. It is reminiscent of the two small fishes and five loaves that fed five thousand. From such humble beginnings, a future world was conceived.

Yet, immediately, we are faced with some questions. What do we mean by intelligence, and can intelligence be artificial at all? Surely you are either truly intelligent, or not? Is artificial intelligence some pretence? Is it like an artificial limb, a dental implant, a lawn made of plastic, a bunch of flowers made from silk stretched over metal wire? Artificial flowers may look very much like the real thing, but they cannot be pollinated, do not produce nectar and will never produce seed. So, is artificial intelligence some sort of mimic, some counterfeit concept, unable to deliver the heavy fruit that only hangs from the bough of the true tree of wisdom?

Also, if it is a mimic, what kind of intelligence should AI be seeking to mimic? Is intelligence a strictly organic property, existing only in living organisms? As organisms increase in complexity, from viruses to bacteria, protists, fungi, plants and animals, at what point does intelligence occur and does it differ at different levels of organization? Can a unicellular organism really be intelligent, or do you need to be multicellular? Vegetarians think that animals are more 'special' than plants and should not be eaten, because they feel pain and are sentient. Yet a plant is an extremely complex and, many would say, intelligent sort of organism. Do you have to have a brain, or at least a proto-brain, or can plants be intelligent?

Was there a moment in the evolution of life on Earth when the first intelligent being suddenly popped up, and proclaimed "*you guys, you just don't get it, do you? 'Cause you just don't have it up here, right*"? If not, then either everything is intelligent or nothing is intelligent.

The same goes for species. Imagine lining up all of your ancestors in a row. First would be your mother and father, then, behind them, your grandparents. Next the great-grandparents and so on. You then step out and start walking along the line that stretches far into the distance. I've often thought about this. As you walk along, the fashions would change over time. Eventually, your distant relatives would be dressed in animal furs, or, further back, nothing at all (an awkward moment). Houses would become more primitive, until they were living in caves. The landscape might well change as you discover that many of your ancestors came from other countries or continents. Further back still, you start feeling the chill of an ice age, or the thirst of one of the great droughts that destroyed the Mayan civilization, the Mesopotamian Akkadians or the Egyptian Old Kingdom.

At what point in this walk, as you feel more and more distant from the individuals you are encountering, would one of your relatives try to eat you? At what point would you decide that they were no longer the same species as you? Putting it another way, did a non-human give birth to a human at some point, or a non-lion to a lion? When does a species become a species? This is very hard to grasp. Yet if we tie intelligence to being human, then at that very moment, when the first human was born, and realized its parents were a different species, we would have to say that at that moment too, human intelligence was also born.

These are difficult issues. Yet they are part of a whole set of similar questions. Biology deems that viruses are not living, but that bacteria are alive. To be in the club of life, science deems that you have to pass a number of tests. These are, in no particular order, the ability to reproduce, responsiveness to the environment, homeostasis, trait inheritance, metabolism, development and cellularity. A mule is the offspring of a male donkey and a female horse. Yet a mule is not able to reproduce. If it stares at you, or kicks you, it seems very much alive. However, according to our definition of life, it isn't alive because it cannot reproduce. Yet no one would say that mules aren't alive. Putting this another way, if you kill a mule, have you really killed at all? It seems obvious that you have killed the mule but, strictly speaking, this definition of life might say otherwise.

Viruses are tripped up on their way into the club of life by the cellularity issue in particular. Yet viruses do very well for themselves, and have outwitted us on many occasions. They use our cells, so don't need to be cellular. The Spanish Flu killed more people at the end of the First World War than all of the conflict of the previous four years. Viruses have been around for hundreds of millions of years, surviving mass extinctions and the rise and fall of countless species of cellular organisms. Surely to claim they are not alive is a bit churlish?

I suspect that the real reason behind this is that we don't like to have anything too different from us in the club. At least bacteria are cellular, just like us. If we include viruses, then we might have to include prions. Prions are short strands of protein that cause diseases such as bovine spongiform encephalopathy (mad cow disease). Then we really would have trouble, because proteins are just molecules, and so where would it stop?

Some turn to *panspermia*, the theory that life came from another planet. It kicks the can down the road a bit more. But at some point, we have to face up to it. At some point, life started, and if it didn't, then is there such a thing as life at all? We like to think we know things, but the beginnings are much more difficult that we realize. We conceptualize the Big Bang as the start of everything, but what happened before that?

Even with the big bang, big shrink model, wherein the universe has continuously expanded (the big bang phase) until it cannot support itself and has collapsed (the big shrink), which came first, and what came before that? This is cosmology at its most bewildering.

If life emerged on the planet from biochemical events, at what moment was the first living organism actually born, and how? We don't like continua, because we like a beginning and an end. We like cause and effect, but imagining how the first cause came into existence to have the first effect is far beyond us. And it's the same with intelligence. What was the first intelligent organism and when did it arise?

A fairly typical definition of AI is as follows: "*Artificial intelligence (AI) is a variety of human intelligent behaviors, such as perception, memory, emotion, judgment, reasoning, proof, recognition, understanding, communication, design, thinking, learning, forgetting, creating, and so on, which can be realized artificially by machine, system, or network*" (Li and Du, 2017). Mikolov et al. (2016) set out four basic requirements for AI:

1. Interactive communication using natural language;
2. Channelling of non-linguistic information such as sensory perception;
3. The capacity to learn;
4. The capacity for motivation.

Driving the rapid progress towards these goals are the huge amounts of data (*big data*) streaming in from the internet of things and cheaper processing and storage facilities.

We'll examine what it means to be intelligent, and what kinds of intelligence exist in Section III, when we question the choice of human intelligence as the foundation of AI. But let's examine AI in a bit more detail first, because within it lie a number of definitional complexities that we need to be clear on.

I.4. Getting to grips with the jargon: Symbolic and non-symbolic AI

Fundamentally there are two forms of AI: the autonomous, thinking, feeling, sentient genius that is still the stuff of science fiction, and the more mundane, but still very impressive current models like IBM's chess-playing wizard, *Big Blue*, and IBM's quiz mastermind, *Watson*. These two categories of AI are extremely different beasts.

The first type, the autonomous thinkers and doers, are examples of *strong AI* or *artificial general intelligence (AGI)*. The later, *weak AI* (also called *narrow* or *applied AI*), are focused on single narrow tasks, like winning a game show or a chess match or diagnosing diseases from symptoms.

Another important differentiation is between *symbolic* and *non-symbolic* AI. This lies at the heart of the history of AI. Symbolic AI is all about facts. These facts are written in code and placed within the memory of the AI. When needed, they are recalled. It works on an if-then basis, where if a certain question is asked, then a stored answer is given. Symbolic AI is also called *classical AI* or *good old-fashioned AI* (GOFAI). It is a top-down approach, wherein the software controls what happens, like a director guiding his actors to produce a movie of the director's imaginings. Symbolic AI relies on a giant library of books, with all the answers to whatever the AI unit is tackling. For example, a symbolic AI programmed for fishing would have books on what weight of line should be used, what type of fly is best, a taxonomy of fish, where to fish, how to cast, how to net the fish and so on.

My mother is a brilliant oil painter (not too shabby with acrylics, pastels and watercolors either). She has sold her paintings around the world. I have always wanted to paint like her. I have attended countless art classes in search of the muse. I've read all the self-help, learn-to-paint-in-three-weeks books. I've followed the television programmes. I've hung out with artists, and even teach a course (on ecology and design) in an art college each year. But the muse just turned her back on me and walked away into the ethereal mist. It's not that the muse doesn't like me - muses are above all that nonsense. Rather, I am invisible to her. She is a spirit, and I am mere non-artistic flesh. I just can't paint. I don't have 'it', whatever it is. It's an inherent, non-symbolic thing, not an if/then thing. You can't reduce love, art or life to if/then responses.

Symbolic AI has strengths. Its decisions can be easily understood by humans, as it follows processes laid down in comprehensible steps. It is rule-based, and we make the rules. With the exception of the earliest years and the most recent period, symbolic approaches have dominated AI research. However, it has well-recognized issues. Firstly, in a dynamic, changing, emergent world, the paths of if-then logic can quickly become inadequate. You would need to programme vast amounts of pathways to cover all eventualities. In other words, you have to map out the journey for the AI unit with step-by-step directions, and this can become difficult in new and unfamiliar landscapes. Trade-offs, lack of complete knowledge, rapid changes and stochastic events all mean that the AI may be left high and dry, with inadequate flexibility, crying out "*But you didn't tell me it would be like this!*"

Another big problem with symbolic AI is what has become known as the *common sense knowledge problem*. This problem relates to the fact that our intelligence is not only based on available, explicit facts, but on a whole host of other influences, including cultural, social, historical and philosophical. When we start trying to convert these into if-then statements, we run into an ocean of problems. Common sense is not easily programmed.

Indeed, the relationship between the conscious mind and the subconscious mind is still a subject for conjecture. And in terms of what the subconscious mind actually is, there is very little firm evidence. Symbolic subconscious software doesn't exist. You can't map what you have never seen. Symbolic AI also does not do well dealing with variations in light, pressure, texture and sound.

The symbolic approach became viewed as being too brittle and rigid to fully deliver the kind of intelligence that programmers dream of. And it was not only the dynamic nature of the world that proved problematic, requiring extremely complex calculations to model it, but also the ambiguity. The world is a functioning system, and therefore is an emergent entity. It is non-linear and full of surprises, unless you possess omniscience. While we may feel like gods, making AI in our own image, we lack a number of deific qualities, one of which is omniscience.

Of course, symbolic AI can be very good at some things. Deep Blue, who beat Gary Kasparov in that historic chess contest in May 1997, used symbolic AI. It could consider and assess 200 million moves per second. Deep Blue started life at Carnegie Mellon University in Pittsburgh in 1985, where that smart soda machine lived. Indeed, the graduate students who first worked on Deep Blue in its earliest stages of development, Feng-Hsiung Hsu, Thomas Anantharaman and Murray Campbell, may well have used that very soda machine. Maybe there was something in the soda, but that's a whole other discussion. Later developed at IBM, what would become Deep Blue represented the zenith of AI at the time.

Non-symbolic AI (also referred to as *data-based AI*, *connectionist AI* or *empirical AI*) operates in a completely different way (note that you will also encounter *sub-symbolic AI*, which many authors equate with non-symbolic AI, though some define as an intermediate technology between symbolic and non-symbolic AI. For example, Willshaw (1994) describes sub-symbolic AI as being built up of entities called sub-symbols, which are the activities of individual processing units in a neural network. We will use the term non-symbolic AI.

The aim of non-symbolic AI is to mimic the human nervous system, and in particular the neural activity of the human brain. This is a bottom-up approach, designing the hardware and then letting it produce the outcomes. The AI unit is fed a stream of raw data, and it then identifies patterns within that data. From this analysis it builds a representation of its world. The best fishermen don't rely on a library of facts alone. In fact, many have never read a single book on fishing. Rather, they have gained insight from countless fishing trips to all sorts of different aquatic landscapes. Big ponds, lakes, oceans, small rivers, swamps – they've done it all. Hot days, wet days, stormy nights and ice have all been faced. Fundamentally they have been exposed to lots of different learning experiences and have learnt how to fish. This doesn't rely on symbols, libraries or if/then decisions. It's non-symbolic.

Non-symbolic AI learns, rather than regurgitates, and excels at pattern recognition and classification. As amazing as this sounds, the approach requires huge amounts of data in order to approach an accurate representation of a real-world problem. Non-symbolic AI is more to do with hardware than software. It uses a series of artificial neurons, connected together, to mimic the structure of the brain. This is called the *connectionist approach* because of the importance of the connections.

Structurally, it consists of large interconnected networks of simple processors that run in parallel. At any one point, a given processor may be delivering input or output, and what emerges is knowledge in the form of a pattern composed of the strengths of connections between these processors. These patterns are products of the activation states of each processor and of the linkage strengths between each processor. This neural network is like a black box. You feed lots of data into it, and something comes out, but you can't follow the thinking and can't be sure how it ended up with the outputs that it did.

An example of a neural network is the *AlphaGo* AI from Alphabet's DeepMind. AlphaGo was the first computer to defeat a professional human Go player, in October, 2015. Go is an ancient Chinese strategy game and is thought to be much more difficult to play than chess, because of the larger board (19 × 19 squares as opposed to 8 × 8 squares in chess) and the greater variety of possible moves. The game is played by two players, one with black stones, who starts, and one with white stones. Players can place their stones on any free point (where two lines cross), but if one of your stones is surrounded by the opposite color on all sides, the stone is removed. The winner is the one who surrounds the most territory and captures the most stones.

At the heart of non-symbolic AI lies *machine learning*. Machine learning represents a watershed in AI development. The AIs work out their own mappings, rather than depending on programmers, using the huge data sets flowing in to them. This has led to major increases in cognitive and perceptive powers. AlphaGo used machine learning, improving each time it played. It was given training data and played against humans and other computers. This ability to learn and to improve is viewed as an important technological advance. It also sets non-symbolic AI at the forefront of technologies.

However, whether it is Go, checkers or chess, these games are examples of what are called *perfect-information games*, where both the human and AI unit know the exact state of the game at any given point. More challenging are *imperfect-information games*, such as poker, where a player can hide cards and bluff. Such games are very much like the real world, where we ultimately want our AI units to function. In the real world there is a lot of hidden information and complexity, partly because of the emergent qualities of complex systems, and partly because of deliberate subterfuge. Such hidden-information games pose useful testing grounds for AI. Recently for the first time, an AI called *Libratus* has beaten top professionals in two-player Texas hold'em poker, the most popular form of the game in the world (Brown and Sandholm, 2018).

The dream of machine learning has been anticipated for many years. Remember the proclamation of Marvin Minsky in 1970 that "*In from three to eight years we will have a machine with the general intelligence of the average human being*" (in Darrach, 1970) or the even earlier proclamation by Herbert Simon (from Carnegie-Mellon University again) that "*machines will be capable, within twenty years, of doing any work a man can do*" (Simon, 1965)".

And the hyperbole surrounding the field in the 1960s has returned. Grace et al. (2017) claim that: "*AI will outperform humans in many activities in the next ten years, such as translating languages (by 2024), writing high-school essays (by 2026), driving a truck (by*

2027), working in retail (by 2031), writing a bestselling book (by 2049), and working as a surgeon (by 2053). Researchers believe there is a 50% chance of AI outperforming humans in all tasks in 45 years and of automating all human jobs in 120 years." Futurist Raymond Kurzweil predicts a moment of technological singularity, where AI will overcome human capabilities, ushering in a utopia, as predicted by some philosophers in the Enlightenment (Kurzweil, 2005). Robert Geraci (2008) looks forward to a "*victory of intelligent computation over the forces of ignorance and inefficiency.*"

However, weaknesses still exist, particularly in less clear-cut areas of application. For example, the training sets of data can have biases within them. To rule out such bias, a huge amount of training data would need to be gathered. Yet if such biases have gone unrecognized and are fundamental to a particular area of study, this could be difficult to rule against. Also, a game of Go is not real life. There is a fixed set of rules, a finite, albeit large, playing area and a rich data set of previous games to learn from. Novel, dynamic and emergent situations may not be so easy to deal with, even for AlphaGo.

Perhaps the most serious issue with non-symbolic AI is the reliance upon mathematical optimization. The solution chosen is the most optimal one. This is problematic in terms of trade-offs and systems theory, as we shall explore later. At this point, let's just say that sub-optimality is generally the rule in systems, and if you try to optimize any given component, this will ultimately lead to system collapse. Imagine a squirrel, who optimizes ways to find all of the nuts it buried. The outcome would mean there were no nut stashes left, and so there would be no new trees growing from these forgotten stashes. Over the years, this would lead to the disappearance of the very forests that produced the nuts in the first place, and the squirrels would starve. Hence, optimizing for the needs of a hungry, greedy squirrel would not be sustainable.

An interesting way to explore the difference between symbolic and non-symbolic approaches is in the way children learn language. For the first few years, this is an informal process, where children learn through playing a linguistic game with their parents, without any real rules and a constant data stream. Connection strengths develop between neurons. It is a black box approach, where language emerges from the grunts and laughs of a baby, to spoken words, then couplets and finally sentences. This is akin to non-symbolic AI. Once the child starts school, this changes, with rules of grammar now directing things. Spelling, writing and syntax follow. This is akin to symbolic AI.

What is interesting is that both approaches combine to deliver our linguistic ability. In AI, there is now also a move to combine symbolic and non-symbolic approaches. Here, symbolism operates within an overall non-symbolic network (Gärdenfors, 2000; Chella, 2004; Forth et al., 2010; Öztürk and Tidemann, 2014). In what has been called *perceptual anchoring*, the symbols associated with a particular object can be anchored to the sensory data coming from the outside world, in space and time, allowing correspondence between what is sensed in the outside world and what is symbolized.

Here, the concepts related to an object (the qualitative, symbol level) can be combined with quantitative data flowing from sensory perception of the outside world.

The respective strengths of each approach are envisaged to combine together, and diminish the respective weaknesses. Rhett D'Souza puts it this way: "*the non-symbolic representation-based system can act as the eyes (with the visual cortex) and the symbolic system can act as the logical, problem-solving part of the human brain*" (D'Souza, 2018).

Historically, both concepts had been developed in the 1950's. William Grey Walter's turtle robots, built in 1951, were early non-symbolic machines. However, the symbolic approach came to dominate in the 1970's and 1980's, receiving most of the funding, and taking the lion's share of research journal pages, graduate programmes and conferences.

The whole field of AI took a huge hit from a report compiled by Professor Sir James Lighthill, entitled *Artificial Intelligence: A General Survey*, published in 1973. This demonstrated that it would be impossible for AI to solve anything but the most basic problems because of the combinatorial explosion or rapid increase in complexity, when operating in the real world. A debate in the Royal Institution in London following the report brought together some of the great minds to discuss the issues raised. It makes fascinating viewing and is available on You Tube (Lighthill Debate, 1973).

There are so many variables and dynamics, that the possible outcomes are close to infinite. As Carl Sagan said, referring to bacterial exponential growth, "*Exponentials can't go on forever, because they will gobble up everything*" (Sagan, 1997). The same applies to computer memory. Lighthill's attack dramatically impacted research funding globally, with governments withdrawing support on both sides of the Atlantic, leading to what was described as the *AI winter* that would last from 1974 to 1980. However, even during the AI winter, the non-symbolic, connectionist approach was viewed as less worthy than the symbolic approach. This was because the non-symbolic approach was thought to need far more calculation and computer power than the symbolic approach.

A re-framing of AI took place towards the end of the 1970's. Instead of promising strong AI, that could ultimately think and plan beyond the capabilities of the human mind, AI was instead marketed to businesses as a tool to help with specific problems, and these tools started to be created and sold. AI became a practical engineering project rather than a theoretical one. The aim was no longer to change the world but to sell some useful software to companies and make a living.

As the vision for AI became tamed, so too were many of the myths slain. AI was no longer a threat that would take over the world and destroy humankind, like some alien, or the dogs of war unleashed. Instead it was that plastic and metal box of tricks, sitting in the workplace, helping with mundane office life. How had the mighty dreams faded? Deep Blue restored some pride, but while being a chess champion was good, it was far from artificial general intelligence.

But, as we know, hope is the last to die, and adherents to the non-symbolic church of the latter-day AI regrouped. Towards the end of the Eighties, a shift began, noted by Paul Smolensky (1988), who wrote "*In the past half-decade the connectionist* [*non-symbolic*] *approach to cognitive modelling has grown from an obscure cult claiming a few true believers to a movement so vigorous that recent meetings of the Cognitive Science Society have begun to look like connectionist pep rallies.*" And with its resurgence, non-symbolic AI re-awakened the mythologies of fear and fantasy.

One of the leading advocates was Rodney Brooks, the Australian roboticist, who founded *iRobots*. In a seminal paper written in 1990, entitled *Elephants don't play chess* (this line of attack continued as late as 2010, with the Kelley and Long (2010) paper, *Deep Blue cannot play checkers*), he opened with the unambiguous criticism that "*Artificial Intelligence research has foundered in a sea of incrementalism*", going on to argue that "*the symbol system hypothesis upon which classical AI is base*[*d*] *is fundamentally flawed, and as such imposes severe limitations on the fitness of its progeny*" (Brooks, 1990).

Brooks set out a brave new world, where rather than building a representation of the real world as an extremely complex model of overwhelming immensity, you take the approach that this world is its own best model. It is totally up-to-date, and so all you have to do is sense it often enough and in an appropriate way. By tapping into the here-and-now, with a constant data flow, rather than trying to reconstruct it with symbols, you overcome the risks of your model quickly becoming out-of-date and cumbersome.

When we combine the non-symbolic approach with the IoT, we end up with *cloud robotics*, where each individual AI unit is receiving data from their surroundings, learning, and then sharing the learning with all of the other units. This flow of raw data and information has the potential to be extremely powerful.

The symbolic, software approach relied on the assumption that everything was programmable, but it became clear that the real world was much more complicated and ambiguous than had been initially thought. To know anything, the machine first had to know everything. Limited scripts, such as how to win a chess game, did not approach true intelligence. Not only are we unsure how the brain works, but it has become clear that every brain works slightly differently, based on its journey through space and time.

In other words, we possess *mindful brains*, as Edelman and Mountcastle (1978) discussed in their book of the same title. And these mindful brains are emergent outcomes of our social and environmental realities. Furthermore, it is the "unknowingness" that defines the mind, whereas knowingness defines the brain, and cognition becomes a process of filling in the gaps based on our unique experiences. As Wolfe explains (1991), "*Winning games is something our brains do; playing them is something our minds do.*"

And so, when it comes to understanding what AI is, we need to embrace both the breadth of intelligence that exists in Nature (due to a lack of species barriers), the

Index

A

Acknowledged dependence, 170
Actor-network theory, 188, 219
Affective processing, 39
Agrarian age, 208, 219
Agricultural revolution, 86, 96, 99, 159
Agroforestry, 159
AI
  narrow, 12, 49, 54
  non-symbolic, 11-17, 21, 24, 33, 63, 162, 225, 229, 231, 232
  strong, 12, 17, 19, 23, 28, 41, 45, 64, 219, 231
  symbolic, 11-13, 16, 20, 126, 225, 231, 232
  weak, 12, 23, 41, 65, 233
AI4All, 36
Aichi Biodiversity Targets, 136
AI rights, 30
AI winter, 1, 16, 53
Algorithm, 23, 25, 27-29, 31, 34, 36, 38-42, 49, 51, 54, 55, 68, 69, 107, 120, 121, 123, 124, 128, 131, 132, 139, 190, 219, 222
Algorithm-based prejudice, 41
AlphaGo, 14, 15
Ancient brains, 199, 200
Anthropocene, 93, 219
ARPANET, 6
Artificial general intelligence (AGI), 12, 17, 19, 23, 24, 63, 219, 231
Artificial intelligence (AI), 1-4, 9, 11, 16-18, 25, 33, 35, 37, 46, 58, 59, 63, 66, 67, 73, 77, 78, 88, 123, 128, 130, 137, 154, 158, 216, 219, 221, 222, 225
Artificial Intelligence and Ecosystem Services (ARIES), 137
Artificial Intelligence for Disaster Response (AIDR), 130
Artificial social intelligence, 128

Asimov's laws, 21, 219
Associative learning, 74, 75, 78, 220
Atmosphere, 25, 78, 89, 93, 99, 103-106, 134, 148, 151, 220, 221, 223
Atmospheric pollution, 103, 178
Autism, 77, 129
Automaton, 2, 34
Autonomous vehicle, 120, 124, 130

B

Ball and stick model, 127
Barriers
    philosophical, 192, 196, 199, 200, 203
    psychological, 192, 199, 228
    structural, 192, 196-199, 204
Beaver reintroduction, 146
Behavioural momentum, 201, 205, 220
Big bang, 5, 11
Big data, 11, 24, 25, 27, 33-35, 37, 43, 50, 76, 120, 128, 129, 131, 132, 140, 220
Big nudging, 131, 220, 229
Big society, 214
Biochemical mind, 18
Biofuel, 6, 104, 183, 220
Biomimicry, 69, 70, 141, 143, 191, 193, 220
Bio-participation, 70, 220
Biosphere, 44, 58, 70, 78, 93, 96, 98, 103, 185, 191, 220, 223
Blackhouses, 173
Bletchley House, 124
Bostrom, Tim, 48
Bottom-up approach, 13, 41
Brain chauvinism, 73
Broken window fallacy, 123
Brooks, Rodney, 17, 67
Buenvivir, 94-96, 145, 166, 169, 191, 220

C

Capital
    Human, 220
    Human-made, 221

Natural, 221
Social, 221
Capitalism, 44, 95, 177, 214
Carbon footprint, 160, 195, 198, 221
Carbon trading, 182, 221
Carrying capacity, 86-88, 157, 221
Catherine the Great, 7
Cerebrocentrism, 73, 221
Chatbox, 41
Child AI, 48
Chlorofluorocarbons (CFCs), 104
Circular economy, 148, 212, 221
Citizen science movement, 138, 221
Civilization, 10, 46, 58, 76, 81, 94, 96, 99, 155, 157, 174, 175, 200, 209, 219
Civil society model, 212, 221
Cloud computing, 124, 222
Cloud robotics, 17, 222
Co-design, 216, 222
Co-dissemination, 216, 222
Collective intelligence, 28, 49, 131
Colonialism, 44, 177
Colonization, 142, 168, 179
Common sense knowledge problem, 12, 222
COMPAS, 40
Complementarity, 44
Condorcet, Marquis de, 58
Connectionist approach, 14, 16
Conflict theory, 127
Contingency theory, 22, 222, 230
Co-production, 216, 222
Correspondence, 16, 30, 44
Court jester hypothesis, 22, 222, 230
Cultural issues, 36, 43
Cultural services, 98
Cyanobacteria, 101
Cybersecurity, 8

D

Darwin, Charles, 72
Data-driven modeling, 140, 222
Data ethics, 25, 26, 222
Dead zones, 101, 102, 145, 151, 153, 192
Deep Blue, 13, 17, 182
Deep learning, 41, 129, 139, 222
Delusion of conservatism, 134
Deontology, 20, 22, 29, 38, 222, 233
Development, 1, 4, 9, 10, 13, 14, 19, 23-25, 30, 31, 34, 44, 50, 52, 53, 59, 64, 65, 71, 76-78, 92, 93, 96, 107, 110-112, 118, 123, 124, 139-144, 155, 165, 169-171, 208, 209, 215, 222, 225, 229-232
Direct feedback, 88
Discredence, 53, 202
Dispersed intelligence, 70, 72, 73, 78, 223
Disruptive civil technologies, 6
Distributive justice, 45
Diversity, 1 21, 29, 33, 42-44, 51, 52, 63, 87, 90, 93, 96, 98, 102, 103, 107, 108, 121, 131, 135, 140, 142, 158, 168, 179, 180, 187, 204, 213, 217, 225, 226
Division of labour, 73, 96, 97
Double devolution, 214
Dunbar's number, 130
Dust bowl, 99, 100
Dynamic equilibrium, 78, 82, 90, 210, 211, 223

E

Earth Overshoot Day, 156
Earth system, 45, 65, 67, 70, 77-79, 82, 83, 86, 88, 90, 92, 98, 103, 106, 113, 133, 137, 138, 144, 148, 149, 154, 157, 161-165, 169, 170, 172, 178, 181, 187-190, 193-196, 198, 200, 201, 203-206, 209-212, 214-218, 223, 226-228
Eco-actualization, 170
Ecology, 1, 12, 44, 78, 87, 136, 138, 146, 147, 154, 166, 171, 181, 191, 217, 218, 220, 223, 225, 226, 230
Ecological footprint, 160, 182, 184-187, 195, 198, 199, 209, 215, 223, 228, 232
Ecological intelligence, 91, 94, 155, 167-169, 171, 172, 199, 204, 223
Ecological modernization theory, 110, 111, 144, 223
Economics, 7, 44, 93, 96, 97, 109-111, 114, 115, 121, 122, 125, 143, 145, 148, 154-157, 168, 169, 171, 177, 189, 190, 198, 200, 215, 219, 222-224, 226, 232
Ecosystem, 44, 52, 61, 66, 70, 73, 77, 80, 84, 85, 87, 90, 92, 96, 98, 99, 102-105, 117, 134-141, 143, 148, 150, 160, 164, 166, 168, 169, 178, 180, 181, 188, 194, 211, 223-225, 227, 229-