Health Care in the Information Society

VOL. 1

FROM ADVENTURE OF IDEAS TO ANARCHY OF TRANSITION

David Ingram

https://www.openbookpublishers.com

David Ingram, *Health Care in the Information Society, Vol. 1: From Adventure of Ideas to Anarchy of Transition*. Cambridge, UK: Open Book Publishers, 2023, https://doi.org/10.11647/OBP.0335

Digital material and resources associated with this volume are available at https://doi.org/10.11647/OBP.0335#resources

ISBN Paperback: 978-1-80064-952-1
ISBN Hardback: 978-1-80064-953-8
ISBN Digital (PDF): 978-1-80064-954-5
ISBN Digital ebook (EPUB): 978-1-80064-955-2
ISBN XML: 978-1-80064-957-6
ISBN HTML: 978-1-80064-958-3

DOI: 10.11647/OBP.0335

Cover image: D koi, *White H2O Molecules* (27 June 2022), https://unsplash.com/photos/5nI9N2wNcBU
Cover design: Jeevanjot Kaur Nagpal

2. Knowledge, Language and Reason–From Ancient Times to the Information Age

The story starts long ago, with the gradual conceptualization of knowledge as an encyclopaedia–a circle of learning. This chapter traces a path from the invention of medicine in classical times, through philosophy, language and logic, and through mathematics, natural science and computer science into the modern era of information technology and health care. It follows the librarian's dilemma over the ages–discovering how best to position books and documents within collections and search them in pursuit of learning.

The chapter proceeds to consider languages as expressions of knowledge, and the different forms they take–spoken, written, artistic, mathematical, logical and computational. This sets the scene for introducing computational discipline that grew from endeavours to formulate rigorous logical foundations of mathematics, in earlier times, and the development of formal logic in support of rigorous reasoning. From there, the computer has become integral to how we express and reason with knowledge, and to problem solving and the discovery of new knowledge. These are twenty-first-century frontiers of machine learning and artificial intelligence.

Moving to the complex world of medical language and terminology, used in representing knowledge about medicine and health care, the chapter discusses difficulties faced in evolving their *corpora* of terms and classifications, from pragmatic organizations into reliably computable forms. Notable pioneering initiatives and their leaders are profiled, highlighting some ideas that have acquired staying power and others that have not, looking for patterns of success and failure.

Finally, the chapter moves to a discussion of some pioneering computer-based systems for capturing, storing and reasoning with medical knowledge, such as for guiding the prescription of antimicrobial drugs. It closes with a light-hearted take on how we use the terms knowledge, information and data, and a reflection on the traction that is needed in the unfolding of new knowledge and its application in practical contexts.

 https://doi.org/10.11647/OBP.0335.02

> It seems like we all seem to know what 'to know' means. It is one of only 100 or so words that have a comparable translation in every language on earth.
>
> –Marcus du Sautoy[1]

The story begins with knowledge, but how to begin with knowledge? It is an elusive idea. Maybe a picture is best (see Figure 2.1).

Fig. 2.1 Knowledge as earth, air, fire and water–and the passage of time. Photo by Andranik Sargsyan (2020), Pexels, https://www.pexels.com/photo/silhouette-of-unrecognizable-woman-jumping-above-sea-beach-at-sunset-4149949/

In ancient times, humans experienced and came to know earth, air, fire and water, and had a sense of passing time. The concept of a 'beginning' was expressed through words, and through a belief in many and powerful gods. In the classical Greek and Roman sense, encyclopaedia denoted a circle of learning. Learning begins with awareness and experience of reality. In our sense of beginning, there is the hypothesis of science known as the Big Bang, which gives the universe a history and challenges belief. Humans find inspiration and illumination in science and art, and in culture and belief.

Maybe we could begin with one grain of sand–from the opening lines of a poem by William Blake (1757–1827), loved by my dad:

1 M. du Sautoy, *What We Cannot Know: Explorations at the Edge of Knowledge* (London: Fourth Estate, 2016), p. 413.

> To see a World in a Grain of Sand
> And a Heaven in a Wild Flower
> Hold Infinity in the palm of your hand
> And Eternity in an hour[2]

There are less lyrical perspectives of knowledge, also based on sand–reality as an arid desert landscape and shifting sands of time and experience. We talk of structures built on sand.

In the works of T. S. Eliot (1888–1965), we find poetic lament of pointless knowledge:

> The endless cycle of idea and action,
> Endless invention, endless experiment,
> Brings knowledge of motion, but not of stillness;
> Knowledge of speech, but not of silence;
> [...] Where is the wisdom we have lost in knowledge?
> Where is the knowledge we have lost in information?[3]

Human knowledge is much discussed. It connects and communicates throughout the worlds of experience, learning, imagination and action. The book's Introduction started with a quotation about connection. We connect, communicate about and reason with knowledge through language and logic. And in the Information Age, machines connect and communicate through language and logic, too–machine and program languages and formal logics.

Human knowledge and learning, and machine knowledge and learning, which we now also speak of, evolve in time. They may cooperate and they may diverge, in different dimensions, on different scales, and according to different values. David Deutsch has described knowledge as information that has causal power. Charles West Churchman (1913–2004) and his doctoral student Russell Ackoff (1919–2009) expressed the relationship the other way around, describing information as knowledge for the purpose of taking effective action. Information, too, is a central but elusive idea. Knowledge and information pervade the eternal world of ideas about the nature of reality and the everyday world of appraisal, decision and action.

Information technology also relates to sand–silicon crystals and embedded impurities at the semiconducting interface, transistors that shape and direct electron flow, and the electrical circuitry of computers. Proton gradients energize and shape electron flow across membrane interfaces and along the pathways of chemical reactions and molecular transport enacted

2 'Auguries of Innocence' (c. 1803), ll. 1–4.
3 'The Rock' (1934), ll. 6–9, 15–16.

in living systems. And information technology connects and communicates information flow at the interface of machine and human worlds.

Other luminaries provide further avenues into knowing and knowledge:

> I neither know nor think that I know.[4]

> *Nosce te ipsum* [know yourself].[5]

> Knowledge is always accompanied with accessories of emotion and purpose.[6]

> Pure logical thinking cannot yield us any knowledge of the empirical world; all knowledge of reality starts from experience and ends in it.[7]

It is as du Sautoy says: everyone knows about knowledge! One sometimes resonates with several lines recorded in *The Rubáiyát of Omar Khayyam*–I have my grandfather's 1928 illustrated edition of this gem close-by on my office bookshelf. Omar Khayyam (1048–1131) was born in eleventh-century Persia, and, so the story goes, the Vizier to the Sultan Alp-Arslan granted him 'sufficient provision to enable him to devote himself to the pursuit of knowledge'.[8] Despite his achievements in the science of astronomy and as a poet, and his high-level patronage, his worldly concerns in those pursuits ran into opposition from the mystics of the day. He records his frustration with their philosophy, thus: 'Myself when young did eagerly frequent, doctor and saint and heard great argument, about it and about: but ever more, came out by the same door wherein I went'.[9]

I used to use this quotation as a joke in talks about information technology and health care–substituting 'it' with 'IT'! My colleague, Thomas Beale, pointed out the coincidence that Khayyam was alive around the time of the founding of St Bartholomew's Hospital (1123), when we first worked together (in 1992).

Here is my own, not very serious, starter about knowledge:[10]

> A is for: Alphabet, Appearance, Attribute, Alpha, Acquisition, Antiquity, Authority, Alchemy, Astronomy, Astrology, Aristocracy, Abstraction, Argumentation, Abduction, Acceleration and Age. Knowledge is an **A**lphabet of learning, distilled from a kaleidoscope of **A**ppearances. It has

4 Socrates (470 BCE–399 BCE), recorded in Plato, *Apology*, section 21d.
5 Carl Linnaeus (1707–1778), describing humans in *Systema Naturae* (1735).
6 A. N. Whitehead, *Adventures of Ideas* (New York: Macmillan, 1933), p. 12.
7 A. Einstein, *Ideas and Opinions* (New York: Crown Publishers, 1954), p. 271.
8 O. Khayyam, *The Rubaiyat of Omar Khayyam*, trans. E. Fitzgerald (London: George G. Harrap & Co. Ltd, 1928), p. 8.
9 Ibid., p. 41.
10 I had the idea of starting each succeeding chapter with a paragraph like this, based on succeeding letters of the alphabet. But it didn't work out—this is a one off!

> an infinity of **A**ttributes. It starts from **A**lpha and stops short of omeg**A**. Stories of the pursuit and **A**cquisition of knowledge date from **A**ntiquity. They demonstrate knowledge as both manifestation and instrument of **A**uthority. Science in its earliest days was embroiled with mythology and mysticism–early chemistry with **A**lchemy, and **A**stronomy with **A**strology. The pursuit of knowledge engaged the rulers and **A**ristocracy of church and state. Theory of knowledge, embodying processes of **A**bstraction, **A**rgumentation, induction, deduction, **A**bduction and more, has evolved over centuries, as it challenged the shaping confines of mysticism and religious belief. This unending process has seen huge **A**cceleration in the Information **A**ge.

Concepts of knowledge and truth seem destined always to be strongly coupled together in the human mind. Discussion of knowledge is encircled by perspectives of discipline, and discussion of truth by schools of philosophical thought. There is much diversity in these discussions and debates, which engage the finest of minds. Karl Popper had little time for their deliberations. At his ninetieth birthday party, he is quoted, loose-lipped, as follows(!):

> I think so badly of philosophy that I don't like to talk about it. [...] I do not want to say anything bad about my dear colleagues, but the profession of teacher of philosophy is a ridiculous one. We don't need a thousand of trained, and badly trained, philosophers–it is very silly. Actually, most of them have nothing to say.[11]

Ridiculous or not, and very much not, I think, philosophy endures. The issues debated are defined and refined in the languages, logics and contexts of their times. Language has extended from the grammar, meaning and symbolism of the spoken and written word to the language of mathematics and computer science. It has entered the complex world of information systems that represent and reason about health care. Differences of perspective can then become embedded and buried within software. The use of this software connects back into human experience of the material world. We must keep our wits about us, lest we fall, unsafely and unaware, under the whim of computer and virtual reality.

I have neither the credentials, nor strong interest, to engage in the detail of philosophical debates, but have observed and followed some of their progression, and read many reams of printed pages through which they have played out. This chapter is built from and expresses that personal experience–no more than that.

11 E. Y.-C. Ho, 'At 90, and Still Dynamic: Revisiting Sir Karl Popper and Attending His Birthday Party', *The Karl Popper Web* (29 January 1997), http://www.tkpw.net/hk-ies/n23a/

When we talk of knowing, we debate theory in the realm of epistemology ('The theory of knowledge and understanding, especially with regard to its methods, validity, and scope, and the distinction between justified belief and opinion').[12] When we talk of truth, we jostle over belief within the realm of ontology ('The nature of reality and being'). With this linguistic etymology, science (*scire* [to know]) is bounded within the domain of epistemology. Epistemology seeks definition and is evolutionary. Ontology seeks truth and is a historical and intellectual battleground of competing perspectives and beliefs.

The nature of truth has been much debated by philosophers of logic and science. Colleagues tell me that: 'Is this an ontological question?' has become ever harder to answer in the Information Age, as we grapple with the space between the material reality of human experience and the virtual reality of the computer. The word ontology has become more narrowly appropriated to mean a set of concepts and categories in a subject area or domain that shows their properties and the relations between them. Concerning ontology, the Oxford Dictionary of Philosophy, 2nd Edition 2008, has this to say:

> Philosophers characteristically charge each other with reifying things improperly, and in the history of philosophy every kind of thing will at one time or another have been thought to be the fictitious result of an ontological mistake.[13]

Brief and to the point! It brings to mind the send up of modern-day political gobbledygook in Michael Dobbs's novel, *House of Cards*: 'You might very well think that. I couldn't possibly comment' (the catchphrase of the character Francis Urquhart).[14]

Francis Bacon (1561–1626), an English philosopher who wrote influentially about scientific method, developed the idea that knowledge should be comprehensively classified and universally shared, for the greater good of humanity. He has a close connection with my home city of St Albans, the location of the Roman city of Verulamium. He is remembered in the phrase: 'Knowledge itself is power', from *Meditationes Sacrae* (1597).

Bacon was devoutly religious, believing that knowledge is the rich 'storehouse for the glory of the Creator and the relief of man's estate'.[15]

---

12 'Epistemology', *Oxford English Dictionary*, https://www.oed.com/

13 S. Blackburn, *The Oxford Dictionary of Philosophy*, 2nd ed. (Oxford: Oxford University Press), p. 261.

14 M. Dobbs, *House of Cards* (London: Harper Collins, 1990).

15 F. Bacon, *Advancement of Learning*, ed. J. Devey (New York: P. F. Collier, 1901), p. 66.

He sought scientific method rooted in his Christian faith. This shaped his idea that philosophy and nature must be studied and reasoned about by process of induction, founded on unquestioned divine revelation, but illuminated through observation and experiment and leading incrementally to the steady accumulation of knowledge. That was his theory, and it suited orthodox thinking of his time.

Theory of Knowledge

I would be foolish and will certainly not attempt a comprehensive survey or summary of the theory of knowledge over the ages. My purpose, here, is a limited one. I seek to highlight its historical development from earliest times, alongside language, logic, mathematics and computer science, and its interface with the organization of books and documents, and now computerized databases and knowledge bases. I extend from this to its contemporary interface with life and medical science, health care and the computational problems these have encountered in their transition into the Information Age. The problems reflect, and reflect in, language and terminology descriptive of these domains and methods for representing and reasoning with medical knowledge, more generally. By analogy with databases that store and manage data and records, computer systems that represent and reason with knowledge have been termed knowledge bases.[16]

Theory and practice intertwine. Theory involves abstraction and simplification, consistent with the domain and purpose it serves. Leonardo da Vinci (1452–1519) wrote that:

> Those who are enamoured of practice without theory are like a pilot who goes into a ship without rudder or compass and never has any certainty where he is going. Practice should always be based upon a sound knowledge of theory.[17]

I am not sure if he was a sailor–possibly he was, given his wide range of talents and interests–but he would surely also have known that there is a lot more to sailing safely than knowledge about winds, headings, charts and weather forecasts. These must be experienced–the knowledgeable pilot on board is more oracle than mariner in good sailing and survival at sea.

16 A. Rector et al., 'On Beyond Gruber: "Ontologies" in Today's Biomedical Information Systems and the Limits of OWL', *Journal of Biomedical Informatics: X, 2* (2019), 100002, https://doi.org/10.1016/j.yjbinx.2019.100002.

17 Leonardo da Vinci, *Notebooks*, comp. I. Richter (Oxford: Oxford University Press, 2008), p. 212.

That said, Whitehead makes these points about theory of knowledge and method in matters of science and philosophy:

> A great deal of confused philosophical thought has its origin in obliviousness to the fact that the relevance of evidence is dictated by theory. For you cannot prove a theory by evidence which that theory dismisses as irrelevant. This is also the reason that in any science which has failed to produce any theory with a sufficient scope of application, progress is necessarily very slow. It is impossible to know what to look for and how to connect the sporadic observations. [...]
>
> No systematic thought has made progress apart from some adequately general working hypothesis, adapted to its special topic. Such an hypothesis directs observation, and decides upon the mutual relevance of various types of evidence. In short, it prescribes method. [...]
>
> A method is a way of dealing with data, with evidence. [...]
>
> Every method is a happy simplification. But only truths of a congenial type can be investigated by any one method or stated in the terms dictated by the method. For every simplification is an oversimplification. Thus, the criticism of theory does not start with the question, True or False? It consists in noting its scope of useful application and its failure beyond that scope. It is an unguarded statement of a partial truth. Some of its terms embody a general notion with a mistaken specialization, and others of its terms are too general and require discrimination of their possibilities of specialization. [...]
>
> In the preliminary stages of knowledge, a haphazard criterion is all that is possible. Progress is then very slow, and most of the effort is wasted. Even an inadequate working hypothesis with some confirmation to fact is better than nothing. It coordinates procedure.[18]

So much has changed in the one hundred years since Whitehead wrote these words about science and philosophy. Perspectives on the relationship of theory, practice and meaning have unfolded within new domains of knowledge, such as psychology, anthropology and behavioural science, forming fresh connections between science and society, politics and economy, and lifestyle and ecology. In human affairs, theory and practice coexist within the context of the art of the possible, and programmes for reform inherit from adventures of ideas. The story of the computer in medicine and health care is characterized by counterflows of knowledge and meaning, in both theory and practice. There is an adventure of ideas in theory extending towards a programme for reform of practice, and there

18 *Adventures of Ideas*, pp. 213–14.

is experience of meaning in practice, feeding back to illuminate and help improve both theory and practice.

The struggle to match theory, evidence and practice in medicine has been a recurrent and haphazard thread running through the Information Age, as it has explored and struggled with its roles and credentials that lie midway between life science and health care practice. This anarchic zone has been further highlighted and amplified by the haphazard nature of forays to computerize, as the capacity to make measurements has accelerated beyond imaginable bounds.

The first part of this book is mainly concerned with the adventure of ideas. The middle section highlights the anarchy that the computer has unleashed in the haphazard collision of new ideas with current practices. Programme for reform takes centre stage in the final part of the book, aiming and pointing towards a fruitful and meaningful experience and practice of personal health care for the individual citizen, drawn together around a more principled theory and practice of information that connects health care services with individual citizens, in their home settings and within populations.

Returning now to theory of knowledge, it is impossible for me to think, let alone start to write about knowledge, without feeling a sense of inadequacy, and awe of Bertrand Russell (1872–1970), who, with Alfred North Whitehead, did much to establish logical foundations of pure mathematics, as set out in 1910–13, in their *Principia Mathematica*.[19] My wider admiration of Whitehead will already be clear. Russell, too, was an extraordinary person, combining aristocratic demeanour, ruthless logical argumentation and loudly expressed pacifist belief.[20]

In 1946, Russell published *History of Western Philosophy*,[21] of which the historian G. M. Trevelyan (1876–1962) wrote in his review (as quoted in the cover notes for the seventh imprint): 'It may be one of the most valuable books of our age'. I have a copy from 1963 on my desk as I write, presented to me as the Bishop Burroughs' Prize for Science, and inscribed by my eccentric, classics enthusiast headteacher of the Bristol Cathedral School. It is one of my early inukbooks, spell binding in its incisive clarity, traversing two thousand years of Western philosophical thought.

19 A. N. Whitehead and B. Russell, *Principia Mathematica* (Cambridge, UK: Cambridge University Press, 1925–27).

20 K. Willis, 'Russell and His Obituaries', *Russell: The Journal of Bertrand Russell Studies*, 26 (2006) 5−54, https://doi.org/10.15173/russell.v26i1.2091

21 B. Russell, *History of Western Philosophy: Collectors Edition* (New York: Routledge, 2013).

Of course, there has been much else of philosophy argued through the seventy-five years since it was written, but this book is a marker, published at the very start of my life, and it seems appropriate to start from there, leaving perspectives arising from more recent relevant traditions of thought, such as philosophy of mind, to be introduced later in the book.

Russell writes: 'To teach how to live without certainty, and yet without being paralysed by hesitation, is perhaps the chief thing that philosophy, in our age, can still do for those who study it'.[22] For Russell, that aim is principally served through clarity and logical precision of thought and reasoning. He was not beyond the occasional incisive moral judgement, though! In the Preface to his *magnum opus*, apologizing for writing a book covering such a wide field, Russell says, 'If there is any unity in the movement of history, if there is any intimate relation between what goes before and what comes later, it is necessary, for setting this forth, that earlier and later periods should be synthesized in a single mind'.[23] That belief is to be pondered in our era of information explosion, where Deutsch speaks of knowledge as information with causative power. Is humankind any longer capable of such synthesis of knowledge, and thus awareness of its causative potential?[24]

Setting out his stall further, Russell writes:

> Philosophy, as I shall understand the word, is something intermediate between theology and science. Like theology, it consists of speculations on matters as to which definite knowledge has, so far, not been ascertainable; but like science, it appeals to human reason rather than to authority, whether that of tradition or that of revelation. All definite knowledge–so I should contend–belongs to science; all dogma as to what surpasses *definite* knowledge belongs to theology. But beyond theology and science there is a 'No Man's Land' exposed to attack from both sides; this No Man's Land is philosophy.[25]

22 Russell, *History of Western Philosophy*, p. 14.

23 Ibid., p. 7.

24 Advances in machine learning have caused the champions of the games of chess and Go to be deposed by machines, guided by the brilliant, Turing-like minds of those such as the inventor Demis Hassabis. As I revise this chapter, the company DeepMind has just triumphed again with AlphaFold, working from the known nucleic acid sequence of many millions of proteins to derive their three-dimensional topology. Of course, there are more dimensions of protein function still to unfold. Having accepted defeat to the machine in many areas, wise synthesis of different threads of knowledge is the next humanly-perceived boundary that is claimed to distinguish humans from machines—a noble thought somewhat thrown into question by the pervasive political turmoil of the summer of 2020!

25 Russell, *History of Western Philosophy*, p. 13.

According to Russell, philosophers through the ages have divided into disciplinarians–characterized as advocating some form of old or new dogma, that could not be proved empirically, and thus tending to be hostile to science–and libertarians–characterized as 'scientific, utilitarian, rationalistic and hostile to violent passion and profound forms of religion'. Russell's personal conclusion, in reviewing the ebb and flow of this philosophical debate was sombre, observing 'endless oscillation' between these 'partly right and partly wrong' parties.[26]

From its origins alongside those of science, in the philosophy of Thales of Miletus (*c.* 624 BCE–548BCE), Western philosophical thought has been dominated by a succession of 'ologies', 'doxies' and 'isms' (and 'isn'tisms'!), characterizing perennial dispute! Religious doctrine became dominant up to and through the Middle Ages, and then, through Renaissance, Reformation and Enlightenment times, science and logical thought became more dominant. Nowadays, philosophical debate is also framed within the scientific and technological contexts of the Information Age. As Russell remarks, philosophical arguments seldom achieve resolution. Some may, in time, come to be conceived as wrong-headed or logically unsound–he is cholerically vituperative about what he sees as non-sense, or ideas contrary to common-sense–but may, even so, re-emerge, re-framed and re-expressed, and argued anew, in contemporary contexts.

In thirty-one chapters comprising over eight hundred and forty-two pages, Russell situates each philosopher as the product of the *milieu* of their times, in whom 'crystallized and concentrated thoughts and feelings which, in a vague and diffused form, were common to the community of which he was a part'.[27] In turn, he observes 'a reciprocal causation: the circumstances of men's lives do much to determine their philosophy, but, conversely, their philosophy does much to determine their circumstances'.[28]

Much philosophical argument centres on the use of words and language. In Grecian times, philosophy was propagated in open debate, and by people moving from city to city. In increasingly stark contrast, much knowledge and theory of knowledge in the age of the computer (such as knowledge entailed in artificial intelligence algorithms) is moving out of reach of communities of scholarship, to be embedded in software and owned or appropriated as property of commercial organizations with global reach. This is a new structure of scientific revolution, with considerable practical and philosophical implication for the balance of public and private enterprise in society. At many points as the storyline of the book develops,

26 Ibid., p. 22.
27 Ibid., p. 7
28 Ibid., p. 14.

it highlights contemporary movements to fulfil Bacon's vision of open knowledge within Popper's vision of Open Society, expressed variously as open access to knowledge, open data and open-source software. The key to this future, if that is the way society evolves, is in the three taken together. The expression of knowledge, the data on which it rests and the software used to reason about and utilize the two, must demonstrate their mutual connections, to justify their correctness and useful application.

Consideration of the nature of knowledge devolves into that of truth and belief. Chapter thirty of Russell's book is devoted to John Dewey (1859–1952), a New Englander most known for his book *The School and Society*, and much admired by Russell, personally and professionally.[29] Dewey was the leading proponent of instrumentalism and a critic of traditional notions of truth as 'static and final; perfect and eternal'. Russell writes:

> Since Pythagoras, and still more since Plato, mathematics has been linked with theology, and has profoundly influenced the theory of knowledge of most professional philosophers. Dewey's interests were organic rather than mathematical, and he conceives thought as an evolutionary process, and human knowledge 'as an organic whole, gradually growing in every part, and not perfect in any part till the whole is perfect'.[30]

Dewey maintained that inquiry is the fundamental concept of logic and theory of knowledge. Russell writes that Dewey defined this as follows: 'inquiry is the controlled or directed transformation of an indeterminate situation into one that is so determinant in its constituent distinctions and relations as to convert the elements of the original situation into a unified whole'; and later, 'unified wholes are to be the outcome of inquiries'.[31] Russell took issue with Dewey over the primacy of inquiry, substituting his own concept of truth as the yardstick of logic and theory of knowledge. This got him into further mathematical and logical complexity, in tackling so-called Russell Paradox that his ideas led him to puzzle over. The instrumentalist view was also criticized by George Santayana (1863–1952), who Russell quotes as saying: 'In Dewey, as in current science and ethics, there is a pervasive [...] tendency to dissolve the individual into his social functions, as well as everything substantial and actual into something relative and transitional'.[32]

29 J. Dewey, *The School and Society and the Child and the Curriculum* (Chicago, IL: University of Chicago Press, 2013).
30 Russell, *History of Western Philosophy*, p. 775.
31 Ibid., p. 778.
32 Ibid., p. 781.

As we delve further into the nature of truth and falsity of knowledge, we encounter debates about language, and words and terms, such as: belief, meaning, significance, thing, fact, common sense and reality. Profusion of differently understood terms also characterizes the discussion of knowledge in the context of computation and I return to this problem in the discussion of computer knowledge bases, later in the chapter.

In 1912, Russell advanced the correspondence theory of truth, writing that: 'truth is understood in terms of the way reality is described by our beliefs. A belief is false when it does not reflect states-of-affairs, events, or things accurately. In order for our beliefs to be true, our beliefs must agree with what is real'.[33] In *History of Western Philosophy*, Chapter XXX, he produced examples whereby Dewey's reasoning leads to positions he believed absurd, based on what he saw as common sense, writing that:

> Dewey's divergence from what has hitherto been regarded as common sense is due to his refusal to admit 'facts' into his metaphysic, in a sense in which 'facts' are stubborn and cannot be manipulated. In this, it may be that common sense is changing, and that his view will not seem contrary to what common sense is becoming.[34]

These words chime with the buzzwords of recent months–'alternative facts' and 'fake news'.

Chapter thirty-one is devoted to the philosophy of logical analysis, in which school of thought Russell was a leading figure. It has come to the fore in formal logic of the Information Age, which has been influential in the evolution of knowledge bases. Russell describes its origins in the work of mathematicians such as Karl Weierstrass (1815–97), in placing the infinitesimals of the calculus of Gottfried Wilhelm Leibniz (1646–1716) onto a logically secure foundation, and Georg Cantor (1845–1918), whose theory of continuity and infinite number, he says, put to bed a good deal of 'muddled and mystical philosophical musings' over the ages! He did not mince *his* words!

The historic connections of knowledge with philosophy, language, mathematics, logic and reason are mirrored in the transition into the world of the computer, and its languages and logics. This transition is developed further in sections below, but first I make another historical detour, this time into the world of librarianship and the age-old struggle to organize the increasing scope and proliferation of books and documents. There are illuminating parallels between this story and the quest to curate, classify

33 B. Russell, *The Problems of Philosophy* (Oxford: Oxford University Press, 1912), Chapter XXXI (n.p.).

34 Russell, *History of Western Philosophy*, p. 780.

and reason with knowledge, using computers in the Information Age, very much a story playing out in medicine and health care of the Information Age.

Libraries and their Classification of Knowledge

> A book is a fragile creature, it suffers the wear of time, it fears rodents, the elements and clumsy hands. So, the librarian protects the books not only against mankind but also against nature and devotes his life to this war with the forces of oblivion.[35]

Librarians provide access to knowledge and maintain order in and between bodies of knowledge and the communities of their producers and users, for research, education and practice. Libraries are also often custodians of historically important collections of books and other artefacts. Wear of time, rodents, elements and clumsy hands all have their correlates in the computer age! The forces of oblivion do not change much; ultimately, they are just physics! The protection of knowledge remains hard and complex, spanning lifetimes of work.

I have chosen to continue this chapter's exploration of knowledge in the world of libraries, and not just because I like books and spend a lot of time with them. My purpose is to draw out parallels between difficulties faced by librarians over the ages, in organizing the storage and retrieval of the written word, and those faced in the evolving Information Age, in envisioning and creating computerized databases as well as knowledge bases.

The advent of electronic publishing has transformed libraries into organizations that connect continuously, both locally and globally, with many communities and languages of disciplines and users. This has brought new challenges to the standardization of methods, bridging from local to global content and organization. Library information systems are, perhaps, not quite as complex in their semantics as are health care record systems, but they are challenging, nonetheless.

Mine is a small personal collection of books, and libraries have always curated collections specific to the interests and needs of their client users. In medicine, the specialist libraries of the Royal Society of Medicine, the British Medical Association, The Wellcome Trust and the Royal College of Physicians are well-established, located close by to my university (University College London). As with many capital cities, there are myriad such libraries, both specialist and general, in central London. As with any

35 U. Eco, *The Name of the Rose* (London: Pan Books, 1984), p. 38.

university, in ours there are multiple libraries distributed across the separate academic campuses and within the academic departments established there over many decades. They are places to work in as much as to access books. Coherent and standardized information technology (IT) infrastructure and services are now crucial for both purposes. Mergers of previously independent institutions and library teams, changing academic needs over time and the alignment of library systems with IT infrastructure more generally have posed major transitional challenges on many simultaneous fronts, over many years.[36]

The world of books and documents and the world of data have increasingly converged during the Information Age. This convergence requires curation that connects knowledge with the data and methods on which it is based, so that findings can be replicated, and the knowledge created more openly shared and accessible, within and between communities of discipline, practice and use. Science places a high priority on the ability of different teams to replicate one another's experimental findings. A recent study showed that fewer than half the results published in twenty-three highly cited papers in preclinical cancer biology research could be successfully reproduced.[37] Such findings are of increasing concern and reinforce the need for curation of both knowledge and the data on which it is based, when publishing research findings.

The history of libraries deserves to be appreciated and learned from in the interconnected domains of knowledge and practice for health care in the Information Age. Books in libraries are not databases but there are some analogies. Waving hands somewhat: persistence of data is the placing of books on shelves; indexing of data is the tagging or coding of books and documents within the library classification system; inheritance of properties of data is the reuse of patterns of subdivision in library classifications.

36 During the course of my career in London, latterly at UCL from the mid-1990s, I worked in many libraries and with their communities of librarians and users. I observed and helped in the transition of university and national libraries, and those specific to medicine, through the Information Age. In those years, I passed and saw daily the emergence of the magnificent new British Library, close by to UCL and St Pancras station. The head of information systems, there, approached me to join and then chair a group overseeing the development of a new system for the PubMed electronic library in the UK, and also serve on the UK advisory board for a Research Information Network across disciplines. The interface of library IT systems with the wider IT infrastructure of our university was a world I experienced closely as a member of its Information Strategy Committee, where I chaired the Information Infrastructure sub-committee for some years.

37 T. M. Errington et al., 'Reproducibility in Cancer Biology: Challenges for Assessing Replicability in Preclinical Cancer Biology', *Elife*, 10 (2021), e67995, https://doi.org/10.7554/eLife.67995

Library history demonstrates the importance of interpreting knowledge within historical contexts: assumptions made and understandings reached; skill and motivation of writer, curator and reader; resource deployed; power exercised. The efforts made to bring sustainable and useful order to library classifications, notably over the past one hundred and fifty years, parallel attempts over recent decades to systematize the language, methods and procedures used in organizing medical knowledge and information systems supporting medicine and health care. There is a continual interplay of theory, pragmatism and edict in these stories.

It is understandable that such breadth of ambition and quest for generic method, proves too risky and intractable for many to encompass. An often-made pragmatic compromise is to narrow the focus and effort into a domain-by-domain approach. That was the decision of the US Library of Congress a century or so ago, in the world of library classifications. It is often the only practical way forward–grand schemes addressing grand challenges are risky and prone to failure! However, the price of pragmatic simplification in the shorter term can be escalating confusion in the longer term. The underlying problem does not go away–it is ignored or deferred down the road until another day. And of course, that day will bring new contexts–including new kinds of problems.

The experience gained and limitations encountered when attempting to formulate and refine a useful and applicable general method, by conducting experiments that implement and use the proposed method in practice, play out in context of the motivation and capability of the participants in the experiment and the availability of resource. The outcomes inform choices that are made about the standardization of method. The twentieth century history of the computer system and its designers and users, in creating and operating health care-related knowledge bases and maintaining records of practice, has interesting historical parallels in library science. It reflects issues of scope, rigour, flexibility to change, cost, utility and governance. In the Information Age, librarianship has extended to open curation, access to and governance of electronic sources of knowledge and data, for example in the arenas of Creative Commons and open-access publications and data. Citizens, in both their personal and professional roles, can now more readily, and usefully, take part as creators, reviewers, curators and governors of these resources. Citizen science is emergent in many fields.

In the next section, I trace the story of the organization of knowledge within libraries, as book and document stores stretching back over more than two thousand years. I do this with Edward Gibbon (1737–94) as my early guide, and later draw inspiration from conversations I had some decades ago with a friend who played a leading role in the UNESCO project entitled the Broad System of Ordering (BSO). This initiative had, and retains, a close link with my alma mater, University College London (UCL), to which its copyright is now assigned.

Historical Origins

Gibbon's one thousand-page history, written from 1783–88, has been my introductory inukbook and source for the following survey.[38] His writing is often richly polemical and controversial, but some of the ancient sentiments expressed amuse, more than offend, and some ring true in the frustrations of our age!

In Gibbon's discussion of progress in the sciences, he adds, dismissively, that:

> The libraries of the Arabians, as with those of Europe, were possessed only of local value, or imaginary merit [...] The shelves were crowded with orators and poets, whose style was adapted to the taste and manners of their countrymen; with general and partial histories, which each revolving generation supplied with a new harvest of persons and events; with codes and commentaries of jurisprudence, which derived their authority from the law of the prophet; with the interpreters of the Koran, and orthodox tradition; and with the whole theological tribe, polemics, mystics, scholastics and moralists, the first or the last of writers, according to the different estimate of sceptics or believers.
>
> The physics, both of the Academy and the Lycaeum, as they are built, not on observation but on argument, have retarded the progress of real knowledge. [...] the human faculties are fortified by the art and practice of dialectics; the 10 predicaments of Aristotle collect and methodise our ideas, and his syllogism is the keenest weapon of dispute. It was dexterously wielded in the schools of the Saracens, but as it is more effectual for the detection of error than for the investigation of truth, it is not surprising that new generations of masters and disciples should still revolve in the same circle of logical argument. The mathematics are distinguished by a peculiar privilege, that, in the course of ages, they may always advance and can never recede.[39]

There have been efforts to classify writing, in the form of clay tablets, papyrus manuscripts, documents and books, from earliest times. As such archives grew, this became a significant challenge, and the library and librarian profession were born. No small undertaking–essentially any

38 E. Gibbon, *The History of the Decline and Fall of the Roman Empire* (London: Strahan and Cadell, 1788).

39 Ibid., p. 982.

writing, from any domain of scholarship, be it in a narrow discipline or multi-disciplinary, needed to be placed somewhere, and known about so it could be discovered, retrieved and used.

Where there is a relatively small set of items to store, one can stack them in piles–just as the numerous source books are stacked as I write and refer to them, here, and rely on the human eye to find and retrieve them. Maybe one can group them in some way, again in piles, as I have done for each chapter, or along shelves in order of size or date of publication, or alphabetically by the name of the author. What about finding the book according to the name of the second or subsequent authors, though, and how should they be managed when they cross different zones of classification? Creating lists and indexes to keep track of library contents became a necessity.

Bibliographies and indexes have grown in scale and detail as human knowledge has grown. Modern day data processing has brought new opportunity and likewise also become a necessity in keeping track and enabling access. Searching recent medical literature just fifty years ago was a laborious and time-consuming process, involving scanning of the microscopic print of hefty annual indexes with a magnifying glass, on frequent treks to libraries.

Clay tablets in the royal archives of the Assyrian king, Ashurbanipal (685 BCE–631 BCE), were organized in a catalogue divided into classes of grammar, history, law, natural history, geography, mathematics, astronomy, magic and religious legends, each divided into subclasses. The poet and scholar, Callimachus (*c.* 310 BCE–240 BCE), is reported as having organized the Great Library of Alexandria, using a classification of poets, law makers, philosophers, historians, rhetoricians and miscellaneous writers, subdivided by form, subject and time. According to Gibbon, the library was described by the Roman historian and philosopher, Livy (*c.* 59 BCE–17 CE] as *elegentiae regum curaeque egregium opus* (Google translates this as 'The elegance and care of kings, an excellent work'), an encomium with which the Stoic philosopher Seneca (*c.* 1 BCE–65 CE) disagreed. That is the way with philosophers! Gibbon has a waspish turn of phrase, here, criticizing that their 'wisdom, on this occasion, deviates into nonsense'![40]

Liu Xiang (77 BCE–6 BCE, a Chinese astronomer, poet, politician, historian, librarian and writer) and his son Liu Xin (*c.* 50 BCE–23 CE, a Chinese astronomer, mathematician, historian, librarian and politician) devised the first library classification for the Seven Outlines ('Qi Lue') in the Han Dynasty of China. Libraries of China in the West Han period, around the first century CE, used a classification of philosophy, poems and songs, military art, sooth saying and medicine.

40 Ibid., p. 956.

Gibbon gives a lively story of the historical context of those times, particularly the Saracen invasion of Egypt in 638 CE, led by Amrou (Amr ibn al-As al-Sahmi, *c.* 573 CE–664 CE), the defeat and retreat of the Greek rulers, and the demise of the Great Library of Alexandria. This history has been substantially pawed over and re-written in modern times, but Gibbon records that Amrou discussed the status of the library with John Philoponus (490 CE–570 CE), a Byzantine Alexandrian philologist, Aristotelian commentator, and Christian theologian, famed for his 'laborious studies of grammar and philosophy':

> Was the library an inestimable gift, in the eyes of the Greek, or contemptible, in keeping with the contempt for idols of their conquering successors? The Caliph Omar was consulted and gave the opinion that 'if these writings of the Greeks agree with the book of God, they are useless and need not be preserved: if they disagree, they are pernicious and need to be destroyed'.[41]

Magisterial having and eating of cakes! As the story goes–which Gibbon doubted, but it is a good story–the paper and parchment was distributed to four thousand baths in the city and fueled their heating for six months! One wonders how many baths the hot air associated with today's Cloud-based knowledge stores could heat!

Culture, belief and learning went hand in hand and stirrings of science occupied a lowly position in the order of things. The early astronomers maintained their credibility by combining their observations with astrology and the mysticism of the Zodiac. They were careful not to challenge the abstract geometry of the heavens described by Claudius Ptolemy (*c.* 100 CE–170 CE). Discovery of the solar system rested in wait until the Copernican revolution broke the mould of established doctrine many centuries later. The books and documents were collected and organized in line with the predilections and pragmatic choices of their times; the common goal was to give each item a place in a collection.

Around the turn of the sixteenth century, in the England of Elizabeth I (1533–1603), the classification of knowledge proposed by Bacon sought to organize all types of knowledge into groupings of history, poetry and philosophy. He understood information to be processed through human memory, imagination and reason, but his methods for the categorization of knowledge were based on inductive principles of experiment and reasoning, proceeding from divine revelation, which he insisted on as the basis of his scientific method.

41 Ibid., p. 952.

William Torrey Harris (1835–1909) built on Bacon's ideas about knowledge structure and scientific method. He created a library catalogue for the St. Louis Public Library School and his ideas were widely influential. He proposed a practical system of rules for the classification, ranging from the generic to the specific. There were main divisions, ultimate divisions, appendices and hybrids. Bacon's approach sought to define all knowledge within a predetermined structure of classification. Harris used generic main divisions to provide a 'guiding principle' of the form of knowledge and dealt with the detail of knowledge content more flexibly in minor divisions and sections.

Enumerative and Faceted Classification

In this section, I explore the origins of more formal library classification systems–how they describe the subject matter of books and documents such that they can be placed efficiently within a library collection and readily discovered there by its enquiring users. In passing, we might note that the computer has a similar problem to solve with data, as we come to consider in Chapter Five. In that context, the computer, as data librarian, must be able to decide how and where to store data as efficiently as possible, in different data storage media, such that they can be efficiently managed and retrieved, as required, for use in the computations specified by its programs, as the data library users.

Enumerative classification focuses on the place for the book. Faceted classification focuses on the content of the book.[42] The section draws on historical detail from *The Encyclopaedia of Library and Information Sciences*;[43] a more recent appraisal of theory and practice of library classification schemes is provided by S. Batley.[44]

First, a simple and rather fanciful example to set the scene. Imagine you are a librarian, and you receive a shipment of new books for the library from a publisher, packed in a large box. You open the box, take out a book and

42 In library classification, 'facet' refers to a particular aspect of a subject or train of characteristics—e.g., in literature, there may be four facets: language, form, author and work. An enumerative classification contains a full set of entries covering all defined concepts. A faceted classification uses a set of semantically cohesive categories that are combined as needed to create an expression of a concept. In this way, the faceted classification is not limited to already defined concepts. Wikipedia contributors, 'Faceted Classification', *Wikipedia, The Free Encyclopedia* (24 May 2023), https://en.wikipedia.org/wiki/Faceted_classification

43 R. Wedgeworth, ed., *World Encyclopedia of Library and Information Services* (Chicago, IL: American Library Association, 1993).

44 S. Batley, *Classification in Theory and Practice* (Oxford: Chandos Publishing, 2014).

glance at the title and brief description of the work on its cover, and perhaps the table of contents. It is a book about Euclidean geometry, say–geometry as a subdomain of mathematics. It might devote some of its content to the history of mathematics in classical times, in Greece. Leaving aside all the administrative steps in registering the book in a library catalogue, your imagined task is to take the book up staircases, along corridors and into aisles of shelving, to place it within the library. It must go somewhere. There is a limited number of spaces available.

One can imagine several possible procedures–the first, a purely pragmatic one. You and your librarian colleagues have previously put your heads together and decided on the layout of the library, dividing it into a fixed number of zones for the principal subject domains: mathematics, literature, science, history, technology etc. These zones may be associated with separate buildings, floors or rooms of the library. In each zone, you have subdivided the space available: mathematics might encompass subsections of algebra, probability, numerical methods, combinatorics and geometry etc., again each with a limited and fixed capacity to house books. Likewise in the history zone, its space has subsections organized by time, region of the world, and kinds of history: social, military, economic etc.

Further choices have been made to define the structure of this very inflexible imaginary library, allocating a numeric code to each book position in the subsection of the library in which the new book is to reside. Your thinking heads have thought through this conundrum and decided upon a divisional structure to deal with all possibilities, and how many book slots to allocate to each subdivision. It is an 'enumerative coding system' (i.e., based on numbers, but essentially an orderly set of symbols) that covers all the possibilities and expresses the structure of the cascading subdivisions of the library book locations ('pigeonholes').

You opt to classify the new book under history and allocate the code that then guides you through the labyrinth to the fixed slot on the shelf, within the room, on the floor, and within the building that is to be this book's home. And so on with the rest of the box of new accessions to the library. I did say it was fanciful!

Assuming there is a space available, this works for placing the book, but there are further difficulties: what about the library users, who come to the library with a topic in mind in search of relevant books? A history student enters looking for information about Euclid's place in the history of Greece, they browse the history shelves dedicated to history of mathematics in classical times. Had you opted for this book to be housed in the mathematics building, they might not have found it so easily. Likewise for a mathematics student who seeks out a book comparing the pros and cons of Euclidean geometries alongside non-Euclidean geometries–they might not have been so lucky, after trekking expectantly to the mathematics building. It is unlikely that many users will be sufficiently determined and resilient to visit every building and browse the shelves there, according to the possible combinations of topics whereby the sort of book they seek might be located.

Both librarian and student have choices to make; they get harder as the scale and granularity of content in the library collection grows in terms of numbers of books, diversity of subjects covered, and their interconnections–history connecting with mathematics, science connecting with technology, politics and economics connecting with pretty much anything. New subjects arise that undermine the integrity of the structure that has been imposed. If one part of the enumerated code has space for four numeric digits, the ten thousandth book that would legitimately be represented by that code segment will have nowhere to go. A valid book classification has overflowed the fixed number of slots available for books thus classified. The material needing to be positioned cannot be accommodated within the structure of permissible classifications of content. The library may yet have spare slots elsewhere–perhaps books about performance of Beethoven symphonies in nineteenth-century Tbilisi have not yet filled their allocated slots!

In theory, you and your librarian colleagues could juggle the enumerated structure of the classification to use the available number of slots more efficiently, moving the books around accordingly. But a 'general post' of books, relocating them to different shelves of a dynamically evolving library, to provide valid positions for new books, is not an attractive option. The frequent users of a particular reference book would be unhappy for it to be moved from place to place. You fiddle with the scheme, and over time it becomes untidy and does a poorer and poorer job.

An alternative strategy might be to disconnect the problem of classifying the book content from that of allocation of slots on bookshelves. Each book is to be uniquely classified according to different facets of that content. In principle, such flexibility should allow for the addition of new divisions and subdivisions of content of the book. But, as ever, the devil is in the detail; the choice of available facets and how they are combined become rather fundamental issues.

A cut diamond is structured with many facets (faces) to reflect and channel incident light that passes through and issues from the whole diamond, in different ways and directions. The quality and appeal of the diamond is expressed and perceived via the cutting of its facets. The content of the book is expressed in a common language of facets. The user can interrogate the facet-based classification and the language can extend in time to introduce new domains of content, and connections between them, that the enumerative procedure is not equipped to handle. We still have the problem of optimal physical placement of books within library collections, but this can now become a separable concern, no longer a tail wagging the dog of a satisfactory system of content classification. An automated book storage and retrieval system sounds a good idea–well, the database engineer and Amazon warehouse manager think that, too, and it is a more realizable one in their worlds of data storage persistence and warehousing, where there are no browsing users who persist in liking to experience the look and feel of books!

This has been a rather artificial and hand-waving introduction to a complex field with a complicated and anarchic history. The librarians' problem over the ages has emerged from and intertwined with the unfolding world of knowledge and technology which they struggle to curate and manage. In passing, we might note analogy with the computer's problem, as knowledge and data librarian, emerging from and intertwined with that same world. There is a difference, of course. The book librarian is not responsible for the problem they confront, but, in the Information Age, the computer is closely implicated in the problem–itself integral to the creation of the exponentially growing body of knowledge and data that it struggles to curate and manage. We are sometimes a bit like the latter-day Christopher Marlowe's (1564–93) Dr Faustus: frustrated with the vicissitudes of medicine, law, logic and theology, seeking to acquire magical mastery over the world, accepting, albeit with similar repeating misgivings, the services of a Mephistophelian computer!

Let us not dwell here on Faustian bargain and fate–it does not have to be that way, but we should be aware and beware! The bargain with the computer can penetrate deeply into health care services. The complexity of challenges to their balance, continuity and governance in the Information Age reflects the intertwinement of problems in management of the data explosion created by and with the computer, and the battle for understanding the proliferating detail and nuance of practice that it creates and exposes. Best not to pass that problem back to Deep Mephistopheles for resolution. We need human hands on how we judge and contain the fractal complexity of data and knowledge–otherwise we risk escalating and intractable battle between signal, bias and noise in human judgement, of the kinds that Daniel Kahneman and colleagues are signalling.[45]

The library story illustrates a general tension between an aesthetically appealing, open and theory-based approach, enabling any book or document to be classified as exactly as is desired, and a more tied-down and pragmatic approach that limits classification in a predetermined manner. It affords no wholly satisfactory solution and requires compromise. It remains

45 D. Kahneman, O. Sibony and C. R. Sunstein, *Noise: A Flaw in Human Judgment* (New York: Little, Brown Spark, 2021).

a continually evolving story and how it has played out is best followed along its historical timeline, as, for example, set out in *The Encyclopaedia of Library and Information Sciences*, which I have used, here. There has been a mixture of enumerative and facet methods of classification. Shiyali Ramamrita Ranganathan's (1892–1972) method of colon classification, the subject of its own section below, stayed true to a wholly facet-based approach but proved too challenging to implement at scale in practice.

There are instructive parallels between this story and that of medical language and the terminologies and classifications descriptive of medical knowledge and clinical practice, as told in the succeeding sections. Computer-based knowledge and library management systems have evolved to tackle the limitations of enumerative and facet-based methods, enabling new tools that work better, now, in both library curator and user contexts. In the medical domain, formal logics have emerged to play a new part in taming the complexities of medical language and knowledge bases. This is a topic that joins my storyline as it moves on into the world of mathematical and formal logic. For now, we continue with the history of library classifications.

Melvil Dewey (1851–1931) was a pioneer of educational reform and librarianship. An early trailblazer in library classification systems was the 1873 Dewey Decimal Classification (DDC), originally enumerated in a thousand subdivisions over twelve pages and criticized at the time as overly detailed! Here is how the Encyclopaedia describes it:

> Dewey's innovation was to use numbers (decimal fractions) as subject (class) markers, infinitely expandable in size, over time, within existing classes and with limited ability to expand within a hierarchy of classes, integrating new subjects within a single unified scheme. Further detail could be added in ancillary tables. And some representation of relationships between subjects was provided for by allowing the subdivision of one class with numbers built (inherited) from another.[46]

As with the much earlier examples, the chosen organization reflected the assumptions and outlook of the times. However, the principle adopted–of maintaining integrity of the numbers–prevented restructuring of outdated classification schedules to incorporate new subjects. Cognate areas of knowledge, such as technological applications of basic sciences, were separated in the number schemes. And new ways of providing generic patterns of structure, such as faceted classification and the later contribution of Ranganathan, with his colon coding scheme, were not well

46 Wedgeworth, ed., *World Encyclopedia*, pp. 209–12.

accommodated. The DDC received powerful backing from the US Library of Congress, which established an organization to take the work forward.

In the 1880s, seeking to break away from the pragmatic (enumerative and pre-coordinated) 'pigeon-hole' filling approach of DDC, Charles Ammi Cutter (1837–1903) proposed an Expansive Classification (EC), aiming to reflect 'evolutionary order in nature'. This initiative did not survive his death, but his idea of encompassing a more philosophically enriched, ideal ordering of subjects influenced subsequent policy and developments at the Library of Congress and the work of Henry Evelyn Bliss (1870–1955) in the United Kingdom (UK).

In the 1890s, seeking an international approach, a Universal Decimal Classification (UDC) was proposed under the auspices of the International Federation for Documentation (FID), now the International Federation for Information and Documentation. Substantially but not fully DDC-compatible, this introduced colon notation to link two or more codes. It was championed by the Union of Soviet Socialist Republics (USSR), which made it mandatory in 1963, and had a wide-ranging user base in Eastern Europe, Japan, Brazil and Latin America.

With the pragmatic foundations of the Dewey system proving difficult to maintain and sustain in the changing and rapidly growing libraries of the time, and with Dewey himself unwilling to agree to substantial revision of his scheme, the US Library of Congress Classification (LCC) proceeded to introduce a Federation of twenty-one loosely coordinated classifications. The notation adopted was a mix of letters and numbers for main classes and it left space for expansion. Each classification adopted its own approach to subdivision of classes. Management of the classifications became rather haphazard, with arbitrary use of vacant spaces, deletions and reuse of blocks of allocations and movement of subjects to different schedules.

Bliss devoted his lifetime of work to devising a scheme of bibliographic classification (published from 1935–53) to represent the 'order of things and ideas'. This had twenty-six main classes (A/Z) and anterior classes (1/9). It was flexible in allowing alternative locations or treatments for many subjects. Though considered a significant advance in the underlying principles adopted, the investments of US libraries in DDC or LCC made it infeasible for them to branch their efforts to a new and unproven system. A hundred UK libraries made slow progress with its improvement.

Thus far in the library story, some common themes and stages have started to emerge.

- Pragmatism–every book or document must have a place in the library and the classification system serves the main purpose of defining and providing that place: a set of 'pigeon-holes' such that everything can find its home somewhere;

- Idealism–whereby the method of classification seeks to be configurable and evolvable, providing a coherent description of the content of any book or document, based on an underlying theory;
- Context and choice–decision about what constitutes a logically consistent, practically achievable and useful home, and how that should be coded for within an index, to cater to prevailing needs, ideas and cultures;
- Pattern–in exploring and experimenting with different systems, the possibilities, and their relative strengths and limitations, emerge over time and general patterns crystallize;
- Growing pains–strength of personality and the commitment and staying power of innovators, combined with organizational and national rivalries, assert themselves;
- Power–sponsorship is dominant.

A not dissimilar story to this has played out along my songline, in its encounter with endeavours to formalize the description of medical knowledge and link this with records of health care practice. Progress in such domains is made slowly and then in jumps, as in the Niles Eldredge and Stephen Jay Gould (1941–2002) characterization of 'punctuated equilibrium' in biological evolution.[47] As for the state of the art today, one might remark, as the Chinese leader Zhou Enlai (1898–1976) was said to have done, when asked his opinion of the success of the French Revolution: 'it is too early to tell'![48]

The Colon Classification of Shiyali Ramamrita Ranganathan

Trained as a mathematician, Ranganathan worked as a librarian in Bombay and is recognized as a founding father of modern-day librarianship, which he called library science. He is credited as the last person to single-handedly envision and enact a library classification that was used in practice. He is also credited as the person who broke decisively with the pragmatic tradition of

47 S. J. Gould and N. Eldredge, 'Punctuated Equilibrium Comes of Age', *Nature*, 366.6452 (1993), 223–27, https://doi.org/10.1038/366223a0

48 It was a misunderstanding, apparently, but a good story, nonetheless (see, further, 'Not Letting the Facts Ruin a Good Story', *South China Morning Post*, https://www.scmp.com/article/970657/not-letting-facts-ruin-good-story).

classification in libraries and sought an underlying theory. Encyclopaedia Britannica records his influence thus: 'Perhaps the most important advance in classification theory has been made by the Indian librarian, SR Ranganathan, whose extraordinary output of books and articles has left its mark on the entire range of studies from archival science to information science'.[49]

I discovered an extensive archived collection of his works online at the University of Arizona (search keyword Ranganathan at repository.arizona.edu), but little evidence survives elsewhere today. He introduced his facet-based colon classification of library contents in six editions between 1933 and 1960.[50] In this system, facets describe 'personality' (the most specific subject), matter, energy, space and time (PMEST). These facets are generally associated with every item in a library, and so form a reasonably universal sorting system.[51]

Ranganathan grounded his ideas in what he set out as Five Laws of Library Science: books are for use; books are for all; they should be openly accessible as if in the reader's private library; and organized to protect the reader's time from laborious search. Finally, the library should be seen as a growing organism and thus needs to be organized around strong enough and flexible enough foundations, so that its communications could be complete, concise, considerate, concrete, courteous, clear and correct. He certainly knew the alliterative power of C lists! His sentiments were sound and the organic characterization resonates strongly with similar requirement for life-long digital care records.[52]

In anticipation of the later UNESCO Broad System of Ordering (BSO), he set out to devise a theory-based method for expression of the content of books, and thus of the full range of knowledge contained there. He conceived of a circle of knowledge, which was described to me as a twenty-four-hour clock face, starting with philosophy at midnight and proceeding counterclockwise through successive domains of knowledge. The hour hand moved on through mathematics and sciences devoted to theory and experiment, from the physical world into the living world, and

49 D. J. Foskett, 'The Dewey Decimal System', *Britannica*, https://www.britannica.com/topic/library/The-Dewey-Decimal-system

50 S. R. Ranganathan, *Colon Classification*, 6th ed. (Bangalore: Sarada Ranganathan Endowment, 1989).

51 Wikipedia contributors, 'Colon Classification', *Wikipedia, The Free Encyclopedia* (7 May 2023), https://en.wikipedia.org/wiki/Colon_classification

52 S. R. Ranganathan, *The Five Laws of Library Science* (Bangalore: Sarada Ranganathan Endowment, 1989) https://repository.arizona.edu/handle/10150/105454; S. R. Ranganathan, *Philosophy of Library Classification* (Bangalore: Sarada Ranganathan Endowment, 1989), https://repository.arizona.edu/handle/10150/105278

then to education and social sciences. At about twelve noon it moved into demography, politics and law.

It then proceeded through economics and finance and on to technologies and industries, language and literature, arts and religion, and finally to the occult and mystical, where, at the reverse striking of midnight, it emerged again into the world of philosophy. Here is a slide I constructed to illustrate this ordering; a previous version was used in my lectures of thirty years ago (see Figure 2.2).

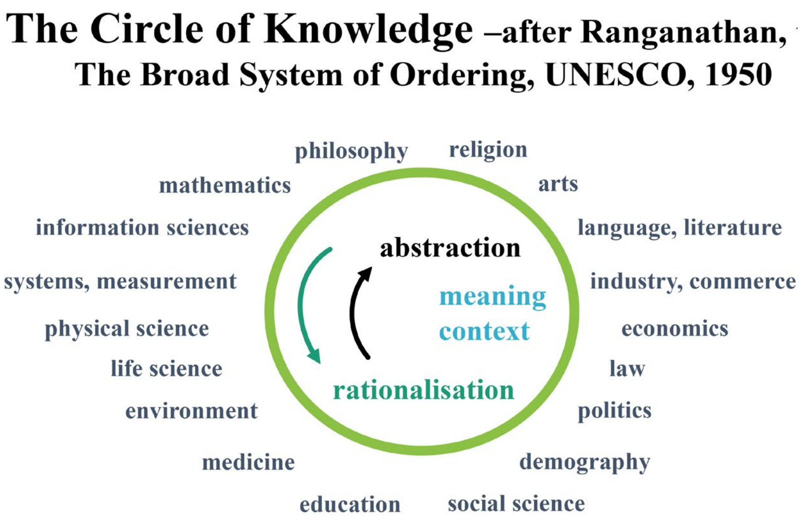

Fig. 2.2 The Ranganathan Circle of Knowledge. Image created by David Ingram (2022), CC BY-NC.

Ranganathan set out to create a set of general principles for an evolving and enduring classification of books and documents. He abandoned the pragmatic approach of enumerating a set of fixed pigeonholes, each for a preconceived and precoordinated class of things and ideas, with ever more detailed subdivisions. By contrast he proposed what he termed an 'analytico-synthetic method of classification', the Colon Classification. Analytico-synthetic scheme, according to Ranganathan, is used 'to denote any scheme in which a compound subject is first analysed into its facets in the idea plane, and later synthesized in the verbal plane and in the notational plane respectively'.[53]

53 S. R. Ranganathan, 'Colon Classification Edition 7 (1971): A Preview', *Library Science with a Slant to Documentation*, 6 (1969), 205.

The term resonates with philosophy, descriptive of human reasoning (particularly that of Immanuel Kant (1724–1804))–involving analysis as a rational process working from *a priori* assumptions, and synthesis as an empirical process working from *a posteriori* evidences. Perhaps Ranganathan saw his ideas as bridging these two. The following quotations are from the American Library Association *Encyclopaedia of Library and Information Services* and Wikipedia description of colon classification, to give the flavour of his approach–its detail is best followed up in the references given.

> Colon Classification lists relatively few and simple objects and ideas as main classes. These are combined, at will, with what were called facets as opposed to precoordinated subdivisions–facets are class markers or tags. These are combined, using a formal punctuation notation, to express exact document subject, from a formal vocabulary and notation capable of unlimited variety of expression and extension. His other main innovation was an alphabetical chain indexing mechanism.
>
> The classification enables 'all possible subjects' to be constructed from a set of standard units covering what Ranganathan calls main classes (mathematics, physics, [...] zoology, [...] medicine, [...] arts, law, economics), common isolate, time isolate, space isolate, language isolate, phase and intra-facet relationship. The colon syntax acts to join these parts together and build a so called analytico-synthetic classification.

This avoided the rigidity of previous systems and gave new flexibility to incorporate new subjects and their various relations over time. Describing his ideas, Ranganathan likened the Colon Classification to Meccano–a favourite construction toy of my childhood and a predecessor of Lego. In this analogy, the classification of a particular book or document is depicted as a Meccano model, nut and bolted together by the colon syntax, from Meccano plates, girders, axles, cogs, wheels and so on, representing the component facets that are the building blocks of the system.

I discovered this Meccano model analogy quite recently, when researching his work, and was immediately struck by its parallel with the Lego model analogy we have used in describing the openEHR methodology for constructing clinical data models. openEHR (discussed in Chapter Eight and a Half) used the analogy to illustrate its compositional method of construction of these models, which provide generic patterns of data entered into electronic care records. The building blocks (Lego blocks) are selected from a set of predefined types and groupings of data, comprising the openEHR Reference Model. These are joined together to create clinical data models (Lego models) known as openEHR archetypes. The record itself is thus akin to a Lego village, such as the Bekonscot Model Village not far from my home, in Beaconsfield, and the one at the home of Lego–Billund, in Denmark.

The UNESCO Broad System of Ordering (BSO) for Documents and Books

In 1951, in pursuit of wider dissemination and adoption of his theory, Ranganathan took on the role of coordinator of the International Federation for Information and Documentation. From this base, his work had a strong influence on future developments in the field. However, his colon classification, and derivatives from it, as adopted in the BSO, ultimately lost momentum some fifty years later, notwithstanding the efforts to support them by UNESCO.[54] I doubt that many now know of its existence, and so I have taken space here to record some of its history and introduce and honour one of its principal adherents and advocates, Eric Coates (1916–2017).[55]

There were important general lessons from the decades-long efforts and final burned-out failure of the BSO. They have parallels in the difficulties and impediments encountered in bringing innovation in information technology to fruition at a comprehensive scale, in support of health care services today. It is a common and costly characteristic of the Information Age that rather than recognizing and learning from past failure, failure is often rationalized, swept under the carpet and forgotten about, as credulous attention switches to new predictions and aspirations.

I knew nothing of Ranganathan until his name came up in a chance discussion with Eric Coates in 1990, where I learned that he had been a key figure in establishing British technical standards in the post-war years and the Director of the British Technology Index publication. I had known Eric as an attender at Quaker meetings for many years before that chance event–a quiet, wiry, civil and slightly austere man, who came each week, never spoke in the meetings and left rather quickly at the end. Eric knew my father well and we talked about him and another young person who my dad knew from a Quaker family of the wartime era, Fred Sanger (1918–2013), who went on to become a double Nobel Laureate, a founding father of molecular biology.[56]

54 For detail of the BSO, see the archive maintained at University College London, *BSO* (2000), https://www.ucl.ac.uk/fatks/bso/

55 For an appreciation of Eric Coates's contributions, see K. Kawamura, 'In Memoriam: Eric Coates, 1916–2017', *Knowledge Organization*, 45.2 (2018), 97–102, https://doi.org/10.5771/0943-7444-2018-2-97

56 Fred Sanger and my father were members of a Quaker community at Spiceland, on the Blackdown Hills in Devon, which sought to establish effective community action in wartime. I remember my dad pointing to its towering chimneys on a distant hilltop, when, years later in my childhood, we drove to Devon from the

At our meetings, I told Eric about a major switch in my career, around 1990, when I set off in a new direction, having just been appointed to the first UK Chair in Medical Informatics, at London University. Among my new directions was the leadership of the European Union AIM research and development initiative to formalize the architecture of electronic health records. This was the GEHR (Good European Health Record) Project, which evolved over the coming decade, through a succession of further projects and collaborations, into openEHR, as described in Chapter Eight and a Half.

Connecting with this topic, Eric told me about Ranganathan and provided me with some early documents–sadly lost in a flood at our home in later years. Eric himself featured prominently in the story of the BSO right through until his retirement, when its copyright was assigned to UCL. The story is recorded by Keiichi Kawamura,[57] who covers two hundred and sixty-two summaries of reports, articles and meetings between 1973 and 2011.[58] I have selected excerpts that illustrate the human dimension of the struggle that Eric led, which I connect with struggles for methods of standardization in the medicine and health care domain.

In Reference 261, the following appears:

> Eric Coates was working as a cataloguer and classifier at the then recently established British National Bibliography. Earnest, sometimes a little severe, transparently sincere, and humane, Eric later became the first editor of the British Technology Index and wrote a book, *Subject Catalogues: Headings and Structure*, much influenced by facet ideas. He has also played a major part in constructing and testing the Broad System of Ordering, a high-level classification system.

---

children's home he and my mother ran in Hampshire, and we stopped for a break and to admire the rhododendrons on the hills. The history of this self-reliant community—men and women, some four hundred strong—is recorded by Stanley Smith, with the amusingly telling subtitle *Cups without Saucers*, nodding towards a rather spartan life (*Spiceland Quaker Training Centre, 1940–46: Cups without Saucers* (York: W. Sessions, 1990)). I have it on the shelf above my desk and see, on p. 35, that Sanger worked at the time as a ward orderly in the local Winford Hospital—quite a modest setting and role, and perhaps a signal of why and how he achieved so much thereafter.

57 K. Kawamura, *BSO−Broad System of Ordering: An International Bibliography* (Koshigaya: K. Kawamura, 2011), https://repository.arizona.edu/handle/10150/129413

58 I was interested to see that a key meeting, among many where BSO struggled to be heard, had been held at Helsingør in Denmark, in 1964. This location (in dramatic full view of Hamlet's castle!) hosted a similar meeting many years later, bringing together leaders of national health IT programmes and the HL7, IHTSDO and openEHR organizations, at which I represented the openEHR Foundation.

In Reference 262, his contribution and staying power, pitted against institutional inertia, obstruction and vested interest, is recognized:

> [...] the works of Eric Coates who put into practice and advanced Ranganathan's thought mainly through the British National Bibliography (BNB), the British Technology Index (BTI) and the Broad System of Ordering (BSO) [...] these three systems demonstrated: (1) how his works are connected with each other, (2) why his achievements should be estimated by a global standard, and (3) which of his contributions will throw light on unsolved problems in knowledge organization. The conclusion is that the underlying conceptual coherence in the work of Coates should be highly regarded as the persistent survival of interest and concern about classification, despite its marginalization.

Eric described himself as 'in favour of a revolution, not of classifications but of the management of classifications' (Reference 1). He was, in essence, seeking a grounded and sustainable method. In his work for the BNB and BTI, he perceived (Reference 62) 'key issues for furthering their work were in distinguishing methods and techniques from sciences and products of human activities, including technologies and religions'. Promoted through the UNESCO UNISIST Programme (United Nations International Scientific Information System), the BSO was conceived as a 'switching system' to enable interconnection and cooperation between information systems, standardizing communication of content among the key classifications of books and documents. The underpinning theory, providing coherence and a *lingua franca* for this communication, was the conceptual framework of Ranganathan's circle of knowledge and the colon classification. Initial efforts were geared towards defining and grouping thousands of subject fields. By 1984, four editions had been published (Reference 78). By 1990, the BSO had expanded to three times the size of the first published form in 1978 (Reference 86).

A review in 1980 (Reference 133) highlighted the problems it faced:

> The progress in library classification has been slow in its long history. But there was a drastic change in the 1960s due to the rapid development of science and technology, an increase in number of publications, and the advanced information processing technology. Looking back on the theoretical studies and practical activities in library classification for the last 10 or 20 years, the following are recognized: (1) the trend towards faceted classification, (2) unified view of classification and indexing, (3) mechanization and automation, and (4) standardization.

This review illustrated these trends with examples of national and international perspectives on the situation. The BSO's aim to provide a generalized 'switching language' for the UNISIST standardization programme implied generalization and was recognized to be at a price in terms of its lack of particular specificity. This was subsequently regarded as a serious defect and the suggestion was made to combine use of BSO with UDC, to attempt to overcome it. The trade-off between general and particular considerations and means to broker coherently between the two has become a widely enduring feature of design and standardization endeavours for health care information systems.[59]

Reference 134 provides a historic flavour of the issues and underlying tensions in play, because of the diversity of disciplines:

> The paper traces the main lines of development of scientific and technical terminology (STT) and the sub-languages of individual scientific and technical subjects. It is emphasized that for a long time virtually every branch of science and technology developed in isolation and has evolved its own closed terms system. In the 20th century, when sciences are interpenetrating on a wider scale and new and promising research trends emerge at discipline interfaces, the interactions of isolated terms systems have wrought havoc in the STT sphere. This tells on the evolution of indexing languages, both classificatory and descriptor ones. The situation is aggravated by the fact that most indexing languages have been developed to serve the needs of one organization or a group of organizations and seldom crossed national boundaries.
>
> Main merits and demerits of the UDC are discussed and an assertion is made that the underlying principles and main scheme of this classification are not consonant with the present-day condition of scientific knowledge. It is pointed out that the Broad System of Ordering (BSO) has a big role to play in perfecting STI exchange processes and in information organization in major information centers, above all international ones. However, BSO is not without fault either, its chief drawbacks being a strong influence of Anglo-American STT, an overly pragmatic nature, the difficulty of classifying multisubject documents, and a potentially strong dependence on political and ideological factors. A crisis of traditional hierarchical classification is postulated. The rapid proliferation of thesauri adversely affects STI exchange and the use of large information networks. It also happens that practically all thesauri,

59 There are interesting parallels in this story with the issues faced and highlighted by Alan Rector and colleagues ('On Beyond Gruber'), seeking viable combinations of 'open world' methods from description logic, 'closed world' frame-based methods, and *ad hoc* annotations, to achieve alignment and harmonization of the important SNOMED and ICD medical terminologies and classifications. These issues of medical language and terminology are discussed in further sections below.

> including international ones, are semantically and logically incompatible even where they refer to one knowledge area. It is suggested that the basic concepts of hierarchical classifications and upper-level descriptors be integrated, and thesaurus systems be established on a common conceptual and logical basis. Latin American nations provide a good testing ground for such a global-scale experiment.

There is a lot to digest and reflect on in that long summary. With Eric already retired, and the failure to replace his breadth of knowledge and engagement as the driving and anchoring force at the heart of its leadership, the BSO lost steam in the standards arena. And FID bowed to the inevitable, passing its copyright to the BSO panel members in 1990 and backing UDC. In 1993, the panel members established a not-for-profit company for distribution of the BSO. In 2000 the copyright was vested in the School of Library and Information Sciences (SLAIS) at UCL.

It is of note that by 2004, the year after the openEHR Foundation was launched, Reference 169 in Kawamura's bibliography records:

> The need for structured machine-readable data and not just 'simple text' and thereby to have a common standardized data model was highlighted. This would enable automatic classification and management and control of concept hierarchies and vocabulary facets and sub-facets. Proprietary data formats in applications are seen as costly and limited. The requirement for independence and integrity of data elements is emphasized.

This motivation has much in common with that addressed in the mission of openEHR, as covered in Chapter Eight and a Half.

The language of these extensive quotations reflects that they derive from an era where manual, facet-based method, which influenced medical terminologies and classifications of the times, was state of the art. Computer science pioneers of the description logics of more recent decades developed theory and method for tackling these goals on a higher level of abstraction. But in seeking to raise endeavours to this level, historic investment in what had become intractable legacy inhibited and contended progress. In turn, the semantic complexities of medical knowledge uncovered in these endeavours challenged the available and tractable methods of description logic. Change of approach required radical new thinking and reform, which further challenged the communities involved with each standard, internationally. This set of related problems was the focus of pioneering research in small, not widely visible, or recognized projects, such as Alan Rector's GALEN project, that grew alongside the GEHR project and its successors in the EU AIM Programme, as discussed later in this chapter, in the section on medical language and computation.

Confusions of the times and struggles for greater clarity are evident throughout the long history I have traced here. There is a sense of an elusive perfection of theory that bedevils practical efforts. Although always imperfect, theoretical models can nonetheless prove valuable, and we seek always to improve those we have and make them more useful–this is a theme developed in Chapter Four. We need a clear sense of why we are building new systems, what our goals are in this, when, where and with whom we are going to tackle them, and, most importantly and most overlooked, the method that embodies how we will achieve these goals. Fulfilling a useful purpose must reign over achievement of perfect execution.

With the benefit of hindsight, one can see in histories like that of the BSO, all these dimensions of challenge playing out in a theatre of life. As William Shakespeare (1564–1616) wrote in *Hamlet*, the play holds a mirror up to nature, and we must learn from the images it provides for us.[60]

As highlighted at the beginning of this chapter, we express, communicate and reason with knowledge through language and logic, and now, in the Information Age, machines do, too, through machine languages and logics. As with many other generic themes encountered along the storyline of the book, language and logic are, in themselves, extremely deep and wide-ranging subjects. My purpose, here, is to emphasize how they have connected with where we have reached in the evolution of computer-based health care records and related knowledge-based information systems.

Languages and their Expression and Communication of Knowledge

For last year's words belong to last year's language
And next year's words await another voice.[61]

In the evolving quest to make and express their sense of the reality of the world, the Greeks became absorbed with metaphysics–*τὰ μετὰ τὰ φυσικά* [*ta meta ta physika*]. *The Metaphysics of Aristotle* (384 BCE–322 BCE)

60 'Suit the action to the word, the word to the action, with this special observance, that you o'erstep not the modesty of nature: for anything so o'erdone is from the purpose of playing, whose end, both at the first and now, was and is, to hold as 'twere the mirror up to nature: to show virtue her feature, scorn her own image, and the very age and body of the time his form and pressure'. *Hamlet*, Act 3, Scene 2.

61 T. S. Eliot, 'Little Gidding', *Four Quartets*, ll. 118–19.

concerned matters 'after the things of nature'. This was a domain envisaged to lie beyond objective study of material reality–idea, doctrine, the nature of reality. Not a domain that the embryonic biological mind had much time or use for. But one of enduring interest and perplexity to the enquiring, embryonic, civilizing mind, concerned with making sense of and interacting with the material and human worlds.

More Greek words came into play–όντως, δόξα, δοκεῖν, ήθος, λόγος, πάθος, ἐπιστήμη, ἐπῐ́στᾰμαι, τέχνη [ontos, doxa, dokein, ethos, logos, pathos, episteme, epístamai, techne]. Here is a rough run down, based on my ten-kilogram Compact Edition of the Oxford English Dictionary (published complete with essential magnifying glass!)–better-informed readers than me may justly roll or avert their eyes.

ontos–being, nature of reality

doxa–opinion or glory

dokein–to seem or to seem good–led to **dogma** and **paradox.** More distantly to **decent** (in connection with seeming 'good')

ethos–custom or habit–connecting with ἠθικός [ethikos]–showing moral character–more generally, characterizing the spirit of a culture, era or community as expressed through its beliefs and aspirations

logos–word, reason, discourse, study–from λέγειν [legein]–to speak

pathos–suffering, experience

episteme–science or knowledge

epístamai–to know, to understand, to be acquainted with–about knowledge of principles

techne–craft, art–making or doing–concrete, variable and context-dependent–also a kind of knowledge

Adding the suffix **λόγια** [***logia***] (plural of logos), describes an associated oral or written expression. Combining with stems leads to 'ologies'.

Οντολογία [***ontologia***]–ontology is about how we answer the question, What is reality? It affects how we approach our subject. It quickly gets complicated and convoluted when we dig deeper into detail. But a key feature of an approach is that it be accepted, appear to be a good one, and be right. In matters beyond the senses this is argued and judged in the realm of belief. It becomes a matter of ownership, assertion, power and persuasion–of rhetoric, logic and reason. **Ontos**, **ethos** and **logos** became central to argument. Knowledge became prescribed and proscribed.

> **Doxology** became connected with praise and glorification, as expressed orally, first appearing in English around 1645. It became liturgy.
>
> **Λειτουργία**, **liturgy**–divine or public service; function, operation, service, working.
>
> **Ορθοδοξία**, **orthodoxy** (**ορθό** [*ortho*], 'right') became authorized or generally accepted theory, doctrine or practice. It extended to 'right' thinking elsewhere–e.g., monetarist orthodoxy.
>
> **Ετεροδοξία**, **heterodoxy** is another doxology (**ηετερο** [***ietero***], 'another'). Someone quipped that orthodoxy is my doxy–yours is heterodoxy (or even heresy)!
>
> ***Επιστημολογία***, **epistemology**–The theory of knowledge; especially pertains to its methods, validity and scope, and the distinction between justified belief and opinion.

Plato (*c.* 428 BCE–348 BCE) contrasted **episteme** with **doxa**–common belief or opinion. The term **episteme** was also distinguished from **techne**–a craft or applied practice.

For Aristotle, **pathos** was a means of awakening people's emotions in order to sway their opinion towards that of the speaker. **Rhetoric** embodied **pathos**, **logos** and **ethos**. Disease was suffering and suffering became disease. *Pathology* became the study of disease.

The ether (or aether) of classical times was the fifth element (quintessence), after earth, air, fire and water. It was a medium filling the universe above the terrestrial sphere. Science later conceived of light propagating in the universe through an ether. We now talk of other universal media–communication systems and the information they transmit. These bear some (pharmacological) resemblance to chemical ether! As Marshall McLuhan (1911–80) wrote in a different context, the medium is the message.

When we reason, express and communicate, we do so through language, of which there are many kinds. Good use of spoken and written language enables and stabilizes communication, as a medium for the expression of thought. It evolves. It reflects and describes domains of knowledge and understanding, and their different jargons. It is conditioned by purpose and context–of culture and practice, geography and time. Interestingly, I've read that some believe that language evolved first as a means of misleading rather than informing, to confuse and discourage potential jungle predators in threatening situations. Information pandemic, a term used to describe confirmatory bias in communications about the Covid pandemic, has born some resemblance to cacophony in a jungle!

Poor use of human language harms communication and risks misunderstanding, distortion and confusion of meaning. Overly elaborate language becomes a linguistic Tower of Babel, overly simplistic expression a Dalek drone, to the human ear. Finding the right words is a struggle for personal understanding as well as for its expression–words may fail us due to overthinking them, as much as not thinking about them enough. It is easy to overthink or be careless with words. Words connect on different levels of meaning and intention: simple and complex, vague and precise, gentle and harsh. Ideas thought through and framed in language, no matter their significance, connect on different levels of expression, too.

Psychologists tell us that the greater part of human communication is non-verbal. Human language has infinitely variable and subjective contexts that impacts its meaning and its integration with wider non-verbal communication. Words toe a line between a defining framework of the language in which they are expressed–we speak of syntax–and communication of their meaning–we speak of semantics. Poetry taps into meanings, feelings and emotions beyond the words and forms expressed. We say that music and art speak to us, and silence speaks volumes.

Human language was born into the cradle of civilization. Philosophy, logic and mathematics evolved from and around natural language. The philosophy and 'term logic' of Aristotle and the philosophy and 'Stoic logic' of Chrysippus (*c.* 279 BCE–206 BCE) were launched, clothed in natural language. They disappeared and resurfaced over the centuries, formalized in new languages of mathematics, science and computation. Now, language of logic permeates and underpins foundations of mathematics and language of mathematics permeates foundations of logic. The *linguae francae* of the world now extend through the languages of mathematics, logic and computation.

The Language of Mathematics, by Frank Land (1911–1990), is the title of my first book prize at school, sitting above me at the far end of the shelf. Over recent centuries, the languages of mathematics, and then computation, have become intertwined with scientific methods. They are now intrinsic to the modelling and analysis of complex systems, and no systems we work with in this way are more complex than living systems. In the Information Age, theory of computation progressed alongside experimentation with novel forms of computer language whereby computing machines could be made to enact human instructions, manipulating first numbers and then data, symbols and reasoning, more widely. Experimental programming languages were in their infancy at the start of my songline. Languages of computation now underpin the rigour, expressiveness, reliability and trust in computer systems and communication.

Fluency in many and diverse languages of communication is increasingly significant for understanding what medicine is and what it does–the science on which it is based and how health care services acquire and use knowledge to investigate, reason, act and communicate. This spectrum of languages is akin to a spectrum of electromagnetic radiation–from the long, medium and short waves of radio transmissions in my childhood, now transmitted and detected at ultrashort wavelength and measured, manipulated and transmitted in digital form, with tools and methods that vary across the spectrum.

When we seek to enhance our lives by intermixing language and communication of machines with human language and communication, we must be careful as this may risk impeding and harming both. As the use of language fails to satisfy and cohere, it is patched by narrow and diverse approximations and appropriations of words and meanings. As Russell wrote, logical formalization of reasoning must be carefully pinned down within its applicable context, as logic grapples with the nature of appearance, reality and truth.

In health care, that is a very hard bar to rise above, and not always or necessarily a useful one. Simpler approximations may suffice and be more effectively enacted and communicated than more complicated and precise ones. As technology advances towards artificial intelligence and the *Novacene* era that James Lovelock describes in his 2019 book, the language of the machine encroaches, superimposes on, and supplants the predominant languages of former eras.[62] Some of that potentially for the good, and some for the not so good. Zobaczymy [we will see]![63]

Facing these numbingly wide-ranging issues, I was undecided about where and how to write about language in the context of health care information systems. Specialists at all points of the circle of knowledge have something to say, splitting the spectrum of languages into all the colours of the rainbow. Spoken, written and machine languages differ and differ differently, in different languages. Language integrates and language differentiates. Language unifies and language divides.

So where to begin? I will cop out and start somewhere else, with some personal, linguistically untutored reflections about my experience of learning two new languages, in later years–the very tricky Polish language, having married into that wonderful country and culture, and the language

62 J. Lovelock, *Novacene: The Coming Age of Hyperintelligence* (Cambridge, MA: MIT Press, 2019).

63 On this Polish expression, see Preface.

of dance, which has preoccupied our home life for two decades, now. Encyclopaedia, in its classical context, was a circle of learning before it became a circle of knowledge. Language, too, starts with learning.

Learning a Language

I learned Polish when marrying again, and the motivation was huge. My newly extended family had few English speakers. Hugs and kisses and the punctuating *dzień dobry*, *dziękuję* and *dowidzenia* pleasantries can only get one so far, as with the broken German on both sides, which is where we started! A brilliant young teacher, Ela Wolk, at the London School of Slavonic and Eastern European Studies, now part of UCL, taught a small group of dedicated learners for three years. We were all seeking out Polish connections and/or heritage in our lives.

On opening one's eyes and ears in a new country, there is a richness of culture to be shared, that would be substantially inaccessible without the language of that country. Ela drilled into us grammar and marked our homework; she took us with our partners to experience Polish theatre and music. One thing we learned quickly about Polish is that you must listen hard to catch the words. Ela used to tell us to pause from analyzing the puzzling, consonant-riddled written forms and consider instead, how does it sound? She would ask 'what do you think it means?', encouraging us to identify what similarities it had with other words we already knew. But, hm…! Here, on the one hand, we were facing the extraordinary complexity of written Polish language, with its somewhat pedantically Latinate formal grammar, while, on the other hand, everywhere in sight the rules were being broken to make pronunciation easier. In principle, the phonetic structure and the spelling go hand in glove in Polish, more than any language I have studied. It requires a special configuration of jaw and tongue to get it right. English does, however, get its revenge–Polish people have difficulty with 'th', just as I do with ę and ą–try them with fifth or thistle!

Languages have spelling and grammar–the structure of sentences parsed into subject and predicate, noun phrase and verb phrase, main and subsidiary clause, noun and adjective, verb and adverb, associated inflexion, mood, gerund and gerundive, pronoun, preposition and the rest–or maybe all new descriptors, now. I had learned Latin and Greek at school. Latin was quite pleasant to my mathematically inclined mind as it kept pretty much to the rules. Greek was a blur as I had a year to study the language and did not quite get there. My scholarly, eccentric head teacher–who persuaded my parents I should study both Latin and Greek, when I would rather have skipped Greek–had lifted me from my surreptitiously-preferred geography class to study Greek with him in a tiny group of four. I am now very glad that he did as it gave me a rusty key to many doors, discovered and walked through years later.

Expression in spoken language is a mix of formal, blurred and broken rules of grammar, and an enveloping contextual and cultural mishmash of associations and meanings. My doctor wife trained as a medical linguist in the UK. She acts as interpreter across health, education, social care and law. The communication she brokers is from different languages, cultures and experiences, of both client and service provider. Sometimes in English, sometimes in Polish, sometimes with native English-speaking professionals, sometimes with professionals from other cultures and tongues, speaking English as a second language. This is the world; and this is health care language and communication.

Esperanto is one approach to overcoming the dividing lines of language–everyone sharing a common auxiliary language–might that be a solution? Incidentally, Esperanto is a nice vignette of Poland and medicine–it was created by the Polish ophthalmologist Ludwik Zamenhof (1859–1917) in 1887. The word translates into English as 'one who hopes'. Clearly its time has not yet come, though hope springs eternal and it is a living movement, still pleasing and enriching to those who keep it alive. Context of language always matters. Languages merge and standardize and then pidgin languages take flight. Now the dangled offering is Google Translate. What cultural implications and perturbations may lie lurking when leaving to the machine, the brokering of human communication?

It is important to remember that human communication is substantially non-verbal. I think of that in the context of dance. We talk about dancing around the point when not quite communicating with one another and being (badly) led a dance. One of the things my wife and I discovered was that we both loved dancing. She expressed it in her nature and drew it out of mine. We did not share a mother tongue and communication through dance proved foundationally important for us in lots of ways and has shaped our lives together amazingly.

One learns that communication on a dance floor is both subtle and dramatic, rife with potential for miscommunication and mistake. Dance is a language of connection and flow. It comprises human form, emotion, fitness and balance, all of which must be nurtured and practised. Dance needs good teachers who love dancing and know how to dance. We have Tom and Ali, who teach teachers of Tango, and Sarah, who performed in the Royal Ballet and toured with the prima ballerina, Darcy Bussell. Dance has a musical context of melody, structure, rhythm and interpretation, and that must be listened to and communicated, sometimes best with eyes shut–more so for the led partner, of course! My wife will often say 'I do not feel that in the music', as I chide her that she is taking the lead when I feel she should not be! Dance may appear unbalanced along gender lines in leading and following, but that is not true for the dancers themselves: gender no longer determines dance partnership, and dance is better for it. Some of the best and most artistically-led partners are the emotionally and physically strongest. And the best leaders know the best dance can only be led, in all moods and stages of the music, when the partner is listened to, given space and with leadership negotiated and flowing, to and fro. If you like jazz, you will probably resonate with tango.

There are several lessons about language that I draw from these two very different learning experiences. Learning language is about listening, experiencing, practising, performing and enjoying. Different languages are not isolated domains, they cross-fertilize in both method and context. The bedrock of fluency in language lies in *how* we learn it, and, specifically, in whether we learn it within its rich cultural contexts. Learning and knowledge go hand in hand. Literacy, in the sense of the effective use of language, is the foundation of knowledge and wisdom. Literature is an account and record of knowledge, the organized expression of thoughts, feelings and ideas–a medium that further connects and flows. In the era of ubiquitous information, we, too easily, talk and write as if we know before we have experienced and learned. We learn from experience and that requires expression and testing of yet unformed ideas.

Some questions then arise. We use words (literately, literacy) and information (informedly, 'informacy'?) well and we use them carelessly and blindly (ignorantly, ignorance). In moving beyond words to information more generally, what are fluency and literacy in the combined use of the many different kinds of language that underpin the Information Age? This ability was once the preserve of the polymath, an expert in many disciplines. But those days have receded beyond the horizons of human capability and capacity. What characterizes a polymath of the Information Age and what distinguishes them from a Jack of all trades? Is the computer fluent and literate? If so, how is it learning languages and what does it know? Will we come to think of it as wise? And how does the information that it communicates connect and flow, in human terms, between machine and human worlds? In 2023, the likes of ChatGPT are dramatically raising the stakes in relation to how we approach these matters.

Language and Machine

In the Information Age, we speak of language that specifies instructions to a computer to execute programs. When the computer reads program instructions from its memory store and follows them to enact the computational steps they specify, there are three different but closely interconnected considerations in play. Taken together, these determine how the enactment plays out, each reflecting and depending on a different language. The first concerns the machine itself and the kinds of operations that it is capable of enacting, commonly referred to as its machine code or machine language–precise and readable within the computer CPU, although opaque and unreadable binary code to most human eyes. The second concerns the programming language chosen to specify the computational process and task to be performed, in a manner intelligible to and reflecting the purposes of the programmer. Precise and readable, here, by the compiler or interpreter program that runs on the machine, first to translate from the programming language into the machine's language, enabling the machine to perform the instructions generated there. The third concerns language descriptive of the data that the program enactment causes to be captured or generated, processed, stored and communicated.

Precision of language matters in all these contexts but is not in itself a guarantee that the machine can or will function as the programmer intends. Machine, program and data descriptive languages each embody and exhibit precise expressive capabilities and rules for how they are used. They also embody precise limitations. The enacted program utilizes these capabilities and must combine them correctly and circumvent their limitations. It needs to be done efficiently. The data provided to the program must exist in a form consistent with the requirements and capabilities of both program and machine.

And for all this to be meaningful, program and data must be jointly expressive and representative of the task the program addresses. This involves considering the relevant context in which the data is generated and processed, as well as ensuring that the results computed are to be properly understood and interpreted. The machine and program must likewise perform together efficiently and acceptably, in terms of the time and resource they require.

Where the program goal served is purely concerned with manipulating data within a particular machine environment, as a closed system (reading a block of data from a disc store, adding up a column of numbers) no wider practical issues of meaning arise. Where the program goal is integral with concerns outside that environment–making a weather forecast, predicting numbers of cases in a global pandemic or results in an election (like today, as I first write this section, with election and viral pandemic both raging on 6 November 2020 in the United States of America)–the question 'What does it mean?' is embedded within a wider context and the answer potentially a matter of human controversy!

The programming of the machine is, in a sense, an art of the possible, seeking to pitch at a sweet spot between the capabilities of the machine and the framing of the nature and requirements of the task being addressed. Sweetness consists in operating within those capabilities and limitations, and with outcomes perceived as useful and meaningful for the task at hand. It can, though, become a sour spot, where delegation to the needs of the machine is a derogation of the human needs that the task addresses. As new boundaries of the possible are approached and explored, the potential for doing harm needs to be understood. This can then be better balanced against the curiosity and excitement about the potential of the new, and the value anticipated from its realization. This is the nature of innovation. The bridge between science and society created by the innovation of information systems and technology is a focus of Chapter Five on information engineering.

What can different computers, or any computer, compute? What can systems of logic or any logic express, prove and decide? What can different languages of computation express and represent? What can be measured and described? These kinds of questions, and the limits they probe, arise within the languages of mathematics, logic, computer science, natural science and engineering. They ramify all around the circle of knowledge, as the Information Age spreads more pervasively into and across human affairs.

Precision of Language

Writing can either be readable or precise, but not at the same time.[64]

There speaks the logician, in his wide-ranging work spanning from precise language of mathematics to logical precision in communication of thought and meaning, connecting with discussion of the nature of truth. On my bookshelf is Russell's, *An Inquiry into Meaning and Truth*.[65] We talk of a computer reading and writing when it transfers information to and from data storage devices. We say that programmers write programs and computers read and execute them. In this transfer of information, precision and readability go hand in hand.

In Russell's world of philosophical logic, consistent, clearly-framed and articulated use of language are the bedrock of thought and reasoning. Expressive range and use of language are also the bedrock of human culture, arts and communication. And in the Information Age, the language

64 Quote attributed to Bertrand Russell, unknown source.

65 B. Russell, *An Inquiry into Meaning and Truth* (New York: Routledge, 2013).

of computation is fundamental to the integrity of data and algorithm, and to the efficiency, effectiveness and robustness of computer systems. All these dimensions matter in the context of health care information systems. How they fit and work together is complicated!

Russell was very much aware of and actively engaged with the social context of the times in which he lived. In his philosophical works, he maintained that logical abstraction and definition must be precisely stated within relevant context. Herein lies a considerable challenge for the domain of health care in the Information Age. How, where and to what extent is logical precision a valid and achievable goal in reasoning with knowledge in this domain? How, where and to what extent does the limitation of logical precision of language matter? Every model of appearances, as with language and logic, is a simplification of the reality it purports to represent. The world of model-based representation is the subject of Chapter Four.

Physics is thinking and discovering its way through a multidimensional maze of experiment, theory and mathematical language, in its quest to describe and understand physical reality in greater detail. Scientifically, this is an exciting pursuit; it focuses research on unsolved but potentially experimentally tractable areas of current unknowing. This is not solely about instruments of ever more precise and specific measurement. Whitehead quotes Jules Henri Poincaré (1854–1912) in pointing out that instruments of precision, used unseasonably, may hinder the advance of science–giving the example that knowledge of the tiny relativistic imperfections of Johannes Kepler's (1571–1630) law of planetary motion might have delayed the imagination by Isaac Newton (1643–1727) of the law of gravitation.[66] Truth must, he says, be seasonable (Microsoft Word does not recognize 'seasonable' and suggests I probably mean 'reasonable'!). That is perhaps another way of saying that meaning and truth must be considered in relevant context.

Formalism of Language

My maternal grandfather, who I never met but whose picture is on the wall to the left of my desk as I write, was an English teacher at Westminster School in London. He was a stickler for grammar and wrote numerous textbooks used for teaching the 'rules'. His writing style in these books and in his articles for the school magazine that he edited reads now as primly pedantic flourish–not very readable although immaculately well-formed. And he, as a teacher then, was always formally dressed, of course! In my student days, too, there was a lot more formal stuff–the balance is better, now. The

66 Whitehead, *Adventures of Ideas*, p. 232.

formalism of a language is meant, here, to cover its definitions and rules of use. Rules of grammar are a formal model of the structure of language. Communication is improved to a point by attention to formal grammar, but when pushed to extremes, this becomes restrictive of expression–words and their usage constantly evolve over time.

English lessons devoted to grammar were an oft-practised routine in my secondary school days–it was a Grammar School! We were given complicated sentences and set to compete, to see who could succeed first in breaking the structure into its different parts of speech and their groupings, linking and labelling them, and arriving at a final standard, hierarchical diagrammatic form. The method for constructing and deconstructing (parsing) natural language sentences might now, I understand, be termed a 'phrase structured grammar following the constituency relation'. It is quite complicated to remember and apply such formalism unless one deals with it on a regular basis! Here, a constituent is 'a word or a group of words that function as a single unit within a hierarchical structure'. The constituency relation comes from the subject-predicate division of sentences, with their clause structure understood in terms of binary division between subject (noun phrase) and predicate (verb phrase). The parts of speech (terms) constituting a complete sentence, and the stepwise reduction of the sentence structure, through binary division of its different kinds of clauses and phrases, maps to a tree structure. The sentence terms and their groupings at successive stages of this reduction appear as nodes in the tree. Such formalism connects with the term logic of Aristotle, as discussed further in the sections below.

Natural language, logic, mathematics and computer science evolved over many centuries, along connected pathways of formal method–mysticism, religion and philosophy initially intertwined. Logic from the time of Aristotle and Chrysippus was expressed in sentences of natural language, and more abstractly characterized and expressed as premise, proposition, predicate and syllogism–a 'sentential logic'. Concept of number, and calculation with numbers, intertwined with philosophy and logic of reasoning and argument, and the glimmerings of science. Mathematics and logic of inference intertwined with philosophy and method of science. Mathematics of infinitesimals and infinities gave birth to calculus. Logic and mathematics intertwined in 'logical calculus', reimagining mathematics in the language of logic.

By the time of Gottlob Frege (1848–1925), mathematics and logic had reached a competing understanding of the logic of sentences. He broke away from Aristotelian logic, built around binary divisions of the structure of the sentence, aiming to replace it with a mathematically rigorous formalism. And two runways of lift-off ensued. The first was directed towards the reinvention of the foundations of logical inference, providing new mathematical reasoning about propositions and predicates. This was no longer formalized in the grammar of sentences comprising words as atomic elements, but in manipulation of mathematically precise constructs that could be seen as either true or false, using logical operations applied to symbols and formulas as the new constituents of a mathematical language of logic. Logic itself became a calculus. The second and parallel runway inherited these ideas and was directed towards unifying mathematics within a framework of logical deduction from a small set of basic axioms.

The language of mathematics and logic of the nineteenth and early twentieth centuries led to and cross-fertilized with language of computation and algorithm, today–a revolution led by mathematicians. Computer technology was a physics and engineering revolution and the Information Age an Industrial Revolution. Medicine of the twentieth century was a scientific, professional and computational revolution, and health care, today, is embroiled in a twenty-first-century cultural and social revolution. I will pick up on the two runways and the post-Frege story of the past one hundred and fifty years in the next section on the language of mathematics and logic, and its embedding within the history of reasoning with knowledge, in the succeeding section. First, I will look briefly at where it led in the formal study of grammar.

The languages used to write computer programs require rules of grammar whereby program sentences (statements) can be rigorously and reliably generated and parsed. Noam Chomsky developed an overarching theoretical foundation for a hierarchy of grammars which underpinned this quest, from the 1950s.[67] In linguistics, Chomsky is known for his theory of universal grammar as an inherited 'hard-wired' human capacity to learn grammar. Wired or not, most people find natural language grammar gets harder as the complexity of meanings expressed increases. Unsurprisingly, the same holds true for the computer and its program languages!

Natural language and programming language share a common feature. Richness of expression correlates with complexity of its analysis. With increased richness and diversity of natural language, comes greater difficulty in mapping or parsing to reveal underlying grammar. And likewise with formal grammar of machine programming languages, grammar with greater power of expression brings harder problems in its parsing. In natural language, the boundaries of correct and incorrect use of grammar

67 Chomsky categorized four types of grammar, with each higher numbered category subsumed within the lower numbered category: regular (Type 3), context-free (Type 2), context-sensitive (Type 1), unrestricted (Type 0). Each had a set of associated rules for generating syntactically correct statements in the language.

are a fuzzy space, and usually not a huge problem, except to the more insistently rigorous of minds. This fuzziness is not possible in the realm of program language and computation, save within a framework that precisely defines what it is to be fuzzy! The programmer may specify valid program code that obeys the grammatical rules of the language used for creation of a program, but the machine used to implement it may, nevertheless, find the task too complex or impossible to execute. The problem of tractability of computation is studied in the machine language of the Turing machine, which features in Chapter Five on information engineering.

Two final questions arise about expressiveness of language in the Information Age. Are we, as some fear, on an *Académie Française-* or Esperanto-like path, risking loss of meaning by overly normalizing and constraining language to a common denominator that serves principally the purposes of machine communication? In refining information systems and infrastructure, accommodating the standardization they require, and adding artificial intelligence overlays, are we also implicitly deskilling and disqualifying humankind, such that it will no longer be able to understand and express itself, and thereby control its own fate? There are no persuasive answers to be had in response to such questions about how emanations from Pandora's box will play out. At best, as yet, both yes and no, potentially, and not necessarily! How we approach them will largely determine events.

We should not look back for answers. As Whitehead wrote a hundred years ago, and it remains true:

> Today the world is passing into a new stage of its existence. New knowledge, and new technologies have altered the proportions of things. The particular example of an ancient society sets too static an ideal and neglects the whole range of opportunity.[68]

We should rather look forward, with somewhere between Barack Obama's audacity of hope and Mervyn King's audacious pessimism. Zobaczymy! First, I will explore how the languages of mathematics, logic and computation came together to advance how we think about thinking and know about knowledge.

The Language of Mathematics and Logic

Leibniz argued that human thinking can be grounded in laws described in the language of mathematics. He was born at a time of turmoil and civil war in England, in the era of Oliver Cromwell (1599–1658), which overturned

68 Ibid., p. 261.

the 'Divine Right of Kings' to govern and established government by Parliament. Those times were described in the eponymous English ballad of that era as the 'world turned upside down'. The term was borrowed by the historian Christopher Hill as the title of his book describing the history of the era.[69] Leibniz was a father of mathematics who overturned the divine nature of inference, turning the world of logic upside down! This was the taxiing zone before the first runway.

Logicism asserts that mathematics is reducible to logic and that pure mathematics can be deduced from a few simple axioms (sometimes called primitive notions) through a process of formal logical argument. Russell wrote that it is the logicist's goal 'to show that all pure mathematics follows from purely logical premises and uses only concepts definable in logical terms'.[70] This was what I described as a second and parallel runway of advance, in the previous section.

There is as much learned dispute among logicians about theory of logic as there is among philosophers about ontology. The topic would quickly move to a level of detail that disrupts the flow of the book, by delving too deeply into what may be unnecessary, distracting or bemusing explanations and examples. It seems better to provide some brief outlines, footnotes and pointers to detail elsewhere. Wikipedia or a good logic primer are reasonable starting points for exploring further.

The domain of mathematical logic, which was also called formal logic, intersected early on with the mathematics of algebra and set theory.[71] In 1847, a self-taught, religiously devout English mathematician, George Boole (1815–64) published an essay entitled *Mathematical Analysis of Logic* that laid the foundations of what became known as Boolean algebra. In this algebra, statements in logic are expressed as algebraic equations. The symbols in the equations represent groups of objects (mathematical sets) and statements in logic. Their algebraic manipulation provides a rigorous method of logical deduction, thus representing logic as algebra. This algebra provides a

69 C. Hill, *The World Turned Upside Down: Radical Ideas During the English Revolution* (New York: Viking Press, 1972). The turmoil of seventeenth-century England has also been described as a reflection of a society challenged in its norms and beliefs as it came to terms with new ways of communicating the printed word and managing the explosion of information that this heralded. This polarized society into civil war between Cavalier and Roundhead armies, although historians differ as to how these loyalties segregated along political, economic, religious and demographic lines. The parallels with our Information Age turmoil, in this case coming to terms with the computer, are tempting to opine!

70 B. Russell, *My Philosophical Development* (London: George Allen and Unwin, 1959), p. 74.

71 Wikipedia contributors, 'Mathematical Logic', *Wikipedia, The Free Encyclopedia* (222 May 2023), https://en.wikipedia.org/wiki/Mathematical_logic

basis for analyzing the validity of logical statements, capturing the binary character of statements that may be either true or false. Boolean logic has been described as akin to a mathematics restricted to the two quantities, 0 and 1.

Boole worked at Cork, in Ireland, a few miles from the location of the Blarney stone. I was taken there once, when visiting to talk at his old University, but was not brave enough to stretch down over the cliff edge, to kiss it! In 1858, Augustus De Morgan (1806–71) was the first to propose the term mathematical logic. He was a close contemporary of Boole, based for many years at London University, the predecessor of my alma mater in London, UCL, which was established in those times as a non-conformist institution, countering the conformist religious regimes of the era, at Oxford and Cambridge.

De Morgan expressed logic in the language of set theory and logical propositions were cast into theorems of mathematical, logical inference. Rules used in translation and reduction of these logical expressions into a standard, not further reducible or simplifiable form, bear the name De Morgan's Laws. The precision this afforded enabled new insight and clarification of principles and methods of logical inference.

The wider application of the ideas of Boole and De Morgan into the realm of reasoning with knowledge, started with John Venn (1834–1923). He proposed the term symbolic logic and is remembered in the Venn diagram, used to represent the overlap of sets of objects arising in logical reasoning. Different kinds of logical formalism evolved over time, specialized for different requirements arising in the representation of verbal logical argument. It is a blurry panorama of separately identified and named branches of logic. I provide here just brief notes and introductory pointers to easily accessible further explanations.

The increasing semantic richness of formalized logic gave rise to different levels of what became known as logic calculus. Boole's work focused on the formal logic and algebraic manipulation of logical statements. There arose what was variously termed propositional logic, propositional calculus, statement logic, sentential calculus, sentential logic.[72] It deals with propositions and relationships between propositions, including the construction of arguments based on them. Frege, with Charles Peirce (1839–1914), made what is seen as the crucial break from the Aristotelian tradition

72 K. Klement, 'Propositional Logic', *Internet Encyclopedia of Philosophy*, https://iep.utm.edu/propositional-logic-sentential-logic/

of logical argument, replacing his term logic with what was called a first-order logic, which became known as the first order predicate calculus.[73]

Seeking greater order and precision of logic languages over subsequent decades, different levels of logical expressiveness became known as zeroth,[74] second-[75] and higher-order logic, such as multi-valued logic (a calculus of propositions permitting of more than two truth values, worked on by Jan Łukasiewicz (1878–1956) and Alfred Tarski (1901–83)).[76] Logical operators, or connectives, have evolved to encompass various sub-specializations, including modal logic (expressing statements about necessity or possibility), temporal logic (expressing quantification over time), deontic logic (expressing obligation and permission) and relevance logic (expressing relevant connection of antecedent with consequent of inference). Each of these languages of logic sought rigorous expressiveness of the subtlety of different strains of natural language and verbal reasoning, pricking up the ears of philosophers defending their own boundaries of discipline. Some eminent philosophers would have none of it–notably Willard Van Orman Quine (1908–2000).

As the foregoing brief account well exemplifies, the names and descriptions of the different kinds of language of logic calculus that have been explored are, in themselves, confusing. Proposition, predicate, statement and sentence mixed with propositional and sentential logic, and with propositional, sentential and predicate calculus. The terminology of this rather chaotic domain became a bit like that of medicine of recent decades! My chief aim here is to illustrate the complexity that mathematization (and then computerization) of commonly expressed human logical argument can foist upon us.[77]

The story continues in the next section on logic and reasoning with knowledge. This section is focused on the language of mathematics and

---

73 See Wikipedia contributors, 'Term Logic', *Wikipedia, The Free Encyclopedia* (19 June 2023), https://en.wikipedia.org/wiki/Term_logic; Wikipedia contributors, 'First-order Logic', *Wikipedia, The Free Encyclopedia* (21 June 2023), https://en.wikipedia.org/wiki/First-order_logic

74 Wikipedia contributors, 'Zeroth-order Logic', *Wikipedia, The Free Encyclopedia* (20 May 2023), https://en.wikipedia.org/wiki/Zeroth-order_logic

75 Wikipedia contributors, 'Second-order Logic', *Wikipedia, The Free Encyclopedia* (28 May 2023), https://en.wikipedia.org/wiki/Second-order_logic

76 Wikipedia contributors, 'Higher-order Logic', *Wikipedia, The Free Encyclopedia* (20 June 2023), https://en.wikipedia.org/wiki/Higher-order_logic

77 I am, at least, in good company in my perplexity: Encyclopedia Britannia likewise throws up its hands over the persistent disagreements among expert logicians over the connections of theory of logic with discourse on language (The Editors of Encyclopedia Britannica, 'Logical Relation', *Encyclopedia Britannica* (20 July 1998), https://www.britannica.com/topic/logical-relation)!

logic, the next on the problems it is used to address. Clearly, these two considerations intertwine, and things then get even more complicated!

I return, now, to the second runway, that I described as following from the seminal contributions of Frege. This is Russell's logicism–the quest for a reformulation of mathematics that proceeds from simple axioms using formal logic. Many well-remembered nineteenth-century mathematicians laid foundations to underpin the development and tools of formal logic of the coming century. Mathematical logic was consolidated by Giuseppe Peano (1858–1932) and later taken up by Russell and Whitehead, in their work towards establishing secure logical foundations for mathematics. By the late nineteenth century, there was broad consensus that a great deal of mathematics could be formally derived in logical progression from a small number of simple axioms. Frege's work set out such a framework, but it was seen not to perform in resolving paradox, including Russell's own eponymous paradox. This inspired Whitehead and Russell to work together to extend their own current thinking, as expressed in their earlier books on these topics: Whitehead's 1898 *A Treatise on Universal Algebra* and Russell's 1903 *The Principles of Mathematics*. Russell worked particularly on the theory of descriptions, and the no-class theory, in which he argued that to be meaningful, set or class terms must be placed in well-defined contexts.

In 1910, 1912 and 1913, Whitehead and Russell published the three volumes of *Principia Mathematica*. The drafts circulated between them and Russell wrote that: 'There is hardly a line in all the three volumes which is not a joint product'.[78] The title echoed Newton's *Philosophiæ Naturalis Principia Mathematica*–also three volumes but in Latin, and first published in 1687. Whitehead and Russell's work posed a different linguistic challenge, introducing new mathematical notation that frustrated colleagues! Nevertheless, these mathematical giants enlivened and recast the foundations of mathematics.

The quest entered new territory over the following decades with the landmark findings of Kurt Gödel (1906–78), who showed that within any consistent formal system of mathematics based on axioms, there are statements that are undecidable–neither provable nor disprovable. The formal system is said to be incomplete. 'Formally consistent' means that within the system there can be no statement such that both the statement and its negation can be proved. This put the cats among the pigeons and the debates trod widely on the toes of philosophy and logic–peering at and disputing language and terminology, as ever, about: assumption, axiom, type, class, category, context, predicate, proposition, fact, meaning,

78 Russell, *My Philosophical Development*, p. 74.

description, appearance, reality, truth, beauty... This era and these debates were the mathematical crucible of computer science, where the now legendary names of John von Neumann (1903–57), Alonzo Church (1903–95) and Alan Turing (1912–54) appeared to heat and stir the molten mix and pour it into new moulds of theory of computation.

Philosophical and mathematical debate extended into the discourse of computer science and its program languages. It seems that the connections of logic, mathematics and computer science are strong but not watertight in the sense of being fully and consistently argued and proven from common axioms. Frege's 1874 *Habilitationsschrift*, a tough read, set the scene for a hundred years of transition in logic, mathematics and computer science. Tractable elements of the first-order predicate calculus have been assimilated within description logics today, such as within the Web Ontology Language (OWL), introduced in the section on computational reasoning, below. The goal of aligning mathematics and formal logic is still worked on. Some of Frege's original constructs, such as his second-order calculus of predicates, proved not to be watertight, and the full richness of his first-order calculus is, in some cases, intrinsically not computable, and, in others, not yet feasibly implementable in computer systems of today.

In many areas of interface of theory and practice, a decision is required about where to pitch pragmatic compromise in the choice of method for tackling a problem. Both are important but create bias in different directions. How far should we simplify and blur precision of theory, as a compromise in favour of tractable practical application? How far should we pursue precision of theory and, by so doing, compromise practical relevance and ability to implement? It is a balance decided in context, and such balance shifts in time, accommodating new considerations of theory and new methods and technologies of implementation. Advancing precision and range of measurement makes possible new balance and interaction of theoretical and experimental physics. New things that can be measured help in refining theories about the phenomena observed, and new theories of these phenomena helps ground further experiments; together advancing the field. Each aspect is provisional and imperfect but in combination they work towards greater insight and capability–searching for a sweet-spot of theory and practice. Such compromise can alight on sour spots where neither theory nor practice advance and may regress. They may become a sweetshop of readily accessible goodies that taste nice but turn out to do harm. Short term sweetness can evolve or turn to sourness as times move on.

The first-order predicate calculus of Frege was a transforming insight–a sweet spot of formal logic and practical reasoning with knowledge. It broke from the classical language of logic to a mathematical language. It advanced the alignment of mathematics with logic and evolved over the next century to set computer science on a pathway to knowledge engineering and then machine learning. In practice, its theoretical potential had to be reined in considerably, to render computable the answers to the questions it asked, and the calculations required.

The debates are far from over, as they never seem to be when they tread deeply into matters of language, logic, ontology and epistemology. Advocacy against formal logic emphasizes its limitations. It asserts that the restrictions implicit in mathematically-precise logic and reasoning do not match well with the variability of detail, context and nuance of human knowledge and affairs, and the ways in which we express and reason about them in words. Quine and others have argued that the limitations lead to ambiguity in both the formal syntax and semantics of knowledge thus represented. Advocates in favour counter by arguing for a narrowed and more precise scope of philosophical discourse, limiting the context of application of formal logic to one based on realism and practicality rather than abstraction. Others argue that all perceptions of reality are subjective appearances, reflecting belief as much as reality. Logicists counter that 'facts', expressed in a precise context, are not matters of appearance–they can be relied on, axiomatically, to underpin unambiguous logical argument. And so on…

Khayyam's eleventh-century verse could readily be updated to Information Age discourse about knowledge! As could Jean-Baptiste Alphonse Karr's (1808–90) much-quoted saying from the January 1849 issue of his journal *Les Guêpes* [The Wasps]: *Plus ça change, plus c'est la même chose!* [The more it changes, the more it's the same thing!] Philosophers might both agree and disagree, I suspect, and perhaps even take this as a profound statement and compliment–it all depends on how you take the meaning of the words.

Whitehead's writing on these matters from a hundred years ago, in *Adventures of Ideas*, rings true to me. He explored the philosophical issues arising in discussion of logical proposition and predicate,[79] in his examination of the appearance, and truth, of reality. Here are some quotations that characterize his approach:

> Truth is a qualification which applies to appearance alone. Reality is just itself, and it is nonsense to ask whether it be true or false. Truth is the conformation of appearance to reality.[80]

> A proposition is a notion about actualities, a suggestion, a theory, of things.[81]

79 Whitehead, *Adventures of Ideas*, p. 234.
80 Ibid., p. 231.
81 Ibid., p. 233.

> No verbal sentence merely enunciates a proposition. It always includes some incitement for the production of an assigned psychological attitude of the proposition indicated. In other words, it endeavours to fix the subjective form which clothes the feeling of the proposition as a datum. There may be an incitement to believe, or to doubt, or to enjoy, or to obey. This incitement is conveyed partly by the grammatical mood and tense of the verb, partly by the whole suggestion of the sentence, partly by the whole content of the book, partly by the material circumstances of the book, including its cover, partly by the names of the author and of the publisher. In the discussion of the nature of a proposition, a great deal of confusion has been introduced by confusing this psychological incitement with the proposition itself.[82]

> The most conspicuous example of truth and falsehood arises in the comparison of existences in the mode of possibility with existences in the mode of actuality.[83]

He is, though, forward looking and I like that, too. He cautions against what he said might be called an 'out with the new, in with the old approach!'

Logic and Reasoning with Knowledge

> Nothing illustrates better the danger of specialist Sciences than the confusion due to handing over propositions for theoretical consideration by logicians, exclusively.[84]

My purpose in this section of the chapter is to introduce how the connected histories of mathematics, logic and computer science, and the methods of formal logic that have evolved from them, have led to new ways of framing and reasoning with knowledge. It introduces how the computer now reasons about real world problem domains. This leads on to sections discussing the application of these methods in medicine and health care.

Improvement in how we think about, express and reason with our knowledge of the world around us is an important and infinite quest. It reflects what we believe, can observe and measure, wish to see and think about, and thus seek to organize, express and systematize. I think Whitehead was saying that logic is important in this broad endeavour, but perhaps not *that* important! Clearly *that* depends on what we are talking about. Errors of machine logic in the processing of electrical signals in a computer chip or

---

82 Ibid.
83 Ibid., p. 234.
84 Whitehead (ibid.) discussing truth and appearance, predicate and proposition.

circuit board, are clearly crucial to put right. Whole worlds would collapse if 2+2 did not make 4!

When we think about and reason with problems and ideas, we say we turn them over in our mind and sometimes that we sleep on them, relegating them to the unconscious mind. When awake we think and reason instinctively and seek to do so, also, logically and with balance and perspective. In Kahneman's terms, this is *Thinking, Fast and Slow*.[85] It is a mix that is matched to our capabilities and the purposes our bodies and minds are addressing. We navigate between the capabilities and limitations of our brains and the languages they employ to represent and reason; likewise, we grapple with the motivations that drive us to understand and act. These languages have extended our canon beyond biological signals and words into mathematics and logic, and into computation.

If mathematics, logic and computation are to be in harmony, rigorous computational methods of mathematical logic that we can depend on (whilst being cognizant of and allowing for their limitations) are very important. As we become increasingly dependent on computer systems that use these methods to reason with clinical knowledge, similar attention to their applicability is needed. Each language or tool has its characteristic applicability and limitations in what it seeks and is able to do in support of thought and reasoning and the solution of problems. The models of the problem domain that they express are representations of reality and must be understood in the context of the purposes they serve and how well they are achieved. And it is thus with computation, as we construct languages and methods and write programs utilizing them to represent and reason with ideas, analyse observations and measurements, and determine, control and regulate actions.

Apart from a passing reference to Aristotle and his Organon in classical history classes at school, I cannot recall the study of logic cropping up anywhere in my education until the names of Boole and De Morgan appeared in the Theory of Computation module of the Masters course in computer science of the University of London, that I attended in 1970. In my school-day mathematics, the good marks came when one had mastered calculus, vector methods and conic geometry. These foundational methods had been envisioned, refined and evolved in the practical context of problem-solving, over centuries. Exposure to the mathematics of symmetry, set and group theory and topology was extremely limited–school curricula were only slowly catching up with these directions of travel and their contributions in physics.

85 D. Kahneman, *Thinking, Fast and Slow* (New York: Macmillan, 2011).

In my university physics days, the mathematical problems got harder and partial differential equations, and matrix and tensor algebra loomed large. Statistics played only limited roles, given the measurements and data we encountered, as then characteristic of the physics domain–five sigma it was not! Mathematics and theoretical physics were closely intertwined. Principles of mathematical symmetry, illuminated and characterized through set and group theory, were emerging into the world of particle physics and field theory. Early progress got even the greatest of minds a bit carried away. My scientific hero, Richard Feynman (1918–88), was quoted as having commented–after the discovery of the omega-minus particle in 1964, which theory had predicted would exist to complete a modelled symmetry of elementary particles–that particle physics would be done and dusted within fifty years!

Theory of computation was nowhere in my education in physics and not very persuasive in the very practically focused Masters course in computer science that I followed at the end of the 1960s. A decade later, in the early and mid-1980s, Feynman envisioned and created one of the best introductions I know to the subject, first for his physics students at the California Institute of Technology (Caltech) and then more widely. He approached it from the physics of computing devices and spread the net wide to theory of computation, information and coding theory, quantum mechanical computers and parallel computing. I review the edited collection of these horizon-scanning lectures in Chapter Six.

Before getting too starry-eyed about the prospects for reasoning based on formal and computational logic, first a story about stars in the sky and how knowledge about them grew two hundred years ago. Perhaps no other endeavour has captivated the human imagination as greatly as the observation, appreciation and quest for knowledge of the night sky.

The Babylonians knew about the revolving planets but not until the late eighteenth century was Uranus added to the list. It had of course been seen but was thought to be a star, until William Herschel (1738–1822) learned to plot its distant orbit. It was through advent of the Newtonian reflecting telescope and its refinement to much greater sensitivity by Herschel (thus conferring ability for observers to reflect more clearly and systematically on what they saw), that the paradigm of astronomy changed. Newton had given up on the challenge of improving the blurring imperfections of the refracting glass lenses of telescopes and substituted mirrors. Herschel conceived of and painstakingly polished single metal reflecting mirrors to a new scale and precision. No one could help him–there was no commercial basis for such production. The measurement of star position for purposes of navigation did not require such precision.

Observation and measurement of the night sky thus moved up an octave and so did reasoning about the observed stars and their constellations. The long observed cloudy nebulae resolved into galaxies of stars. In time, the term nebula came to characterize gas clouds as the factory of stars. Herschel discerned Polaris as a double star–observation over time of their optical parallax revealed their relative distance. In time, binary stars, lacking such observed parallax, were revealed as gravitationally coupled neighbours. He was so far ahead of contemporary observatories that his results were difficult for professional astronomers to confirm. Some thought his ideas 'fit for bedlam'. But he had the credibility of having discovered Uranus as a planet, to protect him against such assault. In using coloured glass filters to process the images, he experienced heat originating from the red end of the spectrum and, by checking with a thermometer, discovered the infrared. He introduced the time dimension of the observed night sky–how long the light took to reach the earth, speculating that it might be millions of years.

Herschel was a gifted amateur astronomer, long earning his living as a musician. Only once famous was he given patronage and honour to be accepted into the brotherhood of astronomy. It was his sister, Caroline Herschel (1750–1848), who made it a sisterhood, too. She had a similarly sharp and focused mind but had struggled to be allowed to branch beyond domestic life, being constrained by the expectations of their parents. She was helped to escape by her brothers, who surreptitiously enabled her to work alongside William in England. She recorded his observations, called down from the pitch blackness enveloping the position in which he worked, at the end of his twenty-foot telescope tube. She took the opportunity to use the telescope independently and discovered comets. She attended to the rigorous documentation of the new more precise observations and found errors in the then classic Flamsteed catalogue. She added five hundred stars to the previously recorded three thousand. It was called a New General Catalogue and I understand that its coding system survives to this day in their naming. She introduced what we might call a new ontology for this record–moving beyond classification within constellations to one based more systematically and painstakingly on geometry of angles.

Much of the reasoning in the foregoing story was based on a combination of measurement and mathematical analysis. No formal logic was needed, which is the way advances in science have long played out. Indeed, the mathematically based methods of formal logic, introduced in the previous section along the timeline of developments in the logical foundations of mathematics, have mainly connected with reasoning about knowledge within the disciplines of mathematics, logic and computer science, themselves. Their application to provide the theoretical underpinnings of life science and medical knowledge bases is growing but has still quite limited and partial practical application in health care delivery. There have been brave attempts along my songline but not many have persisted.

This pattern has started to change in the data intensive world of the computer, where we increasingly depend on formal logic to help us cope and reason with the huge volumes of data, and complexity of their detail, that land on our doorstep. Systematizing and reasoning with domains of knowledge is becoming mainstream computational science. I survey this scene in the concluding sections of this chapter.

I continue, here, with my brief historical overview of logic and reasoning with knowledge. And to clear my head for this, I spent a day collecting and comparing dictionary entries in Oxford Reference, for common terms describing them that have entered different fields of discourse. I interrogated the learned dictionaries of philosophy, mathematics, logic, computer science, psychology and mind, and some from wider domains–religion, biology and medicine. In practice, I could have gone right around Ranganathan's circle of knowledge, but I called a halt at thirty-five pages of comparative study. The hardest to follow tended to be elaborated at the greatest length. Philosophy excelled in length, followed by mathematics and computer science. These disciplines swapped in order, computer science then coming first when I compared the extent to which they used specialized appropriations of the meanings of commonly used terms.

To communicate successfully with the computer about logic and reasoning with knowledge, we clearly have to define the meaning of words used in that discourse more narrowly than when we share them with one another in everyday life. In normal life, it is often rather easier to construct arguments and reach agreement when we can talk across one another a bit, using words loosely to mean different things in different contexts. That is the necessary ambiguity of politics, after all–the art of the possible and making things possible in the human world! Look no further than the tipping point in the dramatic final session at the COP26 United Nations Climate Change Conference, last week as I write, which hinged on one word, 'down' replacing 'out', in relation to the future trajectory of the use of fossil fuel, on which depends reversal of the habitable earth's trajectory towards 'down' *and* 'out'!

I conducted another exercise in relation to the usage of words. I checked with my weighty and now ancient Compact Oxford English Dictionary (OED), to see how it defines some of the words now descriptive of computer-based information systems. These have become blurred, appropriated and fertile ground for misunderstandings in the fragmented discourse of the Information Age. I was interested to see how the classical wordsmiths of OED define some of them–at least I could confidently give them an alphabetical order! I list these words, here, just to illustrate the complexity that arises when mixing natural and specialist languages to express, communicate, and reason with knowledge, and share this, formally, with the computer.

***Class*:** In logic: a class differs from a catalogue by virtue of a common resemblance in the midst of diversity. Way of grade or quality. Grouping by common attributes.

Classify/classification: modern verb/noun–to arrange or distribute in classes, according to a method or system. Especially in relation to general laws or principles. Department of science that consists in or relates to classification. First use in medicine–1799 by Took: the diseases and casualties are not scientifically classified.

Encyclopaedia: Greek: circle of arts and sciences considered to be essential for a liberal education. The circle of learning–a general course of instruction. Word derives from general education. From seventeenth century: a literary work containing extensive information on all branches of knowledge, usually arranged in alphabetical order.

Language: the whole body of words and methods of combination of words used by a nation, people or race, a ‘tongue’.

Nomenclature (*nomenclatura*, Pliny): the act of assigning names to things. A list of collection of names or particulars; a catalogue or register.

***Taxonomy*:** early nineteenth century. *Taxis*: origin Greek: arrangement, order. From *tassein*, to arrange.

***Term*:** A limit in space, duration–that which limits the extent of anything. Or a limit or space in time. Limiting conditions. Uses leading up to the sense of an expression:

- in maths, quantities used;
- in logic, each of the two elements, the subject and predicate, which are connected by the copula;
- in relation to a syllogism, the subject or predicate of any of the propositions composing it and forming one of its three elements: major term, minor term, middle term, each of which occurs twice.

Definite use of word or phrase in a particular subject. Expressing a notion or conception or denoting an object of thought. Manner of expressing oneself. Way of speaking.

Terminology: system of terms belonging to any science or subject; cf. nomenclature.

> ***Type***: (There are seven usages recorded for the noun, four for the verb!) That by which something is symbolized or figured. Distinguishing mark or sign. Characteristic form of a fever. The general form structure or character distinguishing a particular kind group or class of beings or objects. Hence a pattern or model after which something is made.

As will already be clear, and leaving aside the vagaries of terminology, which feature in the section below on medical language and computation, the applications of formal methods of logic are often hard to grasp and understand–in their details, capabilities and limitations. I studied theoretical physics and had testing experience of ideas that were mathematically 'hard' to understand when I came across them, and that was just within the discipline of physics! It is burdensome for a newcomer to read across disciplines, disentangling what is written about logic in the engagements between learned minds, parsing this into one brain, and sharing it with others.

I am labouring this point, here, only to now point out that the computer is also burdened by the effort required to compute with terminology and logic expressed in such convoluted terms; its elaborations blow fuses of computability. There is an evolving ping-pong match of what the human mind envisages and what the machine can recognize and work with. There is experiment and learning in this game, occupying nimble and learned minds, and, when let loose too soon, it risks creating confusion and burden that becomes entrenched in the wider world.

Verbal and Mathematical Reasoning

As we have seen, the search for precision and consistency of thought and verbal reasoning about the material world and human affairs, has long been the domain of philosophy. Historians trace this quest to early civilizations. Hundreds of volumes were written about it in classical times, and some survived to be pored over and debated in succeeding contemporary contexts. How this inheritance has transmogrified into mathematical and computational reasoning with knowledge, today, is, again, as we have already seen, a long and interconnected set of stories. Natural language of speech and writing has gradually assimilated into specialized language of logic, used to reason about appearances in the natural world. It has assimilated into a meta language of abstraction, to generalize ways of describing and reasoning: words about words, data about data, language about languages.

From classical times, observation and measurement of the natural world and the study of paradox evolved and revolved, encompassing new ideas about the nature and meaning of words, numbers and symbols. The study of reasoning evolved and revolved, encompassing new ideas about grammar, statement and argument expressed in natural language. Philosophy evolved and revolved within schools of thought about the nature of reality and its consonance with ways of reasoning about the world. This Pandora's box of challenge and insight–of discipline struggling to be born–had the lid pressed firmly down for many centuries by the force of prior belief, expressed in the language of mysticism and religion. The power of authority and orthodoxy, and fear of the new, are ever with us.

In the history of Greece, the name Parmenides (born *c.* 515 BCE) carries a flag for metaphysics and ontology–about existence, being, becoming and reality. Aristotle and Chrysippus were flag carriers of two emerging schools of logical thought. Aristotelian logic was based on terms and their conjunction in expressing ideas, and syllogism in reasoning with them. Stoic logic was based on proposition. Stoicism has been remembered mainly in context of philosophy of life. Diodorus (*c.* 90 BCE–30 BCE) was also later associated with propositional logic.

In the time of Aristotle, logical reasoning was pursued through verbal argument and seen as the organon (described as 'an instrument of thought, especially a means of reasoning or a system of logic') through which we come to know anything about the world. Aristotle described logical syllogism as 'a discourse in which certain (specific) things having been supposed, something different from the things supposed results of necessity, because these things are so'.[86] The study of logical reasoning requires the parsing of sentences like that! Formal rules of (deductive) reasoning were the basic principles of this logic, which was accepted in Western philosophy until the nineteenth century, when it was notably disrupted by the mathematical advances of Frege.

Stoic propositional logic fell away, to be rediscovered and built on many centuries later, including by the Polish mathematical logician, Łukasiewicz, influential as a historian of logic and in making bridges from mathematics and logic to computer science. He was the originator of Polish Notation of operators and operands, from as early as 1924–in Polish, and so not widely translated or read!

Bacon produced his *Novum Organon* (1620), to replace the Aristotelian deductive reasoning with inductive reasoning. He advocated the eponymous Baconian method for reasoning inductively from observation to abstract concept, through scientific method and experiment. Immanuel Kant promulgated his *Critique of Pure Reason* (1781). Other notable philosophers of the changing times, as ever, had their say.

86 Aristotle, *Prior Analytics*, ed. and trans. by R. Smith (Indianapolis, IN: Hackett Publishing Co., 1989), p. 2.

Patterns of syllogism and the reduction of logical propositions to a set of normal forms were explored extensively in classical times and evolved in stages. They jumped many centuries into mathematical and computational logic, embodied and debated in many systems of logical inference that covered increasingly diverse and semantically rich kinds of logic, reflecting more complex styles of reasoning, such as the modal, temporal, deontic and relevance logic sub-specializations mentioned in the preceding section about the languages of mathematics and logic. Subsets of these were explored, conditioned by practical considerations of feasibility, correctness of enactment and proof of consistency.

Terms like sentence, subject and predicate, proposition, premise, class, category, inference, truth, causality, existence, universality, necessity, possibility, obligation, permission… came to populate philosophy, grammar, linguistics, mathematics, logic, computer science, psychology and religion… through the ages and to this day. They have provided plenty of scope for people to talk past one another, plenty of reasons to be clear and careful, and difficulties aplenty in connecting logical thought and reasoning with the programming of computers.

When feeling bemused by the rapid inflation of the universe of knowledge in the Information Age, it is salutary and steadying to read Whitehead's critique about such adventure. In his extensive review of the classical philosophical and scientific foundations of knowledge and reasoning with ideas, he writes that 'Where Aristotle said "observe" and "classify", the moral of Plato's teaching is the importance of the study of mathematics'.[87] In his *Seventh Epistle*, Plato had opposed the notion that a final system of reasoning could be verbally expressed. His thinking revolved around seven main notions: ideas, physical elements, psyche, eros, harmony, mathematical relations and what he termed the receptacle. Whitehead expressed caution about logic disconnected from mathematics, saying that 'Aristotelian logic, apart from the guardianship of mathematics, is the fertile matrix of fallacies. It deals with propositional forms only adapted for the expression of high abstractions; the sort of abstraction usual in current conversation where the proposed background is ignored'.[88]

He further argues:

> We can never get away from the questions: How much–In what proportions–and, In what pattern of arrangement with other things? The exact laws of chemical proportions make all the difference; CO will kill you, when CO2 will only give you a headache. Also, CO2 is a necessary

87 Whitehead, *Adventures of Ideas*, pp. 137–55 (p. 148).

88 Ibid., p. 150.

> element for the dilution of oxygen in the atmosphere; but too much or too little is equally harmful. Arsenic deals out either health or death, according to its proportions amid a pattern of circumstances. Also, when the health-giving proportion of CO2 to free oxygen has been obtained, a rearrangement of these proportional quantities of carbon and oxygen into carbon monoxide and free oxygen will provide a poisonous mixture.[89]

It is interesting that, as a philosopher and mathematician, Whitehead used an example from physiology to make this point. Similar arguments about the importance of context, when reasoning with knowledge, can be made with even stronger force, using examples from the medicine and health care domain. He also argues that 'the essential connectedness of things can never be safely omitted' and that all languages witness to the 'error' of 'investing each factor in the Universe with an independent individuality'. Here, as well, 'even the appeal to mathematics is too narrow, at least if mathematics is taken to mean those branches hitherto developed' and 'in the absence of some understanding of the final nature of things, and thus of the sorts of backgrounds pre-supposed in such abstract statements, all science suffers from the vice that it may be combining various proportions which tacitly presuppose inconsistent backgrounds'.[90]

This line of reasoning was crystallized in Whitehead's 'process philosophy', described as a philosophy of organism which envisions reality as composed of processes rather than material objects, defined by their mutual relationships. This organic characterization lends support to how I have come to characterize care information as an organic utility, as developed in Part Three of this book. Also, interesting (at least for me!), is the physicist Carlo Rovelli's current pursuit of a similar idea, in proposing theory of relationship as a fundamental unifying ground of theoretical physics. Other physicists are grounding their 'What is reality?' quests in ideas based on theory of information. Theories of information, relationship and process seem to be circling one another in the physical, biological and virtual worlds!

Later in his book, Whitehead writes that 'All knowledge is conscious discrimination of objects experienced'.[91] Here, he sounds to be at one with the Einstein quotation at the beginning of the chapter: 'Pure logical thinking cannot yield us any knowledge of the empirical world; all knowledge of reality starts from experience and ends in it'. Machine intelligence is, though, beginning to experience and learn for itself, about the world and

89 Ibid., p. 149.
90 Ibid., p. 150.
91 Ibid., p. 173.

the problems we, as humans, face in it. Thus far post-Aristotle and post-Frege; henceforward post-artificial intelligence (AI)?

It is apparent when thinking about this substantially and continuingly anarchic scene, that we have much still to learn about logic and reasoning with knowledge. It has made all manner of human reasoning more complicated, as well as offering opportunity, and necessity, to make it better. In the panorama of medicine and health care, the full spectrum of this adventure and the programme for reform it makes necessary, is playing out all around us. How is the computer coming to grips with this scene? The story now moves into the decades during and after which Whitehead was writing, to the rise of computer science, the computing machine and machine intelligence.

Computational Reasoning

The quest to connect logic and reasoning with knowledge, with computationally tractable axioms and proofs, is a complex challenge. It is playing out in the invention of new computational methods that are being explored across multiple problem domains. Given current uncertainties, it seems a reasonable concern about this direction of travel that, even if new mathematics, science and computation do succeed in representing knowledge, measuring phenotype and analyzing data about human physiology and behaviour, to provide tractable new methods of logic and reasoning about human health care, in useful ways, the methods thus employed may ultimately prove beyond the capacity of human minds to understand and regulate. The machine may then be left to establish yet another sense, adding to the many that du Sautoy remarks already exist, of what it is to know, and what it is to reason, and to decide what should be made of and done with that knowledge.

With that caution in mind, I start here by drawing together a historical timeline of the challenge of computational reasoning. In much of scientific endeavour, whole communities share both task and credit. The communities of actors are large and widespread, and progression a complex mix of shared and contested insight and activity, context and staying power. That is not to deny the inuksuk-quality significance of the insights and historical contributions of key individuals, and I have opted to highlight a few.

Charles Babbage (1791–1871) was both an engineer and mathematician and his interest in building his first mechanical computer was as a tool for mechanizing the labour of calculating astronomical tables. A hundred years later, the system of Boolean logic, and theorems based on it, such as De Morgan's laws, became a central motif of circuit design for electronic computers, applying logical functions to binary data and performing binary arithmetic.

Theory of mathematical logic became foundational to and increasingly intertwined with computer science of the twentieth century. As we have seen, this evolution accelerated in the decades either side of the turn of the century, with the landmark contributions from such as Cantor, Frege and David Hilbert (1862–1943) in Germany, Whitehead and Russell in England, Łukasiewicz, Polish-born in the then Austro-Hungarian province of Galicia, and Gödel, from then Czechoslovakia. Such people were giants and recognized the giants on whose shoulders they stood. Modern-day theory of proof treats the reasoning employed in proving mathematical theorems and verifying their correctness within a framework of mathematical logic. It grew from the work of Cantor and Hilbert and connected with theory of computer science as it got to grips with the expression and correctness of algorithm and program.

Church and Turing focused on theory of computation and its mathematical limits. Turing developed the idea of an abstract universal computing machine, on which any computational process could be performed. This gave traction to the study of computability–what computers could and could not, in principle, compute, from a mathematical standpoint. Church invented the lambda calculus as an abstract theory of computation, also drawing on Łukasiewicz's Polish Notation and the mirror Reverse Polish Notation used in describing functions. This part of the story links with Chapter Five on information engineering.

Theory of computation enabled and helped formal logic to grow, and limitations of computer power and computability guided and constrained its applicability as a method for solving problems. Some computational tasks arising within theory of formal logic defied feasible program implementation and available computing resource. Exploratory implementation efforts were restricted to handling subsets of the expressions that the formal logic could, in principle, handle.

The story became more complex and specialized as it extended more widely across cognate disciplines, framed by their different language and discourse. Philosophy of mind debated with neuroscience and computer science, giving new context to the study of logic and reason. Reasoning about knowledge extended in the context of the computer science of program languages and development of early knowledge base systems. The idea of self-referential systems entered discussions of consciousness and mind, as the cognitive psychologist, Douglas Hofstadter explored in his book, *I Am a Strange Loop*, which I introduce in Chapter Six.

This history has, thus far, trended towards philosophical bemusement and computational intractability! The perplexity it presents is reminiscent of René Descartes' (1596–1650) philosophy of mind and body dualism, once memorably batted off with: 'What is mind? No matter: What is matter? Never mind'.[92] And in response to the still baffling physics question, What is reality?, a similar response has been: don't think about it, just solve the Schrödinger equation! Reducing theoretical physics to its simplest axioms has long been a recurrent work in progress and the quest has involved much new mathematics.

Notwithstanding the fractal complexities and contingencies of human affairs, we decide and enforce rules whereby society operates, and individual components are required to function and behave, presided over by judicial systems. These change over time and exceptions are made–the rules are said to be 'defeasible'. Rules that do not admit of such change are 'indefeasible'. Where we seek generality, we focus on the indefeasible. Where we recognize and allow for variability and contingency, we are in the defeasible realm of the particular. The more we seek to allow for and define the contingent and defeasible, the more complex and ungovernable the computational edifice we create. Today's national tax codes that look after every nook and cranny of compliance and default, extend from hundreds to thousands of pages–those of some countries are much more complex than others. Judicial functions are challenged, too, and human judges teeter on the limits of their oracular finesse and prestige. Human judgement admits of increasing noise and bias and, as Kahneman, Sibony and Sunstein suggest, assumes the character of imprecise measurement.[93] Health care knowledge and clinical records populate a highly defeasible domain.

What problems does the ingestion of this noisy and uncertain world of discourse present to the machine? One might say that it tends to exhibit the machine equivalent of indigestion, burping and adding to the noise and general disarray! It needs a digestive system for the data it eats and the information it processes and promulgates. It causes problems with the use of terminologies and their change over time. It creates problems of burdensome legacy as knowledge evolves while knowledge base systems are unable, or fail, to keep track. It causes problems in the navigation of a middle ground that attempts to represent and reason with what we deem indefeasible axioms of knowledge and what is contingent.

Problems of how to embody the generalism of knowledge about the human circulatory system–what it is and how it works–and the particular knowledge about individual patients who, given an underlying disorder, may present with a generalized pattern of symptoms, but not always so. Myocardial infarction generally, but not always, exhibits a pattern of

92 Ibid., p. 173.

93 Kahneman, Sibony and Sunstein, *Noise: A Flaw in Human Judgment*.

radiating right arm pain. Most might imagine that putting a lower bound on blood pressure compatible with life might be framed as an axiom, but someone, somewhere, for sure, will put up their hand, in all seriousness, to ask about cryogenic preservation of frozen bodies–whether they have a blood pressure and in what sense they are still alive!

Computational logic and reasoning lie within evolving, both theoretical and empirical domains of endeavour. Logic is concerned with provable truths. Formal logic has struggled to represent uncertainty and contingency. Computational logic endeavours have come to lean more to the theoretical than the practical side of the purposes they serve. Statistics helps us with uncertainty and Bayesian statistical methods help us with incremental learning to improve reasoning, in the context of evidence elicited and the framework of concepts we use to measure, model and analyse the system being studied or worked with. Computer software sets us free to create and reason with representations of knowledge that are detached from enforced mathematical correctness. That may be a good or a bad idea depending on the situation. This approach is sometimes called heuristic–applying 'rules of thumb'. At least we should be aware when we are usefully employing heuristics and when we are just getting our sums wrong!

All methods of computational logic require wide-ranging appraisal of their rigour and applicability in particular situations: how faithfully can they mirror a desired way of describing a system; are they formally rigorous and consistent in their statements and proofs; and how useful can they be in refining the knowledge they embody, and improving the services that teach, use, sustain, promulgate and update that knowledge?

Machine learning methods, as further discussed in the next section, widen the software scope further, to provide the ability to learn an optimum method for reasoning in a defined context of questions asked or decisions to be made. What is the three-dimensional structure that this DNA sequenced protein molecule will fold to? What is my best next move in a game of chess or Go? The method trains itself in this skill through an experimental process of trial and error, to create its own heuristics, based on historic databases of known and classified cases and other relevant knowledge about rules of the game or domain–rules of chess, a bank of known protein folding structures, scientific knowledge about the domain concerned.

These many and varied exploratory methods butt horns with human discourse. Formal logic has butted horns with philosophy in discourse of ontology. Mathematical and computational formulations of reasoning have butted horns with other communities and their methods of measurement, analysis and reasoned judgement, such as health professionals making judgements about the patients they care for.

In the realm of medical terminology, we are trying to arrive at both expressive and consistent vocabulary that can reliably be used to codify knowledge and pass it for further use and analysis. We must expect changing concepts of health and disease, that develop in line with changing means of measurement and analysis, and new representations of this knowledge in computer systems. A formally rigorous logical representation is required to ensure that consistent and correct inferences can be made when joining together elements from diverse computer representations of knowledge. Combined in this endeavour are validated methods for representation of the modelled reality with mathematically provable methods for inferences then drawn. Medical terminology has proved a fruitful area for exploration of the potential of the branch of formal logic known as 'description logic'. Medical language and terminology and the representations of knowledge about them are further discussed in the section below.

In the wider context of the use of knowledge bases in support of clinical reasoning and decision making, we encounter a mixed world of indefeasible and defeasible (contingent) statements, and consideration of how well and usefully different computationally anchored representations handle them. The methods explored have evolved considerably over at least six decades and I trace some landmark systems in the section on medical knowledge bases. One approach is to operate as a law maker and use the program to express rules to be followed, with embodied reasoning anchored to these rules. Early 'rules-based expert systems' adopted this approach.

Often, a hierarchy of knowledge representation is postulated, seeking to give a sense of the structure of central and subsidiary detail required. Sometimes this is expressed as a span from a background knowledge base about generalities and a foreground database about particulars. In clinical contexts, this database may include care records of patients and other defeasible and contingent knowledge that is relevant to the purposes served. Another consideration when deciding on a formal method for modelling a domain of knowledge, is how we are to interpret and act on results arising from computations involving the knowledge base thus created. This leads us to consideration of 'closed-world' versus 'open-world' assumptions.

The closed-world assumption is more prescriptive–if a statement is not provably satisfiable within the context of the knowledge base under consideration, it and statements that logically follow from it are deemed false. It is sometimes expressed as taking the view that all knowledge relevant to the purposes served by the knowledge base is represented within it. And thus that, if a proposition or statement can be shown not to be satisfied in this knowledge base context (sometimes called a 'world'), it can reasonably be treated as false, and all consequential propositions or statements that follow from it, likewise.

The open-world assumption is more restrictive and sets a higher bar of proof for a statement, requiring that its negation is provably untrue in the context of any possible 'world' consistent with the defined knowledge base framework. Otherwise, the matter remains undecidable. It recognizes that no system can have complete knowledge and that future changes in the knowledge base may render statements that were formerly undecidable, decidable.

Knowledge bases using a mixture of open- and closed-world assumptions have been explored. These assumptions become of great significance when searching for information by interrogating the knowledge base with logically constructed queries. The semantics of such query constructions become all important in guarding against incorrect, misleading, or unintended results.

Two principal approaches have predominated in this exploratory era of computerized knowledge bases–substantially pragmatic approaches called frame logics, which adopt the closed-world assumption, and later, substantially theory-based, approaches drawing on description logics, which follow the open-world assumption. Over time, the two have been admixed, rather as we saw in the world of library classifications. All this quickly becomes a highly context-dependent discussion and the arguments adduced, one way or another, need to be based on a clear statement of goal and definition of method, informed by experiment with implementation. Such uncertain endeavours seem easily to become entrenched in debate over mutual understanding of methodology, between subject domain and logic domain specialists, more than experimental findings in practical implementations.

Wider historic and up-to-date reviews of the field of clinical decision support, have been provided by luminary figures such as Mark Musen, Blackford Middleton, Robert Greenes, Dean Sittig and Adam Wright.[94] The previously cited paper of Rector et al.[95] is essential reading for understanding how this field has played out in attempts to harmonize the knowledge bases of two state-of-the-art terminologies, SNOMED (Systematized Nomenclature of Medicine) and ICD (International Classification of Diseases), as introduced in the section below. This incisive and battle-hardened overview of the state of the art is magnificent and Rector's personal advice has been generously given in helping me make my own, considerably more limited,

94 M. A. Musen, B. Middleton and R. A. Greenes, 'Clinical Decision-Support Systems', in *Biomedical Informatics*, ed. by E. H. Shortliffe and J. J. Cimino (Cham: Springer Nature, 2021), pp. 795–840; B. Middleton, D. F. Sittig and A. Wright, 'Clinical Decision Support: A 25 Year Retrospective and a 25 Year Vision', *Yearbook of Medical Informatics*, 25.S 01 (2016), S103–16.

95 'On Beyond Gruber'.

sense of this domain. I hope I have not departed too far off beam in this. The paper is a tough read but a worthwhile one–as with the Frege papers on predicate calculus it is best approached with a cold towel around the head!

Frame Logic

Frame-based methods emerged as a mechanism for capturing knowledge concepts and relationships among concepts, expressed as descriptive data structures and tools for reasoning with these. They seek to represent what is known and deemed relevant in a particular domain of application. The set of frames thus defined, and connections between them, form the basis of reasoning about the domain of knowledge represented. Frame logic is a pragmatic, cut-and-dried methodology and computer programs can more readily navigate the interface of defeasible and indefeasible domains of knowledge that it may encompass.

Notable tools to support frame-based knowledge systems have been pioneered by Musen at Stanford University. His team's renowned Protégé system became a widely used open-source ontology editor and knowledge management system in the field of biomedicine. In its later evolution, Protégé frames methodology moved on to adopt the OWL (Web Ontology Language) description logic methods alongside the frame structure, enabling the language of OWL also to be used to reason with the knowledge about the domain of application.

Protégé was a generation further on from an earlier pioneering initiative, the MYCIN rule-based expert system, created to represent and reason about antimicrobial therapy in clinical practice.[96] Even earlier initiatives originating from the same team, notably at Stanford and Massachusetts Institute of Technology (MIT), had led to the first expert systems, such as the Heuristic Dendral system of Edward Feigenbaum, known as the father of expert systems, and his team. This system identified chemical structures from their mass spectrometry profiles; one of the pioneering examples of knowledge bases that I introduce in a separate section below.

The evolving Stanford team, environment and field of work has made connections all along the timeline and songline of this book. It illustrates many of the lessons that recur throughout. It has been a sustained source of effort to construct knowledge bases useful for the support and improvement of clinical decision making in the Information Age. These experimental endeavours have connected new and evolving methods of observation and

96 E. Shortliffe, *Computer-Based Medical Consultations: MYCIN* (New York: Elsevier, 2012).

measurement, modelling and engineering (Chapters Three, Four and Five in Part One) with the transition of medicine into the Information Age (Part Two). From the early days, the experiments went hand in hand with the research of computer scientists, such as Ivan Sutherland, exploring new programming languages like John McCarthy's (1927–2011) LISP, tuned to reason with formal expressions of knowledge, just as they now connect with the unfolding field of machine learning.

Description Logic and the Web Ontology Language (OWL)

Description logics are languages designed to express and reason with knowledge about real world entities, utilizing methods of mathematical logic. Statements in these languages are used to represent what is known about a domain of knowledge, in a computable format. They are used to describe the logical structure of the domain and populate, manipulate and interrogate an associated knowledge base with content that fits within this structure. There are many such languages, and they utilize different levels and constraints of mathematical logic.

They are used to check the mutual internal consistency of both the knowledge structure and the entries placed within it. The logic statements expressed in the description logic language are treated as mathematical axioms and used to prove theorems based on them, to assist in reaching conclusions and informing decisions. An underlying inference engine, embodying formal rules for logical reasoning about the knowledge represented, such as *modus ponens*, is used to prove the consistency of the logical statements made and infer further true statements relevant to the purposes the knowledge base system serves. These might be to place the entities being described into a logical structure or taxonomy, showing where they fit within a hierarchy, and how they otherwise interrelate. This reasoning enables the construction of a logical network of relationships.

In knowledge bases programmed using the methods of open-world description logic, the representation of knowledge is in the form of indefeasible statements (axioms) about the things described–these brook no exceptions and are called invariant in formal logic. The computer may understand these axioms, but they may, nevertheless, prove impossible to compute with (be intractable) when attempting to reason about this knowledge representation, for example due to mathematical limitations or the scale of 'combinatorial explosion' (escalation in amount of computation) involved.

A growing range of description logic languages has evolved to represent and reason with knowledge. They occupy a sweet spot between theory and practice, comprising a set of formal logic methods that are capable of computing with the knowledge in the form they represent it, which can be implemented efficiently. In the main, those in use today are said to be more expressive than propositional logic and less expressive than first-order predicate calculus.

Description logic provides for two areas of reasoning and these exhibit different scales of computational complexity. There is T-box reasoning about what are sometimes called 'necessary truths' about classes of descriptive knowledge (for example, axioms about the classes of lungs, pneumonias and bacteria) and A-Box reasoning, which covers assertions made about individuals (for example, London, Manchester and Bristol as examples of English cities). The system might be designed to reason about cities, with some axioms being statements about the characteristics of any city, in T-box style, and some being assertions about London or Manchester, in A-box style.

In the health care domain, description logics have proved useful in furthering the quest to ensure consistent use of terminology for representing, analyzing and communicating its wide-ranging and complex hierarchies of knowledge (sometimes called terminological knowledge). To quote Rector et al. directly:

> In general, reasoning about classes ('T-Box reasoning') can be optimized computationally (although in expressive dialects it is worst-case intractable). However, reasoning about individuals ('A-box reasoning') is much more difficult. Conveniently, it is primarily axioms about classes that are relevant for terminologies.[97]

The Web Ontology Language (OWL) is a family of description logic languages that brought formal logic to the Internet, as a foundation of the Semantic Web.[98] This built on the World Wide Web Consortium (W3C) standards for data description on the web, notably Extensible Markup Language (XML) and the Resource Description Framework (RDF). OWL is for the interchange of knowledge representations, what XML and RDF have been for the interchange of data. OWL has been used as a tool for the formal representation of many different domains of biomedical knowledge.

Computer systems that represent and reason with knowledge about different domains of study, using methods of formal logic, have increasingly found their way into exploratory real-world applications. In medicine, methods of description logic have principally found application in

97 Rector et al., 'On Beyond Gruber', p. 4.

98 Wikipedia contributors, 'Web Ontology Language', *Wikipedia, The Free Encyclopedia* (17 April 2023), https://en.wikipedia.org/wiki/Web_Ontology_Language

endeavours aiming to provide logically consistent and useful representation and management of medical terminologies, such as SNOMED, and in creating models of biomedical structures, such as the Foundational Model of Anatomy (FMA). I follow both in sections below. And in recent times, the advance of machine learning and artificial intelligence, to analyse, reason with and make decisions, based on data collected from widely across the domain under consideration, has created a new buzz. I often tend to elide the two terms as 'machine intelligence'.

How and where different kinds of formal logic and machine intelligence can be applied, and prove reliable and useful, in well-characterized domains, and in relevant practical contexts, is an open question. Time will tell and a lot will change along the way. The hope is that they will enhance human health and wellbeing and improve the ways in which society operates and develops. The concern is for caution, lest this trend prove an overly Faustian bargain, heralding an era of machine intelligence that acquires and maintains momentum in directions that quickly deskill, demotivate and degrade human endeavours, adding to social inequalities and divisions. We must create and navigate this future songline, guided by King's audacious pessimism and Obama's audacity of hope.

Machine Learning and Artificial Intelligence

Machine learning adopts a naive approach to solving problems. In theory, it starts from almost nowhere in terms of knowledge assumed, save for some nascent capabilities for structuring information, using, for example, 'neural network', 'genetic algorithm' and database, with statistical and algorithmic dexterity–a bit like the Chomsky idea of an inbuilt human capacity for grammar. But if the problem posed is to learn to play a game, it seems only fair to let it know and make it follow the rules! And in setting it to tackle any problem, the machine might also, advantageously, be pump-primed in some way with human experts' accumulated knowledge and expertise in solving that problem–track record and strategies acquired in becoming good at winning in a game like chess or GO, for example, or a 'test set' of cases for identifying abnormalities in clinical images, where the answers deemed correct are given, and the problem posed is for the machine to become good at arriving at these answers on its own.

The machine kicks off with a best guess idea of a computational method or solution it seeks for solving the problem posed to it–playing its chess moves, for example–a bit like starting from a prior probability in a Bayesian method of inference. Equipped with its internal computational method for managing and learning from experiment, it tries this idea out, observes the outcome it achieves, revises its guess and tries again, seeking improvement. And it does so again and again, iteratively homing towards better and better method or solution for the problem posed.

Such a quest for improvement can be rapidly played out over many millions of iterations in virtual reality. And such systems consume electricity in gargantuan amounts! It is a method akin to the proof of a pudding being in the eating, as the computer eats the numbers and digests them to become a more useful tool! This tuned and continuously updated resource can then be applied prospectively to newcomers–for example, to play chess for real against an opponent, having tired of playing itself, or to propose a three-dimensional folding structure of a protein, for which the proteome sequence is known but the corresponding structure has not yet been found. The machine learner becomes a skilled pattern recognizer by working out its own patterns. All this is somewhat akin to humans who say that they will sleep on a problem. On waking with an answer, they may struggle to articulate in humanly accessible ways, how they arrived at this solution. Just as skilled clinicians may struggle to explain their immediate ability to recognize and interpret a pattern in what they see, when caring for a patient. The machine, and it is just a machine, may likewise become good, and even better than humans, at discerning patterns in, and interpreting, observations and measurements recorded, as, for example, in medical images and biochemical profiles. And as with any such tool, its results tend to prove better, prospectively, for well-defined and -circumscribed problems, and such constraint may not be easy to pin down in the highly contingent world of clinical practice. At some level of description, each patient is unique, and, in that sense, their associated contextualized data exist in a sample space of one.

The machine learner surveys all manner of recorded appearances in its circumscribed problem domain, or 'world', and searches out more, learning to make its own connections and adapting its learning, guided by the experience gained in trying and failing, as it iterates to improve. In its simplest form, this is an empirical process of learning from experience, naive because childlike–the way the newborn starts to experience, learn and acquire knowledge. This process might now be termed heuristic and I further connect the storyline with pioneers of 'heuristic programming' later in this chapter and in Chapter Five. It is much as the Greek philosophers confronted the foundations of mathematics and logic, learning to reason, faced with paradox and inconsistency revealed in the application of their arguments. It is much as the world learned to become systematic about medicine, in the evolving practice of wandering healers, bridging from concepts of divinely conferred affliction to descriptions and management of treated disorders. The machine-learning algorithm is a naive learner, its neural networks a childlike (but inhuman) learning brain.

At first glance, human factors may not appear to feature in the machine. This level of machine intelligence will not seek to psych an opponent in a game and nor will it get emotionally tired, drained or frustrated, and tip up the board, or become erratic and go home with a headache! One can imagine that machine intelligence may, though, subtly, and perhaps harmfully, distort the human world it samples and learns from, as it evolves. It may break down or get too difficult or expensive to sustain. As it interacts with and influences the wider world, that world will change, and the problems addressed there with machine intelligence will, too. New challenges to human governance will arise.

Digging deeper in one particular 'world' of specialism does not necessarily lead to a 'better' solution, when set against the immense breadth and variability of human phenotype, knowledge and behaviour intrinsic to the wider 'world' of health care. It may create new problem domains where machine intelligence, itself, becomes part of the 'problem'. What might a world look like where a Go or essay-writing competition is battled between AlphaGo- and ChatGPT-like machines? Of course, they will have come up with much more taxing pastimes! Machine fusion, powered by nuclear fusion, resulting in human confusion is probably best avoided! Such are some of the real and imagined issues that look to face us in coming to terms with what the world now recognizes as AI, and these have to be weighed experimentally, in human ways, in the real world.

In September 2019, a thoughtful and fluently interesting article appeared in the *Times* newspaper, characteristic of its well-respected journalist author, David Aaronovitch. It headlined the potential future benefits of AI in medicine, drawing extensively on a visit to see the work of a DeepMind AI company team collaborating with clinicians at my nearby Royal Free Hospital.[99] A patient, doctors and nurses at the hospital were quoted as expressing strong endorsement of the revolutionary benefits flowing from a DeepMind software App called Streams, that they had been piloting on their wards. They described the clinical insight it provided, by monitoring and alerting the team to significant deterioration of a patient's kidney function, in a timely manner that enabled effective intervention to support and stabilize them.

The article made a sharp comparison of the reported frequent failings of the breakdown-prone hospital-wide information system in use, with

99 D. Aaronovitch, 'DeepMind, Artificial Intelligence and the Future of the NHS', *The Times* (14 September 2019), https://www.thetimes.co.uk/article/deepmind-artificial-intelligence-and-the-future-of-the-nhs-r8c28v3j6

the clinical utility of this free-standing App. The App itself accessed a limited subset of the patient record, as stored on the hospital system, and communicated directly with mobile devices carried by clinicians treating the patient concerned, alerting them to imminent and acutely threatening kidney injury (AKI). The article made further connections with pioneering work on AI in the USA, including in the Veterans Administration health system, exploring the efficacy of AI systems to help analyze patient data–for example providing accurate interpretations of clinical images and timely alerts about detected adverse trends–and optimize and streamline workflow for clinical teams.

There are a number of issues here, which illustrate the checkered pathway along which AI is rapidly evolving, today. I consider these further in the context of factors shaping future health care, in Chapter Eight. There, I describe an authoritative 2019 book (surely an inukbook) by the eminent clinician and medical scientist, Eric Topol, in which he reviewed and provided a wide perspective on the field.[100] He quotes what he calls a published 'sharp critique' of the domain and key points made there, including about the lack of transparency in its methods and 'growing (often self-interested) misinformation and mystification of the field'.[101] He was interviewed, and his book cited, for the *Times* article.

When I read the article, I was both impressed and surprised. The message about AI was a good one but the lead-in and main story about a patient and clinical team, which served effectively to humanize and dramatize its impact, was misleading. First, the Streams App reported was apparently a version based on a straightforward calculation using measurements of creatinine level in blood samples, not on any much more complex computation associated with and conjured in the reader's mind by the term artificial intelligence. Thus, the approval reported was not a reflection of AI, but rather of an important much-improved clinical workflow that the App enabled. Apparently, the measurements of creatinine level that the App used were already being collected and filed in the hospital IT system, but their significance was not being detected there, and was thus not alerted to the clinicians overseeing care. The long article did say, later on, that the Streams App was not based on AI. But then, if so, why use it in this way to promote a message about the potential benefits of AI?

In this regard, I recalled a paper from forty years ago, with authors including the current President of the Royal Society, the eminent statistician of Bayesian methods, Adrian Smith. This reported that a Kalman-filtered time

100 E. Topol, Deep Medicine: *How Artificial Intelligence Can Make Healthcare Human Again* (London: Hachette, 2019).

101 Ibid., p. 94.

series of essentially the same creatinine measurements, enabled a similar, up to several days, advance warning of imminent failure of kidney function, relative to that which was otherwise evident to attending clinicians–for renal transplant patients in that case.[102] I remembered collecting the paper and communicating with its authors at the time. A later section of the *Times* article described progress in development of AI methods, achieving similar detection and advance warning of significant clinical problems. It reported a visit to a Veterans Health Administration (VA) hospital in the US, where AI researchers extolled the potential of AI in enabling these benefits.

If this level of clinical impact is now seen as revolutionary, and on the face of it, it certainly appears so, why has it not been pursued with the statistical methods explored and reported forty years ago, and become a recognized standard practice over those intervening decades? Does labelling it as an example of AI somehow give it a magical status and justify its realization now, through AI becoming a priority?

All that said, it is entirely feasible that present-day state-of-the-art machine learning algorithms will progress to make immense contributions in preventive medicine and the further enhancing and streamlining of health care services–whether providing AKI alerts, detecting adverse pathology in radiographic images or retinal investigations, identifying and monitoring population subgroups at specific risk of disease, guiding patient self-care, or the like, and whether used in the patient's home, or a care home or hospital setting. Based on that reported experience at the Royal Free Hospital, it was, though, incautious exaggeration to align AI with the benefits reported–clinically, economically, and in both staff and patient appreciation–since no machine learning method was then involved. Unfortunately, hyped and uncritical projections about the health care benefits of AI are rather common at present. Such Apps as the one described in the article are a front end of the improved information flow on which AI can and will feed. A key observation, also made by Topol in his landmark 2019 review of health care workforce for the NHS, and in his book, is that to work well and beneficially, such new method will require coherent, accessible and timely digital care records, which is a central concern for the patient-centred care information utility advocated in this book.

There is a further relevant and evolving perspective about AI arising from the analysis and prediction of weather systems. Today, as the modelling of these systems combines with machine intelligence, both are proving differently advantageous in forecasting of weather and what is

102 I. M. Trimble, M. West, M. S. Knapp, R. Pownall and A. F. Smith, Detection of Renal Allograft Rejection by Computer', *BMJ*, 286.6379 (1983), 1695–99, https://doi.org/10.1136/bmj.286.6379.1695

being called its 'nowcasting'. For an immediate (now) prediction of local weather trends, machine intelligence can nowcast based on measurements of the current weather, including wind, temperature, cloud cover and time of day, to outperform complex physics-based model predictions. Longer-term and wide-area forecasts are still the preserve of complex models of atmospheric physics.

At this point in the book, and after peering briefly to a future reality of AI-supported health care, it seems appropriate to come back down to earth, in the here and now. The applications of computational logic and reasoning are experimental; however, as their ambition broadens in scope, the feasibility of conducting controlled experiments to study and make informed decisions about them diminishes. Somehow, our age has become of the mind that such applications are not of an experimental kind–often treating them, in embryo, as mature in form. Treating principles of computation as purely technical abstractions has had a harmful influence on the interplay of information technology and medicine. The interaction of science and technology with health care engages with personal and social problems as much as with mathematical and scientific ones, and the field of health informatics has often lacked, and not sufficiently prioritized, practical engineering interface between these related scientific and social domains. The framing of policy for services that bridge them has been characterized as a 'wicked problem'. Such problems feature throughout the book, though I discuss them in the context of health care policy specifically in Chapter Seven.

The interactions of theory with experiment and practice, accumulate as the storyline and chapters of the book moves around the Ranganathan circle, through the worlds of philosophy, mathematics, science and engineering, into the worlds of life science, medicine and health care information systems. The importance of efforts to study and experiment with their interconnections, in the context of the grand challenges they illuminate and address, is common sense. But a relevant community and environment in which to achieve a credible balance of theory with tractable and sustained implementation, where such experiment can proceed and prosper, has proved exceptionally hard to create and manage. There have been notable exceptions and I describe a number in Chapter Eight, and the pioneers who created and led them, and enabled good things to happen there.

The story now moves into the world of medical language and computation. In this, I focus on two examples of knowledge representation taken from the medical and clinical domain–the clinical terminology SNOMED, and the Foundational Model of Anatomy (FMA). They exemplify the state of the art in software whereby knowledge can be represented and used to create programs that organize, search and reason in these domains. They illustrate two general forms for representing hierarchy of knowledge, termed 'compositional containment hierarchy' (as exemplified by the FMA) and 'subsumptive containment hierarchy'(as exemplified now by SNOMED).[103] Here the container aims to provide a full and consistent description, at a chosen level of detail and within a chosen context of use. These initiatives illustrate the evolutionary interplay of knowledge, logic and machine capability in the practical problems tackled. They have been important foundational experiments–SNOMED, over many decades in exploring and developing theory and method in the context of practical health care requirements.

These two examples then bridge to the topics of Chapter Three, connecting from knowledge, language and reason to observation and measurement, and Chapter Four, broadening the discussion to the modelling of different kinds of systems, to represent and reason about their structure and function and the way they work.

Medical Language and Computation

In the storyline thus far, we have visited the historical evolution of library classifications, in their aim to achieve ordering of knowledge, and parallel quests to express, communicate and reason with knowledge, through verbal language, mathematics, logic and then computation. These histories have touched on examples from health care while setting the scene more generally for discussion, now, of the quest of recent decades to tame and computerize medical knowledge.

It has proved a perilous and contested terrain! One on which unsuspected dragons have revealed themselves, to become fired up and magnified in the Information Age! The quest to computerize always brings us face to face with what we know and do not know about a subject, how we express and communicate it, and how we reason with and from it. It challenges us as to why we are doing what we are doing, to what end and how well.

Medicine and health care appeared to offer enticingly rich pickings for computerization, in all these dimensions of expression, communication and reasoning with knowledge. Early endeavours focused on creating dictionaries and databases of medical and health care language and terminology. These evolved into quests for logical rigour and consistency in taming the huge and multi-faceted corpora of terminologies and classifications of knowledge

103 For basic definitions of these subtypes of hierarchy, see Wikipedia contributors, 'Hierarchy', *Wikipedia, The Free Encyclopedia* (21 June 2023), https://en.wikipedia.org/wiki/Hierarchy

that emerged, reflecting different but overlapping purposes, goals, methods and governance. It has been a torturous escapade of discovery, with some, but not much, low-hanging fruit found on these growing computerized trees of knowledge. They were more like a coconut palm–a very tall trunk to climb before reaching the fruit, with the coconuts often falling on the brave climbers' heads before they reached them!

Terminologies used in expressing ideas are noisy and imprecise and so may be the judgements based upon them. A connecting thought caught my attention when re-reading Whitehead's *Adventures of Ideas*, as I was considering how to construct this section of the Chapter. In Chapter XV on Philosophic Method, where he is discussing how we use language to express and generalize experience, he says that what may appear a redundancy of terms used, is in fact required and that 'the words correct each other' in conveying meaning. In this sense, tying expression to a limited vocabulary risks limiting and harming communication of meaning.

On another occasion, preparing for writing Chapter Six, I was re-reading an equally wonderful book, *Feynman Lectures on Computation*, where binary coding and transmission of electrical signals are discussed in relation to the methods used to correct errors generated by the electrical noise experienced during their transmission.[104] Here again, redundancy in the transmitted binary data is key to its accurate communication. Error-correction methods, without which no digital network infrastructure can function, work on the principle of transmitting redundant additional data, which is generated from the signal being encoded and tagged onto it during transmission. The system of coding is designed such that the undamaged transmitted data can be reconstructed, to an extremely high degree of accuracy, from the noise-beset erroneous data received. This is at the expense of transmitting the redundant bits of the digital signal and the encoding and decoding process needed to set up the transmission of messages and detect and correct for any errors encountered.

Communication of meaning about individual health care is prone to all manner of noise, discontinuity and imprecision. During my career, I sat for many years within hearing distance of the lunch-time discussions among clinicians about their patients. It is a very efficient channel but not widely scalable to the high-intensity, multi-faceted, distributed settings of health care today. Where we computerize these records and communications, it is essential that their human meanings are communicated well. Clinical meaning is poorly communicated between today's non-coherent digital care

104 R. P. Feynman, *Feynman Lectures on Computation* (New York: CRC Press, 2018).

record systems. The motivation of openEHR, as discussed in Chapter Eight and a Half, has been foursquare focussed on improving this reality.

Records of health care are both narrative and structured datasets. They connect with the personal story of the individual patient and the professional record of their care. The records of individuals also need to be shared and integrated within the wider health care system. Achieving expressiveness in the natural language of the record and formal structure that facilitates meaningful access, analysis and communication, more widely, places natural language in apposition to language of the computer system. These goals pose complex requirements and the solutions thus far evolved have been a complicated mix.

Natural Language and Medicine

Joseph Weizenbaum (1923–2008) expressed concern about language in relation to medicine and computation. His seminal book, *Computer Power and Human Reason*, features strongly in my Chapter Seven, where I move on to discussion of health care services.[105] He spoke of language used for the organization of facts and assertion of axioms and theorems, in the context of processes.

> Human language in actual use is infinitely more problematical than those aspects of it that are amenable to treatment by information theory [...] language involves the histories of those using it, hence the history of society, indeed, of all humanity generally. And language in human use is not merely functional in the way that computer languages are functional. It does not identify things and words only with immediate goals to be achieved or with objects to be transformed. Human use of language manifests human memory. And that is a quite different thing than the store of the computer, which has been anthropomorphized into 'memory'. The former gives rise to hopes and fears, for example. It is hard to see what it could mean to say the computer hopes.[106]

Notably, he comments:

> Even the kinds of knowledge that appear superficially to be communicable from one human being to another in language alone are in fact not altogether so communicable. Claude Shannon showed that, even in abstract information theory, the 'information content' of a message is not

105 J. Weizenbaum, *Computer Power and Human Reason: From Judgment to Calculation* (Harmondsworth: Penguin Books, 1993).

106 Ibid., p. 209.

> a function of the message alone but depends crucially on the state of knowledge, on the expectations, of the receiver.[107]

The quest to constrain expressiveness by enforced normalization of language has a history. In his book *Medical Nemesis*, Ivan Illich (1926–2002) records that, in 1635, Cardinal Richelieu (1585–1642) set up an Academia of distinguished scholars of French literature, for the purpose of protecting and perfecting the French Language.[108] They imposed and mandated an elite language of the bourgeoisie and made it normative for all social classes. The term 'normal' has history as well. Illich also wrote that in England of the 1830s, normal was a geometrical term for perpendicular, or standing at a right angle. In the 1840s it was generalized to mean conforming to a common type. And around 1840, Auguste Comte (1798–1857), in France, talked of the laws relative to the normal state of an organism as a basis for study of comparative pathology. In that era, pathology was largely used to classify anatomical anomalies, and, towards the end of the century, Claude Bernard (1813–78) started to label and catalogue functions and homeostasis of the body. Clinical normality gradually became associated with wellbeing.

Specializations of medicine have led to ever more extensive vocabulary and complexity of language used to name, group, describe and record thoughts, ideas and reasoning. In the Information Age, these have extended beyond dictionaries to databases. Terms and their interrelationships have been grouped and expressed as hierarchies, enabling computer programs to meaningfully and rigorously handle a vast corpus of terms that would otherwise be unmanageable for humans. Here have arisen clinical terms as codes used to label and integrate medical knowledge with the procedures and reasoning of clinical practice. It is hard for humans to master consistent use of the hundreds of thousands of terms of the language that have arisen in this way. Their selection and use are increasingly enacted using machine software.

Nomenclature and Terminology

There is a story told of a blinded set of observers presented with a huge elephant and each asked to feel its body and describe it in words. One encounters the tail, another the tusks, another the trunk and another the legs. The contrasting descriptions they give are all true, but incomplete and inadequate to describe the elephant. How they connect would be missing

107 Ibid.

108 I. Illich, *Limits to Medicine: Medical Nemesis: The Expropriation of Health* (London: Boyars, 1995), pp. 115−16.

from these descriptions, even if they are simply concatenated together. Neither medicine, nor health care, more generally, nor the individual citizen cared for, are elephants, but they are huge in the sense of the domains of knowledge and data that they encompass. The descriptive range covered in these domains is likewise huge, and prone to intractable problems when seeking to draw together disparate descriptions, connections and communications, to in some way faithfully represent the wholeness of health care, as it affects and is experienced by individual citizens. This has been an elephant in the room, where blinded observers have often flailed and failed to connect reality with the computer, from the outset of the Information Age. That story is told in Part Two of the book.

The decades-long struggle for the standardization of electronic care records serves as a persistent testament to wider endeavours that have involved and implicated the computer in the ways we understand and express our knowledge of what medicine and health care services are, and what they do. Care services require and depend on good records that are faithful to this understanding and kept up to date with relevant new knowledge. Record plays a central role in the connection, communication and application of the wholeness of this knowledge–logically, consistently and ethically. Central to faithful records are the terms and systems of terms (nomenclatures) used, the structure of relevant knowledge they reference and contain, and the logic of the reasoning and decisions they reflect.

It has proven extremely hard, if not as impossible as once it might have seemed, to compartmentalize the personal, professional and organizational content and context of records kept about each citizen and patient. As in the domain of library classifications discussed earlier in the chapter, this has developed into a battleground of rivalries and special interests–national, commercial, institutional and professional–with the patient, who is the focus of the record, spoken of and for, but too often left in the margins of the circle of those speaking. In the Information Society of the future, citizens will reclaim both their personal data and the common ground on which care record systems are centred and shared. That is the vision and perspective informing Part Three of the book, which looks forward to how it can now be created. In the story of how electronic care record standardization has evolved, there are illuminating parallels with the experience of library science in its efforts to move beyond enumerative hierarchies of terms and codes towards a compositional approach, using generic building blocks applicable at different levels of hierarchy and detail of content.

Lancelot Hogben (1895–1975) studied medicine at Trinity College, Cambridge and worked then as an experimental zoologist.[109] His wider interests extended to medical statistics, etymology of language and experimental biology. He recognized the struggle students had in learning the terminology of biology, especially those unfamiliar with etymology and classical language, and worked on an *interglossa*–a language for international communications–believing that eight hundred words constructed with very simple grammar, rooted in Latin and Greek, would suffice as a basic vocabulary across languages! He believed that an international committee could easily take it on and make it happen! It is hard not to smile wryly at this expression of belief, looking back over the years, but he was a great and lauded thinker, and such were the times.

Recent decades have seen rapid escalation in the scale and complexity of nomenclature and classification used in records descriptive of health care practice. The history of medical terminology, the clinical tasks it addresses and the challenge of adopting formal method, were covered in a 1980 review by Roger Côté and Stanley Robboy[110] and later reviews by James Cimino[111] and Cimino and Xinxin Zhu.[112]

There is continuing struggle to tame what might be called the linguistic noise and accumulating entropic disorder present in records, through their incorporation of imprecise and changing terminology and categories and structures of content, over time. If records are to be computable and sustainable throughout patients' lifetimes, consistent use of terms is important. GIGO (Garbage In, Garbage Out) may overstate the impact but NINO (Noise In, Noise Out!) is physical reality. Pretending otherwise is a no-no! One can artfully filter signal from noise, but noise remains noise, at whatever level of resolution–'random' signifies 'don't know'. Leaving aside

109 Lancelot Hogben's book, *Mathematics for the Million* (1936), was described by the historian and imaginer of future worlds, H. G. Wells (1866–1946), as of first-class importance and was praised by Einstein, Russell and Julian Huxley (1877–1975). He was socialist and atheist—in the First World War, he served in the Red Cross and Friends Ambulance Unit and in the second as curator of the army's medical statistics. His work in experimental zoology was recognized by the Royal Society in his fellowship there in 1936.

110 R. A. Côté and S. Robboy, 'Progress in Medical Information Management. Systematized Nomenclature of Medicine (SNOMED),' *JAMA: The Journal of the American Medical Association*, 243.8 (1980), 756–62, https://doi.org/10.1001/jama.243.8.756

111 J. J. Cimino, 'Desiderata for Controlled Medical Vocabularies in the Twenty-First Century,' *Methods of Information in Medicine*, 37.4–5 (1998), 394–403, https://doi.org/10.1055/s-0038-1634558

112 J. J. Cimino and X. Zhu, 'The Practical Impact of Ontologies on Biomedical Informatics', *Yearbook of Medical Informatics*, 15.01 (2006), 124–35, https://doi.org/10.1055/s-0038-1638470

the many frustrations and differences of opinion that have prevailed in this area, they do reflect real, and very difficult, problems experienced at the coalface of health care services.

Multiple perspectives and interests are in play when seeking greater clarity, coherence and usefulness of medical language. The benefits sought, and the burdens imposed in attempting to realize them, are experienced differently, in the contexts where professional practice is performed, and in those where it is managed and regulated. At the coalface of practice, the benefits are often fewer and the burdens often greater than they are perceived and experienced to be when sitting away from that front line, in places where policy is set and designs and plans are made and mandated for implementation.

Safe use of language for coding and classification of care records requires consistent methods and the availability of resources–people, time and money. It can assist rapid browsing and access to the relevant content of the record, during a consultation. It can enable subsequent extraction of data for secondary analysis elsewhere, in support of audit and resource management of clinical services, research and budgetary control. In UK hospitals, dedicated local teams navigate the immense collections of terms that are selected from when making, coding and classifying clinical records. They make decisions about how to use these terms to tag individual records and provide aggregate statistics descriptive of the care processes and outcomes of the services that the institution provides. This can bring further work for the clinicians whose records are being analyzed, drawing them into the checking of judgements made, where ambiguities have arisen. All this embeds closely with cost and remuneration of services, and the pursuit, and sometimes gaming, thereof.

A good way to meet these important but different requirements, in a less burdensome manner, has long been, and remains, urgently needed. Achieving this requires that the data can be drawn together across different component services, simply and accurately, within a coherent information framework. There has been considerable and costly redundancy of efforts in this regard. For example, as mentioned to me by a former UK Chief Medical Officer some years ago, there were then some thirty or more different systems in everyday use across the NHS, for generating the nationally mandated critical incident reports–collecting, structuring and mapping local disparate datasets to a centrally defined report template, as further discussed in Chapter Seven. And I have read that during the Covid pandemic, data from a hundred differently structured spreadsheets were regularly cut and pasted into an aggregated report for senior NHS management purposes.

If we decide to keep trying to improve record keeping–and as Chapter Seven will unfold in detail, over the past fifty years this has been recognized as a *sine qua non* of progress in health informatics–coherent, robust and connected specification and implementation of the related computational methods programmed and deployed, and clarity about their bounded domains of application, are imperative requirements. E. M. Forster's 'only connect', says it all, here, too. Otherwise, disconnection will propagate and the entropy of information and information system legacy will grow, because of the lack of such discipline.

Historically, considerable momentum was injected into coding and classification of medical language, long before rigorous methods of description logic had or could have crystallized. Today's health care information systems remain strongly influenced by the early methods that were used to organize them, substantially through 'pre-coordination' of pragmatically-structured lists of terms. This has inevitably led to a considerable legacy of codes and methods of coding that must now be accommodated within new methods experimented with. The complexity is compounded by the continuous emergence of requirements for new terms and the need to extend or redesign the structures of the systems into which they fit. 'Post-coordination' of terms (creating new terms from combinations of existing ones), though tractable, has been found a difficult and onerous method to implement and sustain rigorously in practice. The result has been shaky edifices that have struggled, after initial enthusiasm and support, to reinforce or underpin their foundations with new more formal methods, and thereby reinvent themselves within changing contexts. However, the legacy of systems in current use was hard to achieve and expensively won, and it is hard to justify and seldom if ever seems timely to change them.

The limitation and Achilles' heel of pragmatic decisions about medical and health care language (focused on the here and now of the domains represented, and adjusted piecemeal to changing requirements over time) is the incremental complexity and blurring of the imposed order of the system of terminology adopted. Entropy accumulates within the system and impedes further evolution and change, requiring continuous work for maintenance of the current system, to keep it in order and functioning. That is physics and that is life!

The promise of a stable and expressive description logic is that its adoption would provide rigorous discipline and tooling in support of this curation task. It would provide a framework whereby the coherence and consistency of the nomenclature could be validated and analyzed, using generic software tools based on the adopted description logic formalism. Once established and well-grounded in custom and practice, this would facilitate smoother evolution alongside changing requirements. Of course, a model of the term set expressed in formal logic is itself a complex design. It requires anchoring expertise and engagement, spanning mathematics, computer science and engineering as well as medicine and health care, and suitable environments in which to draw them together.

In recent years, the SNOMED, ICD and UMLS (Uniform Medical Language System) systems–leading international initiatives that I introduce below–have experimented with description logic for restructuring their extensive and 'multiaxial hierarchies' of terms. The tools available to support this effort have continued to evolve, and, as ever, pragmatic choices have been needed to balance evolving theory of the underpinning computational methods and established practice in their application. The essence of these endeavours has been to represent the system of terms as an ensemble of logical descriptions of their form and content, expressed as a 'subsumption hierarchy'–one that subsumes and generalizes terms within hierarchies of detail, reflecting the knowledge they represent and how it fits together.

The brave and difficult SNOMED International initiative has received multi-government backing, initially orchestrated by my colleague, Martin Severs, as I describe below. Over recent years, merged with the UK NHS Read Codes project, they have taken the plunge and adopted description logic methods. Yet more recently, the ICD, the most venerable, and arguably most successful, of medical terminology initiatives, has also put its toes in the water to weigh up whether to jump in and swim towards the adoption of description logic methods. Another great colleague, Alan Rector, worked with a SNOMED- and ICD-knowledgeable team to confront the methodological issues that would need solutions to update ICD methods in this way, and facilitate alignment with the UMLS.[113] They reviewed the functionalities required and tools available to realize them, and described a combination of state-of-the-art description logic, frame-based logic and ad hoc code, that appeared suitable for the experiment.

In earlier years, Rector had pioneered a green field description logic approach to medical language, encyclopaedia and nomenclature, the 'len' of the 1991 EU GALEN Project, which I also introduce in a separate section below. As with the BSO for classification of books and documents, it failed, but Rector and his team's work has, nonetheless, been widely influential in the field. Failure to recognize and support this pioneering experimental initiative and help it achieve traction, was a highly consequential failure of policy of recent decades. Not that it was necessarily going to prove a successful approach in the short term, but the creation of a well-anchored and competent academic and professional environment and community of

113 UMLS is an initiative of the US National Library of Medicine, seeking to draw together the current medical language standards used in publication and record keeping, within a common framework. It is introduced briefly, below, alongside several other initiatives of note.

endeavour in description logic and medicine would have been a good goal to support, enable and sustain.

It is a struggle to keep one's feet on the ground in these emerging adventures of ideas, but that is what is needed to improve quality and achieve traction in the development of methodology and implementation of electronic care records. Staying power is also needed and this requires sponsorship. Lacking this cohesion, developments typically proceed piecemeal, and pragmatic choices made at each stage easily lead to sub-optimal combinations of the old and the new. These choices are fiercely debated by rival interests and critics of all the parties in play, whose teeth tend not to be biting the bullet of hard work involved in tackling the task at hand, on the ground.

Achieving expressiveness and rigour of computation and applying it efficiently and effectively in diverse health care contexts (while accommodating legacy systems and coming to terms with the opportunities and constraints of newly evolving, but often rapidly obsolescent, computational methods) is a multiplicative set of challenges. These play out alongside the rapid change, intrinsic uncertainty and difference of perspectives, implicit in individual patient care. It has been said that health care is the most fertile domain for proposing problems likely to benefit from the application of information technology, while, at the same time, the most difficult in which to realize the hoped-for benefits.

Experience of three attempts to underpin the concrete foundations of our house here in St Albans, have shown that such efforts can often be only temporary fixes! Problems continue to recur. Building a new house, possibly in a different place, might often prove a better long-term bet, but it costs money, and location matters! In a universal domain like an academic or professional discipline, we often wait for the structure to collapse before being forced to build anew. Thomas Kuhn (1922–96) and Gould have had something to say along these lines, about how paradigms of knowledge change through resistance to change and punctuation of equilibria.[114] Pressure towards a new home sometimes pushes the old one towards tipping over before it is ready to fall. Leaning structures, such as the famous Tower of Pisa, somehow manage to persist against seemingly poor odds, with the burghers of Pisa no doubt very keen and active, behind the scenes and viewing their tourism statistics and cash registers, that it should not fall!

On the following pages, I briefly highlight global initiatives working to standardize the language and nomenclature of medicine and health care,

114 T. S. Kuhn, *The Structure of Scientific Revolutions: 50th Anniversary Edition* (Chicago, IL: University of Chicago Press, 2012); Gould and Eldredge, 'Punctuated Equilibrium Comes of Age'.

drawing heavily on the up-to-date websites describing their work. I focus first on the two that are connected with my close colleagues of the era, Severs and Rector. They have played stellar roles in clinical and scientific leadership along the timelines of the SNOMED and GALEN initiatives.

SNOMED™–Systematized Nomenclature of Medicine[115]

SNOMED started as a library-like enumeration of terms used in records of care, organized and grouped under axes (facets) of topography (anatomic site), morphology (form), aetiology (origin) and function. It has grown into a present-day incarnation of hundreds of thousands of terms, each with unique code and linked in twenty hierarchies, organized as subsumptive containment structures.

SNOMED CT (Clinical Terms) describes itself as a multi-lingual, multinational logic-based health care terminology. It traces its origins to the Systematized Nomenclature of Pathology (SNOP), published in 1965 by the College of American Pathologists (CAP).

Here is the paraphrased story of its subsequent evolution, as told on the International Health Terminology Standards Development Organization (IHTSDO) website, now the SNOMED website. The IHTSDO now owns, develops and maintains it, trading under the name SNOMED International.

> SNOMED I and II were released in 1974 and 1979.
>
> SNOMED-RT (Reference Terminology) was released in 2001.
>
> 1999: Agreement was reached between the NHS and the College of American Pathologists to bring together the NHS Clinical Terms Version 3 (formerly known as the Read codes) and SNOMED-RT, under the umbrella of a new terminology, SNOMED CT. The final product was released in January 2002.
>
> 2003: The National Library of Medicine (NLM), on behalf of the United States Department of Health and Human Services, entered into an agreement with the College of American Pathologists to make SNOMED CT available to U.S. users at no cost, through the National Library of Medicine's Unified Medical Language System (UMLS) Metathesaurus.
>
> 2007: The International Health Terminology Standards Development Organization (IHTSDO) was established as an international SDO. SNOMED CT intellectual property rights were transferred from the

115 SNOMED, https://www.snomed.org

> CAP to the IHTSDO, in order to promote international adoption and use of SNOMED CT. IHTSDO subsequently adopted the trading name, SNOMED International.

Rather than paraphrase, the following is quoted directly from the 2020 website, to give the flavour of the rapidly evolving initiative which now has received powerful and wide-ranging international backing.

> SNOMED RT, with over 120,000 concepts, had wide coverage of medical specialties and was designed to serve as a common reference terminology for the aggregation and retrieval of pathology health care data recorded by multiple organizations and individuals. The strength of CTV3 was its terminologies for general practice. With 200,000 interrelated concepts, it was used for storing structured information about primary care encounters in individual, patient-based records. The January 2020 release of the SNOMED CT International Edition now includes more than 350,000 concepts.
>
> SNOMED CT's primary purpose is to support all health care professionals in their recording and sharing of detailed patient information within Electronic Health Records (EHRs) and across health care communities globally. Its ontological foundations allow SNOMED CT data to support detailed data analytics to meet a variety of use cases from local requirements to population-based analytics. With the help of SNOMED CT, programs can translate different medical terms into an internationally standardized numerical code. In this way, clinical data from different countries can be compared and used for research. This creates the prerequisites for treating diseases more effectively in the future, recognizing them faster and supporting prevention. The networking of routine care data and top medical research has great potential–for better medical treatment and for strengthening business and science.
>
> With a complement of 39 Members, SNOMED CT now represents approximately one third of the global population. Adding to that complement with our affiliate licensees, SNOMED CT is now used in more than ninety countries globally.
>
> SNOMED CT is currently available in American English, British English, Spanish, Danish and Swedish, with other translations underway or nearly completed in French and Dutch.
>
> SNOMED CT cross maps to other terminologies, such as: ICD-9-CM, ICD-10, ICD-O-3, ICD-10-AM, Laboratory LOINC and OPCS-4. It supports ANSI, DICOM, HL7, and ISO standards.

The ontology foundations of SNOMED have evolved considerably over recent years. SNOP, SNOMED-RT, Read and CTV3, as well as LOINC, ICD-11 and UMLS do not employ formal logic at their core, with the inbuilt proof and checking of consistency that this provides. Multiaxial hierarchy and class overlap were introduced within the constraints of the EL++ subset of description logic–which lies between propositional and first-order logic, allowing overlapping of classes, in contrast to ICD-11. It provides for post-coordination of terms, drawing together pre-coordinated terms. Laterality and negation are not currently allowed for.

GALEN–Generalized Architecture for Languages, Encyclopaedias and Nomenclatures in Medicine

The GALEN project arose from the computer science community at the University of University, through the work of Rector's luminary team. Although transformational in potential, it suffered from a lack of parallel, everyday health care service support and grounding.

If there was ever a name to fit the acronym, this must be it! GALEN is remembered for his role in the invention of medicine and contributions to the evolution of Aristotle's logic. The parallel is fitting–Rector, father of this initiative, was a notable pioneer. In the language of Ranganathan, GALEN is an 'analytico-synthetic' method for composing codes. In this respect, its approach, as specifically characterized, below, in bullet point five, bears resemblance to the 'switching language' concept of the BSO.

Again, quoting extensively from the now no-longer hosted GALEN website (as with others mentioned in the following section, like ICPC, this is no longer maintained):

> GALEN is the name given to a technology that is designed to represent clinical information in a new way and is intended to 'put the clinical into the clinical workstation'. GALEN produces a computer-based multilingual coding system for medicine, using a qualitatively different approach from those used in the past. GALEN is attempting to meet five challenges:
>
> - To reconcile diversity of needs for terminology with the requirement to share information
> - To avoid exponentially rising costs for harmonization of variants
> - To facilitate clinical applications
> - To bridge the gap between the detail required for patient care and the abstractions required for statistical, management, and research purposes

- To provide multilingual systems which preserve the underlying meaning and representation

To do so, GALEN advocates five fundamental paradigm shifts to resolve the fundamental dilemmas that face traditional terminology, coding and classification systems:

1. ***In the user interface***, to shift from selecting codes to describing conditions. Interfaces using GALEN technology allow a central concept to be described through simple forms. If required, a precise code for reporting can be generated later automatically.
2. ***In the structure***, to shift from enumerated codes to composite descriptions. Correspondingly, GALEN handles terminology internally analogously to a dictionary and a grammar so that indefinitely many descriptions can be composed from a manageable number of base concepts. Traditional coding systems are more like a phrase book; each sentence must be listed separately. No one would think of trying to list all the possible sentences in any natural language in a phrase book; listing all possible disease or procedure terms in a coding system is equally fruitless.
3. ***In establishing standards***, to shift from a standard coding system to a standard reference model. Existing coding and classifications differ because they are used for different purposes. Finding a single fixed set of codes for all diseases, procedures, etc. which will serve all purposes is a chimera. The GALEN Common Reference Model provides a common means of representing coding and classification systems so that they can be inter-related–a common dictionary and grammar. The project's slogan is 'coherence without uniformity'.
4. ***In delivery***, to shift from static coding systems as data to dynamic terminology services as software. Terminology is now at the clinical software. GALEN originated the idea of a terminology server and is participating actively in the CorbaMed effort at standardizing the software interface.
5. ***In presentation***, to shift from translations of monolingual terminologies to multilingual terminologies. GALEN separates the underlying concepts from the surface natural language that presents them.

The founders established OpenGALEN as a not-for-profit organization to enable the widest possible exploitation of some of the results of the GALEN Programme. GALEN is built around what is called the GALEN Core:

> The GALEN CORE Model for representation of the Common Reference Model for Procedures contains the building blocks for defining procedures–the anatomy, surgical deeds, diseases, and their modifiers used in the definitions of surgical procedures. This document describes the structure of the CORE model and gives a detailed account of its high-level schemata followed by a detailed example of the use of the ontology for a portion of the model of the cardiovascular system and diseases.
>
> The ontology for the GALEN CORE model is designed to be re-usable and application independent. It is intended to serve not only for the classification of surgical procedures but also for a wide variety of other applications–electronic health care records (EHCRs), clinical user interfaces, decision support systems, knowledge access systems, and natural language processing. The ontology is constructed according to carefully selected principles so that the reasons for classification are always explicit within the model and therefore available for processing and analysis by each application. This leads to an ontology in which most information lies in the descriptions and definitions. The hierarchies are built bottom-up automatically based on these definitions.
>
> Note that the word ontology has acquired a range of meanings in various communities. Following the usage of Guarino [Guarino and Giaretta 1995], it is used here with a lowercase ***o*** or in the plural to indicate the set of primitive, high-level categories in a knowledge representation scheme together with any taxonomy which structures those categories.
>
> Quality assurance of the model is an ongoing process. The most important quality assurance of the building blocks comes from the checks on the correct classification things built with them–the model of procedures and the other models for subspecialties being built in collaboration with other projects. Preliminary results from such checks are extremely promising.
>
> The structure of the model is now believed to be complete, but there remain many details of anatomy and diseases to complete for each subspecialty area. Future development of the model is governed by the requirements of the applications and the needs of the centres who are using it to develop classifications of procedures. The next areas to be addressed will be based on the needs of vascular, ENT, orthopaedic, and gynaecological surgery to meet the requirements and priorities of those centres.

Some authors succeed in summarizing a field so well, at a point in time, that it would feel almost an insult to paraphrase them. The seminal paper describing the challenges of aligning ICD-11 within a framework of formal logic, where Rector played a leading role, has been previously cited. The conclusions provide a classic account of the state of the art in 2019. Again, it would be a disservice to paraphrase, albeit that the quotations here abbreviate the authors' deep understanding of the domain, and further insight requires engagement with the paper itself and its references.

> In general, reasoning about classes ('T-Box reasoning') can be optimized computationally (although in expressive dialects it is worst-case intractable). However, reasoning about individuals ('A-box reasoning') is much more difficult. Conveniently, it is primarily axioms about classes that are relevant for terminologies, *e.g.,* axioms about the classes of lungs, pneumonias, bacteria, penicillin preparations, etc.
>
> [...] The reasoning path, or 'justification' leading to inferences in OWL/DLs can be surprisingly difficult to work out by manual inspection. Unlike the closed world reasoning in logic programming methods such as Prolog or MYCIN, there is no way to accumulate a simple explanatory trail in the course of the proof. This has led to a body of research on methods for generating 'justifications' computationally.
>
> [...] Almost all statements about signs and symptoms in medicine–and many characteristics in genetics, genomics, and biomedicine more widely–are subject to exceptions, *i.e.,* they are 'defeasible' or, in our vocabulary, they are 'generalizations'.
>
> [...] To summarize: in closed-world representations, 'not' means 'not found or derived based on the individuals explicitly in the knowledge base'; 'all' means 'all found or derived'.
>
> [...] The Information Retrieval and Librarianship communities have a long history of systematic methods for organizing and classifying terminologies, thesauri and other knowledge artefacts, much of it shared with the linguistic community. Standards based on this work include the work of ISO TC-37, especially the ISO standards 704 and 1087. The most important example in biomedicine is the UMLS Semantic Network.
>
> [...] Description logics were a major advance for expressing invariant/indefeasible statements but are fundamentally unable to express generalizations/defeasible statements. The conflict between expressiveness, logical completeness and computational tractability has been a recurring theme in the knowledge representation research, see for example Doyle and Patil. Regrettably, frames, description logics, and ICD-like structures were often seen as rivals rather than complementary. The ICD-11 project required articulating their complementarities, an instructive example of practical development helping to elucidate theory.
>
> [...] Conclusion

> In summary, OWL has been a major step forward for representing terminologies and the invariant part of background knowledge bases. However, the W3C's standardization on OWL has led to pressure to use OWL, often beyond its limitations, and contributed to confusion over vocabulary and to neglect and misunderstanding of other representations.
>
> The experience in the ICD-11 project is that, although it is not possible to express frames in OWL or OWL in frames, it is possible to combine them in hybrid systems that take advantage of the semantics of each without violating the semantics of either. Likewise, while it is not possible to represent systems such as ICD based on JEPD (Jointly Exhaustive Pairwise Disjoint) mono-hierarchies in either Frames or OWL, it is possible to link them through queries. Associations for navigation and language are needed by most applications but should be distinguished from statements with other semantics. SKOS (note: Simple Knowledge Organization System, is a W3C standard, based on other Semantic Web standards (RDF and OWL)), often provides a useful set of relations for this purpose.
>
> If symbolic knowledge representation is to continue to play a role in biomedical information systems, our experience is that such architectures need to be further developed and standardized, preferably within an integrated environment.[116]

In this last paragraph, the paper connects back to a general theme exemplified in the storyline of this book. One might say, *plus ça change, plus c'est la même chose*! GALEN was a seminal initiative towards modernization of health terminology systems in the Information Age, comparable to the BSO of former decades. And like BSO, it drifted onto the rocks when it failed to connect and develop at scale, within the practical and everyday context of coding and classification of medical records, which it was designed to support.

Further Notable Medical Language and Terminology Initiatives

The history of medical terminologies is a long one, and impossible to encompass in detail within the broad range of this book. As with the forgoing truncation of discussion of formal logic, the following section aims only to give the flavour of key initiatives. It relies on, and quotes from, their current websites, where detailed and up-to-date descriptions can be found.

116 Rector et al., 'On Beyond Gruber', pp. 4–13.

UMLS–Uniform Medical Language System[117]

This initiative of the National Library of Medicine has built on its substantial role in collating, first the paper-based Index Medicus of publications and then the online version known as Medline/Medlars. It 'integrates and distributes key terminology, classification and coding standards, and associated resources to promote creation of more effective and interoperable biomedical information systems and services, including electronic health records.

The component parts of UMLS are:

1. The Specialist Lexicon

The SPECIALIST Lexicon is an English lexicon (dictionary) that includes biomedical terms as well as commonly occurring English words. The lexical entry for each word or term records the following information:

- Syntactic (syntax information)
- Morphological (inflection, derivation, and composition information)
- Orthographic (spelling information)

Currently the SPECIALIST Lexicon contains over 200,000 terms and is used by the lexical tools to aid in Natural Language Processing. Words are selected for entry into the Specialist lexicon from a variety of sources:

- The UMLS Test Collection of MEDLINE abstracts
- *Dorland's Illustrated Medical Dictionary*
- *The American Heritage Word Frequency Book*
- *Longman's Dictionary of Contemporary English*
- Current MEDLINE citation records[118]

2. The Metathesaurus

Over 100 vocabularies, code sets, and thesauri, or 'source vocabularies' are brought together to create the Metathesaurus. Terms from each source vocabulary are organized by meaning and assigned a concept unique identifier (CUI).

117 'Unified Medical Language System', *National Library of Medicine*, https://www.nlm.nih.gov/research/umls/index.html

118 'The SPECIALIST Lexicon', *National Library of Medicine*, https://www.nlm.nih.gov/research/umls/new_users/online_learning/LEX_001.html

Sixty-two percent of the Metathesaurus source vocabularies are in English. However, the Metathesaurus also contains terms from seventeen other languages such as Spanish, French, Dutch, Italian, Japanese, and Portuguese.[119]

3. The Semantic Network

The Semantic Network consists of semantic types and semantic relationships. Semantic types are broad subject categories, like Disease or Syndrome or Clinical Drug. Semantic relationships are useful relationships that exist between semantic types. For example: Clinical Drug treats Disease or Syndrome. The Semantic Network is used in applications to help interpret meaning.[120]

[…] The Semantic Network consists of:

- Semantic types (high level categories)
- Semantic relationships (relationships between semantic types)

The Semantic Network can be used to categorize any medical vocabulary.

There are 133 semantic types in the Semantic Network. Every Metathesaurus concept is assigned at least one semantic type; very few terms are assigned as many as five semantic types. Semantic types are listed in the Metathesaurus file MRSTY.RRF.

Semantic types and semantic relationships create a network that represents the biomedical domain.[121]

119 'The Metathesaurus', *National Library of Medicine*, https://www.nlm.nih.gov/research/umls/new_users/online_learning/Meta_001.html

120 'The Semantic Network', *National Library of Medicine*, https://www.nlm.nih.gov/research/umls/new_users/online_learning/OVR_003.html

121 'The Semantic Network', *National Library of Medicine*, https://www.nlm.nih.gov/research/umls/new_users/online_learning/SEM_001.html

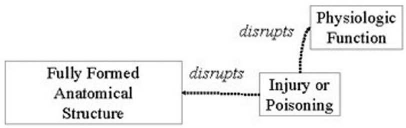

Semantic Type 1	Semantic Relationship	Semantic Type 2
Injury or Poisoning	disrupts	Fully Formed Anatomical Structure
Injury or Poisoning	disrupts	Physiologic Function

Fig. 2.3 An illustration of semantic relations. Image created by David Ingram (2022), based on details from the USA National Library of Medicine website, https://www.nlm.nih.gov/research/umls/new_users/online_learning/SEM_001.html

Semantic types and relationships help with interpreting the meaning that has been assigned to the Metathesaurus concept. This graphic adapted from the website (Figure 2.3) illustrates two semantic relationships. There are currently some fifty-four semantic relations defined. Tools are provided to assist parsing of natural language about any biomedical topic, seeking to recognize words, in whatever form they are expressed or spelt, by reference to the two hundred thousand-strong semantic lexicon. These are then mapped to what is sometimes called a 'controlled vocabulary', in this case that of the vocabulary or concepts of the Metathesaurus. In this usage, the controlled vocabulary being built on is the set of source vocabularies, each of them a thesaurus, embodied in the Metathesaurus. The Metathesaurus is a thesaurus of thesauri! Following on in the website description, 'The UMLS Metathesaurus organizes all of the original data from the source vocabulary including unique identifiers, definitions, or term spelling variants into a common format'.[122]

122 'The Metathesaurus', *National Library of Medicine*, https://www.nlm.nih.gov/research/umls/new_users/online_learning/Meta_001.html

The next step is to take the resulting words–that re-express the originals in the terms of the controlled vocabulary–and map their meaning, in the sense of semantic 'type' and 'relationship', as specified in the semantic network, and the much smaller, but still very wide-ranging model of types and semantic relationships that it provides.

So far and so good. When creating or invoking any model, an important first step is to know why we are doing it and what we plan to use it for. We then come better equipped with a sense of how it will be employed and how it will enable things to be done better. In these respects, the librarian's perspective is, quite rightly, predominant, given the organization (the US National Library of Medicine) that is leading these efforts. The purpose of UMLS parallels the librarians' challenge described earlier in the chapter. It is one of organization of materials, but here the challenge is not where to place the document or book as much as: what is being discussed here and how can I group and generalize all potentially relevant resources within a simplified, useful and computable structure? The driver is one of automation to support management of the exploding contexts of both knowledge semantics and scale of curation required. There is still human control applied in defining the methods adopted. We should be aware that an initiative like this appears to be progressing towards a time when such a purpose will be wholly within scope of machine control, when it will no longer be just human knowledge that is being organized.

The UMLS is reaching its capable arms around the very substantial and still growing challenge of managing medical literature. The integrative methods that it embodies, embracing the many vocabularies, semantic types and relations of medicine and health care, has scaled satisfactorily thus far. It has not dropped the ball, but at some future point, the balance of human and machine effort required may well fail again, and a heavier ball will drop. The benefits to student, researcher and librarian, alike, in managing the new world of literature, are phenomenal, as compared with the time and effort required for the hit-and-miss work involved in reading through and taking notes from Index Medicus of yesteryear, and sending off letters requesting to borrow copies, or receive reprints, of published books and papers. But these very limitations shaped as much as constrained academic discourse of the times. Set free from such limitations, with machine software making decisions required to carry the load, the new methods reflect into a qualitatively different human discourse. Scale and metrication of publication has exploded, bringing new burden and constraint. When we step from the world of education and research into the world of providing and recording services for direct health care, expectations of well-curated documents and records impose new burdens on practitioners, where they require new expertise to interface with an ever-increasing and evolving standardized use of language and codification of knowledge. If the machine takes over that burden in ways that do not chime with human meanings, the real and virtual worlds will further divide.

MeSH–Medical Subject Headings[123]

The MeSH initiative started from the other end of the scale, with the purpose of creating standardized descriptors of a document, to assist in the positioning and retrieval of relevant information from a collection of documents, by giving them a consistent and coherent system of headings. Again, the following rests heavily on up-to-date reference to its website.

> The Medical Subject Headings (MeSH®) thesaurus is a controlled vocabulary produced by the National Library of Medicine and used for indexing, cataloguing, and searching for biomedical and health-related information and documents.
>
> MeSH includes the subject descriptors appearing in MEDLINE®/PubMed® and other NLM databases. MeSH provides a consistent way to find content with different terminology but the same concepts. MeSH organizes its descriptors in a hierarchical structure so that broad searches will find articles indexed more narrowly. This structure also provides an effective way for searchers to browse MeSH in order to find appropriate descriptors.
>
> The MeSH vocabulary is continually updated by subject specialists in various areas. Each year hundreds of new concepts are added, and thousands of modifications are made.
>
> Many synonyms, near-synonyms, and closely related concepts are included as entry terms to help users find the most relevant MeSH descriptor for the concept they are seeking. In NLM's online databases, many terms entered by searchers are automatically mapped to MeSH descriptors to facilitate retrieval of relevant information.[124]

The MeSH website records the history of its evolution since 1954. Many of the issues encountered in the development of the BSO initiative resurfaced there, such as alterative choices of headings and subheadings to group under, depending on topic, and overlap among chosen groupings. It was

123 'The Metathesaurus', *National Library of Medicine*, https://www.nlm.nih.gov/research/umls/new_users/online_learning/Meta_001.html

124 'Medical Subject Headings: Preface', *National Library of Medicine*, https://www.nlm.nih.gov/mesh/intro_preface.html

a cultural change within NLM to bring together what had previously been separately designed and managed methods for compiling book catalogues and periodical article indexes.

> The main heading-topical subheading combination is a pre-coordination of terms, reducing the problem of term permutation, which looms large in most manual retrieval systems in book form.
>
> From its beginning, MeSH was intended to be a dynamic list, with procedures for recommending and examining the need for new headings. The content of the vocabulary related to the usage of terms in the literature itself and evolved to meet new concepts in the field. The use of the computer made revisions more practical and systematic, despite the difficulty in updating printed indexes and card catalogues.
>
> Categorized lists of terms were printed for the first time in the 1963 *Medical Subject Headings* and contained thirteen main categories and a total of fifty-eight separate groups in subcategories and main categories. These categorized lists made it possible for the user to find many more related terms than were in the former cross-reference structure. In 1963, the second edition of *Medical Subject Headings* contained 5,700 descriptors, compared with 4,400 in the 1960 edition. Of the headings used in the 1960 list, 113 were withdrawn in favor of newer terms. In contrast, the 2015 edition of MeSH contains 27,455 descriptors and in 2021 there are 29,917 Descriptors and 270,373 Supplementary Concept Records.
>
> In 1960, medical librarianship was on the cusp of a revolution. The first issue of the new *Index Medicus* series was published. On the horizon was a computerization project undertaken by the National Library of Medicine (NLM) to store and retrieve information. The Medical Literature Analysis and Retrieval System (MEDLARS®) would speed the publication process for bibliographies such as *Index Medicus*, facilitate the expansion of coverage of the literature, and permit searches for individuals upon demand. The new list of subject headings introduced in 1960 was the underpinning of the analysis and retrieval operation. MeSH was a new and thoroughly revised version of lists of subject headings compiled by NLM for its bibliographies and cataloging.[125]

125 Ibid.

ICD–International Classification of Diseases[126]

ICD describes itself thus on its website:[127]

> [as] the foundation for the identification of health trends and statistics globally, and the international standard for reporting diseases and health conditions. It is the diagnostic classification standard for all clinical and research purposes. ICD defines the universe of diseases, disorders, injuries, and other related health conditions, listed in a comprehensive, hierarchical fashion that allows for:
>
> - easy storage, retrieval and analysis of health information for evidenced-based decision-making;
> - sharing and comparing health information between hospitals, regions, settings and countries; and
> - data comparisons in the same location across different time periods.
> - Uses include monitoring of the incidence and prevalence of diseases, observing reimbursements and resource allocation trends, and keeping track of safety and quality guidelines. They also include the counting of deaths as well as diseases, injuries, symptoms, reasons for encounter, factors that influence health status, and external causes of disease.

And:

> The first international classification edition, known as the International List of Causes of Death, was adopted by the International Statistical Institute in 1893.
>
> WHO was entrusted with the ICD at its creation in 1948 and published the 6th version, ICD-6, that incorporated morbidity for the first time. The WHO Nomenclature Regulations, adopted in 1967, stipulated that Member States use the most current ICD revision for mortality and morbidity statistics. The ICD has been revised and published in a series of editions to reflect advances in health and medical science over time.
>
> ICD-10 was endorsed in May 1990 by the Forty-third World Health Assembly. It is cited in more than 20,000 scientific articles and used by more than 100 countries around the world.

126 'International Classification of Diseases (ICD)', *World Health Organization*, https://www.who.int/classifications/icd/en/

127 The text here is from the 2021 version of the ICD website which has now been superseded.

> A version of ICD-11 was released on 18 June 2018 to allow Member States to prepare for implementation, including translating ICD into their national languages [...] Member States will start reporting using ICD-11 on 1 January 2022.

Bringing ICD-11 and SNOMED CT within a unifying framework of description logic has proved challenging. Rector et al. give a forensic review of the internal structure of ICD and the still intractable (because too expensive to rectify) methodological limitations it imposes on its future development.[128] Similar 'too big to fail' limitations came to wider attention when methods of formal logic were deployed to analyze their mutual consistency and coherence.

LOINC–Logical Observation, Identifiers, Names and Codes[129]

The longstanding and influential LOINC resource, first developed in 1994, started as a database and standard for identifying medical laboratory observations. It was created and is maintained by the Regenstrief Institute, a US non-profit medical research organization.

The initiative describes itself as follows:

> [LOINC is] a common language (set of identifiers, names, and codes) for identifying health measurements, observations, and documents. If you think of an observation as a 'question' and the observation result value as an 'answer.' LOINC codes represent the 'question' for a test or measurement.
>
> Where needed, codes from other standards (e.g., SNOMED CT) represent the 'answer.' Of course, you don't always need a code for the result value. For quantitative results, the 'answer' is just the numeric value–with its associated units of measure.
>
> Most laboratory and clinical systems today are sending data out using the HL7 version 2 messaging standard).
>
> The system allows for local coding standards to be applied, as synonyms, within the universal standard it provides.

128 'On Beyond Gruber'.
129 LOINC, https://loinc.org/

ICPC–International Classification of Primary Care[130]

ICPC was developed in stages, from 1987, originally under the name HICPIC. Its last published update was ICPC-2, in 2003, from Oxford University Press. It was recognized by the WHO Family of International Classifications (FIC) as a means for classifying reason for encounter in primary care and general practice, wherever applicable. It takes account of the frequency distribution of problems seen in these domains and allows classification of the patient's reason for encounter (RFE), the problems/diagnosis managed, interventions, and the ordering of these data in an episode of care structure.

The website of the WHO describes the ICPC structure as follows:

> It has a biaxial structure and consists of 17 chapters, each divided into 7 components dealing with symptoms and complaints (component 1), diagnostic, screening and preventive procedures (component 2), medication, treatment and procedures (component 3), test results (component 4), administrative (component 5), referrals and other reasons for encounter (component 6) and diseases (component 7).[131]

There is clearly huge overlap with the services of secondary care and the standards of terminology and classification used there. Although high profile in its time, the lack of update for the past almost twenty years, indicates that it has largely disappeared from use.

Two Illustrious Pioneers

It has been one of my great good fortunes to know and work alongside pioneers of our field, and two such are Martin Severs and Alan Rector. They have been strong and indomitable thinkers about fundamentals, and orchestrators of change at the coalface of practice. Alan has pushed for incorporation of the methods of description logic in organizing medical terminology and knowledge representation. Martin has pushed health care delivery to the centre of national policy for health informatics, breaking down barriers and building international collaboration, notably in establishing the IHTSDO (International Health Terminology and Standards Development Organization) and becoming its founding chair. Their contributions are akin to those of Ranganathan and Coates, whose stories I told in the first half of this chapter.

---

130 'International Classification of Primary Care, 2nd edition (ICPC-2)', *World Health Organization*, https://www.who.int/standards/classifications/other-classifications/international-classification-of-primary-care

131 Ibid.

It is one of the rewards of academic life to supervise doctoral students and it was an illuminating honour for me to be invited to examine some of Alan's great cadre of PhD students, some of whom continued very successfully in the world of health informatics. The role of reviewer of research, especially of large-scale projects bridging across disciplines, organizations and countries, is another such experience. It brings disjoint worlds of endeavour and enterprise together in dialogue among assessors, in hearing and responding to the evolving story and drama of research teams and their work. I will describe more such experiences in the next chapter, exploring biomathematical models of cancer treatment in cancer research. Here, I draw on my personal experience of working with Martin and Alan.

Martin Severs

Fig. 2.4 Martin Severs–clinician and founding father of SNOMED International and medical director of NHS Digital, curating the NHS data forest. Now wielding a chainsaw to curate his own forest near Portsmouth and enjoying the great outdoors. Photograph by Martin Severs, CC BY-NC.

Martin provided leadership, organizational and political skills, clinical expertise and insight, professionally and nationally. He showed extraordinary commitment and staying power in raising SNOMED International into the worldwide position it occupies today. It is a great regret that our combined efforts over two years, some fifteen years ago, to merge the openEHR mission within this organization that he created, were unsuccessful. Some twenty carefully constructed working papers and plans were voted down by its more influential international government representatives, perhaps because the project was seen as too disruptive of contemporary interests and legacy. I have preserved all these working documents in my personal document archive of the field and consulted them when writing here.

I first knew Martin through his several leadership roles: of the Royal College of Physicians of London Computer Committee and afterwards the Medical Informatics Group, then of SNOMED, internationally, and finally as Medical Director of NHS Digital. He immediately stood out in his concern for those at the coalface of health care services, helping them to adapt to and survive the increasing burdens faced in managing information resources and responding to increasing, managerially driven, audit and governance process demands.

He was always a doughty, honest and loyal warrior and friend. He understood that the practical exigencies and burdens of clinical practice imposed exacting limitations of time and capacity on what was achievable there in standardizing the coding of records. Whole departments were being devoted to manual determination of codes descriptive of clinical records. In the electronic capture of records, necessary and sufficient explanatory codes should arise and be recorded integrally with the workflow of care, although this workflow itself, and the skills required by clinical professionals, would necessarily need to evolve, too. He realized that the transition from individual practice to organizations and wider groupings of services would encounter differences of perspective when defining requirements and implementing standards at each level and scale, from local to national and international levels. It was thus a highly political, human and professional endeavour, requiring new organization and clout to cope with and counter operational intractability, inertia and dispute. The same sort of context that Coates had encountered and fought against, for many years, as I described above in the history of the BSO.

Martin's foremost achievement was his leadership of the SNOMED organization for a decade. As a practicing clinician, he had become interested in the coding of clinical practice through observing the inside story of Read Codes, named after James Read, a general practitioner who had composed a loosely structured compendium of the terms used in general practice. From this background, he achieved reputation, trust and sponsorship at a national level, supporting its extension into the wider NHS Clinical Terms project, as a foundational level of recording in electronic patient records.

Established terminology such as the long-established International Classification of Diseases, was focused on what might now be termed static knowledge–although that characterization is, as ever, philosophically contentious. Static tends to mean what is included in and taught from contemporary textbooks of medicine. It consisted of a hierarchy of names of disorders, pragmatically structured and grouped according to discipline, and using the language of the day.

Read and SNOMED terms and codes branched further into new nomenclature for coding medical records, to enable the creation of more coherent and consistent textual accounts–thus more readily and reliably used, cross-referenced, searched and communicated. The terms covered the presenting problems of patients, investigations, clinical reasoning, actions and outcomes–the who did what, when, how and why of health care. They were generalizations and these mushroomed in range and detail, attempting to keep track of the almost infinite variety of patient journeys, treated in context of time, place and person. They started with pragmatic choices in representing this complexity, mirroring the history of library classifications of the content of documents and books, emerging from pragmatic and enumerative approaches towards a theory-based compositional structure of codes. SNOMED metamorphosed into SNOMED-RT to absorb the Read Codes. Read was substantially rewarded for assigning his copyright in these to the NHS.

Under Martin's determined leadership, SNOMED became an internationally backed and fast-growing organization, with United Nations (UN) formula-based national subscription to support its operations and strong defense of its copyright.

Alan Rector

Fig. 2.5 Alan Rector–pioneer of formal logic and its application to the curation of medical records and terminology. Now creating order in his vegetable garden and playing the piano, he tells me. CC BY-NC.

Alan was the intellectual powerhouse behind the GALEN initiative, on which he set to work around 1990. This was, in many ways, an analogous quest to the colon classification initiative of Ranganathan and conducted in very much the same pioneering spirit. Like Octo Barnett (1930–2020), a founding father of medical informatics, Alan brought personal grounding in both medicine and computer science.

I knew about Alan's alma mater, the University of Manchester, through its pioneering contributions in Computer Science and Engineering in the UK, in the days when it was world-leading in semiconductor physics and electronic engineering. Alan led a wide range of projects in electronic health records, most notably in application of description logic to formalizing medical terminology. Also at Manchester were Christopher Taylor, who worked in medical biophysics and imaging science, another pioneer I knew from earliest days when we were both members, with Jo Milan (1942–2018), of the Hospital Physicists' Association Computer Topic Group, subsequently the IPSM (Institute of Physical Sciences in Medicine).

I got to know Alan when we arrived together at a meeting room in Brussels, in 1991, under the auspices of Niels Rossing, the outstanding leader of the EU AIM (Advance Informatics in Medicine) initiative. We had both led successful proposals–he for the GALEN project, described above, and I for the GEHR Project (Good European Health Record), the forerunner of openEHR.

Alan was concerned with the computational integrity and expressiveness of medical knowledge bases. He worked to explore the development and use of description logic as the basis of open-world representations of knowledge–open in the sense of dealing with any possible logically conformant world of content. He proposed new thinking to replace pragmatic rules of enumeration of content with logical models defining content, supporting composition and analysis. This was a close parallel with the transition from enumeration to analytico-synthetic composition exemplified by Ranganathan and Coates's endeavour in the UNESCO Broad System of Ordering, as described earlier in the chapter.

Alan faced an uphill struggle in maintaining and enhancing his position and departmental support in Computer Science, while securing ongoing engagement at the coalface of clinical practice–he did not practice as a doctor when I knew him. At the time, I was in a mirror world, working to establish a position in the clinical environment of a Medical School Department of Medicine, with struggle in the reverse direction, maintaining and enhancing broader connections with computer science and other cognate academic disciplines. He was recognized internationally for his research contribution and looked after a small group of devoted colleagues and students, with whom he maintained contact into his retirement years. His manner had a North American phenotype, from where he came and that I had come to recognize–a sometimes rasping and angular mental toughness, and sometimes rather intimidating! It was a 'no nonsense' and honest style, as was that of Martin, which I admired and appreciated.

Alan brought the greatest of minds to his domain of research, surveying across description logics and their intersection with medicine. He kept abreast of all these tools, alongside others of the era used in natural language processing techniques for analyzing bodies (corpora) of texts descriptive of different domains of medical knowledge. His students' theses moved the field forward. I especially recall examining Nicholas Hardiker's thesis on the use of description logic as a metalanguage for comparing and analyzing several existing and pragmatically constructed terminologies descriptive of nursing care. In this, he concluded that it was impossible to translate meaning from one to the other.

Medical Knowledge Bases

> Knowledge is in part kinaesthetic; its acquisition involves having a hand, to say the very least. There are, in other words, some things humans know by virtue of having a human body.[132]

> Knowledge is only a rumour until it lives in muscle.[133]

As this chapter has unfolded, we have seen how the advent of the computer has brought new perspective and focus to the understanding and communication of human knowledge and reasoning. The Ranganathan circle of knowledge has encircled human reasoning–clockwise in abstraction of theory and counterclockwise in rationalization of practice. Stories of this evolution permeate the book. Joseph Weizenbaum (1923–2008) conjectured about knowledge that is accessible only to humans. We now conjecture about knowledge accessible only to the computer. The Asaro tribe saying connects knowledge and muscle. What is the computer's muscle? What will govern arm wrestling competitions between computer and human muscle?

Earlier sections of the chapter included discussion of the historical development of library classifications and medical terminologies and progressed to describe initiatives reaching from their pragmatic origins towards more rigorous and sustainable theoretical foundations. They tracked the scope and methodology of these evolving endeavours, into the Information Age, and alongside the evolution of formal logic and reasoning embodied in computerized knowledge bases.

132 Weizenbaum, *Computer Power and Human Reason*, pp. 207–08.
133 A saying of the Asaro tribe of Indonesia and New Guinea.

The nature of the book as a whole requires a mixture of organizing principles in its presentation of material, chapter by chapter. These reflect the timeline of events, my personal songline of experience and the thematic development of the ideas covered. The combination of these is a pragmatic one and the coverage of health care knowledge bases is introduced in increments along the way, as methods and applications arise in other contexts. This section is by way of an introduction to this ever-widening domain of experimentation, recently surveyed by Musen and colleagues.[134] It is a field full of promise for the careful and rigorous, and of pitfall for the careless and imprecise. As has so often happened, terminology descriptive of the field has evolved in confusing ways. In their paper, Rector and colleagues made suggestions for how such terminology might be usefully standardized.[135]

As further unfolded in the following chapters, developments in mathematics, formal logic and linguistics, and theory of computation and abstract automata, co-evolved with scientific and technological advance to place computational method on increasingly rigorous foundations of discipline. As discussed in Chapter Five, mathematics and computer science combined to shed new light on computation, through the unconventional genius of Turing and exploration of his eponymous Turing Machine–a conceptual but practically implementable device. Study of the potential solutions of classic problems, such as calculation of the optimum route among clients to be followed by a travelling salesman, revealed new complexity, and led to new analysis, of computational problems and the intrinsic difficulty in solving them–what could or could not, in principle, be solved, and with what resource of time and machinery. Proof of the consistency and correctness of the decision logic representations of knowledge bases mirrored proof of mathematical theorems.

Philosophy has debated 'knowledge that' and 'knowledge how', and the ways they connect–knowing that something is true and knowing how to make or do something. Knowledge bases deal in both these aspects–as a discipline for maintaining good order in the representation of a domain of knowledge and as a tool supporting its effective use in practical contexts.

As with all matters philosophical, there are contrary perspectives, of course. Gilbert Ryle (1900–76), the erudite philosopher of mind who passed me by every day, near my rooms at Magdalen College, considered this a false dichotomy. But humans are fond of arguing in terms of dichotomy, distinguishing supposedly disjoint ideas, and arguing about their connections. Dividing lines between them are often matters of much

134 Musen, Middleton and Greenes, 'Clinical Decision-Support Systems'.
135 Rector et al., 'On Beyond Gruber', p. 3.

musing, and rather less moment. Ontology and epistemology get typecast somewhat in this way, as do health and care, and dare I breathe it, messages and information models in context of meaningful communication of data between health care computer systems! Emoticons are not allowed in the book's text but take it that a smiling one is due from me, here! Some of my colleagues might likely choose to emote in more dramatic moods and modes!

As Robert Oppenheimer (1904–67) argued in his Reith Lectures that I discussed in Chapter One, in relation to scientific understanding, 'both and' can be a more fruitful approach than 'either or', in the pursuit of sound reason, good decision and effective action. 'Either or' often tends towards a playground of zero-sum gamers and their winners and losers.

Many computerized knowledge bases have straddled the 'that/how' dichotomy of knowledge. They have combined formal representation of *that* which is known in a particular domain of knowledge, with methods for reasoning with this knowledge and reaching decisions about *how* to act in its light, to achieve a desired end or solve a problem. The knowledge represented may be of problem, discipline and individual subject–heart attack, cardiology, John Smith presenting with chest pain, for example.

There has been a myriad of such examples, differently specialized as to how they approached and combined their 'that' and 'how' dimensions of knowledge. Among the earliest I worked with were radiotherapy treatment planning programs, which combined the 'that' knowledge of dose-absorption characteristics of ionizing radiation incident on human tissue, with the 'how' knowledge reflecting the state of the art in deciding the alignment and time course of application of the beams of radiation delivered in cobalt and linear accelerator radiotherapy treatment of cancer.

One pioneering medical knowledge base, which has already been mentioned in other contexts, combined knowledge descriptive of antimicrobial therapeutic agents with rules for their use in identifying and treating threatening pathogens encountered in caring for an individual patient. This was the MYCIN system, further described below, that evolved over many years from a team that included luminary figures like Feigenbaum, Bruce Buchanan and Edward Shortliffe, founding fathers of biomedical informatics in the USA. They made foundational contributions to the field of expert systems, a front runner of artificial intelligence.

A common fate of pioneering innovations in the field, such as this one, is that they were gradually disrupted and nudged aside by more generic tools, such as OWL, the OMG-pioneered knowledge representation language. As we have seen, open-world description logics have found fruitful application in the context of organizing medical terminologies, but both these methods and closed-world frame logics, and their combinations alongside other pragmatic software fixes, have struggled to provide and sustain adequate methods for useful and sustainable formalized medical knowledge bases more widely. And more generic tools of artificial intelligence, which are now rapidly coming of age, are poised to disrupt this domain, once again.

As was also introduced above, this machine intelligence has provided a different, and more simplistic and naive, approach to some such problems, which are seemingly now poised to prove more usefully competent there. It largely dispenses with 'knowing that' and restricts itself to 'learning how'. Knowing *that* can be of more limited use than knowing how, in making decisions about something that is required to be made or done. People and machines both learn to play games, and how to play them to win. Broadly speaking, the 'that' knowledge of chess is simply the rules of play with the pieces on the board. And the 'how' knowledge (its 'know-how') is an infinitely more subtle and massive ensemble of pattern, strategy, experience and expertise, which includes weighing up the opponent during the course of play and studying their previous form. Nowadays, the machine can discern what the human does not see in the game, or perhaps cannot articulate. The 'that' and 'how' of the game of Go are yet more simple and yet more subtle, we are told. I do not know, so could not say. Again, the naive but capable machine has scored.

The ideas underpinning medical knowledge bases continue to feature in later sections of the book. To conclude this section, I will briefly introduce some further pioneering examples, whose pioneers I have known and followed in their initiatives to explore clinically relevant knowledge bases in the Information Age.

Pioneering Examples

Heuristic Dendral–Analyzing Mass Spectra[136]

I remember collecting the impressive early work at Stanford of Edward Feigenbaum, from the 1950s, a pioneer in the field of expert systems who coined the term heuristic programming to describe this algorithmic approach

136 B. G. Buchanan, G. Sutherland and E. A. Feigenbaum, *Heuristic DENDRAL: A Program for Generating Explanatory Hypotheses in Organic Chemistry* (Stanford, CA: Stanford University Department of Computer Science, 1968); G. Sutherland, *Heuristic DENDRAL: A Family of LISP Programs* (Stanford, CA: Stanford University Department of Computer Science, 1969); R. K. Lindsay, B. G. Buchanan, E. A. Feigenbaum and J. Lederberg, 'DENDRAL: A Case Study of the First Expert System for Scientific Hypothesis Formation', *Artificial Intelligence*, 61.2 (1993), 209–61.

to reasoning. Written in LISP, a pioneering programming language of the era, Dendral was designed to study hypothesis and discovery in science. The Heuristic Dendral and Meta-Dendral programs reasoned with mass-spectrometer profiles of known molecules, to deduce the composition and structure of unknown molecules, based on their measured mass spectra. Knowing the atomic mass of a compound, a search of known elements was used to create a candidate list of potential composing groupings of atoms. For water, with atomic mass eighteen, this is not a large list, but for larger compounds the list of potential combinations explodes in size and has somehow to be reined into a manageable group for further study. This narrowing down was achieved by applying a set of general rules, based on knowledge of the science governing the formation of compounds. The initiative thus focused on a specific problem domain, and the rules pertaining within that domain, while exploring more widely the general issue of representation and reasoning with data and knowledge.

MYCIN–Prescribing Antimicrobial Therapy[137]

I also read and collected early books on exploratory applications in medicine, including, notably, the rule-based expert system, MYCIN, developed by an early colleague of Feigenbaum, Buchanan, and then spearheaded within the domain of medical informatics over the coming decades by another luminary figure, Shortliffe. MYCIN, one of many so-called rule-based expert systems in medicine, embodied knowledge about microbial diseases and rules governing and justifying an automated reasoning process, to elicit information about the patient, identify the disease and propose the drug treatment that was indicated. MYCIN embodied a database of antimicrobial drugs and their properties and uses. It thus combined knowledge of the clinical and scientific domain with expertise in diagnostic reasoning about individual patients.

Through Musen, Stanford also led the way in evolution from rule- to frame-based knowledge base systems. This generic approach to closed-world knowledge representation and reasoning was disseminated in the Stanford Protégé system, which now combines with the description logic of the OWL language, as introduced above. It populates and edges outwards the current practical limits of computability, implementing tractable elements of predicate calculus.

137 Shortliffe, *Computer-Based Medical Consultations*.

PROforma–Modelling Clinical Decision-making Processes

Closer to home, my colleague and friend, John Fox (1948–2021), who, sadly, died as this book was being written, was a career-long tiger in the jungle of clinical decision support systems, in his career spanning the invention of tools and their testing in clinical applications. He created and ran the Advanced Computation Laboratory of the Imperial Cancer Research Fund at Lincoln's Inn Fields in London, from where he created and led the OpenClinical initiative.[138] He and his team invented PROforma, an agent technology for modelling clinical decision making, and Tallis, a visual design studio for creating such models, and much more.[139] Standout future leaders in health informatics, such as Enrico Coiera and my subsequent colleague at UCL, Paul Taylor, launched their careers in his team. This was a fertile environment bringing together pioneers of the age. Nearby at the Royal College of Surgeons, was the home of a luminary contributor to clinical measurement, Denis Hill, working there with John Bushman and James Payne. He in turn connected with William Mapleson (1926–2018), at Cardiff, a pioneer of the mathematical modelling and control of anaesthesia, who I got to know during the first half of my career, devoted to educational computing and the modelling of human physiology, as described in Chapter Four.

The Foundational Model of Anatomy (FMA)[140]

The Foundational Model of Anatomy is a frame-structured hierarchy and closed-world representation. Human anatomy is complex and so is this model, which provides a computational representation of how an anatomist views the human body.

The architecture of the model is a very large compositional containment hierarchy. As of 2020, it described itself as comprising seventy thousand concepts, one hundred and ten thousand terms, six and a half million instantiations and one hundred and seventy kinds of relation. The hierarchical class relationships descriptive of anatomical structure proceed from anatomical entity and subsume structure, organ and cell classes.

138 *OpenClinical*, http://www.openclinical.net

139 D. R. Sutton and J. Fox, 'The Syntax and Semantics of the PRO Forma Guideline Modeling Language', *Journal of the American Medical Informatics Association*, 10.5 (2003), 433–43.

140 C. Rosse and J. L. V. Mejino Jr., 'A Reference Ontology for Biomedical Informatics: The Foundational Model of Anatomy', *Journal of Biomedical Informatics*, 36.6 (2003), 478–500, https://doi.org/10.1016/j.jbi.2003.11.007

Concepts, classes and instances are represented as frames. Continuity and physical adjacency of entities in the hierarchy are handled as relationships. Connections between frames are handled through slots. One has to see the detail to understand the precise meanings intended by these descriptive terms–a controlled vocabulary of such descriptions, perhaps with the admirable Rector as stern rector, would help a lot!

Overlying the structural hierarchy there are physiological/functional, radiological, surgical and biomechanical system hierarchies. A difficulty faced is how to handle problems of change in the applicable model during ontogeny, as the body develops and grows. Describing the change in lung anatomy and function at birth would be one such transformation.

Knowledge, Information and Data

In the oncoming chapters, the storyline of the book moves from the world of knowledge into one of data and information. It traces this path through successive chapters on observation and measurement, models and simulations, and information engineering. In passing, I interpose a reflection about the oft-enunciated trichotomy of knowledge, information and data.

As the quotations at the head of this chapter signified, it does not pay to apply our mind too seriously to defining what we mean by knowledge, these days. Such attempts often seem self-referential; knowledge is what we know! Hofstadter has gone so far as to suggest that the mind itself is a self-referential system–as discussed in his book, *I Am a Closed Loop*, introduced in Chapter Six. Whitehead wrote that reality just is, and we can connect only with appearances of reality.

The differentiation of knowledge, information and data is another topic that has involved much mental exercise. Again, self-reference abounds–information is how we inform ourselves, to guide our actions, and data are, as Latin indicates (*do*, *dare*, *dedi*, *datum*), 'givens' that we capture, collect and collate. Such abstractions often elude singular definition and are described in multiple ways, each adding to the wholeness of the ideas expressed.

It seems that it does not pay to spend too much time defining mathematics, either. As mentioned earlier in the Introduction, some say that mathematics is what mathematicians do. They explore rigorous tools for reasoning about the manifolds of number, shape and form, that they care, and find it useful, to reason about in that way. They enact connections throughout science, providing a unifying framework and thread. Mathematics of physics connects with physics of chemistry, chemistry of biology, biology of life, and the lives of people, ecosystems and societies.

As with mathematics, it does not pay to spend too much time defining informatics, nowadays, either. We might cop out in a like manner and say that informatics is what informaticians do. They explore rigorous tools that connect *how* we capture, process, codify, analyze, communicate and reason with knowledge and data, at all levels of science and society, in support of *how* we understand their meanings and *what* we make and do. Informatics is to knowledge more broadly, as mathematics is to science, including now computer science. It is elusive to capture and define information as an abstraction, as it is the manifolds that mathematicians imagine and deal with when describing number, shape and form. Some, as does Paul Davies, conjecture that information may be found to be a fundamental component of physical reality and measurable as such–as I discuss in Chapter Six, in the outline of his book, *The Demon in the Machine: How Hidden Webs of Information Are Solving the Mysteries of Life*.

Mathematics and informatics are central to *how* we connect the science and practice of health care with the computer, helping to frame answers to questions about the way things work. Ranganathan placed them side by side, after philosophy, in his circle of knowledge. *How* connects with causal power–information connects knowledge with action.

In the computer era, fragmenting the world into disjoint domains of knowledge, information and data–K, I and D–has sometimes degenerated into kids' play, digging bottomless pits of abstract and unconnected conjecture and distraction! Wholeness reflects the three in concert. If one follows Einstein's thinking, as also highlighted at the beginning of this chapter–'all knowledge of reality starts from experience and ends in it'–and John Archibald Wheeler's (1911–2008) observation, that appears again in Chapter Six–that during his career, 'everything is particles', gave way to 'everything is fields', and then to 'everything is information'–one might assert the primacy of information in characterizing all knowledge.[141] The KID stuff is then all information and informatics is science that seeks rigorous and trusted tools for engaging and reasoning with, and making sense of experience.

A highly abstract and philosophically Olympian perspective, for sure, and speculative, no doubt. But, who knows, in time, in a future Information Society, it may make sense to connect human quests to 'know', in that way. Mathematics connecting domains of number, shape and form, and informatics connecting domains of knowledge, experience and action. The

141 J. A. Wheeler, 'Information, Physics, Quantum: The Search for Links', in *Feynman and Computation*, ed. by A. Hey (Boca Raton, FL: CRC Press, 2018), pp. 309–36, https://doi.org/10.1201/9780429500459-19

'What is life?' and 'What is information?' questions, and what connects their answers, are considered in Chapter Six.

Concern for clarity and consistency of discourse in the languages of philosophy, mathematics and logic have fed into the evolving language and discourse of computer science and medicine. Study of the logically-deduced framework of mathematics proposed in *Principia Mathematica* led to its axioms and reasoning coming under a new spotlight. In similar manner, the evolution of medical language, terminology and classification, and programming languages that manage them in computable form, have brought new focus on the consistent and coherent use of language descriptive of medicine and health care. This has been a slow process of discovery, refinement and standardization, based on how ideas have worked out in practice, in managing information that embodies and connects knowledge and data with logic and algorithm. Along with this must come consistent and coherent information, framed within coherent discipline of informatics.

Progress towards health and care made whole, which will underpin their reinvention in the Information Society, will rest substantially on information and informatics relating to them made whole, which will rest on knowledge and experience of them, made whole. These are, of course, most unlikely ever to be 'made whole' in any absolute sense and must therefore always be treated as provisional and evolving understandings, with recognition of their limits. The computer is both a central ally in this quest and a powerful adversary. Concern for clarity, consistency and transparency in this evolving discourse will become ever more important as reasoning about them increasingly disappears into the inner worlds of ever more capable machines.

Health and care are intrinsically grey and messy areas, where noise is rife and context all important. Will computing machines have the capacity or the 'will' to handle the symbols they work with, cautiously and within a humanly recognizable rationale of provisionality and conditionality? Or will this evolution of the machine lead to assertion and imposition of a more machine-based rationale , and how will that play out? The recent advances in machine learning and the potential of quantum computation to widen capability in formal logic systems portend a long road ahead in which these technologies will play out, from playing with checkers to playing with Go, to engaging with life and health care. Chapter Six charts the connections being made between the science of information and the science of life. These seem destined to play forward into Lovelock's *Novacene* era of machine intelligence. It was interesting to observe the amazement in the world of life science, reacting to the announcement of AlphaFold, as if of magic, as I first wrote this section some weeks ago.

Parenthesis–Traction

Vehicle traction is force applied that achieves grip and creates movement. Travel and steering depend on it. Physical traction depends on connection and synergy. We generalize the idea in mathematics and computation–the intractable problem is one where we lack a viable method for its solution. In health care, traction is more organic in nature, contributed to and experienced by both carer and cared for. This traction is one of human connection and flow, on which tractable balance and continuity of services depend. Information policy for health care services, in general, has become an intractable problem. Here, I reflect on why that is and what might be done to improve matters.

A driver who skids and spins their vehicle on ice has connection but not traction between wheels and road, lacking friction between tyre and ground. Applied torque force, car momentum and direction of steering are out of balance. The car is difficult, if not impossible, to drive and can become dangerously out of control. Anti-skid technology can assist in maintaining traction between car and ice, making its control more tractable. Where traction is elusive or not possible, we speak of an intractable situation. New technology can make the intractable more tractable. What was unimagined or practically impossible can become feasible, albeit sometimes still very hard–perhaps requiring adaptation of purpose, goal and method in what we are attempting to make and do.

In everyday life, we seek traction between our efforts and their desired outcomes. The pursuit of knowledge seeks traction between ideas expressed and the context in which they arise and are applied. In science, this is traction between theory and experiment, and in everyday life between theory and practice. New ideas can enable progress on previously intractable problems and dilemmas. Advance in knowledge increases our sense of the known, unknown and unknowable. And advance in technology increases our sense of the tractable, intractable and impossible.

It is, though, a phenomenon of our times that fundamental underpinning utilities, on which everyday life depends, tend to become noticed only when they fail. We used to know about what was going on under the bonnet of our car, because we had to. But now we do not, and mostly cannot; it is opaque to us, and we have come to rely on services that keep our cars on the road. How far are we content to allow a similar derogation and delegation of power over and governance of personal data about our health care, to what might be, or become, opaque information engines and utilities? Software and data systems and services can also become opaque and inaccessible, or unusable, when they become technologically obsolete. In this sense, advance of technology may render a previously accessible and tractable utility no longer so. This is another kind of slippery surface on which we need to navigate with better traction!

Traction matters when tackling difficult immediate needs and concerns, faced with cost of effort, risk of harm and uncertainty of outcome. Good government requires that policy, goal and implemented plan of action have traction. Information technology has brought revolutionary advances in measurement devices and analysis of data. But the policy level encounter of information technology with health care has struggled to achieve tractable balance, continuity and governance of high quality and affordable services. Considerable applied torque has spun often threadbare wheels, failing to achieve commensurate movement on the ground! The UK NHS is pointed to from across the world, as having been the home of the biggest public sector procurement failure ever, with its early twenty-first century eleven-billion-pound National Programme for IT. Not all bad but not yet good enough.

This chapter has highlighted issues of traction between theory and practice in relation to knowledge and computation. Theory and practice are both implicated, and both lacking, when new ideas unfold in new practices. A sweet spot is needed, on which to concentrate in gaining traction. This is only discoverable with adventurous people and their environments, ideas, luck, hard work and staying power. Successions of hubristically championed and expensive sour spots have impoverished the pathway of much of health care service computerization and made it slippery. The evidence of growing dysfunction in the vehicles navigating this slippery ground is under their bonnets, out of sight, or observed with poor eyesight, in the many string-and-sealing-wax legacy information machines that pertain there. The issue of *New Scientist* this week, as I write, has a telling article on the hidden crisis of legacy software that is facing the world at large. There must be a more tractable way forward. I will describe one that I favour and have experimental and experiential knowledge of its formative stages, in the third part of the book.

Learning about *what* we assert to be true of something and connecting it with *why* and *how* we reason and decide about what to *make* and *do* in the context of this knowledge, has surfaced new issues of computational traction, as knowledge bases have advanced in the Information Age. Learning both *why* and *how* to make and do something involves *experience* (*expiri* [to try]), which comes from a combination of *discovery*, *experimentation* and *practice*.

Making and *doing* things is about *methods* employed. *What* we do and *why* we do it is tightly bound up with *how* we are able and chose to *implement* it. A tractable method must be *grounded* and *understood*, with awareness of both its scope and limitations. Grounding of method requires *consistent theory* matched with *achievable practice*. This depends on resource, community and environment, as well as capability. In that spirit, too narrow a consistency may not be compatible with a useful practice. As Ralph Waldo Emerson (1803–82) famously wrote in *Self-Reliance* (1841), 'a foolish consistency is the hobgoblin of little minds, beloved by philosophers, politicians and divines'! User community support for implementing, testing, refining and sustaining new ways of working, are essential for them to gain and sustain traction. Where this is lacking, innovations (however promising and much needed) may likely prove burdensome and unacceptable to those expected or required to use them.

Experiment and practice have *errors and uncertainties* which must be understood and accommodated. These can, to advantage, be tamed and narrowed, and thus made tractable, but not eliminated. In facing errors and uncertainties, and finding tractable method for coping with them, human beliefs and values are both motivators and demotivators of action–to stay as we are, asserting *status quo,* or to cope with fear of the unknown, innovate and seek to change. An overly idealistic or visionary approach risks obfuscating, as much as helping efforts to improve. An overly cautious and limited approach risks stasis and decay. Neither will likely gain and retain traction.

Theory and practice can quickly detach from one another within unclear and confusing experimental contexts. The story of the decades of implementation experience in the encounter of information technology with the delivery of health care services is one of failure of traction between policy, goal, method, team and environment. Policy must focus practical health care *engagement*, computational *rigour* and public *trust*, as a tripod of supportive legs on which to create good standing of method, team and environment.

The close connection of innovation in computational methods with everyday presence in, and experience of, practical problem-solving environments is exemplified in many of the stories in this book. Turing's foundational contribution in computer science, which I further describe in Chapter Five, related to his success in aligning theory of computation with machinery of computation–an abstract machine but a palpably implementable one. The wartime code-breaking exploits at Bletchley Park exhibited interplay of theory and implementation, to solve the most immediate and practical of problems–breaking the code of the ENIGMA machine. The story of Barnett and the MUMPS language, featured in Chapters Five and Eight, played out in practical context of the limited capabilities of the DEC minicomputer machine environment available to him and the characteristically sparsely filled and continuously changing data ecosystem of everyday clinical practice. More recently, advances in tractable description logic have played out in the context of problems encountered in creating orderly foundations of medical terminology.

This close connection with implementation was why and how my colleague, Jo Milan, was so successful in building coherent and effective information systems for the Royal Marsden Hospital in London, as I describe in Chapter Eight. Lack of awareness of the nature and depth of his contribution, and its significance, revealed harmful and costly disconnectedness of health information policy and policy makers, from such everyday realities on the ground. In failing to detect and understand this quality, although seeing it performing there, in plain sight, they failed to support this contribution and enable it to gain wider traction. Governments need to rethink their constructive roles in such creative endeavours or they will continue to drain away resources and make worse the problems that they believe themselves to be addressing, in what prove serially disjointed policy initiatives.

These many aspects of traction were all contributory reasons that led me to frame the three top priorities of the openEHR mission as implementation, implementation and implementation. This initiative spearheaded a new approach to achievement of a coherent and soundly based architecture for digital care records, as described in Chapter Eight and a Half. I also described the mission, in a letter to my colleague and friend, the chair of SNOMED International in those early times, Martin Severs, as prioritizing 'Little Data' first, ahead of 'Big Data'. This novel methodology–established, iteratively and incrementally, in and from its foundational team and environment–laid the foundations for the subsequent incremental traction of the openEHR mission.

Surveying the field of medical terminology today, SNOMED has established traction, and now occupies a key position in the context of names and structures for coding and mapping care record entries to a description logic style of knowledge base of clinical terms, and related expression language. ICD also has traction. More narrowly focused on classification of diseases, it remains poised between a history of largely pragmatic methodology, and a future moving towards a mixture of OWL-based subsumption and frame hierarchy and annotations, that together capture the richness of what it seeks to describe. The LOINC codes, which focus on clinical laboratory findings, is performing well in its domain of application. For these terminologies and classifications to increase their traction with health care service delivery, they need to be better aligned with coherent reference models that overarch record, analysis and communication of data expressing the wholeness of individual patient care. The ways in which this might best be approached and achieved remain matters of interest, experiment and debate, currently being played out, principally and in principle, among four communities: HL7, IHE, openEHRInternational and SNOMED International.

A good heuristic for achieving and maintaining traction is to keep close to priorities for things that need to be done, and for which we can establish a clear and practical plan for how they should and can be done better. In this we must prioritize learning from successes and failures of traction along the way, and how they came about. The detailed history set out in Chapter Seven shows that coherent, integrated and interoperable digital records of care, providing a comprehensive, balanced, meaningful and ethically acceptable account, and supporting continuity of care, have long been and stubbornly persist at the head of these priorities. Traction in this endeavour requires synergy across communities of health care practice, research, education and industry. It requires teams and environments in which mutual collaboration can be nurtured and sustained.

Ethical traction of science within society is an increasingly complex matter, when seeking to achieve and ensure wise balance, maintain effective continuity and assure trusted governance of services. These three legs of a tripod–balance, continuity and governance–emerge strongly along the storyline of the book, as it evolves towards a perspective of care information as a public utility that connects and flows organically. This utility is envisioned to be centred on individual care records, owned by autonomous and enabled citizens, as a symbol of their ownership of their individual health care needs. It is envisaged to be co-created by them and the professional teams that treat and support them, as a symbol of the essentially connected roles and responsibilities of both parties, and the mutual dependencies and duties that connect them in their pursuit of traction in both personal health care, centred where they live, and public health care services.

The Quaker short book entitled 'Advices and Queries', which I used to read at one time in my life, advises that a useful test of our knowledge is that we should 'know it experimentally'. In the world of artificial intelligence and Big Data, what is the relationship between what the computer gets to know experimentally and what we do? Society today is observing artificial intelligence at a moment of awe in its evolving power, comparable to that when chess playing machines first challenged and defeated human mastery. Health care is not a game, and the human stakes are far higher. Are we going to tip over our King in resignation, as we had to with the chess game? If not, how will we secure traction in the future balance, continuity and governance of health care, such that Big Data and AI are our servant and do not conspire, by default, to impose what I have heard described, in another context, as a 'confederacy of virtual caricatures' that shape human life and experience in unwished for ways? The story of Dr Faustus and Mephistopheles is not without parallel; I certainly do not impute any devilment, but the hype of AI is a bit magical and different incarnations of Mephistopheles may be lurking, or emerge, there! At least doctors today should think about the other doctor, Doctor Faustus, and the warning to him–*homo fuge*! No one would wish for his fate! I reflect further on this issue, a bit more seriously, in Chapter Eight.

The historical axis of this first long chapter of the book has tracked through philosophy, spoken language, mathematics and logic. It introduced Ranganathan's circle of knowledge, and it is further round this circle, into the natural and clinical sciences, that the storyline of the book now progresses. The next chapter enters the world of observation and measurement, where the Information Age has seen transformational change in science and technology. Medicine and health care have been enticing and sometimes controversial domains in which this evolution has played out. What and how we can now measure has extended almost beyond bounds. Why and to what ends we measure, and how we share and regulate measurement, and records of measurement, have become of increasing ethical concern as we employ ever more powerful and autonomous computational methods.