

# Knowledge Infrastructure

Gordon Rugg  
Hyde & Rugg Associates  
Market Drayton, UK  
gordon@hydeandrugg.com

Colin Rigby  
Keele University  
Keele, UK  
c.a.rigby@keele.ac.uk

Sue Gerrard  
Hyde & Rugg Associates  
Market Drayton, UK  
sue@hydeandrugg.com

Amy Martin  
A.G. Telehelmets  
Market Drayton, UK  
amy@agtelehelmets.com

Jennifer Skillen  
Keele University  
Keele, UK  
j.d.skillen@keele.ac.uk

Emma Bonfiglio  
Keele University  
Keele, UK  
e.bonfiglio@keele.ac.uk

Allison Gardner  
Keele University  
Keele, UK  
a.gardner@keele.ac.uk

Yuxin Guo  
Honeywell UK  
Horsham, UK  
yuxin.guo@hotmail.com

Shailey Minocha  
Open University  
Milton Keynes, UK  
s.minocha@open.ac.uk

Gavin Taylor  
Weavers Academy  
Wellingborough, UK  
gavint321@gmail.com

***In this article, we outline the concept of knowledge infrastructure and describe how it differs from Information Technology (IT) infrastructure, with particular regard to the implications for education theory, practice and policy. We examine the inherent limits to growth for attempts to handle knowledge, as opposed to information, via the types of software and hardware likely to be available in the next few decades. We show how a simple process model can be used to identify pinch points where knowledge, as opposed to information, is a bottleneck. We also show how a simple model of knowledge types and knowledge locations can be combined with the process model to remove those bottlenecks via existing low-cost technology and a more efficient use of existing human expertise. We conclude that a minimal investment in knowledge infrastructure would provide significant human, social and economic benefits, by creating major added value from existing digital and organisational infrastructure.***

*Keywords: tacit knowledge, semi-tacit knowledge, information, knowledge, sociotechnical.*

## 1. INTRODUCTION

In this article, we argue that the very visible success of Information Technology has masked an underlying lack of progress in handling of *knowledge*, as opposed to handling of *information*. We argue that the implications of this distinction have been seriously underestimated.

We discuss the differences between information and knowledge in terms both of the underlying computational issues and of the practical implications. These issues are well known in technology-based communities of discourse and

practice, but these issues and their implications are not widely known in socio-political communities of discourse and practice which deal with public policy.

We then use worked examples to illustrate the issues involved, and to illustrate how they can be handled.

Finally, we discuss the implications for building a knowledge infrastructure to complement the existing IT infrastructure.

## **The growth and the limits of Information Technology**

An issue which has been treated as central in computing for decades is the distinction between data, information, and knowledge.

We argue that this issue now has major implications for how society handles the infrastructure both for Information Technology, and for the topics that Information Technology is unable to handle.

We begin by looking at the historical context of IT development. We then look at the nature of data, information and of knowledge, via definitions from computing and knowledge representation which provide an operationalised foundation for the subsequent discussion.

Although IT in general and the Internet in particular are often portrayed as a single revolutionary and recent innovation, there is a strong case for viewing them instead as a scaling up of numerous technologies, many of which, such as telecommunications networks, date back to the early twentieth century or the mid nineteenth century (Standage 1998). There is also a strong case for arguing that radical innovation in most of those technologies ended decades ago, and that those technologies have now peaked qualitatively, in a phase of improvement more by quantitative incremental refinements than by major qualitative breakthroughs.

An article in the Harvard Business Review provides an example (Davenport and Fitts, 2021). The authors give a good overview of the potential for business development via cutting edge technologies from Artificial Intelligence (AI). Within AI, however, all of the main approaches used today were well established fields of research by the mid 1980s, and have been routinely described for decades as promising technologies on the verge of transforming the world.

The MYCIN system, widely viewed as the first expert system, was developed in the early 1970s (Shortliffe and Buchanan, 1975). Genetic Algorithms (GAs) were well established by the mid 1980s (e.g. Forsyth, 1981) as was Machine Learning (e.g. Quinlan, 1986). Artificial Neural Networks (ANNs) were also an established technology by that time (e.g. Werbos, 1982). The same is true of the standard current types of application software, such as word processing, spreadsheets, and databases, together with browsers, search engines, Graphical User Interfaces (GUIs) and email.

The successes of these innovations are well recognised. What has received less widespread attention, however, is that the inherent limitations of current IT approaches were identified decades ago,

but have been largely masked by the very visible successes of these approaches.

For example, autonomous cars, in the sense of cars driven by computers, are near to passing regulatory hurdles at the time of writing (2022). However, this visible progress is based on software and hardware that use very different mechanisms from those used by humans, because those human mechanisms (e.g. identification of objects from visual input) are still beyond the reach of current computing, even after more than half a century of research. Software systems can perform well in identifying a limited, closed, set of objects in tasks such as industrial inspection (e.g. Steger et al, 2018) but do not perform so well when identifying objects in unconstrained environments; an example of the latter is autonomous cars mistaking the moon for an amber traffic light.

This combination of inherent limitations masked by visible successes brings risks. Some of these risks involve being increasingly constrained over time by initial decisions that become increasingly costly to change, as in the classic case of the QWERTY keyboard (e.g. Norman, 2013). Another set of risks involves missed opportunities. These are much harder to spot, but can be more costly than visible constraints.

We will revisit these issues at the end of this article. First, though, we will examine the underpinning concepts of information technology, and the inherent limitations of those concepts, with regard to knowledge infrastructure as opposed to information infrastructure.

## **2. INFORMATION VERSUS KNOWLEDGE: THE ISSUES**

In traditional philosophy, definitions of “knowledge” have usually focused on the justification and the truth of a proposition. In education theory, there is a long-established distinction between the abstract knowledge learned in establishments such as grammar schools and universities, and the specific information taught in e.g. technical schools and traditional polytechnics. In computing, the term “knowledge” tends to be used with a more formal definition, as the top tier of a three layer classification (e.g. Stair and Reynolds, 1998).

If we look at the successes of current digital technology, they are predominantly in the areas of data processing and information processing. The knowledge level has proved much harder to handle via digital technology.

**Table 1: Knowledge, information and data**

Level	Description, and implications
Knowledge	Information about information, e.g. "British houses quite often have names rather than numbers"
Information	Structured data, e.g. "House number = 8" or "Postcode area = TF1"
Data	Raw numbers and alphanumeric values, e.g. "8" or "TF1".

The history of digital technology can be broadly described in terms of two large waves of successful innovation relating to data and information, and a smaller wave of partially successful innovation relating to knowledge.

***The First Wave: Automating existing processes***

Many real-world problems in an industrialised society involve data and information, and are highly amenable to computerisation; for example, payrolls and commercial transactions. Scaling up regular, predictable record keeping and transaction recording was hugely successful, and became a textbook example of information systems success (e.g. Avison and Fitzgerald, 2003).

This could be called the *First Wave* of IT innovation; using computers to perform already-existing processes.

***The Second Wave: Information-level innovation***

The *Second Wave* is qualitatively different from the first wave. It involves the introduction of new types of information-level hardware and software that change the processes in use, whether by introducing new processes, or making earlier processes obsolete, or by altering earlier processes. Examples include e-commerce, email, and mobile phones, all of which have transformed business and social processes via types of information technology that were invented decades ago.

All these examples involve information-level technologies; for instance, mobile phones per se are simply a way of transferring digital information. Although the mobility of mobile phones has had far-reaching social implications, the technology itself is still information-level, rather than knowledge-level.

***Wave 2.2: Tackling knowledge problems via information-level technology***

This wave is also qualitatively different from the first wave. It involves using information-level technology

to tackle problems that are difficult for humans, often by using different methods from the ones that humans would use. The results have been mixed, in various ways.

One example is chess. Chess had long been viewed as a complex game that required intelligence and the ability to plan multiple steps ahead, in ways that would be extremely difficult to automate. When the 18<sup>th</sup> century chess player Philidor played three simultaneous blindfold games, this was viewed at the time as an astonishing intellectual feat. However, empirical investigation of skilled chess players found instead that a key part of their expertise consisted of knowing huge numbers of configurations each involving only a small subset of pieces (de Groot, 1965). Now, within a few decades of those early systems, the top chess rankings in the world are held by software systems (French, 2012; Hassabis, 2017).

Similarly, early medical expert systems such as MYCIN performed better than human experts in the relevant domain. However, medical expert systems have not been widely adopted. Fault diagnosis systems of similar complexity, in contrast, have been widely adopted, usually as embedded software within devices such as engines (Lo Bello et al, 2019).

An important feature of successful expert systems and of successful fault diagnosis systems is that they used alphanumeric input, either from a human intermediary in the case of expert systems, or via a human intermediary and/or sensor readings in the case of fault diagnosis systems. Attempts to move beyond these input methods soon encountered serious problems from two sources. One was object identification; the other was real world knowledge. In this article, we focus on the latter and its implications.

Real world knowledge is a long-standing problem for AI. The classic attempt to handle real world knowledge in software, the CYC project, dates from the 1980s (e.g. Lenat et al, 1986), and brought home a realisation of how huge the body of facts is within a typical human brain. After thirty years of development, CYC contained over 1.5 million pieces of information, but was still only effective within tightly bounded domains (Knight, 2016). The human brain, in comparison, contains tens of billions of neurons, each with about a thousand connections to other neurons (French, 2012). The massively connectionist nature of the human brain makes it possible to learn huge volumes of real world knowledge with only limited instruction, though it typically takes about twenty years before a human learns enough to be accepted as a full member of society.

Recent advances in IT have been achieved by working round this problem, rather than solving it.

In this article, we describe a way of partitioning tasks between humans and software so that both can play to their strengths, rather than trying to solve problems via IT alone.

### 3. WORKED EXAMPLES

We use two worked examples. The first describes a small scale problem that illustrates some key issues. The second describes a larger scale framework that uses a process model to partition tasks within a given process between humans alone, IT alone, and a combination of humans and IT.

#### 3.1. EXAMPLE 1: SEMI-TACIT KNOWLEDGE

Our first worked example involves support in writing commercial business cases. This is a non trivial problem that is difficult because of the nature of human real world knowledge.

A core problem in this and related tasks involves eliciting a person's goals and values and other semi-tacit knowledge. This is a difficult problem that has been handled within requirements engineering and market research by laddering (Hinkle, 1965; Reynolds and Gutman, 1988; Rugg and McGeorge, 1995). Laddering is an elicitation method that uses a deliberately small set of verbal probes to elicit goals and values in a structured way, whether as a hierarchy or as a network in terms of graph theory.

The typical form of laddering used in this context involves presenting the participant with two options, then asking which option the participant would prefer and why.

Laddering can be automated to at least some extent, though it is difficult to automate well, because the process is highly interactive and depends on rapid judgments by the elicitor that depend on real world knowledge (Rugg & McGeorge, 1995).

Another possible way of eliciting goals and values that has better chances of working via software is chatbots. These are software systems that can conduct natural language conversations, often via Artificial Intelligence (Eeuwen, 2017; Griol et al., 2013; Shawar and Atwell, 2007). The underlying concept traces back to the ELIZA program (Weizenbaum, 1976) which mimicked interaction types such as Rogerian non-directive therapy. A significant advantage of this underlying framework from a knowledge infrastructure viewpoint is that it does not depend on as rigorous a representational infrastructure as laddering, so the software can more easily sidestep potentially problematic issues.

The screenshot and dialogue sample below show an example of such a system, developed as part of a student project by one of the authors (Guo, 2019).

The system was designed to help business start-ups work through their business plan. Business plans are required by most funding and start-up support organisations (Anon, 2021). Most start-ups find it difficult to write a good business plan, particularly with regard to articulating clearly the key features of what they are selling.

The system shown here used the AIML Extensible Markup Language on the pandorabots open source platform. It asked a limited set of open-ended questions that were triggered by keywords in text from the user.

Figure 1 shows the appearance of the interface. A sample section of interaction between the system and a user is shown below. The system's language is designed to mimic human language, following widely used methods such as minor punctuation errors (e.g. French, 2012).




Figure 1: The chatbot interface

### Sample dialogue

SYSTEM: Right, what type of individuals and businesses are your target audience

USER: people who like to go to the cinema so a range of people eg 16 to 50

SYSTEM: Tell me more about the characteristics of this audience

USER: outgoing and young

SYSTEM: What will you do to make sure that you meet the people who like to go to the cinema so a range of people eg 16 to 50's needs?

USER: give a range of services like food and drinks as well as a variety of films to watch

In knowledge infrastructure terms, software methods such as the chatbot above can be used for an initial pass, helping of users to do a first iteration to help them clarify and articulate their semi-tacit goals and aims. Using software enables the first iteration to handle large numbers of users simultaneously at low cost.

For many users, the first pass will be enough. For others, a second pass will be involved. This is the point where human experts can be brought in. The number of users in the second pass will be lower, making it possible for the human experts to spend more time with each user, and to use that time more efficiently, since the users at this point will have a better idea of what they want.

This example appears simple, but it skirts round the edges of problems that IT is unable to handle properly. For example, if the user wanted an in-depth explanation of why the chatbot was asking a particular question, a proper response would involve levels of natural language processing and real world knowledge that are far beyond the ability of any current technology. As part of a bigger process, this approach could work for a specified task, but only within the limitations of what IT can currently achieve. In the next section, we look at how to structure a bigger process systematically.

### 3.2. WORKED EXAMPLE 2: A PROCESS MODEL

Our second worked example involves choices during an undergraduate degree. We begin with choice of modules within the course, and then consider the broader question of choice of career after the degree.

At first sight, choice of modules looks like a straightforward IT problem, involving putting the relevant information, including help files, online, having human advisors available, and then letting students select modules online.

In reality, the problem is deeper, and not properly solvable via IT alone, for reasons identified by Maiden & Rugg (1996). In brief, these involve

various types of tacit and semi-tacit knowledge, which the individual cannot access via introspection, but which can to some extent be accessed via appropriate choice of elicitation methods. The issues are summarised in Table 2.

The Maiden & Rugg framework maps each of these types of knowledge onto appropriate elicitation methods. For example, the contents of Short Term Memory can be accessed via concurrent think-aloud; Taken For Granted knowledge can be accessed via direct observation or via downward laddering; back versions can be accessed via indirect observation or projective methods.

**Table 2: Types of knowledge**

Category	Description, and implications
Future systems knowledge	Knowledge about a product or system that has not yet been developed. Since it has not yet been developed, the "knowledge" here is prediction rather than knowledge, and may be faulty.
Explicit knowledge	Knowledge that the individual can access directly and at will (e.g. naming the capital of France)
Semi-tacit knowledge	Knowledge that the individual can access only in some circumstances, such as recognising the names of French towns that they could not consciously recall otherwise
Strictly tacit knowledge	Knowledge that the individual can use, but that they cannot consciously access, such as "gut feel" knowledge, or "muscle memory" skills.

With this framework of knowledge types, it is possible to identify types of knowledge involved in choice of module, and then partition them between humans and technologies, as shown in Table 3.

For brevity, we have shown how this can be done for three types of semi-tacit knowledge, to illustrate the principle. Other types of semi-tacit and tacit knowledge can also be handled in the same way.


A key point is that the relevant knowledge can often be handled with very little effort or cost. For example, an image on promotional material or a wall poster can show a lot of taken for granted or not worth mentioning knowledge that the student can assimilate via incidental learning, without necessarily being consciously aware that they are doing so.

**Table 3: Implications of knowledge types**

Semi-tacit knowledge type	Implications for module choice
<i>Taken for granted knowledge:</i>	The official description of the module may take for granted and therefore not mention factors that are important to a student (e.g. <i>The field trips on this module involve air travel to other countries</i> ).
<i>Not worth mentioning knowledge:</i>	The official description of the module may view some knowledge as too trivial to be worth mentioning (e.g. <i>This module will involve giving a short presentation to other students</i> ).
<i>Front and back versions:</i>	Student choice of modules may be swayed by e.g. the personality of the staff member delivering the module, or the module's reputation for being well delivered. Such factors would only appear as back versions in off the record conversations; they would not appear in the online module information.

Figure 2 shows how this approach can be mapped onto a process model, kept deliberately simple for clarity, and extended to handle career choice. In this model, the options are shown in rounded rectangles, and the decision points in diamonds. For clarity, we have not shown the decision routes from points 3 and 4. We will first describe the basic process shown in the diagram, and then discuss questions arising from attempting to represent this model rigorously in a diagram.

The diagram starts with *School*. At the end of the individual's school education, they reach *Decision Point 1*. The choice here is between *College*, *University* and *Job*. This is a point where knowledge needs to be available. The framework above is one way of checking that all the relevant types of knowledge are provided by one or more routes (e.g. human, or human plus software).



**Figure 2: Decision points for career choices**

The *College* and *University* options both have a fixed duration. At the end of that time, the students are then faced with the same set of options as before, namely *College*, *University* or *Job*. In the diagram above, we have shown these options from the end of *College*, in *Decision Point 2*.

The table below shows a hypothetical example of the table for *Decision Point 1*. We use “World 3” in Popper’s sense of knowledge recorded in text etc as opposed to knowledge in the brain (Boyd, 2016).

**Table 4: Location of knowledge for Point 1**

	Location of knowledge		
	User	Others	World 3
Future systems	Chatbot	Careers support	
Explicit			Online text
Front/back		Student guides	
Semi-tacit		Admissions tutor	“Day in life” online
Strictly tacit			Online multimedia

The *Future Systems* knowledge about the student’s goals and aims in this example is handled by a combination of chatbot and Careers Support. The chatbot is used for a first pass, helping the student to clarify their own goals and aims at their own pace; the Careers Support staff then provide more focused support, using methods specifically chosen to handle the semi-tacit and tacit knowledge problems involved.

This arrangement means that a lot of the work is being done by chatbot, cheaply and swiftly, but with human beings to handle the issues that the chatbot cannot handle. A key point is that the humans are not only using their real world knowledge; they are also using appropriate methods for the various knowledge types involved.

The various types of knowledge in this model are handled in several ways. In addition to the Careers Support team, there are students acting as guides. The student guides provide incidental learning and access to back versions via their shared experience as fellow students, as well as using methods such as laddering to clarify the goals of the students choosing modules. There is also an admissions tutor trained in the same methods, with different experiential knowledge from the student guides as a source of insight and support.

This arrangement provides support in depth, so if the Careers Support team provide low quality

support, the student has a second strand of support via the students acting as guides, and a third strand via the Admissions officer. Although support in depth has advantages, it also has well known disadvantages, such as the risk of contradictory advice from the different strands, and the risk of key knowledge falling between the cracks because each strand thinks that a particular piece of knowledge is being handled by another strand. The tabular representation helps identify where such risks occur, so that the system designers can make appropriate provision.

The act of producing a visible representation of the process can itself produce significant insights to the system designers. For instance, in the diagram above, the decision points are all shown immediately to the right of an option, implying that they occur after that option finishes. However, other designs are possible, such as having the decision point occur during the option (e.g. students making the decision early in their final semester) or overlapping with it.

Similarly, in this simple model, at the end of the College option, Decision Point 2 offers the same set of options as Decision Point 1 (College, University or Job). One system design choice would be to use the same matrix of knowledge type/location as at Decision Point 1, since the options available from both are the same. This would have the advantages of economy and consistency, by re-using existing resources. However, another possible choice would be to use a new matrix for Decision Point 2, since students finishing college will be different from students finishing school, in age, in real-world experience, and in qualifications or significant absence of qualifications.

The representations in these examples do not in themselves provide the answers, but they do help identify questions that might otherwise easily be overlooked, and they do make it easier to plan a viable, effective system.

Our experience with this approach indicates that it can make a significant difference at minimal cost.

An example involves a student clarifying his career options with one of the authors (Rugg) via upward laddering. The student was trying to decide between two management-track jobs that he had seen advertised. The first laddering probe, *Which would you prefer and why?* elicited that he would prefer one of them because *it paid more*. Laddering up from that value elicited the response that he would prefer the job that paid more because *it would let me travel*. At that point the student realised that he had been going for well paid jobs as a means to an end, and that he could go straight to that end by finding a job that involved travel, which he would much prefer to a management job.

This type of response is common in the laddering literature, and in our experience. It means that a very swift, minimal-cost intervention can significantly improve a person's career choice process by clarifying key issues that were lurking in semi-tacit knowledge.

#### 4. IMPLICATIONS

The framework above provides a practical, low-cost way of handling knowledge-level problems via a systematic use of existing technology and human expertise combined with appropriate choice of methods for eliciting and representing knowledge. A key feature of this framework is its grounding in the research literature about human knowledge, and the implications for eliciting knowledge from humans.

Although the worked example above may look superficially the same as what is currently done in many institutions (e.g. "We already have someone whose role is to help students decide on their options") a key, major difference is that this framework explicitly provides ways of accessing key semi-tacit and tacit knowledge that would otherwise be missed.

This framework also provides an understanding of the underlying problems that limit the growth of information-level technologies. Those limits to growth are already becoming apparent, and need to be tackled soon, for several reasons.

One reason involves the risk of becoming locked in to a solution which works well at the time of adoption, but whose limitations become more costly with time, as in the case of the QWERTY keyboard, which was a reasonable choice in the days of manual typewriters, but which is less efficient than several other layouts in the days of electronic keyboards (Norman, 2013). A current example is the trend towards using Internet plus mobile phone as the default way of handling information-level problems, which is efficient for most users, but which can easily marginalise or exclude entire communities (e.g. Ramsetty and Adams, 2020).

A related risk involves logical incrementalism, where a series of decisions occur over time, each of which is sensible in terms of immediate context, but where the cumulative result is far from optimal, such as the layout of many old European towns. It is easy to identify examples from early digital technology, such as old ways of tackling the explosive growth of telegraph and telephone. It is less easy to identify this in present-day technology, because we can't know what we don't know; there is often no way of knowing for certain whether a given situation is the best outcome that could have been achieved given the complexities of real world systems. However, it is possible to make predictions based on best evidence, and to make

provision for reviewing evidence after time and to build in points where decisions can be changed before the process has gone too far.

A third and particularly interesting risk involves missed opportunities. Again, it is easy to identify historical examples; for instance, if a viable system of patent law had been invented earlier, it would almost certainly have increased the speed of technological innovation. Again, it is less easy to identify present-day examples. However, the framework described here provides a tool for identifying opportunities for making significant improvements to knowledge infrastructure swiftly and easily.

A fourth risk relates to the nature of education. Information Technology is well suited to making information available online. The temptation is to treat education as information, and to follow the easy and obvious route of packaging education into online multimedia modules using the same business model as in online entertainment. Education, however, is not solely about information; it is also about knowledge, and the nature of knowledge means that specific types of teaching and learning are required for different types and different components of knowledge. For handling knowledge, information-based approaches are simply not enough; a systematic suite of other approaches need to be used to complement the IT.

## 5. CONCLUSION

Information Technology has transformed the world, but has inherent limitations. This article describes issues involved in handling *knowledge* as opposed to *information*. These issues can be handled efficiently, simply, and at low cost via a systematic combination of humans and technology, guided by a knowledge infrastructure framework such as the one above.

We plan to apply this approach in a range of case studies. One involves guidance on choice of career, helping people to clarify their semi-tacit goals and aspirations via methods such as card sorts and upward laddering, so that they can make better use of existing career guidance facilities.

Another involves development of improved communication between medical professionals and the public, including patients, via representations that reduce common sources of misunderstanding and error. This builds on work by one of the authors (Skillen 2017, Skillen 2019) that extends previous work by e.g. Kahneman et al (1982) and by Gigerenzer (2008) by drawing on a wider range of elicitation and representation types to clarify key points.

## 6. REFERENCES

- Anon (2021) *Write a business plan*. Available at: <https://www.gov.uk/write-business-plan> (accessed 16 May, 2021).
- Avison, D. and Fitzgerald, G. (2003) *Information Systems Development: Methodologies, Techniques and Tools* (3<sup>rd</sup> edition). McGraw-Hill, London.
- Boyd, B. (2016) Popper's World 3: Origins, Progress, and Import. *Philosophy of the Social Sciences* 46(3): 221–241.
- Dampney, C.N.G, Busch, P. and Richards, D (2002) The meaning of tacit knowledge. *Australasian Journal of Information Systems*, Special Issue, January 2002. Cited in: Stair, R.M. and Reynolds, G.W. (1998). *Principles of Information Systems: A Managerial Approach* (3rd edition). Boyd and Fraser, Danvers, Mass.
- Davenport, T. and Fitts, J. (2021) AI Can Help Companies Tap New Sources of Data for Analytics. *Harvard Business Review*, March 19, 2021.
- de Groot, A.D. (1965) *Thought and Choice in Chess*. Mouton, the Hague.
- Eeuwen, M.V. (2017) *Mobile conversational commerce: messenger chatbots as the next interface between businesses and consumers* (Master's thesis, University of Twente).
- Forsyth, R. (1981) BEAGLE: A Darwinian Approach to Pattern Recognition. *Kybernetes* 10: 159–166.
- French, R.M. (2012) Moving beyond the Turing Test. *Communications of the ACM*, 55:12.
- Gigerenzer, G. (2008) *Gut Feelings: Short Cuts to Better Decision Making*. Penguin, London.
- Griol, D., Carbó, J. and Molina, J.M. (2013) An automatic dialog simulation technique to develop and evaluate interactive conversational agents. *Applied Artificial Intelligence*, 27(9), pp. 759–780.
- Guo, Y. (2019) *A business evaluation chatbot (bebot) acting as an interactive virtual agent prototype*. Unpublished undergraduate dissertation, Keele University, UK.
- Hassabis, D. (2017) Artificial Intelligence: Chess match of the century. *Nature* 544: 413–414.
- Hinkle, D. (1965) *The change of personal constructs from the viewpoint of a theory of construct implications*. PhD Thesis, Ohio State University, USA.
- Kahneman, D., Slovic, P. & Tversky, A. (Eds.) (1982) *Judgement under uncertainty: Heuristics*

- and Biases*. Cambridge University Press, Cambridge.
- Knight, W. (2016) An AI with 30 years worth of knowledge finally goes to work.  
<https://www.technologyreview.com/2016/03/14/108873/an-ai-with-30-years-worth-of-knowledge-finally-goes-to-work/>  
(accessed May 2021).
- Lenat, D., Prakash, M., and Shepherd, M. (1986) CYC: Using Common Sense Knowledge to Overcome Brittleness and Knowledge Acquisition Bottlenecks. *AI Magazine* 6(4): 65–85.
- Lo Bello, L. Mariani, R., Mubeen, S. and Saponara, S. (2019) Recent Advances and Trends in On-Board Embedded and Networked Automotive Systems. *IEEE Transactions on industrial informatics*, 15(2) 10038-10051.
- Maiden, N.A.M. and Rugg, G. (1996) ACRE: a framework for acquisition of requirements. *Software Engineering Journal* 11, 183–192.
- Norman, D. (2013) *The Design of Everyday Things (revised and expanded edition)*. The MIT Press, London.
- Polanyi, M. (1967) *The Tacit Dimension*. Routledge and Kegan Paul, London.
- Quinlan, J.R. (1986) Induction of Decision Trees. *Machine Learning* 1(1): 81–106.
- Ramsetty, A. and Adams, C. (2020) Impact of the digital divide in the age of COVID-19. *Journal of the American Medical Informatics Association*, 27(7): 1147–1148.
- Reynolds, T.J. and Gutman, J. (1988) Laddering theory, method, analysis, and interpretation. *Journal of Advertising Research* 28: 11–31.
- Rugg, G. and McGeorge, P. (1995) Laddering. *Expert Systems* 12: 339–346.
- Shawar, B.A. and Atwell, E. (2007) Chatbots: are they really useful?. *Ldv Forum* 22(1) pp. 29-49).
- Shortliffe, E.H. and Buchanan, B.G. (1975) A model of inexact reasoning in medicine. *Mathematical Biosciences* 23(3–4): 351–379.
- Skillen, J.D. (2017) The symptom matrix: Using a formalism-based approach to address complex syndromes systematically. *Musculoskeletal Care* 15, 253–256.
- Skillen, J.D. (2019) The Friends and Family Test: From card sorts to control charts. *Methodological Innovations*. 12(2).
- Standage, T. (1998) *The Victorian Internet: The Remarkable Story of the Telegraph and the Nineteenth Century's On-Line Pioneers*. Walker and Co, London.
- Steger, C, Ulrich, M and Wiedemann, C (2018) *Machine Vision Algorithms and Applications* (2nd ed.). Weinheim: Wiley-VCH.
- Weizenbaum, J. (1976). *Computer Power and Human Reason: From Judgment to Calculation*. W. H. Freeman and Company, New York.
- Werbos, P. (1982) Applications of advances in nonlinear sensitivity analysis. In: Drenick, R.F., Kozin, F. (eds) *System modeling and optimization*. Springer, London pp 762–770.  
<http://werbos.com/Neural/SensitivityIFIPSeptember1981.pdf>  
(accessed 2 March, 2020).
- Wootton, D. (2015) *The invention of science: A new history of the scientific revolution*. Penguin, Harmondsworth.