Chapter 3
In Humans We Trust: rules, algorithms and judgment

Hamid R. Ekbia

1. Introduction

Modern humanity is facing numerous challenges in all arenas of life – from the economy, politics, culture and the environment to law, healthcare and governance. Cracks are appearing on the seemingly consistent surface of social reality, revealing its irremediable fragility and evoking the image of an ice sheet during a thaw. Many of these have escalated to the level of crisis in the sense that we experience a constant presence of negativity in our lives. Humanity has accomplished much, but despite – or, depending on your outlook, perhaps because of – this, it has left many problems untouched. Not only that, every step of the way it has added new predicaments for human beings, for other species and for our globe not least through the reckless hubris emanating from its marvellous prowess in scientific and technical discovery. Individuals and communities face this situation in different ways, listening to experts and commentators for guidance. It is in this environment that AI has emerged as either the *cause célèbre* of these crises or as the solution to them, depending on whose views, whether of AI detractors or AI prophets, you take seriously. Both of these miss the mark by large margins because AI does not cause these issues, nor can it cure them on its own.

The detractors, those who see in AI the root cause of our predicaments, recommend withdrawal from technology. Their idea is that we can choose to disconnect from technology at will with little or no social cost or, even more naively, that taking leave from technology can help us deal with the growing number of challenges that we face as individuals and communities by falling back on our 'good old ways'. Little do the proponents of such ideas understand that withdrawal is an unrealistic option and that taking nostalgic refuge in a bygone era will not save the day for us. There was no such thing as an 'ideal' past where things were presumably in apple pie order, nor could the old ways work for our current issues, even to the extent that they did work in the past.

The prophets and devout advocates of AI, on the other hand, see in it an invisible hand, all too ready to be extended outward, with a purported magic ability to make things work. You have too many road casualties? Self-driving cars will reduce or eliminate them. Mysterious diseases? AI diagnostics can reveal them. Too many applications for jobs, admissions, benefits or parole? AI can screen them. Too many items on your daily agenda? AI assistants can sort them. Too long a history of legal cases for litigation? Computerised clerks can filter them. Fraudulent actors? AI detection systems can reveal their hand. Looking for voters to support your political campaign? AI tools can convince them. Yearning for romance? AI can find the perfect partner, even giving you the bots and robots that would play that role themselves – satisfaction guaranteed! Perceived in this way, AI appears to be much more effective than its modern predecessor, the miraculous market, to take care of our issues. If that invisible hand could solve just one problem – coordinating buyers and sellers of commodities – AI can solve a whole slew of social, economic, political, professional and personal problems: essentially, anything under the sun that would lend itself to 'pattern recognition' and the logic of computing.

A case in point is the management and allocation of social benefits to different individuals and social groups. Local and national governments increasingly rely on so-called 'smart algorithms' to do this for them. A recent report by the Electronic Privacy Information Center, for instance, revealed the extensive use of algorithmic decision-making by government agencies in Washington DC, adding up to 29 automated systems used by 20 agencies. At the same time, the report highlights numerous loopholes and mistakes present in these systems – for instance, a 93% error rate in an unemployment fraud detection system that led to 40,000 false cases and 1.1 million false flags (Johnson 2022). Washington DC is just one example among many; software systems are currently used for all kinds of purposes by government agencies on different levels in the name of efficiency and objectivity.

Government agencies are not alone in their penchant for efficiency. Similar other reasons are invoked for the increasing application of AI and algorithms in other domains of social life. Prominent among these are accuracy, reliability, neutrality and transparency, which are often played, implicitly or explicitly, against such human fallibilities as social and cultural bias, cognitive limitations, emotional caprice and ethical evasiveness. Road accidents, for instance, kill 1.35 million people around the globe every year (CDC 2023) because drivers are distracted, exhausted, unskilled, under the influence or they simply make mistakes; medical errors due to misdiagnosis, drug interactions or sheer neglect annihilate another 2.6 million people every year in low and middle-income countries around the globe (WHO 2019), more than 250,000 in the US alone (Anderson and Abrahamson 2017); even judges, who are expected to be 'objective' in their views, are shown to be subject to 'judicial temperament', ruling differently at different times of the day because of their mood which might, in turn, depend on their physiological state (Danziger et al. 2011). Human beings are indeed fallible, biased, lousy, clumsy, capricious and evasive – and hence they are unreliable.

It is against the backdrop of such observations that advocates of AI propose technical alternatives as solutions to the increasing predicaments of modern societies. Given the numerous sources of human fallibility, they ask, why should we trust human beings with our social, personal and professional issues and not AI-enabled technologies that are not only devoid of bias and emotion but also more exact and efficient? In asking this question, such proponents are comparing, directly or indirectly, the reliability of rule-based algorithms with the unruliness of human judgment. Rules, in other words, play a central role in this comparison. But rules come in different varieties. In the examples mentioned earlier, the rules of traffic are different from the rules of medical practice, and both are yet more different from the rules of judicial decision-making. This brings up an important question: what are 'rules' and how have they come to play such a central role in human affairs?

2. Varieties of rules

Rules come in many varieties from rules of games, grammars and poetry to the etiquettes of behaviour in a school, a bus or an office, and from government regulations, laws of science and statutes of war to cooking recipes, design guidelines and computer algorithms. Behind this diversity of forms and contexts, there is a long history with discernible continuities and discontinuities.

In her book, the science historian Lorraine Daston (2022) identifies three broad clusters or kinds of rules that have operated throughout history across cultures: tools of measurement and calculation; models and paradigms; and laws. The first cluster includes rules followed over the centuries in the arts and crafts such as painting, cooking and baking, and even the medieval practices of artillery and fortification. The second cluster spans the traditional rules of conduct according to exemplary behaviours of specific individuals all the way to the more recent notion of 'paradigm' in the historical development of modern science (Kuhn 1962). This was, for instance, how rules were understood in the monastic Rule of St. Benedict, where the authority to apply the codes of conduct was vested in the abbot who was considered a rule, or a 'model', to be emulated by others. The third cluster covers the laws and regulations that govern social behaviour based on sanctions, for example the tax laws of ancient cities and medieval fiefdoms or the traffic laws of modern societies.

The historical development of rules can be understood as the interplay between these clusters on three oppositional dimensions which Daston describes as thin-thick, flexible-rigid and general-specific. Thin rules are concise descriptions and 'commands' that can be followed to the letter, without the need for intelligent interpretation and adaptation to context. They are, as such, 'exceptionless and infallible codes of conduct' (Chirimuuta 2023). Thick rules, on the other hand, are context-dependent and require extra capacity to see when and how best to apply them. But then that same capacity should be governed by a higher set of rules that would determine what type of context we are dealing with and this, in turn, would demand yet another rule which itself begs a rule of higher order … ad infinitum. This chain of rules leads to an endless regress and a paradox that has been known to philosophers for a very long time. In the eighteenth century, Immanuel Kant gave voice to this in the following manner:

> If the understanding in general is explained as the faculty of rules [Regeln], then the power of judgment is the faculty of subsuming under rules, i.e., of determining whether something stands under a given rule (*casus datae legis*) or not. General logic contains no precepts at all for the power of judgment, and moreover cannot contain them. For since it abstracts from all content of cognition, nothing remains to it but the business of analytically dividing the mere form of cognition into concepts, judgments, and inferences, and thereby achieving formal rules for all use of the understanding. Now if it wanted to show generally how one ought to subsume under these rules, i.e., distinguish whether something stands under them or not, this could not happen except once again through a rule. But just because this is a rule, it would demand another instruction for the power of judgment, and so it becomes clear that although the understanding is certainly capable of being

> instructed and equipped through rules, the power of judgment is a special talent that cannot be taught but only practiced. (Guyer and Wood 1998)

Kant then talks about a physician, a judge or a statesman who can have many fine pathological, juridical or political rules in their head but who can easily stumble in their application. He describes this lack of power as that which is properly called stupidity, reminding us that such a failing is not to be helped. What, then, distinguishes between a competent physician, judge or statesman and someone who is a mere stockpile or kitbag of rules? Or, by the same token, what distinguishes between a skilled craftsperson such as a master chef and a newbie who has never touched a spatula, between a master and a novice chess player, and between a good and a clumsy musician? Or, even more fundamentally, what is the difference between the conduct of an ordinary human being who adjusts to the demands of situations as they arise, with or without explicit knowledge of the 'rules', and someone who behaves 'mechanically', clueless 'like a robot' – to use modern parlance – in describing a rigid adherence to rules?

On the surface, these are different questions that speak to various domains of human behaviour – that is, expertise, skill and common sense. Deep down, however, they are related questions having to do with how some humans are better than others at applying, adjusting and following rules, not just having a knowledge of them. Ludwig Wittgenstein made an important distinction between following a rule and merely acting according to a rule. Rules, he argued, are not just causes that determine our behaviour in the same way that the law of gravity determines the behaviour of a planet around a star. Rather, they have a normative force because they provide measures of correct and incorrect behaviour. They are not simply dispositions that can drive behaviour like a mechanism; they have to be practised through experience. That is why we can invoke rules to explain and justify ourselves: 'To obey a rule, to make a report, to give an order, to play a game of chess, are customs (uses, institutions)' (Wittgenstein 1953: § 199).

This is how rules-as-models have been understood for an exceptionally long stretch of human history. The model as exemplar for following the rule could be a person, a work of art or simply a well-chosen example in grammar or algebra (Daston 2022: 8). Putting a high premium on experience and practice, such rules brought the head and the hand together, creating a balance between strict discipline and adherence to rules, on the one hand, and, on the other, the creativity and flexibility called for in applying them. A cooking recipe, for instance, leaves room for human experience in judging when the batter is thick, the flour is cooked enough or the cheese spread evenly on the crust.

To talk about these aspects of rules, Daston (2022: 36-38) draws on the notion of 'discretion' as a form of judgment 'which embraces not only knowing when to temper the rigor of rules but also matters of taste, prudence, and insight into how the world works, including the human psyche'. Discretion, as such, 'combines intellectual and moral cognition… but [it] goes beyond cognition… [It] is a matter of the will as well as the mind… Cognitive discretion without executive discretion is impotent; executive discretion without cognitive discretion is arbitrary'.

This way of thinking about rules-as-models, involving the moral capacity to exercise discretion, has been gradually edged out in modern times in favour of the more calculative notion of rules-as-algorithms. This shift started with, among other things, the introduction throughout the nineteenth century of mechanical machines that could perform some of the tasks of human beings and which culminated in the digital automation techniques of the late twentieth century. But it was not only the meaning of 'rules' that underwent a shift. In the process, the meaning of a whole slew of other related terms also underwent significant change. The concept of 'algorithm' was one such notion.

3. Varieties of algorithms

The term 'algorithm' goes back to the ninth century CE when the Persian mathematician, astronomer and geographer, Musa Al-Kharizmi, wrote a treatise on algebra, Indian numerals and astronomical tables. This text, which is considered the founding document of modern algebra, was translated into Latin in the twelfth century, and the earliest surviving manuscript in the Latin text starts with the words 'Dixit Algorizmi' ('Thus spoke Algorizmi'). That's how Al-Kharizmi's name was transformed into the notion of algorithm.

Early on, algorithms had to do with the solution of specific problems, not abstract generalities – for instance, calculating the length of the lunar month or the square and cubic roots of integers. All such calculations were carried out by people, of course, using pencil and paper (or something equivalent). Of special significance here is that the head and the hand were both involved in doing the calculations. In the industrial era, too, almost all the way to the mid-twentieth century, algorithms were applied by astronomers, insurance companies, census bureaus and weapons projects on an industrial scale, and this was done by humans and machines working in tandem.

It was only in the second half of the twentieth century when algorithms became 'automated' and understood as similar to recipes and procedures that can be followed through failsafe computation. That is how Donald Knuth, who wrote the bible of computer science, described them except that he added five features for computer algorithms: 'finiteness, definiteness, input, output, effectiveness' (Knuth 1973). Without getting into the details of these terms, essentially an algorithm is understood as an effective method that can be expressed and executed in a finite time and space. By adding these requirements, Knuth highlighted that the medium of computation matters. For, in principle, the medium on which an algorithm is implemented should not be relevant – one can carry out the same algorithm on paper, on a mechanical device or a digital computer, even with a set of wooden tokens – except that, for most practical and interesting purposes, the finite-time-and-space constraint would exclude all but the digital alternative (and any other that would surpass their speed and efficiency – for example, quantum computing).[1] In short, only algorithms that can produce a result in finite time and space are practically useful.

This unique advantage of the digital medium enabled the automation of algorithms in ways that were inconceivable earlier, launching a new era that has brought us to the current moment of algorithmic decision-making which seeks to replace not only human judgment in the sense of discretion but also the laws and regulations that govern our social relations. ‘By driving the exercise of discretion underground’, Daston argues, ‘rules-as-algorithms blow up the bridges that connected universals to particulars in rules-as-models’ (2022: 21). The question facing us is what bridges are blown up now that rules-as-algorithms are driving underground in addition the exercise of social norms and laws?

To address this question, we need to step back and take a closer look at the development of algorithms in recent decades.

3.1 Expert systems

The automation of algorithms in early computing is best exemplified in the development of ‘expert systems’ in AI. By the late 1960s and early 1970s, a decline of interest in AI research on the part of funding agencies had led practitioners to look for practical problems to solve.

As knowledge was conceived to be the key to such endeavour, a new class of artefacts called expert systems then appeared on the scene. Mycin, for instance, was one of the first expert systems for the diagnosis of infectious blood diseases. It used ‘production systems’ – that is, a set of ‘if then’ rules, like the following, along with the mechanisms for deciding when and how to apply the rules separated from the set of rules themselves.

IF
the site of the culture is blood, AND the gram strain is positive, AND
the portal of entry is gastrointestinal tract, AND
the abdomen is the locus of infection, OR
the pelvis is the locus of infection

THEN
there is strongly suggestive evidence that Enterobacteriaceae is the
class of organisms which therapy should cover.

1. Quantum computing uses the principles of quantum mechanics to perform calculations. Specifically, it uses specialised hardware to manipulate ‘qubits’ — the equivalent of a ‘bit’ in digital computing. Like a bit, a qubit can be in one of two states but, unlike bits, qubits can also exist in superpositions of those states. Furthermore, they can also be entangled with other qubits. Roughly speaking, entanglement is the result of non-local correlations among the parts of a quantum system. What this means is that we cannot fully understand the state of the system by dividing it up into parts and studying the separate states of the parts. Information can be encoded in non-local correlations among the parts of the system. Much of the art of designing quantum algorithms involves finding ways to make efficient use of these non-local correlations. Superposition and entanglement are the sources of power of quantum computing because, unlike the classical case, a qubit is equivalent to a vector in a two-dimensional space of real numbers (Ekbia 2008: 70).

This kind of linear reasoning borrows strong elements from logic, which relies on the explicit expression of facts and a formal modelling of the rules. The origins of this approach go back to rationalism and developments in mathematical logic in the last three hundred years (Ekbia 2008). Systems such as Mycin represent what is commonly referred to as 'symbolic AI'; symbolic because they use the formalism of symbolic logic (predicate calculus) to represent the world as a set of objects, predicates (attributes, properties) and relations – that is, as a set of symbols. The first proposition (sentence) in the above chain of reasoning, for instance, will be represented in the following way:

Box 1 The use of symbols in logic

Let P(x) be 'x is the site of the culture' and Q(x) be 'x is blood'. Then the proposition can be represented as: $\exists x\,(P(x) \wedge Q(x))$

This can be read as 'there exists an x such that x is the site of the culture and x is blood.'

Similarly, to represent the proposition 'the gram strain is positive' in predicate calculus, we can use a predicate symbol and a variable to represent the subject of the proposition. Using the predicate symbol 'P(x)' to represent 'x has a positive gram strain' gives 'x' as a variable that can take on different values depending on the context.

Therefore, the proposition 'the gram strain is positive' can be represented in predicate calculus as: P(x)

where 'x' refers to the subject being discussed such as a bacterium, a sample of biological material or a patient.

The whole production rule (chain of reasoning), for instance, might be represented as in Box 2.

Box 2 Symbolic representations of early AI

$(P \wedge Q \wedge R \wedge (S \vee T)) \rightarrow U$

where:*

P: the site of the culture is blood

Q: the gram strain is positive

R: the portal of entry is the gastrointestinal tract

S: the abdomen is the locus of infection

T: the pelvis is the locus of infection

U: therapy should cover Enterobacteriaceae

$\wedge$: AND

$\vee$: OR

$\rightarrow$: implication

* The symbol → in logic is called 'implication' or 'conditional'. It is used to denote a logical relationship between two statements, where the first statement (the antecedent) implies the truth of the second statement (the consequent). The implication symbol is often read as 'if then' or 'implies'.
Source: author's own elaboration.

Expert systems with such attributes – deep and linear inferences drawn from a small amount of information captured in a few variables – had some early success, bestowing on AI the respect it was longing for not only in academia but also in business, where billions of dollars were invested in expert systems for manufacturing, financial services, machinery diagnosis and so on. But this success was limited because the competence of such systems was restricted to very narrow domains. Two of the most notorious examples of such limitations come from a medical diagnosis program that, given the data for the case of a 1969 Chevrolet having reddish-brown spots on its body diagnosed it as suffering from measles; and a car loan authorisation program that approved a loan to a teenager who had claimed to have worked at the same job for twenty years (Ekbia 2008: 96-97). The problem, of course, was that the builders of the system had failed to include certain facts in the knowledge base: that the given Chevrolet is not a human being; and that one cannot have work experience that is longer than one's age. Another program, given the case of a patient with a kidney infection, prescribed boiling the kidney in water – a good remedy against infection, but a terrible failure to understand the basics of what was going on (Ekbia 2008: 96-97).

To state the obvious, these programs seemed to lack an important feature of intelligence: despite apparent sophistication and expertise in specialised areas like medicine, they demonstrated a clear lack of understanding of very basic facts that a human being takes for granted. They either had to be given the minutest details, as in the first two cases, or else they seemed to be missing fundamental knowledge (e.g. about living organisms), as in the third.

3.2 Algorithms in machine learning

AI researchers and practitioners tried to overcome these issues in various ways but, by and large, they failed until big data and machine learning came to their 'rescue', although with limitations of their own. Machine learning algorithms build on earlier techniques in AI and computing, such as neural networks, while breaking away from many aspects of symbolic AI.

Smith (2019: 47) describes machine learning as a suite of statistical techniques for the statistical classification and prediction of patterns, based on large quantities of data, using an interconnected fabric of processors that are arranged in multiple layers. What makes machine learning distinct from earlier techniques, such as earlier AIs, Smith adds, is its reliance on (a) shallow (few step) inference; (b) by a massively parallel process; using (c) massive amounts of information; and (d) involving a very large number of (e) weakly correlated variables.

In some rough sense, like their neural networks predecessors, machine learning systems are brain-like structures made up of heavily interconnected nodes ('neurons') that run in parallel, finding patterns in the vast amount of data that is fed into them from various sources (sensors, cellphones, social media, location tracking devices, etc.). Broadly speaking, such patterns are statistical correlations that are implicitly captured in connections between the nodes. Because they are implicit, it is hard, if not impossible, for a human observer to recognise and understand what the correlational patterns mean. Unlike the explicit propositions of symbolic AI, which were legible to a human with a basic knowledge of symbolic logic, these patterns are encoded in a 'language' that is not accessible to human beings. This, along with the distributed (multiprocessor) character of current AI systems, is the source of the opacity of machine learning algorithms, making it hard for human beings to understand how they arrive at decisions. That is why you'd be doomed if your social benefits, as a resident of Washington DC, are cut off on the basis of the patterns that a machine learning system has identified in your data. Neither you as the beneficiary, nor the government workers who use them, have access to the 'logic' behind the system's decision.

4. Varieties of Judgment

This shift in techniques from symbolic AI to machine learning, from algorithms that reasoned explicitly to black-boxed systems that we don't understand, has further complicated our relationships with AI systems, putting more power in the hands of those who design and market them and chipping away from the rights and recourses that were available to the rest of us. To see how, we need to examine the principles on which such systems are built, three of which – about social reality, about human behaviour and about our relationship with modern technology – have particular resonance.

The first principle can be described as the principle of regularity. As discussed earlier, modern times have witnessed a gradual transition from rules-as-models to rules-as-algorithms, or from thick rules to thin rules. Computer algorithms are the thinnest of rules not because they are short or simple – which they are obviously not – but because they are built on the assumption of the regularity of the outside world and, more specifically, the uniformity of the social world. This uniformity manifests itself in various ways: in the way similar algorithms are applied to various domains of life, including online trading and online dating, traffic regulation and job, loan or university entrance applications; in the way people are pigeonholed into categories based on criteria deemed relevant to powerful players such as large corporations (and sometimes governments); and in the way statistical generalisations paper over meaningful differences among personal, cultural and economic backgrounds, assuming homogeneity within each category of people. Old statistical techniques failed to capture nuance and difference because they were coarse-grained. Machine learning, equipped with the immense computational power of current technology, has refined those techniques but it has not overcome these basic problems. When it comes to human differences, fine-grained statistics fare no better than their coarse-grained predecessors.

The second principle is closely related to the first, having to do with the malleability of the human condition and of human behaviour. The origins of this go back to the behavioural psychology of the early twentieth century and its emphasis on conditioning. The cognitive revolution of the mid-century, of which AI is an intellectual heir, dislodged the behaviourist emphasis on the observability of the conditioning mechanism, pushing it to the 'subconscious' and the inner workings of the mind. But it maintained the earlier belief in the explainable character of human behaviour from an intentional standpoint

– that is, in terms of beliefs, desires and intentions. In and of itself, that belief does not incur harm to the social fabric. The issue is how current systems enabled by AI techniques seek to mould human behaviour in the image forged by corporations. Daston (2022: 5) captures this issue succinctly:

> An island of stability and predictability in a tumultuous world, no matter what the epoch or locale, is the arduous and always fragile achievement of political will, technological infrastructure, and internalized norms.

Our epoch, it seems, is increasingly resorting to technology to create not islands but continents of stability and predictability, and not on a local but on a global scale. It is not hard to see the fragility of these arrangements in a growingly tumultuous world.

Third, and perhaps most relevant to the present discussion, is the binary principle of comparing humans and machines, often in terms of the 'competitive advantage' of machines over human beings. This false binarism is behind many of the claims about the alleged superiority of AI systems in regard to accuracy, reliability, objectivity, etc. In reality, however, these systems are ultimately the product of the embedding social, economic and cultural environment in which they are developed, designed and deployed. Separating them from this environment under the rubric of 'autonomy' is as misguided as it is to think of human beings in isolation from the social environment. That environment includes, among other things, the very technologies that human societies have created throughout their history and continue to create at this particular juncture. From this perspective, there is no 'pure' human detached from technology and no technology detached from humanity. By the same token, to speak of putting humans 'back in the loop' sounds somewhat like a tautology because humans already are, and will continue to stay, in the loop in the foreseeable future. It is ultimately technologically-enabled humans who make decisions, who act on those decisions and who bear their consequences. Someday 'smart' machines may be able to go their own way, built on thick rules that can emerge from their deep embeddedness in a social environment. Until then, human judgment remains our last resort, as flawed and fallible as it might often be. In humans we should trust!

References

Anderson J.G. and Abrahamson K. (2017) Your health care may kill you: Medical errors, Studies in Health Technology and Informatics, 234, 13–17.

CDC (2023) Road traffic injuries and deaths – A global problem, Center for Disease Control and Prevention. https://www.cdc.gov/injury/features/global-road-safety/index.html#print

Chirimuuta M. (2023) Rules, judgment, and mechanization, Journal of Cross-Disciplinary Research in Computational Law. https://journalcrcl.org/crcl/article/download/22/12

Danziger S., Levay J. and Avnaim-Pesso L. (2011) Extraneous factors in judicial decisions, Proceedings of the National Academy of Sciences, 108 (17), 6889–6892. https://doi.org/10.1073/pnas.1018033108

Daston D. (2022) Rules: A short history of what we live by, Princeton University Press.

Ekbia H.R. (2008) Artificial dreams: The quest for non-biological intelligence, Cambridge University Press.
Guyer P. and Wood A.W. (eds.) (1998) Immanuel Kant, Critique of pure reason, Cambridge University Press.
Johnson K. (2022) Algorithms quietly run the city of DC — and maybe your hometown, Wired, 3 November 2022. https://www.wired.com/story/algorithms-quietly-run-the-city-of-dc-and-maybe-your-hometown/
Knuth D.E. (1973) The art of computer programming, Volume 1: Fundamental algorithms, 2nd ed., Addison-Wesley.
Kuhn T. (1962) The structure of scientific revolutions, University of Chicago Press.
Smith B.C. (2019) The promise of artificial intelligence: Reckoning and judgment, MIT Press.
WHO (2019) WHO calls for urgent action to reduce patient harm in healthcare, News, 13 September 2019. https://www.who.int/news/item/13-09-2019-who-calls-for-urgent-action-to-reduce-patient-harm-in-healthcare
Wittgenstein L. (1953) Philosophical investigations, Blackwell.

All links were checked on 23.01.2024.

Cite this chapter: Ekbia H.R. (2024) In Humans We Trust: rules, algorithms and judgment, in Ponce del Castillo A. (ed.) Artificial intelligence, labour and society, ETUI.