



The necessity of judgment

Jeff Malpas¹

Received: 15 July 2020 / Accepted: 16 July 2020 / Published online: 24 July 2020
© Springer-Verlag London Ltd., part of Springer Nature 2020

In 2016, the Australian Government launched an automated debt recovery system through Centrelink—its Department of Human Services. The system, which came to be known as ‘Robodebt’, matched the tax records of welfare recipients with their declared incomes as held by the Department and then sent out debt notices to recipients demanding payment. The entire system was computerized, and many of those receiving debt notices complained that the demands for repayment they received were false or inaccurate as well as unreasonable—all the more so given that those being targeted were, almost by definition, those in already vulnerable circumstances. The system provoked enormous public outrage, was subjected to successful legal challenge, and after being declared unlawful, the Government paid back all of the payments that had been received, and eventually, after much prompting, issued an apology.

The Robodebt affair is characteristic of a more general tendency to shift to systems of automated decision-making across both the public and the private sector and to do so even when those systems are flawed and known to be so. On the face of it, this shift is driven by the belief that automated systems have the capacity to deliver greater efficiencies and economies—in the Robodebt case, to reduce costs by recouping and reducing social welfare payments. In fact, the shift is characteristic of a particular alliance between digital technology and a certain form of contemporary bureaucratised capitalism. In the case of the automated systems we see in governmental and corporate contexts—and in many large organisations—automation is a result both of the desire on the part of software, IT, and consultancy firms to increase their customer base as well as expand the scope of their products and sales, and of the desire on the part of governments and organisations to increase control at the same time as they reduce their reliance on human

judgment and capacity. The fact is, such systems seldom deliver the efficiencies or economies they are assumed to bring, and they also give rise to significant additional costs in terms of their broader impact and consequences, but the imperatives of sales and seemingly increased control (as well as an irrational belief in the benefits of technological solutions) over-ride any other consideration. The turn towards automated systems like Robodebt is, as is now widely recognised, a common feature of contemporary society. To look to a completely different domain, new military technologies are being developed to provide drone weapon systems with the capacity to identify potential threats and defend themselves against them. The development is spawning a whole new field of military ethics-based entirely around the putative ‘right to self-defence’ of automated weapon systems.

In both cases, the drone weapon system and Robodebt, we have instances of the development of automated systems that seem to allow for a form of ‘judgment’ that appears to operate independently of human judgment—hence the emphasis on these systems as autonomous. One might argue—and typically it is so argued—that any flaws that such systems currently present can be overcome either through the provision of more accurate information or through the development of more complex forms of artificial intelligence.

Within AI research itself, there has long been a debate around the extent to which human judgment, which is to say human thought processes, can be duplicated by artificial systems. For many years this debate focussed on whether such judgmental or human capacities could be duplicated using essentially computational or calculative processes. Although that debate has largely been resolved in the negative—so that contemporary AI no longer looks to model human cognition on computation, but rather to model artificial cognition on human cognitive structures—there is still a tendency, outside of the AI field, to assume that judgment can be understood as essentially reducible to a calculative or quantitative process. One version of this is found in the contemporary obsession with the reduction of judgment to the operation of systems of rules—whether as part of quality assurance mechanisms, audit processes, automatized

✉ Jeff Malpas
jeff.malpas@utas.edu.au

¹ Emeritus Distinguished Professor, University of Tasmania, Hobart, Australia

approval or compliance mechanisms, or even some form of so-called ‘evidence-based’ decision-making (the latter is something of a misnomer since what is at issue is not ‘evidence’ as such, but *kinds* of evidence). One might say that this is partly a notion that derives from the Cartesian emphasis on number and quantity as at the heart of scientific thinking, although the most extreme version of this tendency is probably to be found, not in physics, but in economics, where the market itself becomes the pure model of rational judgment—objective, unbiased, and capable of resolving computational problems beyond the capacity of any human thinker, or so it is supposed.

What is enshrined in this conception of ‘nonhuman’ judgment is just the idea that judgment is itself a matter of computation or calculation as it operates over quantitative values. Yet several problems seem to affect this conception of judgment—widespread though it may be—of which two seem to be the most significant. First, if judgment always relies ultimately on already given values (even if those ‘values’ are understood simply as modes of salience—so that a value involves a particular situational orientation), then it cannot be assumed that such values will always be identical with or reducible to quantitative values—some values, perhaps the most important, are qualitative and resist quantitative reduction. This is especially so, in spite of utilitarian claims, with respect of ethical judgment. Second, one cannot derive any action-guiding judgment, that is, any imperatival, judgment, merely from an accumulation of facts, information, or data. This is not quite the same as the idea that one cannot get an ‘ought’ from an ‘is’, but rather derives from the fact that facts alone do not in themselves bring with them any particular mode of orientation towards those facts.

To put this latter point another way, facts always require interpretation—which is to say that facts themselves require that they be taken up into judgment. There are, after all, a plethora of facts, and what facts are relevant, what they mean, and how they should guide us are all judgments and as such stand apart from those same ‘facts’. This point applies, not only to facts, but even to rules and procedures—there is no rule or procedure that is self-determining or self-interpreting as to its application (a point Kant makes, but which is also central in Wittgenstein, and, one might even say, is suggested by various formal results in logic and mathematics including Gödel’s Incompleteness Theorem). This means that judgment has an inescapable indeterminacy about it—there is always more than one way of judging that is supported by the evidence available. Judgment is not reducible to calculation, computation, to algorithm or rule.

The conclusion of all of this is that judgment is indeed indeterminate, but also that it is ubiquitous and essential. We might even say that seemingly automated systems of ‘nonhuman’ judgment only function as systems of judgment in as much as they are themselves derivative of the judgments we make that allows such systems to operate in the first place. Judgment is

ubiquitous and essential, but it is also fundamental. One cannot escape judgment—nor can one escape the responsibility that goes with judgment.

Indeed, one of the great dangers of automated decision-making systems is precisely that they seem to present the possibility of judgment without responsibility. The drone weapons system makes judgments that may involve the taking of human life and yet the system cannot itself be held accountable nor is any notion of responsibility attached to the system itself. Responsibility has to rest, in such cases, with those who design and implement those systems, and yet since they themselves are typically detached from the process and do not exercise any judgment with respect to specific cases, they may well view themselves as standing apart from the judgments actually made—those judgments belonging to the system and not to them. This is itself part of the implicit attraction of such automated systems—it is likely part of what underlay the reluctance of the Australian Government to make any apology for the Robodebt affair until they were effectively forced to do so.

The divorce of judgment from the responsibility that automation thus achieves is one of its dangers, but it is certainly not the only one. Equally important is the loss of a sense of judgment as itself inescapable—judgment and the burden of judgment is at the very heart of human life. The desire to escape that burden is itself representative of a desire to escape from our own humanity. In our current situation, in the face of the COVID-19 pandemic, looming economic disaster, and the ever-increasing threat of climate catastrophe, a recovery of our humanity, and so of the necessity and responsibility of judgment, is perhaps more important than ever.

Curmudgeon Corner Curmudgeon Corner is a short opinionated column on trends in technology, arts, science and society, commenting on issues of concern to the research community and wider society. Whilst the drive for super-human intelligence promotes potential benefits to wider society, it also raises deep concerns of existential risk, thereby highlighting the need for an ongoing conversation between technology and society. At the core of Curmudgeon concern is the question: What is it to be human in the age of the AI machine? –Editor.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.