

# The answers a machine will never have

*"I suddenly feel kinda guilty... This is not a great way to wake up!"*

This was how my colleague ended an email he sent to my co-founder and me a little less than a month ago, letting us know that he had been diagnosed with leukemia while on a trip to California. He is 29 years old, and this was a life-shattering surprise: to him and, in turn, to our startup team of six people.

Within half an hour, I was walking to the train station with my co-founder, digesting this news on the way to a major negotiation. We agreed quickly and without debate that we would pay our colleague for as long as we were solvent, and that if this was what bankrupted our company, we felt at peace about it. The company is just over a year old, with enough cash to pay everyone, in a worst case scenario, for a few months. We did not care (or check) what his employment contract said. This was the second time in the past year that I've taken Peter Drucker's advice,<sup>1</sup> writing down my predictions so that I can compare them later with what happens. In nine months, I predicted (or hoped), my colleague would be better, and the team would be stronger.

When I phoned our other three colleagues to tell them the news, between data blackout patches on the train, stifled tears on the other end of the phone, and sympathetic looks from fellow passengers, I could also give them the first really hard piece of evidence that our company means what we say about putting people first.

Digesting it a few weeks later, our response was rash, emotional, and risky. And I am deeply proud of it. It passes Drucker's test on managing oneself: it makes me the kind of person I want to see in the mirror.

I also note an irony, because our company specialises in advising organisations on opportunities in artificial intelligence (AI), and an algorithm produced by machine learning could have done a calculation like this for us. It might have weighed up all kinds of data points - our colleague's age and his time with us (eight months), his billable hours and charge-out rate, perhaps even a numerical score that explains what a good-natured, cheerful, and kind person our colleague is, and how much his energy is the heartbeat of our little company. It would also have weighed our other obligations far more objectively than we did that morning: our directors' duties to our shareholders (happily, ourselves - so

---

<sup>1</sup> Advice from P. Drucker "Manage yourself, and then your company" (1996)

not much conflict there), our duties as employers to keep cash in the business for our other team members, and our drive to grow the company.

I expect that an algorithm would have advised a different answer from the one we chose. Perhaps, weighing his importance to the company, it would have advised us to give him a little extra sick leave - say, an extra four weeks. It is unlikely a machine would have advised us to make a decision that could bankrupt us.

This question - why did we deliberately make a subjective decision when we might have made a more objective one? - got me thinking about bigger questions. What kinds of decisions can be made by AI? What kinds of decisions should always be made by people, and not by machines? And how can AI support the modern knowledge worker?

## Artificial intelligence and the knowledge worker

“In a knowledge society”, Drucker writes, “the employees... own the tools of production”.<sup>2</sup> Compared with the typical worker of the industrial age, completing a repetitive task using the machinery owned by his employer, a knowledge worker has enormous freedom. She has inalienable ownership of every thought - every application of her knowledge - and also has discretion about how she applies this to support the organisation that employs her. She receives satisfaction and a sense of ownership when she uses her knowledge to help achieve an organisation’s goal - one which, ideally, she believes in and, even better, may have helped to devise.

Artificial intelligence programmes, like other tools, will have varying impacts based on how people design and use them. In Drucker’s words, writing years before modern advances in machine learning, “[w]ithout [human knowledge], the machines, no matter how advanced and sophisticated, are unproductive”.<sup>3</sup> This human role is one that we need to grip firmly, especially as machine learning algorithms become more sophisticated, and harder to unpick. We need to design organisations now in which AI applications will make our lives better, and not worse. This means asking the right questions: not what AI can do, but what AI *should* do. There are questions we should never ask a machine, because we already know that a machine’s answers - no matter how cleverly designed, mathematically complex, or well-tested - can never be good enough.

---

<sup>2</sup> From P. Drucker “The knowledge worker, and the knowledge society”, Edwin L. Godkin lecture, Harvard University, 1994.

<sup>3</sup> Ibid.

## The questions we should never ask a machine

The approach to artificial intelligence that has technology enthusiasts, business leaders and newspaper journalists most excited is machine learning. A machine learning algorithm is trained on huge quantities of data to spot patterns, which it then uses to make predictions, estimates or even (a loaded word) judgments. When a person has watched thirty films on an internet streaming service, for example, it is very good at guessing - based on millions of pieces of data it has from other streaming users - what kind of film he might like to watch next. This might be quite a nuanced suggestion, based on the fact that it is Friday night, and also that his many of his recent viewing choices have been war dramas.

There are many things the streaming service does not know. It does not know (we hope) whether the viewer is sitting with his girlfriend, or with a large group of male friends. It does not know whether he had a good or a bad day at work. It may know whether he is planning a holiday, or whether his girlfriend is pregnant, because this kind of data is frequently shared between services and can be easy to spot in a user's browsing history. But perhaps the most reassuring thing about the advice the viewer receives from his streaming service is that it is very easy to ignore. The suggestion is for Saving Private Ryan, and he searches instead for a documentary his friend told him about today about life in North Korea.

This example highlights criteria that characterise problems that are solved well by machine learning algorithms. First, they should be problems with a clear, definable answer or set of answers. The streaming service knows quickly whether its algorithm has succeeded by testing whether its viewers follow the recommendations it provides. Second, they should be problems about which there are large volumes of available, accurate data. A complete database of what every viewer has watched so far is large enough to be able to draw conclusions about people's behaviour. This is rare: many of the biggest problems we face, as people and as leaders in organisations, have no such dataset to draw on. In making a hiring decision, for example, I might find data from other companies about the resumés of candidates who have proven successes or failures, but my own company is too small to have reliable data of its own - and the resumé of the person sitting in front of me is, in important ways, unique.

The third point is more controversial. I think machine learning algorithms should only be used to give answers that either do not matter very much, or which are recommendations rather than final decisions.

Why should we consign machine learning applications to the trivial, or limit them to recommendations? We should do so because there is a particular human specialty for dealing with novel situations, and failing to deal properly with the unexpected can have disastrous consequences. Early malfunctions of self-driving cars, for example, have occurred when a car comes across a situation it has not dealt with before - such as a woman walking a bicycle.<sup>4</sup> Returning to the hiring example, a machine learning algorithm may not spot (let alone weigh in its algorithm) a brief reference in a cover letter that the candidate's mother died during his exams.

People, on the other hand, are very good at quickly assessing and handling situations they have never come across before. Knowledge workers are especially good at this, because their training - being taught to think and to reason through years of school, university and workplace training - prepares them for it.

When I ask even my most junior colleague to manage a project, I do not want her to follow exactly the same procedure she used last time. I want her to spot the many ways in which this project will be different, to find new ways of doing things which we might apply to other projects, to anticipate costs that might be higher this time, and ideally to think of ways we can sell additional services. The more initiative she shows, the better - which is the opposite of what I want from a machine. This is what makes her a knowledge worker: she understands what she needs to achieve, but is not told how I expect her to achieve it. The space in which she gets to decide is part of the satisfaction of work for her. Drucker believed that managing a knowledge worker does not involve close supervision, but enabling and inspiring.<sup>5</sup> In this way, a good manager of a knowledge worker makes that space as big as possible - but not so big that she fails, or that the organisations fails.

Ensuring that our knowledge workers have good and interesting work makes them happier, but it also makes them more successful. If we allocate too much of our decision-making to machines (rather than limiting their inputs to recommendations), not only will our organisations fail more often, but we will deprive our employees of the satisfaction of their work: and in turn, I suspect, of their motivation.

---

<sup>4</sup> The first fatal accident of a self-driving vehicle involving a pedestrian occurred in Arizona in March 2018 when an Uber equipped with self-driving technologies failed to see (or react appropriately to) a woman who was walking with her bicycle.

<sup>5</sup> See for example P. Drucker *The Effective Executive* (1966), Chapter III.

## When AI can make humans better

It is an interesting feature of AI - and perhaps a predictable one - that the things it is not very good at (creativity, empathetic decision-making, social interactions) are some of those things that we think of as the most human of skills. Restricting the tasks of machine learning algorithms to other tasks - data crunching, pattern spotting, systems analysis - enables AI to be a humanising force, rather than a stifling (or a dangerous) one.

When an organisation has thought about the questions that a machine learning algorithm should not answer, it may have protected itself from the worst excesses of AI - but it has not benefited from it. The next step should be to find areas where the extraordinary abilities of a machine can be used to free up human minds for more interesting and more rewarding work. Indeed, while automation threatens some of the tasks that make up our jobs, they do not tend to be the tasks that are the highlights of a worker's day. An economist searching for clues in a large data set, for example, could use an algorithm to more quickly clean the data that he is using and eliminate outliers; but will still need to decide for himself which factors he thinks are most likely to explain a rise in inflation.

The most talented of knowledge workers will continue to be more valuable than machines, both because they can handle unfamiliar scenarios, and because they can process emotional inputs. We are a very long way from machines being able to manage people, and this is partly because machines are not good at the kind of thinking required to help get the best out of a knowledge worker. A knowledge worker may have a brilliant idea that she never shares – either because she does not want to, or because she does not know it is a brilliant idea. Returning to the kinds of questions that machine learning algorithms are good at answering, we do not have data on the many times a good idea has not been shared. Nor is there a clear answer on which ideas are brilliant: it will depend on what the company wants, the risk appetite of its managers, and the intangible variable of its culture.

Rather than spending our time trying to quantify the unquantifiable, managers should use the capabilities of machine learning as an opportunity to return to our greatest strengths and specialties - putting ourselves and our colleagues, as Drucker argued, in the positions where we can best succeed.<sup>6</sup> Machine learning tools are at their best when they free up management time, so that we spend more time with our colleagues, and less time staring at screens and rows of data.

---

<sup>6</sup> P. Drucker "Manage yourself, and then your company" (1996)

Vacuum cleaners, dishwashers and washing machines replaced very human work that had been part of our lives for centuries. Their advent did not make us less human, but gave us more leisure, and led to us spending more time with one another.<sup>7</sup> Machine learning technologies can offer the same. They may appear very technical and a lot more threatening than a washing machine, with the very term 'artificial' suggesting a retreat from our humanity. But AI tools are human creations and should be used to strengthen the best human traits.

Every manager is responsible for using technologies, as they develop, to strengthen the skills that make us human. The reason why decisions about sick leave should be made by a human, and not a machine, is because a machine cannot truly weigh the importance to people of helping their colleague and friend to get better.

The future of artificial intelligence, and how it affects the way we live and work, is one we are creating day by day. Every decision about which algorithms we design, and what we use them for, will be made by a human. It is essential that we decide well. We should not let algorithms displace our judgment or our empathy. Instead, algorithms that do some of our work for us should provide the gift of time. When we spend less time trying to think like a machine, we can spend more time thinking like a person. Time is what can help us to understand one another better, to be better managers, and perhaps even that loftiest and most human of goals: to be better people.

### **A postscript**

As I conclude this essay, my colleague is still waiting on a final diagnosis for his cancer. He is back at work while he waits. In an act of daily courage that I find inspiring, he has found reserves of human resilience and strength of character to make him even more cheerful and hopeful than before he got sick.

*Word count: 2460*

---

<sup>7</sup> The Economist Data Team "Parents now spend twice as much time with their parents as they did 50 years ago"