CQI | IRCA

ARTIFICIAL INTELLIGENCE AND THE QUALITY PROFESSIONAL

Artificial intelligence is becoming more prevalent in the era of Industry 4.0. Although AI is bringing a wealth of benefits to businesses who have adopted it successfully, there is still some work to be done in the quality profession to bring quality professionals up to speed and help them reap the benefits of this emerging technology

Words: Carol Long

Artificial intelligence (AI) is changing our world and the ways in which our businesses operate. Although it is bringing some major benefits to organisations who are adopting this tech, AI system developers can sometimes accidentally implement questionable moral standards, encouraging a wealth of negative press. There are examples of digitised identity recognition systems that distinguish between skin colour, fitness monitors that discriminate between men and women owing to physical differences, and recruitment systems that embed biased hiring decisions that favours recruits from the university that current employees attended.

What does a quality professional need to consider to avoid bad publicity? What skills do we already have that can support us in the current era of Industry 4.0? How can we get a head start in understanding what is possible with AI and how we can develop our careers? ▶

AI explained

In layperson's terms, AI is computer software that creates algorithms to make decisions that approximate how a human might deal with that data. AI techniques were developed in the 1980s, and those AI systems required programmers to take human knowledge and convert it into a set of rules or logical statements that can be programmed into the code. The system then used those rules and combined them to make decisions based on data. Sometimes the AI would encounter complexity in the data that led to it giving unexpected answers because the model algorithms based on the rules were not comprehensive enough to reflect the range of real-world data. Developing those systems took a considerable amount of time and effort from experts to uncover their implicit and explicit working knowledge. That effort took them away from their usual roles.

Rule-based machine learning was and is a pragmatic option for healthcare professionals, because their activity followed logical steps from diagnosis to treatment. These medical steps were easily translated into a set of rules that could be programmed to look at medical data and draw conclusions. Some early applications of these rule-based systems were used in oncology for that reason.

> "There is potential for quality professionals to use AI to look for anomalies or recurring themes in the data they use in their own process monitoring"

Other professions have less structured decision-making processes, and the expert would need to understand why they made the decisions in the way that they did before their rules could be defined. Drawing out assumptions and implicit knowledge was significantly harder in those contexts. When the AI produced unexpected answers in testing, it was useful in helping the experts recognise that they hadn't specified some of the decision rules that they would instinctively apply in more unusual situations. These rule-based AI systems were expensive to develop because of the efforts required.

A parallel development path proved more successful. Taking the data that humans had already worked on, AI could find patterns in that data that would lead to the same result as humans. Instead of following the experts' methods, the software found its own. Using this training, the AI model could process larger amounts of data than the human expert and to a similar accuracy. This is then tested in a safe environment under human supervision, before being used in less controlled environments. One example is the joint work between Oxford University Hospitals and University Hospitals Birmingham, who are developing a screening process for a biochemical and physiological signature of Covid-19 in patients' blood samples.

Prevalent technology

After many years of development, AI tools are now usable for many businesses. One of the most common is an automated chat functionality – designed to help website visitors navigate information quickly and efficiently. This system takes the form of a 'bot' (instead of a human), and it is used on the website to analyse what people are interested in and show them similar items. For example, Facebook uses AI to learn what individuals are interested in so that more about that particular subject can be included in that Facebook user's feeds.

This technology is also used inside business systems to analyse customer transactions and requests, using patterns found in previous customers' data to predict potentially dissatisfied customers and pre-empt complaints with more appropriate customer care actions, or to encourage the customer to engage in alternative services better suited to them. Credit card processing has started to use AI to identify transactions outside the customer's usual behaviour patterns to avert fraudulent transactions.

The development of AI tools has three areas of cost:

- developing the toolset
- collecting suitable data to train the tools
- specific application developments based on that training.

From a business perspective, AI is now accessible technology because a large portion of the initial costs of developing the toolsets have been taken on by larger organisations, such as IBM, Google, Amazon and Microsoft. All these organisations are now keen to involve businesses in the use of their tools and target some of their marketing budgets to promote their successes.

Along with successes, there have also been some major failures. This includes:

- Microsoft's conversational chatbot function which became antisemitic, misogynistic and foul-mouthed within 24 hours after it interacted with certain users.
- IBM's Watson for Oncology miscalculations on treatment paths, suggesting potentially dangerous treatments for cancer patients.
- Amazon wanted its recruitment algorithm to take bias into consideration when finding new hires. Instead, the algorithm favoured male recruits because their existing engineering teams were predominately men.
- Google indexed racist material and presented this as reliable information.

These examples of failures started out as well-intentioned AI automations.

They encountered conditions that challenged the original intent when processing complex data. Google did not intend to present racist material when a user searched for organisations that protested against racism. The Google AI simply learned that those offensive terms were associated with anti-racism sites in the same way that ingredients are associated with a cooking recipe. The complexity in that data was beyond the intelligence of the search tool's AI.

AI is not the only technology that is emerging from specialist areas and moving into more general applications. Other examples are:

- process automation
- Internet of Things (IoT) and the associated sensors
- robotics
- wearable devices
- cloud computing
- mobile data and unified communications
- electronic currencies (not only cryptocurrencies, but also cashless transactions).

These technologies have been evolving since the 1990s.

Unfortunately, none of them were initially developed with a social agenda in mind. This leads to complexity in integrating them, causing interesting and unexpected developments. There are likely to be some unpleasant surprises, but some eventualities can be predicted. For example, we are becoming increasingly aware of the interaction between big data, the IoT, and AI, and its capacity to embed insidious biases in future wearable devices.

The quality professional's role

Quality professionals have a role to play in the development and implementation of AI. Professionals who review business cases for operational improvements will need to be ready for the increasing uptake of AI. Acquiring a basic understanding of how AI works will help quality professionals assess whether or not the claims of advantages made for the developments are overblown, or if the ambition of a system design is overreaching. There is a lot of hype around AI, and early adopters of this technology tend to be over-enthusiastic and optimistic.

Those who spend more time looking at legal and compliance risk will recognise that innovation can cause complications. For example, a person with a limited understanding of their organisation's digital AI tools and AI's risks and capabilities could see an increase in liability issues. In this scenario, the person in charge of an implementation could be held responsible if an AI issue takes place on their work premises.

Data protection

AI will need appropriate technical and management safeguards around its use. This includes making sure that people have given their consent before using or analysing their data for a particular situation.

There will also need to be frameworks and guidelines for appropriate and ethical use, the management of artificial intelligence-related intellectual property, and a careful watch on the interaction between AI applications and other technologies that may cause unexpected results. For example, AI monitoring production may misinterpret an unexpected reading from a IoT sensor, calling it a major failure and forcing a factory shutdown, whereas mature integration would raise a maintenance request for the sensor and an advisory note for those reviewing the process.

Data management and data quality is the biggest risk area. There are two main aspects to this; the first is the quality of data within the dataset that will be used to train AI. For this, the standard questions around currency and validation are only part of the areas to be investigated. It is vital that any operational data that is used to train AI is evaluated in a wider context. For example, using customer data to identify the characteristics of good customers is selective in that it only looks at your existing customers and it cannot tell you about the characteristics of potential customers who may be important to your future business.

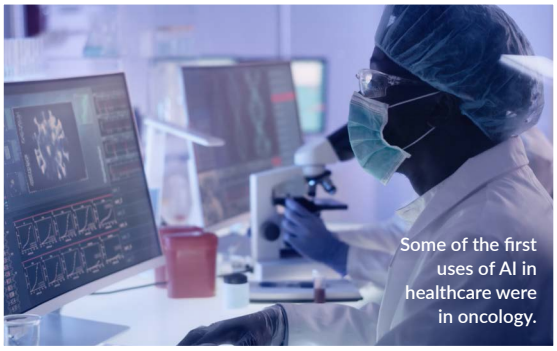

Some of the first uses of AI in healthcare were in oncology.

> "Quality professionals have a role to play in the development and implementation of AI"

The second aspect is as AI applications extend based on current regulations and practices and start to draw data in from elsewhere (eg, pooling customers' social media accounts to assess creditworthiness), the AI's models may be applied in areas for which they were not designed. This means that consumer protection, compliance with the General Data Protection Regulation (GDPR) or questions about ethical business practices could become a concern. ▶

Critical thinking

Perhaps the biggest area where quality professionals will find their expertise being valued is in critical thinking. We are well-trained in recognising the difference between cause and effect and that there is a difference between causality, coincidence, and correlation. By providing constructive feedback, we can help educate other colleagues about the logic in AI and those businesses who are developing it.

It would be quite easy for the technical teams involved to become intensely focused on the work to be done, and not have the wider viewpoint needed to validate the ideas and question the scope of them. Modularisation, software objects and reuse are so embedded into development practices that we often encourage developers to adapt existing tools, protocols and code as a way of speeding up and improving product development.

> "There is potential for quality professionals to use AI to look for anomalies or recurring themes in the data they use in their own process monitoring"

We do not stop to consider the history of the tools we use. In the context of AI , this could be problematic. If a data tool analyses medical indicators for a sports team, reusing it to monitor the health of elderly patients may be entirely inappropriate, even if the same factors need to be modelled. The issue with this is that what is normal for a highly trained athlete could be a danger signal in an elderly person. Understanding the source and provenance of the AI tools being deployed for this and similar situations will become an important part of long-term maintenance and management of an organisation's AI assets.

When data mining and analytics were first used, a pattern was detected where Sunday newspapers were often bought by people who went for a walk. Was this a coincidence simply because both Sunday newspapers and walking only happened together on Sundays and not on other days of the week? Was this causal, in that going to get a Sunday paper meant that people had to walk because of the lack of public transport? Or was it causal because walking to the shops on a Sunday meant people coincidentally bought a newspaper? All those questions said something specific about the social structure of those who bought Sunday papers. At the time, these were simple marketing questions based on patterns that had been noticed in customer data. Using AI technology, such patterns could be used to shape business decisions and communication with consumers on a global scale.

Similarly, false positives and false negatives are familiar concepts to some quality professionals, and they know that these are not equally important in some situations. In Covid-19 testing, a false positive will inconvenience one person and their close contacts, who will need to quarantine. In contrast, a false negative could cause a person to be released with an unrealistic confidence and risk, which may result in that person infecting vulnerable people, thus spreading the virus.

During any development, a quality professional could be an independent voice that considers the balance between striving for the perfect model in the AI algorithm and having something that is effective without taking on unnecessary risks. This independent voice is something quality professionals do in other contexts.

AI systems for rapid one-hour Covid-19 screening in hospital emergency departments are being further trialled in Oxford and Birmingham, after promising initial results.

Auditing with AI

There is potential for quality professionals to use AI to look for anomalies or recurring themes in the data they use in their own process monitoring.

Something similar has been done in the accountancy world, where auditors have produced transaction analytics for their client organisation (the auditee). In the past, such transaction analytics would highlight many false positives as unusual behaviour in the accounts. Applying AI can reduce the volume of anomalies to be reviewed by the auditor by assessing other patterns within the transaction record to evaluate if the unusual behaviour is a coherent action within context of the whole set of accounts. This is a valuable time saver for both the auditor and auditee, allowing efforts to be focused on transactions or patterns of transactions that might reveal more risky behaviour, such as potential expense frauds or internal revenue recognition within a group or overextending cashflows.

AI systems, like all other IT systems, will require supervision and

maintenance, which will fall into the following categories:

- Systems where the model will be fixed and training changed only when retraining is deliberately applied (eg, testing a traditional application before an upgrade).
- Systems which continue to evolve and learn based on transactions and data with which they interact (the organisation cannot exercise change control).

The second of these will require more overt monitoring. In the 1990s, it became a common understanding that every IT system needed monitoring and maintenance. In the 2000s, management became accustomed to their systems being relatively stable, as demonstrated by features such as Plug and Play becoming ubiquitous.

Their experience of AI will be more like the systems of the 1990s or plant and machinery that needs to be monitored to ensure that these systems stay within their intended performance limits. Some quality professionals know about the common and special causes of exceptions from training in statistical process control (SPC). This knowledge will need to be adapted for AI systems to be able to determine when the AI is truly broken and when it is within scope but has found an unusual situation.

Governance and leadership

Governance will continue to be an area where thought is needed. If decisions are automated, who is responsible and who holds the accountability? We are still in the early stages of this technology's use outside of research facilities. The transparency of "explainable AI" extensions helps here. These are relatively new extensions to the AI building toolkits. Those using AI for business purposes are now able to have some record of why a decision is made by the AI. How the decision is made is recorded, showing what data drove the weighting for each part of the decision. This is useful if a decision is challenged: the objectivity of the system is laid bare. Legislation has prompted and supported this approach (eg, GDPR article 22 guards against profiling that can't be explained).

Standards development has and will continue for AI. The first standard initiative was started by the Institute of Electrical and Electronics Engineers (IEEE) five years ago (IEEE 7000 family) and standards for AI and robotics have been published or are currently in final drafting stages. Other countries are now following the UK Government's lead in providing support for AI development for their key industries.

The National University of Singapore developed a tool that uses AI to customise training content to individuals providing accelerated learning.

Similar systems could be used to train people within organisations and perhaps by quality professionals themselves in ways that avoid 'sheep-dip' process training tactics.

Professional development

There are a growing number of AI courses available to bring quality individuals and organisations up to speed. The National Careers Service (bit.ly/3byPQkI) has some basic skills courses, which explain some of the fundamentals to help individuals acquire new skills and support their continuing professional development.

A guide to using artificial intelligence in the public sector, published by the UK's Government Digital Service (GDS) and the Office for Artificial Intelligence (OAI), reminds us that AI can be applied in many sectors and that there is considerable commercial reach. The guide has case studies from banks, market research, transport, utilities and the public sector. Reviewing this document in your organisation's context might be useful to help you start thinking about what would be applicable and where some of the risks might arise in your business.

> "Governance will continue to be an area where thought is needed"

There are also useful intermediate AI courses being run by the Microsoft AI School, Google AI Education, IBM Watson Academy, Amazon Machine Learning University and Facebook AI, to name a few. In addition, more in-depth courses are now available from FutureLearn (in conjunction with the Open University and the Institute of Coding, some of which are free) and there are now universities offering postgraduate courses in AI technology, including Oxford and Cambridge Universities, as part of their computer science programmes.

AI is software engineering, data and probability applied to test a theory of the way the world works, and is a technology that will increase in use over the next decade. There will be new automations using AI implemented in areas not yet imagined. Quality professionals are well placed to move into this arena. This is because most quality professionals are educated in applicable skills, such as statistics, testing, process development, risk management and engineering. They can extend their skills to develop their careers while helping organisations navigate getting the best from AI technology without stumbling. ■

To learn about the effects of digitalisation in the workplace, read the CQI's report *The Future of Work*, available at: **quality.org/future-of-work**