tsg insights
October 2021

Introduction to the testing of artificial intelligence (AI) and machine learning (ML)

This whitepaper is the second in a series which act as companion pieces:

- An overview of artificial intelligence (AI) and machine learning (ML).
- Introduction to the testing of AI and ML.
- Testing of AI (artificial intelligence) and ML (machine learning) – supervised learning.
- Testing of AI (artificial intelligence) and ML (machine learning) – unsupervised learning.
- Testing of AI (artificial intelligence) and ML (machine learning) – reinforcement learning.
- AI and machine learning – algorithmic bias – the cruel mirror AI and ML reflects back at us.

In this paper, we look at recent developments in AI and ML, together with the long-standing challenges involved and the different types of testing that result from these challenges.

Current AI and ML developments

AI and ML will undoubtfully play more important parts in everyone's lives, yet it's difficult to speculate even about the medium-term future.

The compound power that AI and blockchain can provide has been a major topic for years, and the potential for blockchain to add trust and privacy is being explored across multiple sectors.

Along with the challenge of algorithmic bias, the ability of AI and ML to cope with major "before and after" events such as Brexit and COVID-19 will continue to be examined. Will pre-event learning be able to cope with tomorrow's post-event new realities?

In some areas, the answer is a qualified "yes". In others, the jury is still out:

Healthcare – it was partly thanks to new AI models that vaccines were developed, tested and rolled out so quickly in response to the COVID-19 pandemic. What once took years or even decades was condensed into months. Research now is looking at using AI to prevent future outbreaks.

Autonomous driving – the world is getting a lot closer to widespread use of fully automated driving capabilities. In China, a Robotaxi service was recently launched, with the aid of advanced AI systems to safely control vehicles in any environment and conditions.

Facial recognition technology – this is developing apace, and not just in the human sphere. Wildlife conservationists can now use AI to recognise and track individual animals such as chimpanzees.

Financial services – most major companies in this sector are now using ML to improve customer experience and to drive revenue increase. In this sector, however, the joint potential for AI and blockchain continues to be unresolved.

Robotics – applications here range from robots on the ocean floor all the way up to satellite servicing in space, as well as using industrial robots to clean up chemical spills on land via 4G/5G networks.

To date, results remain mixed:

- Closed systems –AI and ML have already objectively been proven to supersede human ability in just about any closed system, including board games like chess and computer games such as Starcraft II.
- Chatbots – here the record is less good. Failures include Microsoft's Chatbot Tay (which exposure to Twitter caused to became racist) and its successor Zo, which went the other way to become overly politically correct and judgmental. In addition, Amazon tried to use AI for recruitment, eventually admitting failure and reverting to the human approach after four years.

Below, we look at some of the reasons why there's still a long way to go.

AI and ML challenges - There are many testing everyday situations that AI and ML continue to face

Common-sense knowledge

When humans make decisions, they not only take the immediate facts into account, they also apply common-sense filters. These tend to be intuitive and not definable by rules. However, ask a computer to generate a sentence using the words "dog, ball, catch, throw" and it could come up with "The dogs are throwing balls at each other". We know that's nonsensical, but the computer doesn't.

Qualification problem

This is closely related to common sense knowledge. Related to this is the qualification problem – the impossibility of listing all preconditions required for a real-world action to have its intended effect. Unexpected circumstances may prevent something from happening – a human would see this, or at least understand why, but a computer probably wouldn't.

Combinatorial explosion

This refers to the exponential growth rate at which search problems grow. Chess is a good example here. A computer may be able to beat even a chess grandmaster, but what it can't do is completely analyse a game from start to finish. There are simply far too many variables involved. The focus now is on heuristic search to effectively reduce the search parameters and make them manageable.

Word-Sense Ambiguity

At least 10 English words have hundreds of definitions each – for example, "go" and "put". The word "set" has 430 senses listed in the current Oxford English Dictionary, but that is likely to be overtaken in the next edition by the word "run", with over 600 different senses. The challenge for AI is trying to understand the context of a word and use it correctly.

Planning

Complex planning scenarios have been handled by computers for decades – for example transport and staff scheduling. However, working out how efficient plans are, or if a new plan will work at all, still has some way to go. The AI and ML challenge is to improve feedback loops plus checks and balances to ensure that planning becomes more efficient over time.

Learning

The issue here is two-fold: learning the wrong things and failing to learn the right ones. A good example is the Tay chatbot mentioned above. Related to this is the issue of knowing for sure that improvements are occurring – that the learning is having the right results.

Types of AI learning

Issues such as the above ensure that testing remains a key requirement for AI and ML into the foreseeable future. This testing must be tailored to each learning type, with the three main types being supervised learning, unsupervised learning and reinforced learning.

- Supervised learning – this is example-based. A computer algorithm is trained on input data labelled for a particular output. Similar to test automation, the test cases have expected results.
- Unsupervised learning – simply presenting data and allowing the computer to look for similarities of its own accord, somewhat like data mining. This can work well, but also throw up false positives, where correlation does not mean causation.
- Reinforcement learning – exposure to an environment such as a simulator, where the computer can explore strategies and tactics to earn "rewards" that guide it in the right direction.

TSG provides expert guidance on AI and ML, as well as assurance and testing services. We make change happen, safely and predictably. If you have any question about issues covered in this whitepaper or would like to know more about how we can help you, please contact us now. **Call: +44 (0) 207 469 1500 Email: info@tsgconsulting.co.uk www.tsgconsulting.co.uk**