AI Whitepaper

JCC AI Expertise as applied to fleet management

Version: 2.0
Prepared for Gareth Roberts and Hampshire County Council
Produced by JCC’s Artificial Intelligence Team

Introduction to AI

What is AI?

Artificial intelligence (AI) is when machines mimic cognitive functions humans associate with intelligence, such as learning, thinking, achieving specific goals or solving specific problems. AI generally excels in areas where humans would quickly become overwhelmed, given the amount of data and/or the repetitive nature of the task. AI is also great at bringing disparate data types together to create a unique insight that humans would likely miss.

It's hard defining a precise date where AI was born but recent availability of labeled data in large scale, the increase in computational processing capabilities and the decrease in its cost explain why we see "AI" in a lot of industries and domains. AI spans from computer vision to speech recognition, language understanding to robotics and Self Driving cars — just to name a few!

Machine Learning and Deep Learning

Machine learning is a class of algorithms that use statistical techniques to learn. By showing a machine learning system many examples of data, it will progressively improve performance at a task using that data (i.e. learn) and will abstract those lessons to repeat the task on unseen examples of the data. The pinnacle of machine learning's ambition at present is SAE Level 5 autonomous cars in which an AI trained with millions of miles of driving sensor data (from cameras, LIDAR, etc.) will be able to navigate through previously unseen streets in any environmental conditions.

Deep learning is a specific approach used in Machine learning for building and training neural networks where Deep learning models learn representations of data with multiple levels of abstraction. They are composed of deep neural networks, which are trained by an algorithm using the backward propagation of errors or backpropagation for short. Each layer of a deep architecture constitutes a different representation of the data. The two most popular deep learning architectures are deep convolutional nets and deep recurrent nets. The former architecture has revolutionized speech recognition, image recognition, object detection, and many other domains. The latter one has dramatically improved the state-of-the-art in problems related to sequential data such as speech-to-text, translation, time-series prediction, etc.

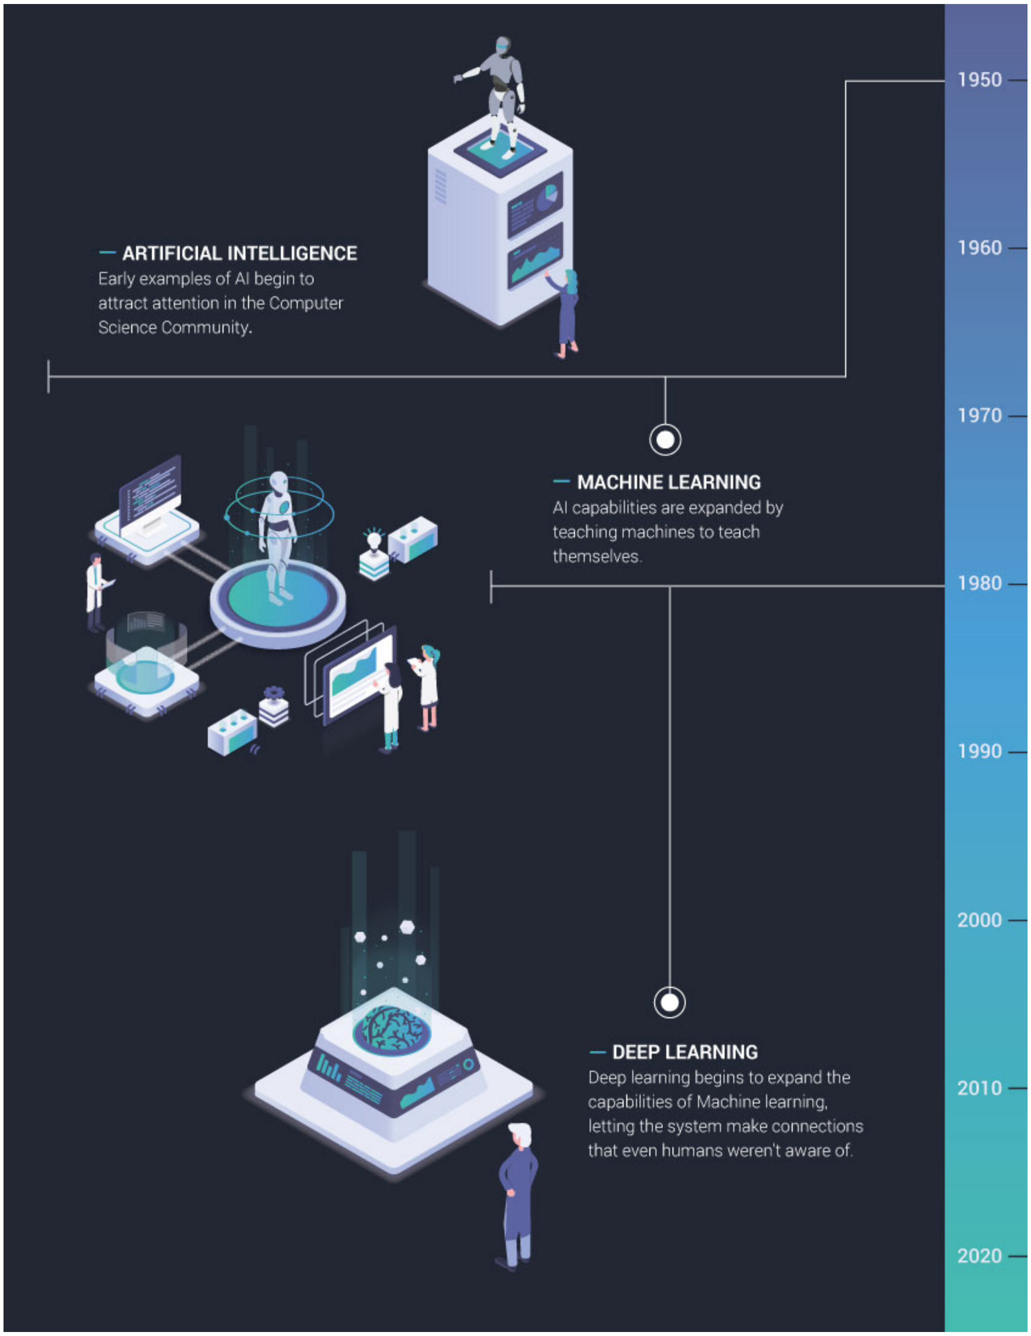

Machine learning may be supervised or unsupervised. Imagine showing an AI machine a set of photographs of cats and dogs. In unsupervised learning, the machine may learn to split the photos into two groups by species. Alternately, the system may learn to group animals by colour, pose, or some other detail.

The results of unsupervised learning will reflect (and provide insight into) the dominant characteristics of the data shown to the machine.

Conversely, supervised learning is where a 'training set' of photographs are tagged as 'cat' and 'dog' so the machine's learning is directed towards the goal of being able to recognise and classify pets in new photographs that it hasn't seen before.

Finally, some ML algorithms use reinforcement learning. This is a style of learning where feedback from the environment and rewards process drives learning similarly to training a dog and rewarding it with food.

Recently Google's AlphaGo AI played itself at millions of games of Go to continually improve, culminating in beating the world's best human player.

Strengths of Machine Learning

Most software systems rely on human intelligence codified as a set of rules telling the system how to behave. In this case, the computer scientists are the limitation as it is exceptionally challenging to codify all the rules for how a complex system should behave in any circumstance. Even systems with well-bounded rules and states such as games like Chess or Go are too complex for humans to analyse and codify every possibility.

Machine learning, in contrast, can take advantages of huge volumes of labelled data to infer rules or mathematical models that would have been too complex to code by humans.

Limitations of Machine Learning

Machine learning algorithms are fundamentally statistical analysis machines and, as such, are subject to the laws of probability. This means that while they are very good at finding patterns and relationships, they struggle with novelty.

Imagine again training an AI with photos, but this time to recognise humans, animals and other objects. Imagine you have all the photos people upload from their phones to train your AI. For all the many photos you have of people there are only a relatively tiny number of photos of gorillas. This remains a problem for Google who has constrained their classifier not to attempt to recognise gorillas because of offensive mistakes.

AIs are not infallible and can present some limitation. The performance of any classifier can be described by two metrics: sensitivity is how often it correctly classifies a thing (a true positive, as if to say this is a dog), and specificity is how often it correctly classifies not-thing (a true negative, as if to say this is not a dog). There are two types of classification errors: classifying a thing as not-thing (this dog is not a dog) and classifying not-thing as thing (this cat is a dog).

Very often, a classifier will be deliberately biased towards making fewer of one type of error in exchange for more than the other. This is common in medicine, where an AI detecting cancer will be tuned to minimise false negatives even if that results in more false-positives, as it is better to send people without cancer for unnecessary tests than to miss a genuine diagnosis.

Artificial intelligence can only currently mimic some human cognitive functions. It cannot yet offer true intuition or creativity, only inference and extrapolation. AI art may be selling at Christie's, but the AI's creator acknowledged only that "algorithms are able to emulate creativity".

AI Applied to Fleet

Fleet Maintenance

All fleet owners carry out reactive and preventative maintenance. Reactive maintenance is expensive for several reasons:

- A failing component can cause additional damage and in the worst case a catastrophic accident
- Any unplanned downtime of a vehicle will trigger recovery, repair, logistical, administrative and contractual costs, with a knock-on impact on other fleet operations

Preventative maintenance seeks to reduce incidences of reactive maintenance.

It is the regular:

1. Replacement of parts with a clear cost-benefit. For example, oil costs little to replace with the benefits of reducing risk of engine failure and increasing the service life of expensive parts.

2. Inspection of parts to replace them before they either fail, wear out or cause more expensive failures. Reading logs and diagnostic trouble codes (DTCs) from the onboard diagnostic (OBD) port enables inspection of inaccessible and complex components too.

Predictive maintenance is enabled by inspecting more frequently and looking for trends that indicate a component is degrading. Ideally, inspection is continuous, occurring all the time the vehicle is in service and in tandem with highly automated decision-making. It involves sensors reporting on the operation of vehicle systems and the detection of deviation from normal operating parameters. At its simplest, it is the logs and DTCs from the OBD reported to the Cloud via Telematics, but it is JCC's belief that simple DTCs (based on sensor thresholds for single components) are an oversimplification that will be blind to many predictive maintenance opportunities.

Cobalt Auta is an IoT device that integrates with in-vehicle CAN buses to report a data stream that is richer and more complex than is available from OBD-based loggers. After smart filtering and dimensionality reduction to extract valuable information, Cobalt Auta uploads this data continuously to the JCC Cloud AI API.

JCC's Cloud AI uses Machine Learning to generate health statuses for vehicle components and detect abnormal behaviour early before a failure happens. Learning is a continuous collaborative process between the fleet owner, JCC's AI, and domain experts, yielding increasingly powerful results over time:

- Unsupervised learning finds anomalies within a fleet of otherwise similar vehicles. JCC's automotive experts will hypothesise as to the causes of such anomalies and what risks they might pose to the fleet owner and advise accordingly.
- The fleet owner uploads service and maintenance records to JCC, including issues identified by manual inspection during regular services and any breakdowns triggering reactive maintenance. JCC uses these time-stamped events in combination with the archive of telematics data to create training sets for supervised learning with the goal of early recognition of component degradation.

AI reports health status data to the fleet owner through JCC's Cloud API. The API allows easy integration with whatever operations and scheduling software the fleet owner uses to manage their business. The component health score can be used to schedule a workshop visit for inspection or replacement depending on the specific circumstances of the vehicle, such as the cost of replacement, the likelihood of false positive, and the consequence(s) of failure such as downtime.

Possible AI Approaches

For proper fleet maintenance, the following machine learning/deep learning approaches can be used:

- **Anomaly/novelty detection — this is supervised and assigns a class/category to data.**
- **Classification — supervised and assigns a class/category to data.**
- **Clustering — this is unsupervised and separates data into groups.**
- **Forecasting — Supervised prediction of specific values of data, e.g. predicting the temperature of the engine.**

Use Cases for Fleet maintenance

- **Predict best time for regular service**
- **Early fault detection of a mechanical subcomponent**
- **Estimate cost of maintenance**
- **Estimate risk of a failure and emergency**

Fleet Management

Many SatNav systems use crowdsourced data from other users to infer real-time traffic congestion and reroute their users to avoid it if faster routes are available. While impressive, this is a reactive rather than predictive technology, and current SatNavs tend to provide little in the way of contextual information when planning a journey. The best they do is use their historical data to calculate the average weight of traffic by day and time in order to try to predict longer journey times during the rush hour.

There are many criteria to balance when planning a route. For example, should a route minimize time, distance, or fuel consumption? Should it be optimized around arrival time? Fortunately, the application of AI can reduce the risk of getting caught in an accident or traffic jam.

There are many factors that contribute to accidents, and some are genuinely random. Some factors though, such as road visibility and tyre grip, can be exacerbated by any number of pile-on factors such as weather, the curve of the road, time of day, and time of year. AIs are good at sifting through all of these disparate datasets to determine the through line. If one stretch of road has been accident-prone in certain conditions, the AI can infer that similar roads under similar conditions will also be risky.

Similarly, AI can parse out whether a particular condition is in play. For example, our AI could determine that a particular stretch of road is dangerous in the rain (e.g. lane lines are harder to see) but actually safe in snow (e.g. one of the first roads to get ploughed) — and adjusted the route accordingly.

This is just one example of how JCC's AI Expertise can be leveraged for Route Optimization, taking destination, route times, and environmental factors into consideration, stitching together data from a variety of sources such as the accident archive, maps, and weather forecasts to return a route that balances risk with efficiency.

The Ways AI Can Help Fleet Management —

- Optimization methods
- Graph-based search
- Machine learning/deep learning
- Mathematical programming

Use Cases

- **Best routes in terms of cost, profit, customer experience**
- **Shortest route**
- **Congestion management for urban EV charging systems**

Driver Assistance

Whenever taking their car drivers may encounter all sort of unplanned situations, again AI can help and assist drivers on a day to day basis to increase their driving experience and here is how:

Possible AI Approaches

For driver assistance the most common approaches are machine learning and deep learning, depending on the problem and the available data.

Use Cases

- **Notifications of emergency events**
- **Assist driver under difficult/extreme conditions (e.g. fog, lowlight conditions, heavy raining, lack of signs)**
- **Classify driver behavior**
- **Enhance passengers' security**

About JCC

JCC brings "Connected, Intelligent, and Autonomous" technologies to a variety of industries, including automotive, aerospace, smart city, freight and agriculture. This breadth of engagement enables JCC to cross-pollinate ideas and best practices, bringing knowledge and experience gained in one industry to bear on the challenges of another.

JCC engages in two phases: innovation and adoption. JCC's data science experts offer consultancy services for new challenges in any of these industries, innovating by gathering data, building training sets, and creating AI solutions. JCC then exposes these AI models through APIs, lowering the barriers to mass-market adoption and accelerating JCC's mission to transform the way people live and move throughout the world.