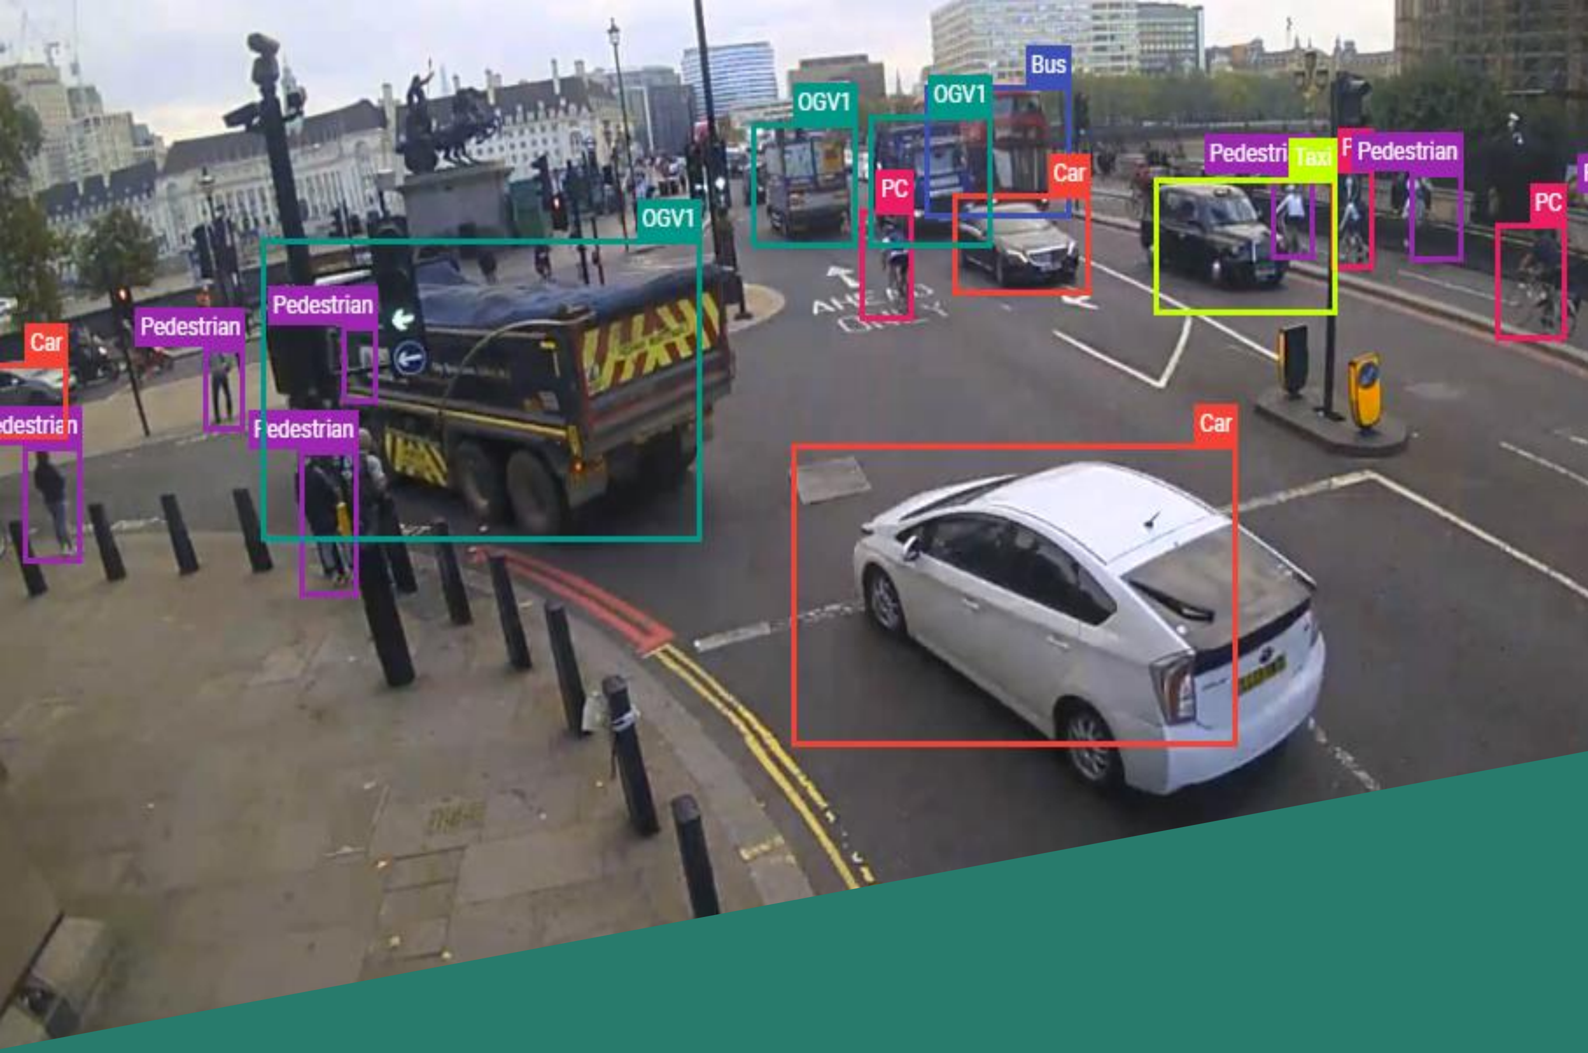

ARTIFICIAL INTELLIGENCE FOR SIGNAL CONTROL

WORKING TOWARDS ROLLOUT IN GREATER MANCHESTER

ARTIFICIAL INTELLIGENCE FOR SIGNAL CONTROL

WORKING TOWARDS ROLLOUT IN GREATER MANCHESTER

1. EXECUTIVE SUMMARY

This 3-year, £3.5m programme is exploring how AI can be used to improve signal control, which will result in a rollout in ~20 signal-controlled junctions in Greater Manchester, ultimately targeting a new paradigm for signal control. Our key objectives are:

- **Multi-modal optimisation**, using sensors able to detect 9 road user classes
- **Fully automated calibration,** eliminating costly calibration cycles
- **Dynamic optimisation**, adapting to coordinated or single-junction control, and to live random events such as lane closure

So far, 1 year through the programme, we have demonstrated:

- AI can **exceed vehicle-actuated performance by up to 20%** in a range of simulated scenarios on a single junction – demonstrating comparable performance to SCOOT already, just 1/3 through our research
- **Full integration** of our AI algorithms with UTC in a safe, bench environment
- A **modern, low-latency OTU (outstation transmission unit)** with **prototype web-based front-end**

Over the coming year, we will deploy initially on one junction and then scale up in Greater Manchester, and continue to demonstrate AI performance, with the ambition **to present real-world improvements vs SCOOT due to AI** this time next year.

2. INTRODUCTION

Over the past 30 years, traffic signal control algorithms in the UK have evolved incrementally, adding new features onto the core of MOVA or SCOOT. These algorithms were designed to optimise overall congestion, and still do a capable job of this, particularly given the computational constraints of the 1980s when they were designed.

However, during this period, the needs of local authorities have evolved. Modal shift and sustainable travel have increasingly become top priorities. Budgets have been cut, resulting in pressure on expensive maintenance and calibration of existing systems. Air quality in British cities is becoming a major item on the news and thus on council agendas.

During this same period, beyond traffic infrastructure, huge shifts have occurred. A new wave of technology, from the internet to microprocessors, has unlocked new paradigms in every industry. Possible optimisation approaches in particular have evolved dramatically with the rise of Artificial Intelligence (AI).

Figure 1: Project Partners

Through this £3.5m InnovateUK-funded project in Greater Manchester, Vivacity is looking to combine these insights. We are leveraging the best of modern technology to control traffic signals, using deep AI trained in simulation with our partner, Immense Simulations. Our ambition is to introduce revolutionise signal optimisation by providing traffic engineers with a radically improved set of tools.

3. WHAT IS "DEEP AI"?

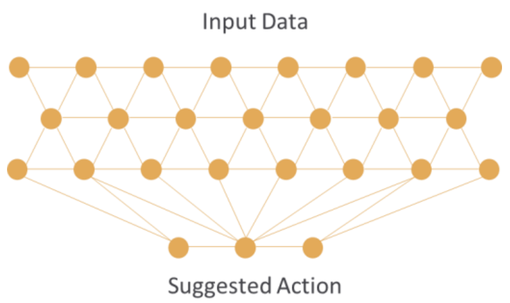

Figure 2: Simplified Neural Network Diagram

Deep Artificial Intelligence, based on deep neural networks, is based on the concept that the best way to solve a problem is to have an extremely flexible algorithm and tune the parameters to fit available datasets. Neural networks are arguably some of the most flexible algorithms ever devised, with deep neural networks often being comprised of millions of parameters in a multi-layered structure inspired by the biological connections in the brain.

However, in order to make use of deep neural networks, these millions of parameters need to be adjusted so that they can model a good solution to the problem. This is done through a training process, where they are shown a very large number of inputs. For each input, the neural network calculates an output, which is then scored, either vs the true answer or against a set of optimisation criteria. The parameters in the neural network are then adjusted to improve this score, and the process is repeated many thousands of times until the neural network produces high-quality outputs for any given input.

Deep AI has been demonstrated to outperform traditional algorithms across a huge range of fields, from image recognition to voice control, from identifying human preferences to driving a car. We are leveraging results from these works to inform the algorithms we deploy for signal control.

4. HOW ARE WE USING DEEP AI FOR SIGNAL CONTROL?

The deep AI closest to traffic signal control is probably the works of Google's DeepMind, who have been training AI to beat computer games. Some of these optimisation problems have similarities with signal optimisation– the player (control system) may have limited information about the world in which it operates, it has a small number of moves available (stage changes) but where each move has ramifications in a much more complex environment, and it is trying to optimise its score (network performance). In many of these examples, DeepMind has beaten top world players, and has demonstrated moves not previously known to the professional community – showcasing not just learning from other experts, but also creativity in designing an optimal solution.

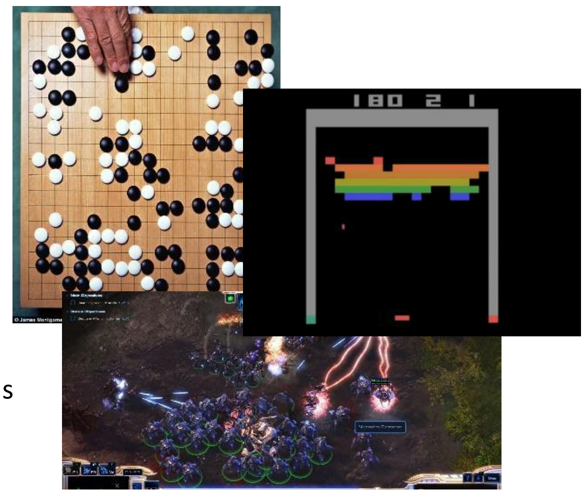

Figure 3: Games that DeepMind has excelled at include the ancient Chinese game of Go, Atari's Breakout, and the strategy game Starcraft

In order to use this type of AI, we need to be able to train it, calibrating the parameters mentioned earlier. When they have not had very much training, AIs typically perform extremely badly, as the flexibility of the algorithm means that it is essentially random until the training process enforces sensible parameters. As such, these algorithms cannot be put in the real world for signal control until they have had some training on realistic scenarios. To do this, we have been working with Immense Simulations, who are providing an accurate microsimulation of a busy section of Greater Manchester's road network. This microsimulation is being calibrated and validated using Vivacity's data, derived from our sensors. Using this microsimulation, we can then run a large number of training scenarios, and calibrate the neural network to provide a very effective solution to controlling traffic at that junction.

Once trained, we are deploying these AIs on an advanced Vivacity outstation to feed commands directly to the controller. Each junction has a number of Vivacity sensors deployed around it, providing the AI with detailed data about the junction. The AI is also fed data about nearby junctions in order to enable cooperative behaviour. The outstations have an intuitive web interface, which links to Vivacity's sensors, providing a range of outputs including images from the cameras on request.

5. WHAT ARE VIVACITY SENSORS?

Over the past few years, Vivacity has developed sensors using AI on video feeds in order to gather data about the road. Each sensor provides anonymised details of the location, classification and speed of each road user (including cyclists & pedestrians), along with vehicle ANPR data for journey times and origin-destinations. These sensors are deployed in cities across the UK to gather offline data, and underpin our signal control work.

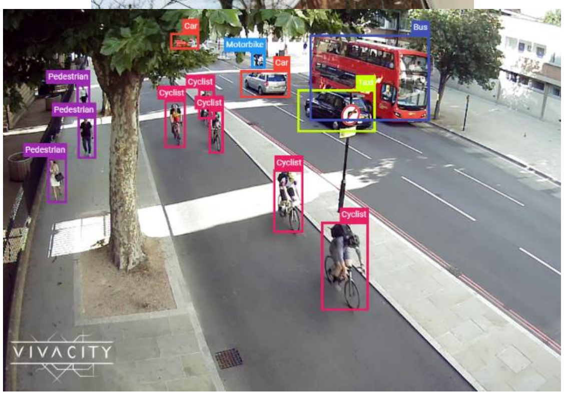

Figure 4: Sensor & Sensor Video Analytics visualisation

Our sensors have been demonstrated to have 99%+ count accuracy and 97%+ classification accuracy in a wide range of trials, including with TfL and Highways England, both of whom are now following up with larger-scale projects. There are now over 1,000 sensors across the UK, providing a range of the capabilities outlined above for different transport authorities.

These sensors provide the ability to gather very accurate granular data, providing the critical breakthrough for creating accurate simulations, as well as giving the AI the insight to choose the best action at a junction for the benefit of all road users - including pedestrians and cyclists

6. WHAT PROGRESS HAS BEEN DEMONSTRATED SO FAR?

Over the first year of this 3-year programme, we have:

- Demonstrated that machine learning approaches can **exceed the performance of vehicle actuated algorithms by 20%+ (average delay time) in a range of scenarios,** in a simulation on a single junction – although there are still some scenarios where AI is worse, which require further training. This demonstrates comparable gains to SCOOT, despite only being 1/3 through our programme, giving us confidence in the overall potential performance.
- Implemented a **safe methodology to integrate AI with the existing UTC** in Greater Manchester, and demonstrated this in an offline environment, with associated **safety features**
- Built a **low-latency, modern OTU** for the AI, including a proof of concept **front-end** for TfGM to manage the AI
- Installed sensors, calibrated simulations, and used these to test a wide variety of machine learning approaches, allowing us to understand the **full costs of an adaptive system optimised for multi-modal control.** Note that these sensors also contribute to huge improvements in data collection in installed authorities due to the wide range of accurate data they gather.

We are now preparing for a single junction AI installation in H2 2019, and scale-up in H1 2020, with substantial work going into research now that we have laid the groundwork. As such, we expect to be able to report on real-world results this time next year, as well as the outcome from a much more substantial research programme. This will include real-world results vs a SCOOT system, and simulated results vs a range of algorithms in a variety of networks.

7. HOW DO WE BELIEVE THIS WILL COMPARE TO MOVA & SCOOT?

We will be looking to achieve a number of key benefits vs MOVA and SCOOT:

- **Dynamic shifting between coordination and individual junction optimisation:** Instead of having to choose between MOVA and SCOOT, or trying to deploy time-based variants, the AI will learn when to coordinate junctions, and when to control them independently. There will be no strict concept of a SCOOT region, but effective zones of coordination will be able to shift dynamically as required by the control system, in a much more flexible manner than existing multi-control-algorithm approaches.
- **Auto-calibration:** Machine learning systems, by definition, improve with experience. Junctions will not need lengthy initial manual calibration periods at installation; signal control algorithm calibration can be executed automatically and continuously through a combination of simulation and real-world learning.
- **Modal & Air Quality Optimisation:** The AI can learn for a variety of different optimisation parameters. We have high-quality modal data already from the sensors, and can prioritise cyclist or bus journey times at will. If appropriate, we can also train the AI to optimise for a proxy for air quality, looking at speed / acceleration by vehicle type.
- **Rapid Response to changing environments:** Today, if a lane is closed or a football match finishes, existing algorithms struggle to adapt appropriately, and typically need manual intervention. However, by including these scenarios within the training set, we can give the AI the opportunity to experiment with different solutions and embed that into the AI's

memory, ensuring that the AI automatically puts in place the right response to these scenarios when they occur in the real world.

8. HOW CAN I FIND OUT MORE?

We will be running a User Group for transport authorities on 14th October, aiming to discuss our work in more detail and explore some of your objections and challenges to deployment. To find out more, or to request an invitation for the User Group, please get in touch with Mark at mark@vivacitylabs.com .

Similarly, if you would like to talk to us about providing further test and demonstration sites, or would like to request more information about our sensors, please get in touch.