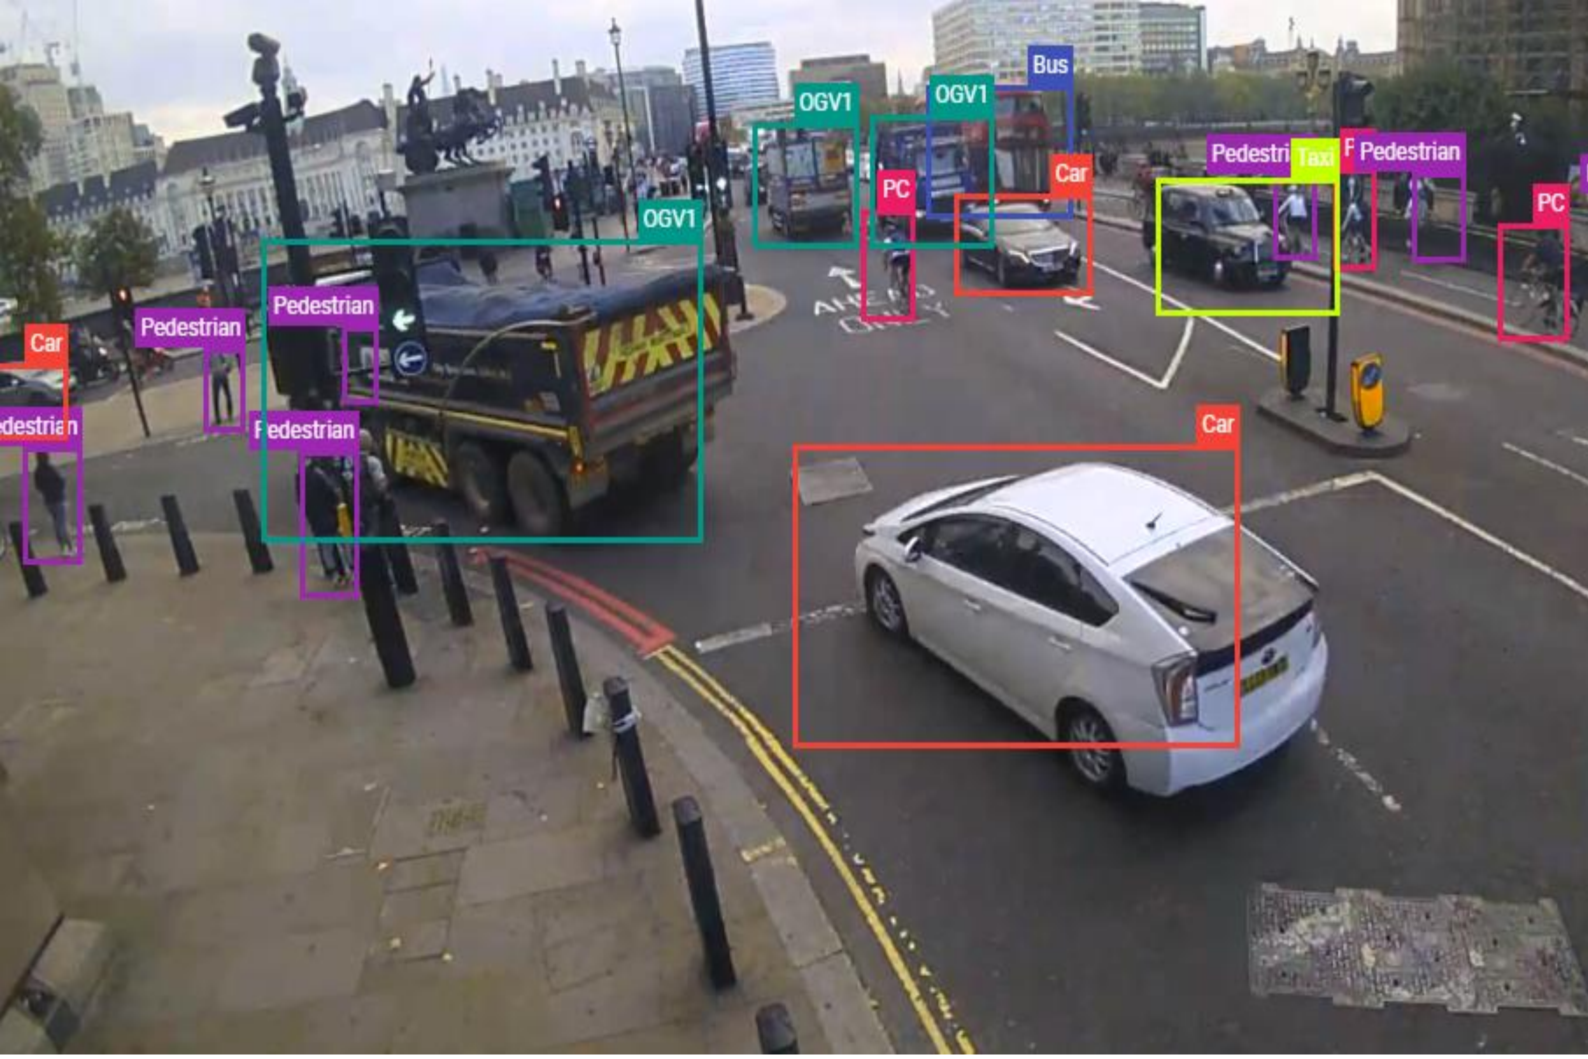

Smart Junctions

Transforming Traffic in Greater Manchester with Artificial Intelligence

1. Executive Summary

Smart Junctions is a 3-year, Innovate UK co-funded programme exploring how Artificial Intelligence can be used to improve traffic signal control. This programme will result in a rollout across a region of circa 20 signal-controlled junctions in Greater Manchester, ultimately targeting a new paradigm for signal control. The key objectives are:

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

Vivacity Labs have partnered with Transport for Greater Manchester (TfGM) and Immense Simulations to use AI to optimise traffic networks. The new system will give transport authorities unprecedented ability to efficiently implement new strategies such as:

a) Focus on **promoting active travel** by prioritising cyclists and pedestrians, thereby improving journey quality and safety

b) Prioritising **air quality** by reducing the amount of breaking, accelerating, and idling of high-emissions vehicles

Authorities will be able to seamlessly shift between strategies in different areas or times of the day to suit their needs and priorities; for example, a focus on **congestion reduction** along main arteries during peak commuting hours, active travel during the weekend, and air quality in off-peak periods.

In the second year of our 3-year programme, we have achieved the following key milestones:

- The **first live trial on Thursday, January 30th, 2020**; successfully controlled the traffic signals at a junction using Artificial Intelligence.
- Continuing to control trial junction, expanding hours of control to include morning and evening peaks
- Initial results saw an improvement on average delays at this junction during the first control.
- **Controlled a total of three junctions independently by July 2020**

Over the coming year, we will:

- Simultaneously control **three junctions over the next month** ahead of the JCT symposium
- Extend the focus on pedestrians and cyclists, expanding to junctions along **Deansgate, an area with high pedestrian traffic**
- Demonstrate coordination and impact in the real world in the form of improved journey quality for all road users in this region.
- Scale up to an area of 20 junctions in Manchester by the end of 2021

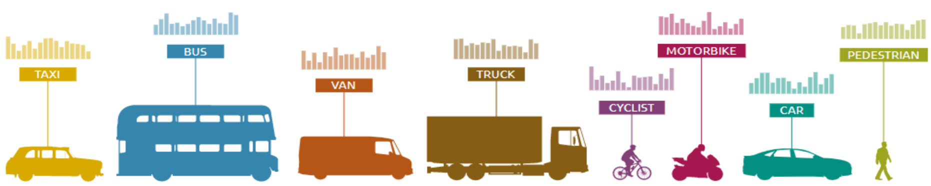

To our knowledge, this is the first application of Reinforcement Learning in traffic signal control in the UK, with one of the largest, real-world trials of its kind world-wide.
In this paper we are going to present the journey from test bench prototype to real world trial, giving further details about the solution that has been successfully implemented in Greater Manchester.

This is a continuation of last year's paper and presentation titled "Artificial Intelligence for Signal Control - working towards rollout in Greater Manchester" in which we presented the solution and early research results.

2. Project overview and Year 1 Recap

Challenges with existing solutions

Many transport authorities are trying to move beyond simply reducing traffic congestion. Optimising air quality and prioritising active travel and public transport reliability are now critical to transport policies. SCOOT and MOVA have dominated traffic signal control in the UK for the last decades and while both have scenarios in which they work effectively, reducing congestion through coordination of multiple junctions (SCOOT) or through adaptively clearing queues and growing cycle times (MOVA), they have both struggled to implement air quality or multi-modal optimisation effectively.

Air quality optimisation with SCOOT has been trialled, but not rolled out at any scale. Bus priority in SCOOT, while well established, is a relatively blunt instrument, overriding optimisation for any other mode to provide late buses with green lights, and thus degrading overall system performance. Meaningful prioritisation for other key modes, such as cyclists, is not widely available. Meanwhile, it is well known that performance of SCOOT degrades over time, often by up to 30% - but recalibration is manual and expensive and often not viable for many authorities.

Vivacity Smart Junctions Solution

At Vivacity, we are addressing all of these issues. Vivacity is using cutting-edge Reinforcement Learning, a branch of machine learning, to develop an algorithm which is able to adapt quickly to changing traffic conditions and efficiently implement high-level strategies at both local and city-wide scales. The system takes full advantage of the unique capabilities of Vivacity's existing sensors to detect and classify 9 different types of road users. This accurate, real-time multi-modal data makes it possible for the algorithm to prioritise different modes of travel, thereby giving councils the unprecedented tools to effectively prioritise active travel, public transportation, and air quality. Crucially, we can also optimise for congestion - allowing transport authorities to choose a balance of prioritisation between air quality, particular transport modes, or underlying congestion. Finally, AI has the potential to self-calibrate, maintaining performance indefinitely as the system self-improves and retrains following changes in the network.

A year ago:

At last year's JCT we presented our idea of applying artificial intelligence to traffic signal control and described why AI and, more specifically, reinforcement learning are well suited to traffic problems. One year ago, we had successfully completed a test bench demonstration of the system working and, perhaps more importantly, proving we had designed and created a safety-critical system. After this key milestone, we were ready for real world trials. The key new work presented in this paper is therefore the deployment in the real world, and scaling towards controlling 3 junctions simultaneously.

3. Reinforcement Learning Research

Reinforcement learning (RL) is an optimisation technique of learning from experience. The RL algorithm or "agent" begins by choosing random traffic stages, and then by using the inputs from the Vivacity sensors (such as positions, speeds, and waiting times of different road users), it can see what the traffic looked like before and after each stage. This way, the agent can learn how 'good' it is to make each decision in a given situation. Over many iterations the system can, therefore, improve itself and choose the best stage at any point in time, to optimise performance according to a council's priorities.

To allow the system to safely make poor decisions as it begins this learning process, we have worked closely with Immense Simulations to build accurate and fast microsimulations. The overall workflow of commissioning a smart junction is shown in Figure 1, and a still image of the simulation can be seen in Figure 2. The dataset available from the sensors has made it cost-effective and scalable to do this very accurately, by automating significant portions of this process. We've now run over 8 million simulated hours to train the system: the equivalent of observing and controlling a junction for a millennium!

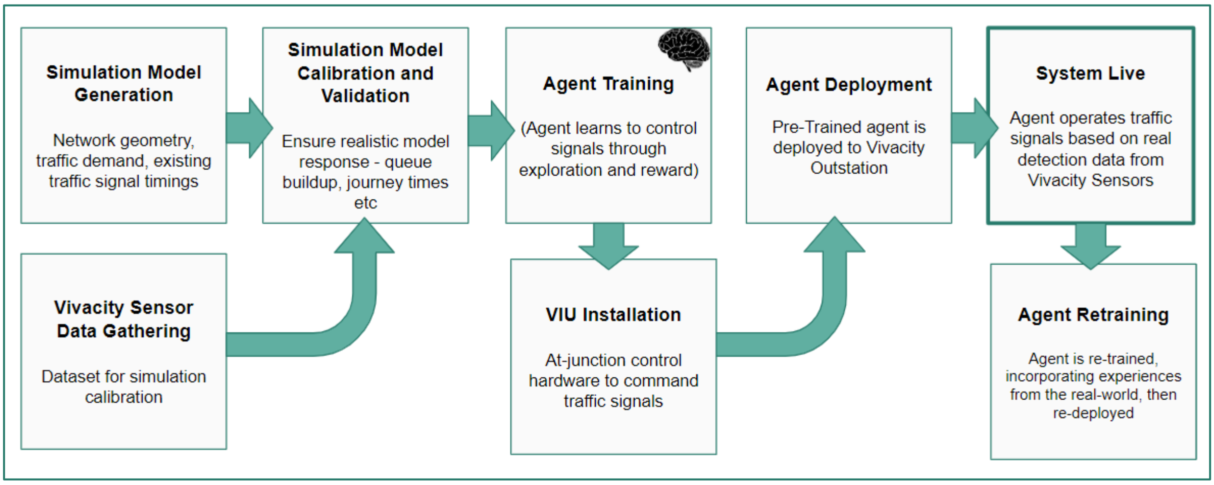

Figure 1*: Overall workflow towards a trained and deployed reinforcement learning system for traffic signal control*

Figure 2: Simulation of a junction in Salford in 3D environment

Reinforcement learning (RL) as a field has a history of learning advanced skills and even innovating new strategies. For example, Google Deepmind [recently developed](#) an RL system which reached the top league of StarCraft 2; a game which involves very quick and complex coordination amongst multiple decision-makers. We are using similar approaches to optimise and coordinate across traffic signals, where several deep machine learning models learn and interact in a structured manner to collectively optimise the overall performance of the system.

Once the system has been trained in simulation, we evaluate its performance in simulation through a series of tests. Again, by leveraging automation we're able to do this at a very large scale, resulting in distributions of performance which are more reliable than an individual simulation. For each training run of the system, we examine its performance across a range of demand levels, and for each demand level we evaluate the system against up to 100 scenarios, allowing a robust comparison of performance envelopes with benchmark systems.

One simple benchmark system we've compared with in simulation is vehicle actuated (VA), or System D. We have seen a 22% reduction of waiting time for vehicles and pedestrians for a typical week, based on observed demand levels at the first trial junction in Salford. The improvement is biggest during high demand, and a simple average across the demand levels tested results in an improvement of 34%. Figure 3 shows how this improvement is distributed across demand levels for both the RL system and VA.

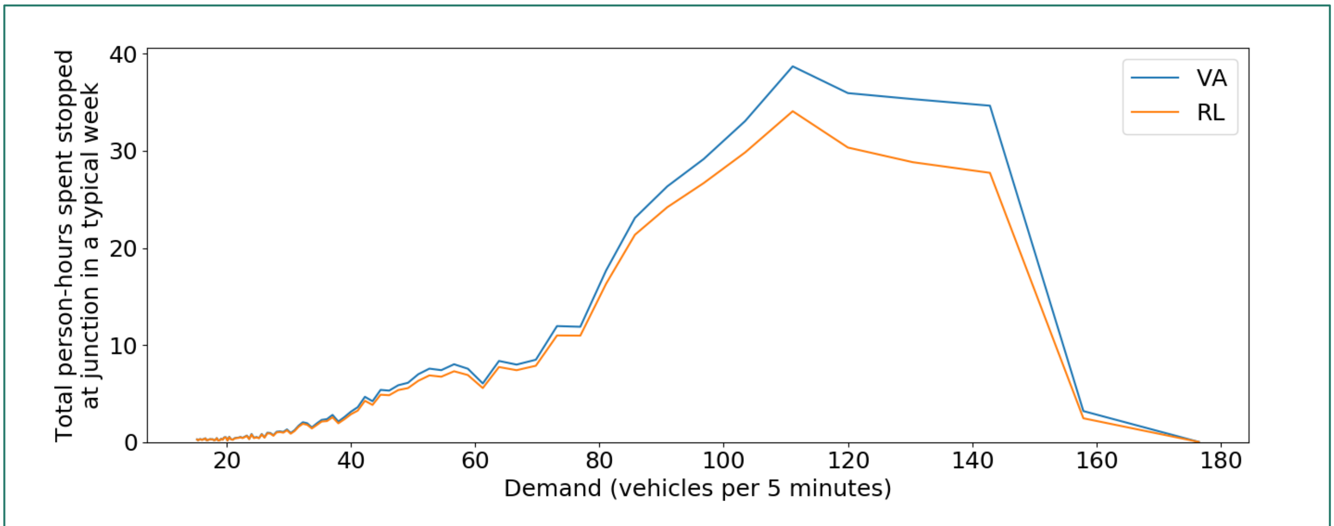

Figure 3: Simulated comparison of total time spent stopped per week under VA vs RL

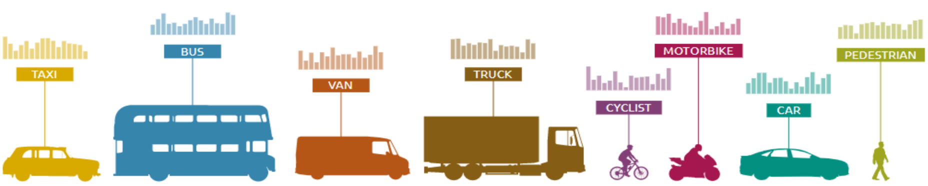

4. PATH TO REAL WORLD CONTROL

The first phase of the project focused on building an offline test bench system (shown in Figure 4) to demonstrate core functionality and system safety. We developed a custom outstation: the Vivacity Intelligence Unit (VIU), which is installed in the controller cabinet and interfaces with the traffic signal controller.

The VIU contains a powerful GPU-based compute unit which receives live data from nearby Vivacity sensors, hosts the RL control agent, and issues requests (force bits) to the controller, as well as streaming data to our cloud-based monitoring systems.

We worked with TfGM to design the system architecture and VIU itself which were engineered according to various real-world requirements, constraints, and assumptions:

- The Vivacity system shall revert control to legacy system as required (TfGM)
- Failures of prototype system shall not cause safety hazard to signal operation
- The system shall be able to be integrated with standard signal control system
- Assumed local network required for low latency sensor to VIU comms

Figure 4: *The test bench included 14 sensors. a Siemens controller with a modified configuration (not pictured) and the VIU installed in the same cabinet*

The test bench demonstrated:

- **A functioning system:** Agent taking action based on observed demand from the sensors and the VIU sending commands to the controller
- **Constraints obeyed**: The system served pedestrian stage within a certain time of the ped button being pressed
- **Controller hierarchy obeyed:** Controller only obeyed commands from the system when an SF bit is set to TRUE

Once this system safety and core functionality was proven on the test bench, we could shift focus towards real-world deployment. We worked closely with TfGM to select the initial pilot site and subsequent trial region in the Salford region of Greater Manchester which can be seen in Figure 5 below.

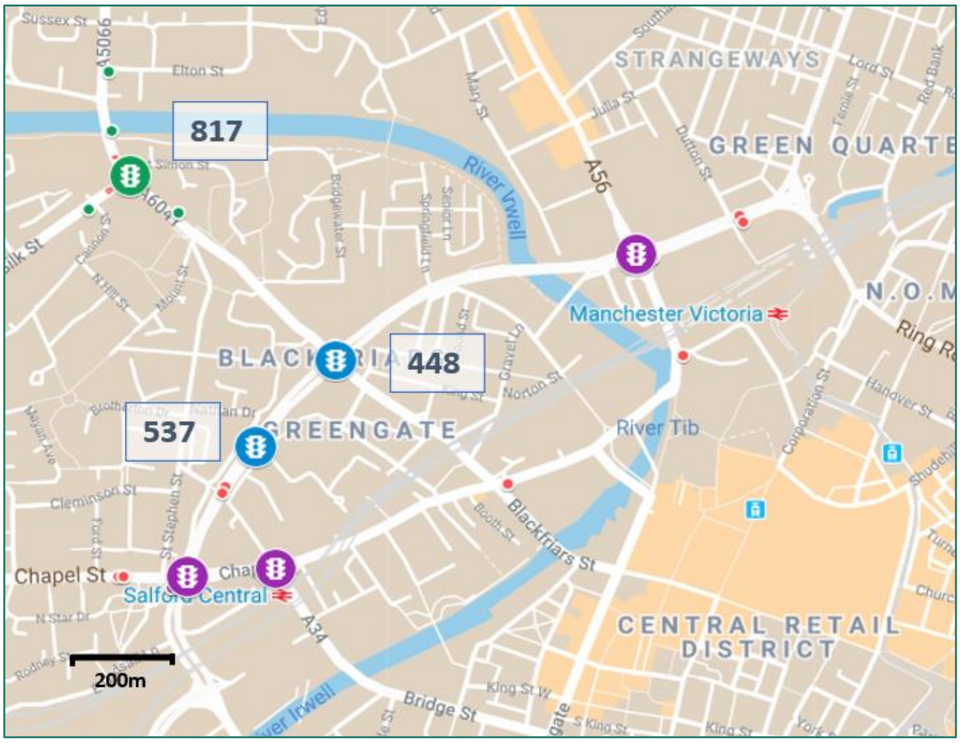

***Figure 5:** Initial trial region in Salford. The traffic light symbols represent the junctions we will be controlling this year and the small dots are the sensors that are currently installed in and around the junctions. First pilot site (junction 817) in green, next two junctions (448 and 537) in blue and subsequent three junctions in purple.*

5. First Control with AI

The first pilot site was junction 817, the intersection of Blackfriars Road and Silk St. This junction has four stages and is currently under SCOOT control and was last re-validated on 19/12/2016. Figure 6 below is an aerial view of the junction and stage diagram.

***Figure 6:** Junction 817, the first pilot site, with representative occupancy zones and stage diagram*

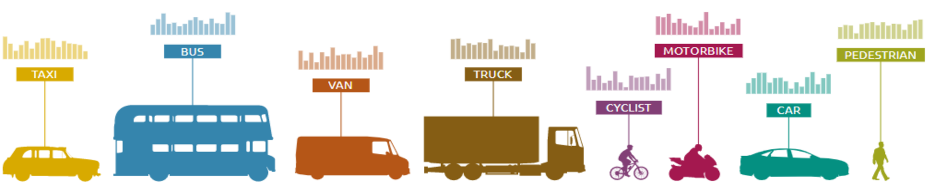

In order to commission the whole system at the first pilot site, and to ensure all core system components were functioning as intended, we employed a staged approach. We gradually increased the complexity of the control algorithm controlling the traffic signals:

1. **Fixed time:** A simple, fixed time algorithm that did not rely on sensor data, but running on the Vivacity system. This proved that the VIU system was fundamentally stable and capable of controlling the traffic signals.
2. **Max occupancy:** This algorithm brought in real-time detection data from the Vivacity sensors and closed the control loop around queue lengths. This deterministic algorithm works to simply serve the stage with the highest demand, as measured by the vehicle "occupancy" of notional queuing zones on the approaches to the junction. The algorithm also employs some backstop constraints to avoid undesirable bad traffic outcomes, such as always ensuring that pedestrian demands are served within a given time, and that all stages are at least served even if their demand is generally low.
3. **RL agent:** The first live trial of the full AI system; **successfully controlled the traffic signals at the pilot junction using Artificial Intelligence on Thursday, January 30th, 2020** from 10:00 to 14:00.

Performance of first deployment

Initial analysis on the performance of the algorithm in the real world saw a reduction of average delays when compared to the performance of the existing system at this first trial junction. Qualitatively, the agent displayed generally sensible behaviour with well-timed stage closing when queues were cleared and longer stage times when queues were longer. Pedestrians were served promptly and the much busier Blackfriars Road got overall higher priority than Silk Street. We also observed some less than ideal behaviour, such as stage 4 not being served quickly enough when only one car was queuing, or pedestrians served too eagerly when clearing a platoon of vehicles would have been a better option. These shortcomings were used to further refine the RL agent, as described below.

Continuing to iterate

After this initial control we continued to control the junction periodically, extending the hours that we were in control into the shoulder periods. Throughout these deployments we improved the system in the following ways:

- **Completely remote deployments:** The first control was done with the team physically onsite. Subsequently we improved the control monitoring interface to allow us to make updates to the system and control the traffic junction remotely, facilitating more frequent control.
- **Algorithm performance improvements**: continuing RL research, we introduced new agents that each performed markedly better than the previous in simulation. Figure 7 shows the performance improvement on average waiting time in simulation of each agent that was used on junction 817.
- **New behaviour tests**: During training of the RL agents, we incorporated certain behaviour unit tests where there is a known "right" answer, for example only a single vehicle waiting.

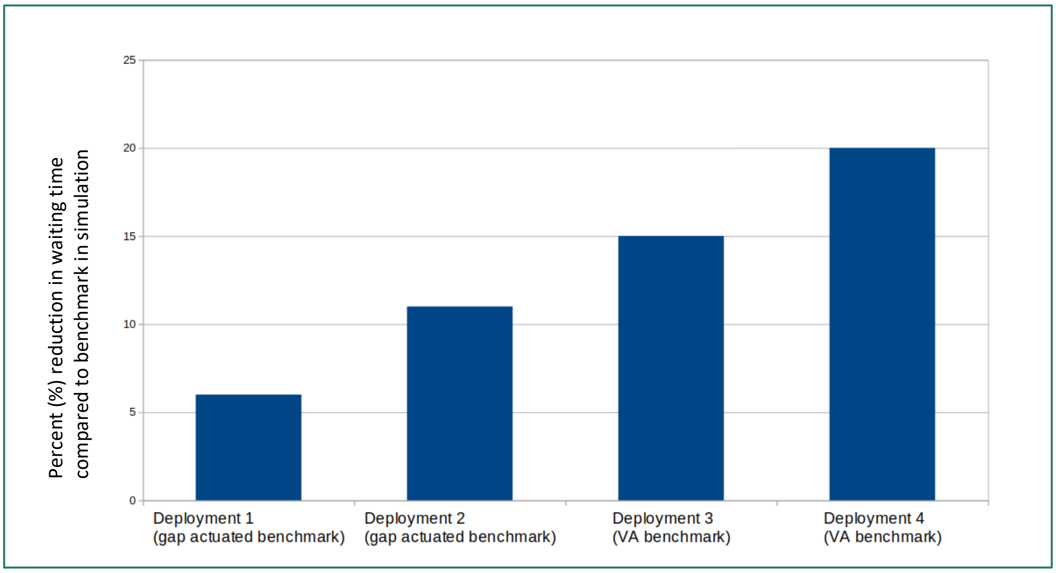

***Figure 7**: Comparison of the first 4 agents we deployed in the real world. This is their performance in simulation compared to the baseline algorithm (gap actuated and then vehicle actuated).*

Unfortunately doing a robust performance analysis on these subsequent deployments and getting a like-for-like real world performance comparisons are currently not possible due to the current pandemic and nationwide lockdown which have severely impacted traffic conditions across the country. Therefore, ongoing performance analysis of the system is being conducted as traffic conditions continue to change.

6. Towards multi-junction control

With consistent, safe control of junction 817, we set our sights on to the next big milestone: multi-junction control. After an 18-month journey to controlling the first junction, we controlled the second junction (junction 537) 5 months later in July 2020. Then, the third junction, 448, just 2 weeks after that.

Scaling to new junctions required a new calibrated simulation model of this larger region as well as improvements to the software stack to generalise it for new junctions. On the research side, bigger junctions and the challenge of coordination has introduced new complexity. We are exploring different approaches to multi-junction and coordinated control, such as comparing the performance of individual agents at each junction versus one global agent controlling an entire region.

This research has already yielded promising results in simulation and we have started to see emergent coordination of the traffic signals. Three agents working side-by-side to control three neighbouring junctions in simulation has produced a 49% reduction in waiting times for vehicles compared to VA. While this is a strong start, this result is not unexpected since VA is not optimised for coordination between nearby junctions. Going forward, we are further developing the benchmarking suite to be able to robustly compare the system's performance against industry standard algorithms.

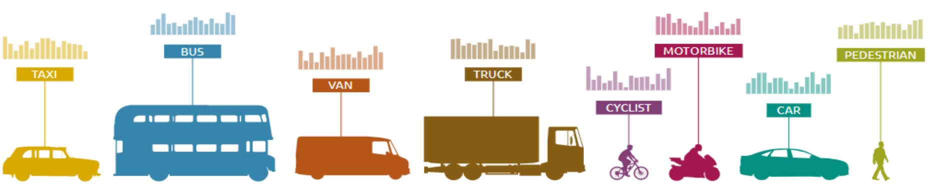

7. What is next?

Simultaneous control

Having already controlled three junctions independently, we are going to be simultaneously controlling all three junctions in early September this year, ahead of the JCT Symposium. Similarly, to the approach during the trial junction, we are taking learnings from these initial controls and improving the agents ahead of simultaneous control. During this next phase we will be looking to gather more real-world performance data in order to demonstrate real world performance improvements such as journey time reductions, as well as working towards better coordination across this three-junction region.

Large-region control

Over the coming year we will continue to scale and expand into larger control regions. More than 50 sensors are already in place on and around a 7-junction area (Figure 5) and we are currently installing sensors on 6 more junctions along Deansgate in Manchester, a heavily pedestrianised area. As we scale, we will be continuing research into regional coordinated control and also exploring multi-modal optimisation further. The expansion into Deansgate will give us an excellent opportunity to this, particularly as this is a region that does not currently use adaptive control.

Expand functionality in the Vivacity Dashboard

We have already begun to incorporate real-time monitoring features into the Vivacity Dashboard, our UI product for data visualisation. We've worked closely with TfGM and other authorities to design new features and over the coming year, we are looking to expand this functionality to create a user interface for junction and control monitoring. The current "junction view" shows the active stage as well as control and reply bits and we will be adding additional real-time data.

8. How do we believe this will compare to MOVA and SCOOT?

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

- **Ability to do more with less:** We are looking to address the budget cuts and skills shortages that many councils and authorities have faced in recent years by creating a system that allows transport engineers to spend less time handling the minutiae of the current, immediate problems and more time thinking about policy and transportation strategy.
- **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 Reinforcement Learning allows us to avoid the concept of a strict region through more adaptive junction groupings, in a much more flexible manner than existing multi-control-algorithm approaches.

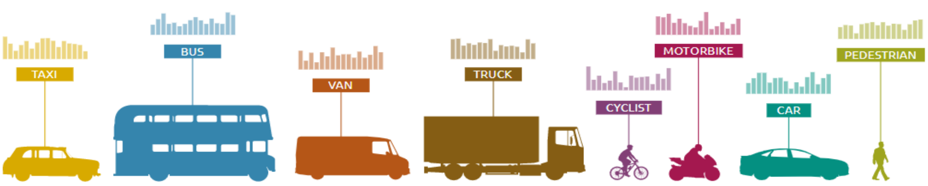

- **Auto-calibration:** Machine learning systems, by definition, improve with experience. Junctions will not need lengthy initial manual calibration periods at installation; signal control algorithm calibration can be executed automatically and continuously through a combination of simulation and real-world learning.
- **Modal and Air Quality Optimisation:** The AI can learn for a variety of different optimisation goals. 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 prioritise air quality, such as by training the AI using a proxy such as acceleration profiles or number of stops.
- **Rapid Response to changing environments:** Today, if a lane is closed or when a football match finishes, existing algorithms struggle to adapt appropriately, and typically need timetabled or 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 system automatically puts in place the right response to these scenarios when they occur in the real world.

9. HOW CAN I FIND OUT MORE?

We will be running a User Group for transport authorities on Tuesday, October 13th, 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 your regular Business Development or Customer Success Manager or reach out to Raquel at Raquel.velasco@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.