

# Smart Junctions

## Three years of developing AI for traffic signals - how far have we come, and what next?

### Overview

Three years ago, we embarked on a three-year programme to develop an AI-based traffic signals optimisation system. We, Vivacity Labs, partnered with Transport for Greater Manchester (TfGM) and Immense Simulations for this Innovate UK co-funded project to build and trial a solution which uses AI to optimise traffic networks.

Along the way, we have presented our progress at the past two JCT conferences. Last year, we were excited to share that we had successfully deployed the system to three junctions in Greater Manchester. This year, we are rounding out this three-part series by sharing some of our key achievements of the past year, summarising our overall learnings, and discussing our further development and commercialisation plans.

Our results from real-world trials have proven that AI can be used to improve on existing systems by **up to 30%**. We are scaling up these demonstrations, and expect to quantify this more precisely and thoroughly in future, but these initial results are extremely promising.

### Key Challenges:

For transport authorities today, congestion is only one of many different priorities. Improving air quality, prioritising active travel, and improving public transport reliability/uptake are all at the top of the agenda. SCOOT and MOVA have dominated traffic signal control in the UK for the last few decades, and while both have scenarios in which they work effectively, reducing congestion through coordination of multiple junctions (SCOOT) or optimising individual junctions (MOVA), they have both struggled optimising signal timings to improve air quality or to help other modes of transport.

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 running buses with green signals, 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 thus not viable for many authorities.

In this context, it is vital that transport authorities get access to better systems, which allow authorities to prioritise any objective according to their local policies.

## System Overview

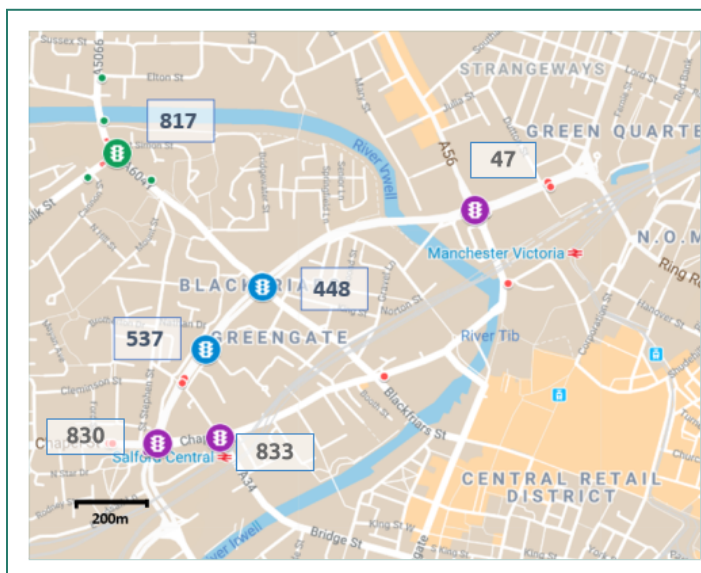
Our system is comprised of several key components:

- Vivacity’s **sensors**, which can detect and classify 9 different types of road user, providing accurate, real-time data
- A **microsimulation model**, using historic data from the sensors to achieve a very precise calibration
- Our **optimisation algorithm**, which uses cutting-edge AI based Reinforcement Learning, a branch of machine learning / artificial intelligence. This is trained in the microsimulation, to work out how best to perform in the real world, and then uses real-time data from the sensors in order to execute an optimal control strategy, adapting quickly to changing traffic conditions and optimising at both local and city-wide scale

## Key Milestones over the past year

### Simultaneous Control of Five Junctions

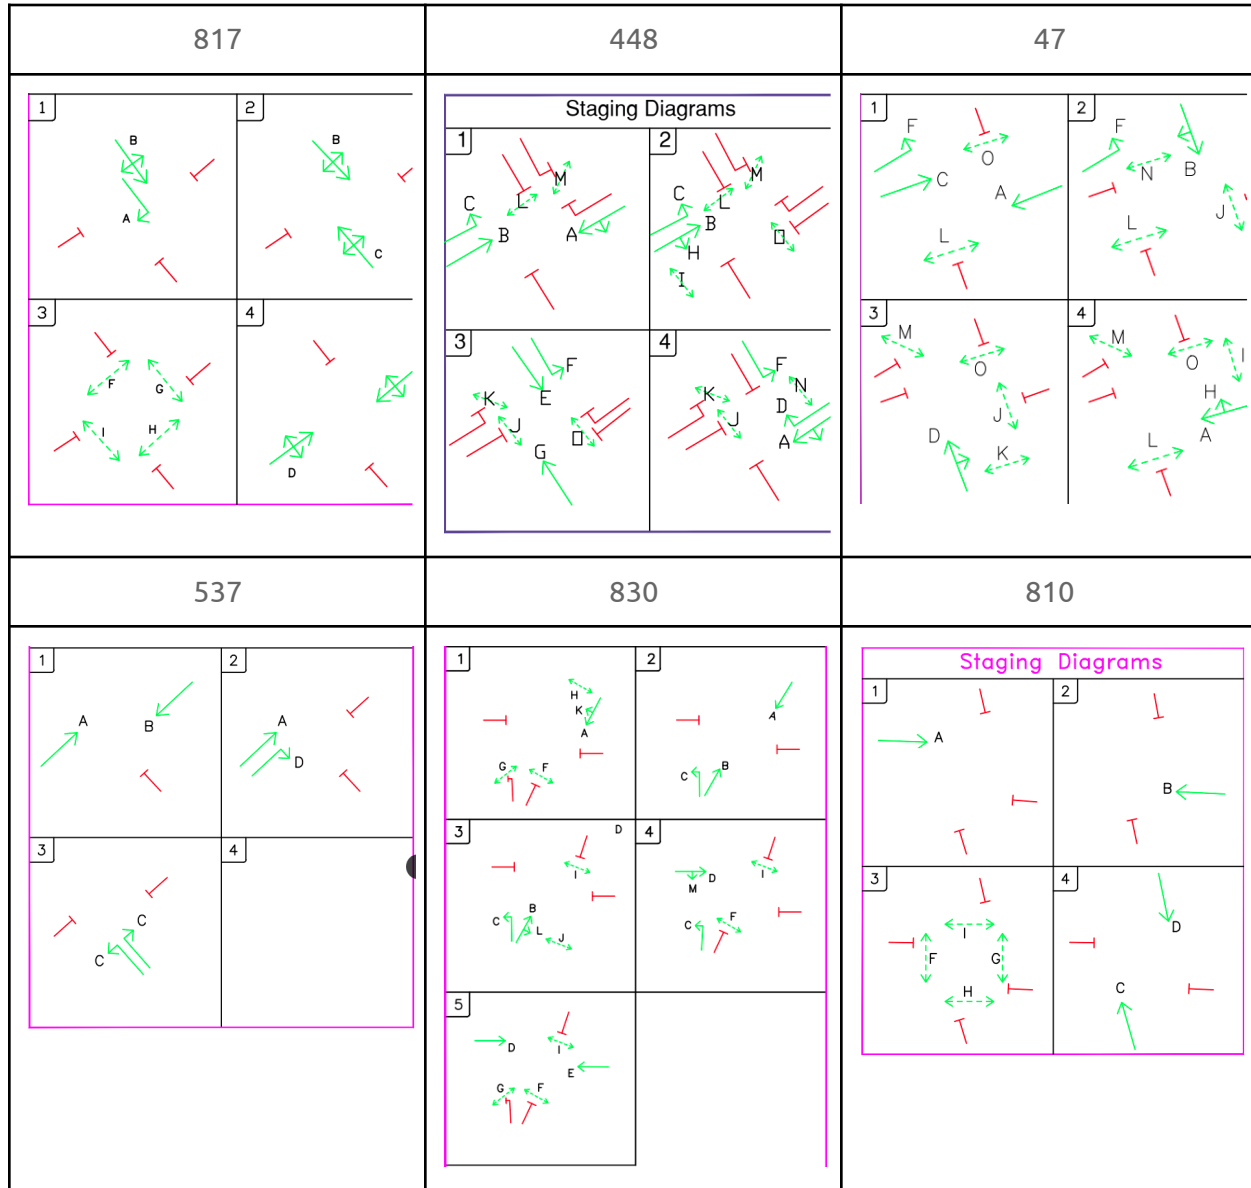
After independently controlling three neighbouring junctions last year, we scaled the system to control a whole key corridor on the main orbital route on the Manchester / Salford border. Figure 1 below shows the location of each of these junctions.



**Figure 1:** Initial trial region in Salford. The traffic light symbols represent the junctions we controlled this past year. First pilot site (junction 817) in green, next two junctions (448 and 537) in blue and subsequent three junctions in purple.

3

These junctions had a range of different layouts, characteristics, and challenges; the stage diagrams for these can be seen in Figure 2 below. We were able to achieve coordination along the corridor, with a single AI agent or algorithm, which had been trained on a large microsimulation of this region, making decisions for all 5 junctions (analogous to a SCOOT region).



**Figure 2:** Stage diagrams for the six junctions in Greater Manchester where the system has been deployed. This highlights the range of complexity of junctions that we are controlling.

## Continuous, unsupervised control

In order to gather significant performance data and to allow us to scale more effectively, our next development target was to improve the robustness of the system to allow us to leave the control pipeline running for long periods of time without the need for constant, human monitoring. In order to achieve this we developed a set of alerts to notify the users (Vivacity internal team initially) about any system or performance issues. We tuned the alerting parameters to adjust the agent configuration and the alert thresholds and build confidence in the completeness of the alerts. The key outcome is that we are currently running unmonitored, A/B testing at one junction in Manchester, Trinity Way – Great Ducie Street(x47), controlling from 7am to 7pm, and adapting automatically to changes in demand level and turning proportions (i.e. no need to timetable plans as in incumbent systems).

## Control Algorithm Improvements

Over the past year, we have carried out a detailed research programme to improve the operation of Reinforcement Learning for signal control, building on our work last year. Combining the latest academic research with innovations of our own, we have seen substantial improvements in underlying performance, with the latest results in simulation showing a 45% reduction in average vehicle waiting times across a 6-junction region when compared against vehicle actuated (VA) control or System D control.

## Scaling the system to new cities

As well as growing the number of junctions we are deploying our AI-system control system to within Greater Manchester, we have also partnered with Cambridgeshire County Council and Peterborough City Council to roll out the system to five more junctions across the UK.

The specific sites that were chosen have significant congestion challenges, made more complex by the significant number of pedestrians and cyclists that also use those junctions.

As a result, the key objectives for these new sites were to improve traffic congestion and then, to proactively promote sustainable transport over car use. Our hope is that we will be able to quantify the impact of our AI-based system on traffic through these trial projects to aid in building a business case for subsequent roll outs of this technology across both regions.

Some of the key things we've done in order to expand to these new cities:

- **Simulating a wider range of junction layouts and features**, including cycle early release stages, crossings controlled by UTC pelican bits, phases which terminate after their minimum green; Multi-stream controllers, etc
- **Integration with Dynniq controllers**: parse their configurations, emulate their operation in simulation and interface with them in the real world, as we already can with Siemens controllers

5

- **System able to support different traffic signal control policies across different cities.** For example, Manchester permit us to use a variable Stage Order, whereas Cambridge and Peterborough require a fixed, predictable cyclical stage order
- **Manage multiple cities concurrently;** improved the architecture of our system in order to monitor performance and roll-out improvements to each independently

## System Improvements this year

### Improved process for end-to-end system deployment

We have streamlined our initial site scoping process, with a team now able to accurately determine the number of sensors required in a single site visit and find the optimal balance between sensor coverage for optimal performance and cost. Increased automation means we have reduced the time required to configure, calibrate and validate a simulation by more than half after sensor installation.

We are able to commission the control hardware and test our control pipeline (end-to-end sensor to instation to controller) for a new site by deploying a simple deterministic algorithm prior to AI being trained, allowing us to investigate any communication or network challenges prior to AI takeover. The initial AI deployment is monitored by both our team and the customer to ensure the agent is behaving as expected and the algorithm meets performance expectations for performance whilst respecting safety constraints (i.e. respecting desired stage order). Post-commissioning, the system can be left unsupervised, whilst performance is continuously measured and monitored. This is followed by regular updates to the AI agents as we gather more data and add new features and functionality.

### System and Performance Monitoring tools

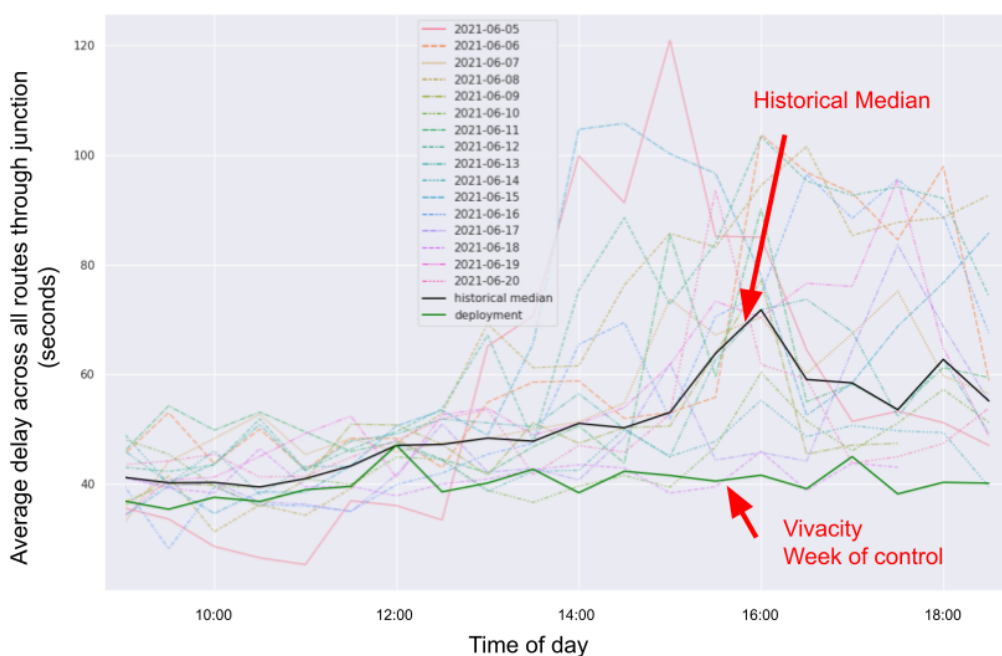
To effectively monitor, improve and respond to issues we have developed monitoring and alerting systems to allow us to react to the following:

- **Traffic conditions** - by consuming data from our sensors we can visualise traffic conditions in real time to validate agent actions and to determine the best course of action should a traffic incident arise. We can utilise our vision-based sensors to take low-latency, blurred images to investigate traffic incidents.
- **Infrastructure performance** - by monitoring our data pipeline, end-to-end, we monitor and react to infrastructure issues in **Network performance** (4G, broadband, etc.) and **System performance**
- **Algorithm performance** - by performing continuous A/B testing we are able to measure the performance of our agent against incumbent control strategies by comparing journey time data which is available from our sensors. This will also expand to include delay for different classes.

- **Alerting** - as we shift from manually supervised control to unsupervised control and beyond we have built an alerting framework which delivers categorized alerts to different channels and operators. This provides both the notification of potential incidents and contextual information linked back to our monitoring dashboards which enables operators to make informed interventions where necessary.

## Current performance headlines

Early results from continuous control have shown **reduction in average journey times of 30% across key routes** during control periods. This can be seen in Figure 3 below.



**Figure 3:** This figure illustrates the performance of our system across a full week of 7am-7pm control, comparing the average delay of each vehicle through a single junction (x47) for multiple days (dashed lines), the median of these historical values (black line), and the average delay while our Smart Junctions system was in control (green line). It is clear that we had a noticeable impact on average delay through this junction. We recognise that SCOOT is designed for multi-junction control, and so a truly representative trial needs to compare vs a full SCOOT region; see below for more on our future plans.

## Future Plans

### Research Focus

While this paper outlines the work that we have done to date on traffic signal control research, we are still investing heavily into further research in this field. In particular, we have a strong emphasis on multi-modal control, and optimising for the wide variety of scenarios that authorities may implement in future. We will continue using delay (calculated from journey times for all routes through a region) as the main metric for performance evaluation of vehicle congestion reduction, while adding stopped time for pedestrians and cyclists as a second measure later this year. To complement our own performance numbers, we will endeavour to have an independent party evaluate our performance as well.

### Real-time monitoring Dashboard

As part of the full product offering, we are doing further development work to enable three key workflows for Network Managers from within the Vivacity Dashboard: Live Junction Monitoring, Automated Congestion Alerts, and Performance evaluation

- We are taking our internal **monitoring tools** and exposing them to our clients. This will complement the standard offering for our sensors.
- We are developing and releasing **additional data features** which will be accessible through the dashboard and API such as Speed, cumulative stopped time, alerts, and zonal occupancy.
- **Analysis tools** to enable performance comparisons and in-depth evaluation of the wide range of different data sources.
- **Notifications** in the Vivacity Dashboard as well as integration with other systems

### Further roll outs

In the coming months, we will be rolling out the system across new regions, including a larger region of Manchester and Salford.

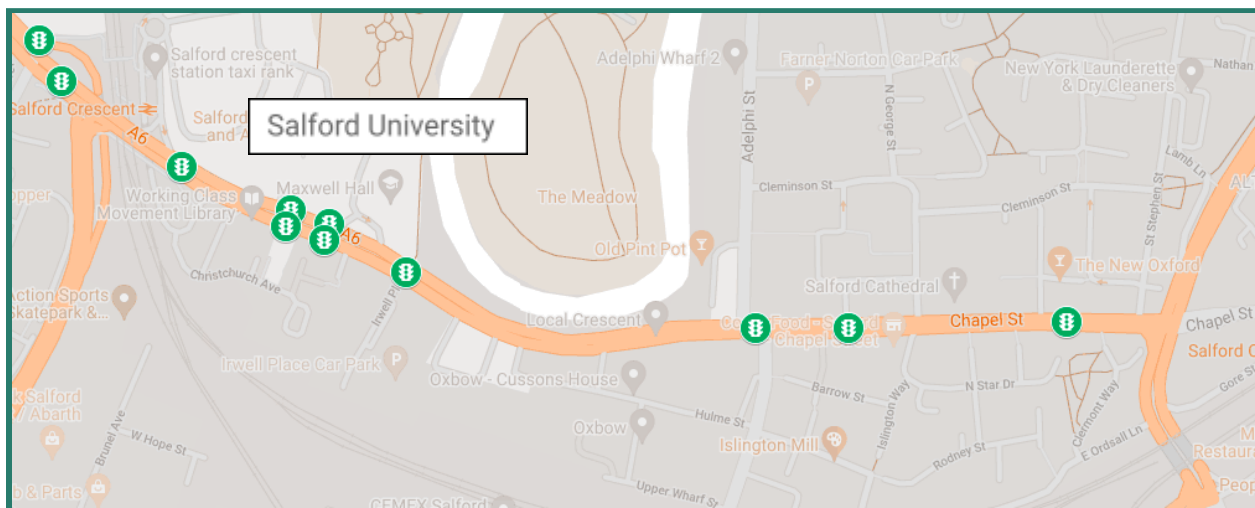
Initially, we will be controlling several junctions on **Deansgate in Central Manchester**. This is a high street environment with a high proportion of active travel modes – pedestrians and cyclists. We will be exploring how to balance the competing demands of different road users, finding the trade-offs between promoting one or more modes whilst keeping others within acceptable delay limits. We will be implementing intelligent multimodal optimisation, rather than a brute-force, modal override such as bus priority.

In addition to the initial IUK programme in Greater Manchester, this project has secured additional investment to expand throughout the trial region via the Department for Digital, Culture, Media, and Sport's '5G Create fund' which was announced in July 2020. We will be scaling the Smart Junctions network up to an area of 20 junctions in Greater Manchester by



the end of this year through this 5G project. Its aims are to demonstrate impact in the real world of 5G in the form of improved journey quality for all road users in this region.

The new junctions are along a **major arterial route linking Salford and Manchester**, which exhibit significant tidal traffic flow patterns (Figure 4 below). The use of 5G will minimise the latency of the system, allowing the system to receive data, make decisions and send instructions more quickly.



*Figure 4: For the Smart Junctions 5G project, we will be controlling this region along the A6.*

### Commercial Release

In early 2022, we will fully launch our first product for signal control, which takes our learning to date and is focused on congestion reduction. This system will make use of our camera sensors and provide our full, rich data to the system. This data will also be available to clients via the dashboard and API.

Clients will be able to monitor the junction(s) behaviour and performance through the Vivacity dashboard. The dashboard will include tools and panels to view real-time, low latency data on controller state, algorithm actions, and demand of each stage. Additionally, users will receive alerts on congestion incidents and will be able to use the camera sensors to validate and investigate incidents.

As well as real time monitoring and alerting capabilities, the dashboard will have enhanced data analysis tools to allow users to quantify performance of the system and compare journey times and dwell times of different control methods.