30 cars, figure of 8, 1 show: large scale proving ground experiments to investigate junction control

Simon Box, John Lees-Miller, James Snowdon, James Hammond
Andy Hamilton, Shashank Gupta, R. Eddie Wilson, Ben Waterson

January 9, 2013

Abstract

An experiment was conducted using the InnovITS proving ground in Nuneaton. Thirty cars with volunteer drivers were asked to drive around a tight closed road circuit causing them to pass repeatedly through a cross-roads junction from all directions. The junction was signalized. In different test-runs of the experiment the traffic lights were controlled by either an automated system or by a human using remote control. All vehicles in the test were instrumented using GPS and bluetooth. Video footage from two cameras was also recorded.

The goal of the experiment was to collect data to validate the results of earlier work carried out in computer simulation. This earlier work indicated that human controllers could outperform commonly used automated systems.

This paper examines some of the issues that arise when trying to simulate an urban road junction in this manner. For example results are presented indicating differences in network performance depending on whether the drivers were instructed to follow a fixed route or a random route of their choice. Thus providing some guidance for maximising the fidelity of this type of simulation in the future.

The paper also presents a detailed analysis of the sensor data and video footage to measure the performance of the junction under the different modes of control. The results support the conclusion that humans can be effective junction controllers.

1 Introduction

1.1 Background

In previous research humans have been asked to control traffic junctions within microsimulations via a computer game interface, like the one shown in Figure 1 [Box, 2011]. Video footage of the game being played can be found here: `http://youtu.be/dB3zoGJiODk`. The results of this research have indicated that players of the game *can* be very good junction controllers and even outperform some of the market leading automated junction control systems, such as MOVA [Vincent and Peirce, 1988] and SCOOT [Hunt et al., 1982]. This particular result has prompted

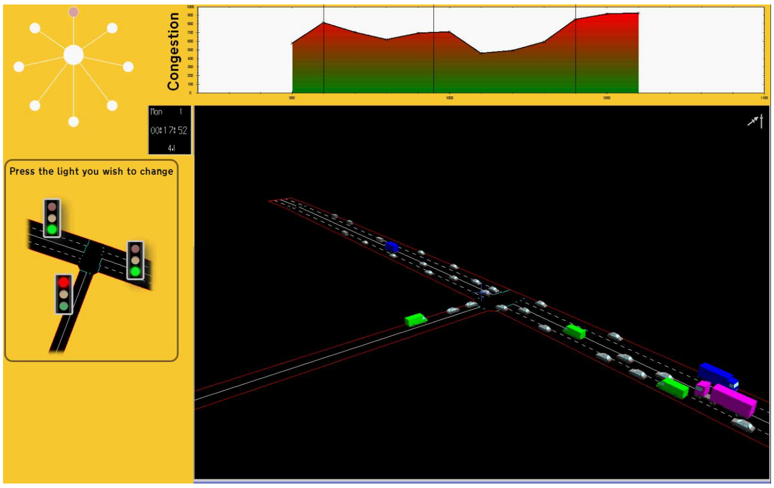

Figure 1: Screenshot of junction control game.

the development of machine learning junction control systems that can learn control strategies from human players of the game and from experience [Box and Waterson, 2012a,b].

1.2 BBC One Show Opportunity

Since obtaining the above results using computer games the authors have been keen to test the hypothesis that humans can be good junction controllers in the real world. An opportunity to do this came about when the authors were approached by the BBC to be involved in the making of an article on traffic control for the television programme "The One Show". The concept of an experiment where a human controls a signalised junction was proposed and the BBC commissioned the article. This approach to funding a scientific experiment has both benefits and drawbacks. The chief benefit was that the time scale to obtain funding and perform the experiment was *much* shorter than one can typically expect from research council funding. The drawback is that the dual purpose of the event meant that compromises had to be made regarding the scientific method of the experiment in order to accommodate the filming. These will be discussed throughout the paper.

1.3 Problems of the Experiment

There are two ways that the proposed experiment could be performed. Either a human could take charge of a real road junction, or a junction in the controlled environment of a proving ground. While the former was investigated, obtaining approval to do this from the appropriate local authorities proved impossible, so the proving ground was the only available option in this case.

There are several advantages to using a proving ground for this type of experiment. It is possible to instrument all the vehicles involved, providing rich data with which to evaluate the experiment. It is also possible to control the demand on the junction and standardise this between repeated tests. It is also possible to issue drivers with behavioural instructions.

However there is still the fact that a proving ground experiment is a *simulation* of a real world junction and while much closer to the real thing than a computer simulation there are still some differing characteristics that must be considered. For example, in order to limit the number of drivers and vehicles that are required for the experiment it is desirable to maximise the use of each vehicle by causing them to drive in a closed loop course and pass repeatedly through the junction. To quantify this problem consider that in the experiment discussed in this paper thirty vehicles were used and in each *run* of the experiment, each vehicle passed through the junction on average twelve times. To achieve the same flow of vehicles without a closed loop would require three hundred and sixty vehicles and drivers.

The closed loop course may reduce the fidelity of this simulation, for example it could lead to feedback effects. When the junction controller releases a vehicle the timing of this decision (weakly) determines when that vehicle will arrive at the junction again. Feedback effects could manifest in a number of ways, including causing vehicles to become 'synchronised' with the junction and encouraging 'platoon' formation. Such effects are difficult to detect and measure conclusively but efforts were made and these are discussed in Section 3.

It is possible that closed loop effects can be managed (or exacerbated) by the instructions that are given to drivers. Specifically, whether drivers are asked to drive a prescribed route or asked to choose their route 'at random'. Both approaches were investigated in this experiment as described in Section 2. Analysis of the difference between these approaches is given in Section 3.

2 Method

The experiment took place at the InnovITS "city circuit" test track (`http://www.innovits.com/`), which is part of the MIRA facility (`http://www.mira.co.uk/`) in Nuneaton, UK. Figure 2 shows a map of the city circuit. Cones were placed around the track in order to restrict the vehicles to within the area marked by the boundary shown in Figure 2. Thus a 1-km long figure-of-eight track was formed with a signalised junction at the central node. Cones were also placed where necessary to restrict the width of the road to a single lane in each direction. Thirty volunteer drivers and vehicles were available for the day of the experiment; the mix of vehicles was 25 cars, 2 motorcycles and 3 minibus vans.

2.1 Track Topology

As discussed in Section 1.3, during some of the tests performed we needed to ask drivers to follow a prescribed route around the track. It was considered important to give the divers clear, easy to follow instructions. So each driver's route had to be represented by a single line drawn on a map of the track that does not go over any stretch of road more than once **in either direction**.

It was also important that the set of routes included all possible turning movements through the signalised junction in the centre of the figure of eight. This raises the following question:

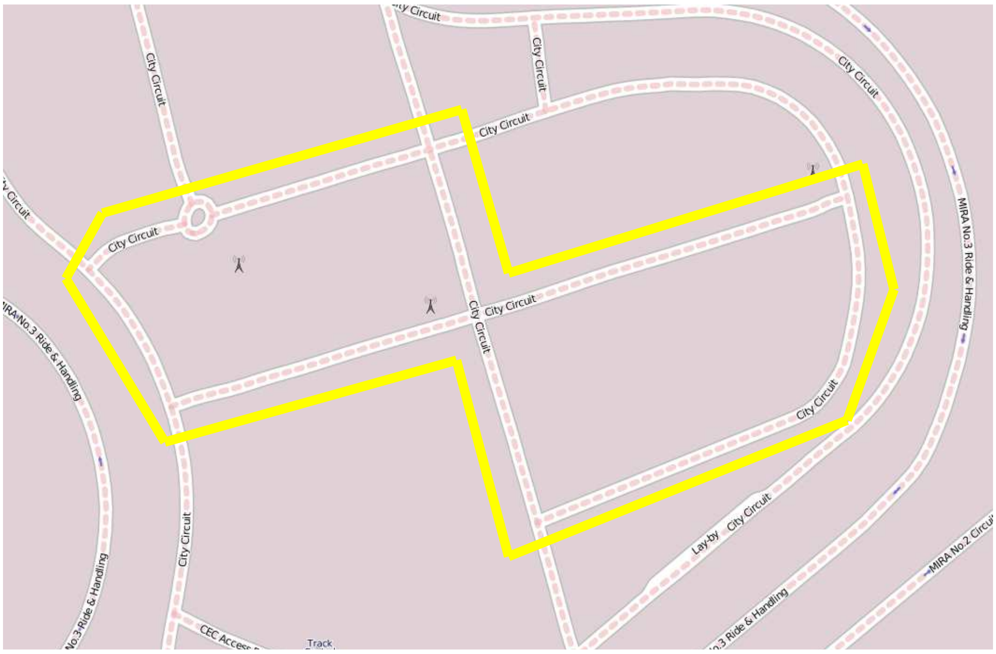

Figure 2: Map of the test track layout.

> "What is the minimum number of routes that include all turning movements while satisfying the constraint that routes must not cover the same stretch of road more than once in either direction?"

Figure 3 shows the topology of the track. On the left there is a graph of the track which has a single node over the junction and four edges that correspond to the clockwise and anti-clockwise routes around each loop in the figure of eight. On the right is an inversion of this graph where the edges in the track graph are now nodes and the node in the track graph is now represented by edges, one for each turning movement through the junction. Our routes can be seen on this turning-movement graph as routes which do not pass through any node labelled **A** or **B** more than once. In this case the routes are represented by the eight simple loops in the graph.

Figure 4 shows the eight loops in the turning-movement graph drawn as track graphs (starting from the top left loop and proceeding clockwise around the turning-movement graph). These eight routes were distributed evenly (as far as possible) among the drivers participating in the experiment. As a point of interest only: if we relax the "either direction" constraint and allow routes to cover the same road section provided they are going in different directions then the number of routes reduces to six, because routes 2 and 6 in Figure 4 can be combined as can routes 4 and 8. On the turning-movement graph these correspond to the paths **A,B,A',B'** and **A,B',A',B**.

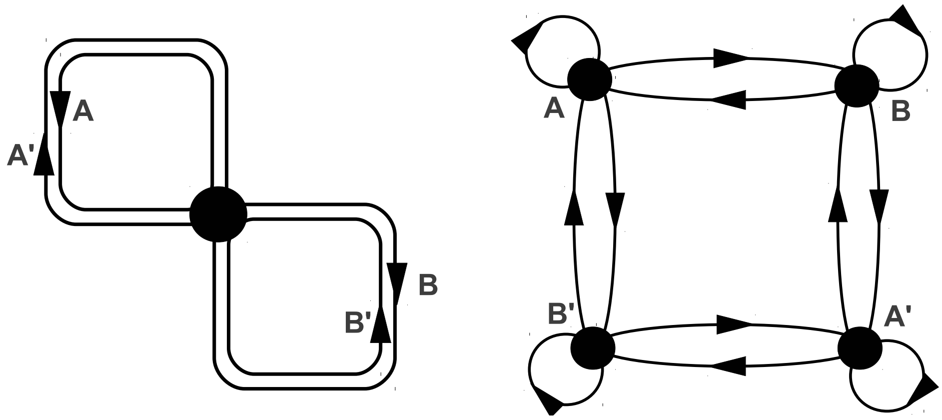

Figure 3: Graph of track topology.

2.1.1 Control Methods

The goal of the experiment was to collect evidence on how well a human can control a signalised junction therefore two control methods were required, a remote control method allowing a human to control the traffic lights and an automatic control method to act as the *Control* (as in baseline) case.

Automatic control The signalised junction at the test track was equipped with the MOVA [Vincent and Peirce, 1988] automatic control system, which is an adaptive control system that uses data from inductive loop sensors under the track to inform its control strategy. However during the first few minutes of the pre-experiment warm up tests the junction control computer suffered a hardware failure and defaulted to an optimized, fixed time control strategy. This was quite unfortunate as this resulted in a lower performance baseline than desired. Due to the constraints of the filming we had to perform all our testing in one day and did not have to option to repeat or reschedule the experiment.

Human control The signalised junction at the test track is a two stage junction that gives the green light either to vehicles on the North-South running road or the East-West running road. The human controller was presented with a very simple control interface on a laptop computer, shown in Figure 5. This has two large buttons allowing the human controller to select the two stages. Also after a stage is selected both buttons are greyed out for 10 seconds. This was a safety measure to prevent the human controller switching stages too frequently. The human controller was located on a gantry above the junction giving them a good aerial view of the scene (See Figure 6).

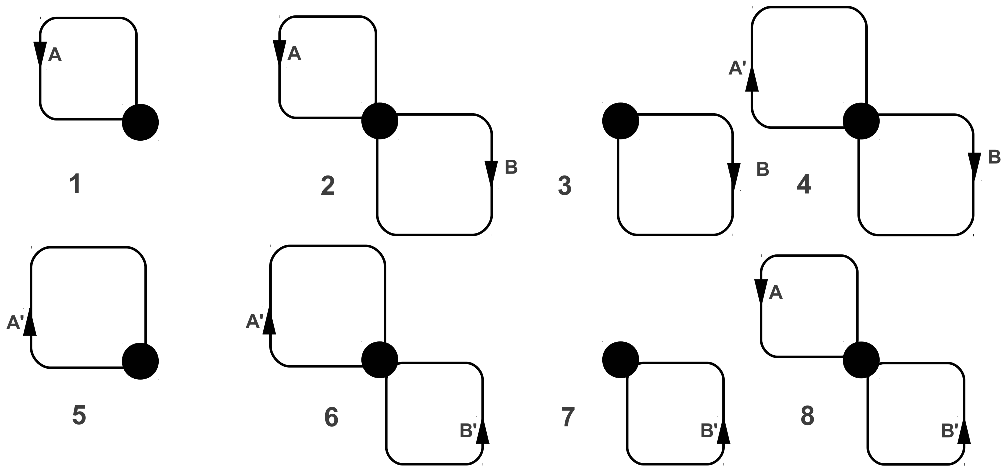

Figure 4: The eight routes provided to drivers during the test.

2.2 Testing Plan

The time available on the day of the experiment allowed for us to give the drivers a short warm-up period where they could drive 4-5 laps of the track and get used to the circuit followed by four 15 minute experimental test-runs with rest breaks in between.

The original plan for the four test-runs was as follows: The first two runs were to be under automatic signal control, where during one run the drivers would be asked to follow a fixed route and during the other the drivers would be asked to pick their own route at random. We were then to analyse the results gathered during the first two runs in order to decide which approach was a better simulation of real world traffic. The third test was to be a repeat test with the chosen instruction method, under automatic signal control. The fourth and final test was to be conducted using human control, and the chosen instruction method.

However on the day, when the results of tests 1 and 2 were analysed, there was evidence of unrealistic behaviour in *both* tests. We postpone detailed discussion on these behaviours until Section 3, but in short: when drivers were asked to follow fixed routes there was evidence of feedback effects, including platooning and synchronisation; when drivers were asked to pick routes at random there was evidence that some drivers (deliberately or subconsciously) chose to take the least congested routes, leading an unrealistically homogeneous distribution of vehicles and free flowing conditions.

At this point the authors suspected that the two instruction sets we had devised may in fact lead to extreme “corner-cases” of behaviour and that a more realistic approach would be a hybrid strategy where some drivers are asked to drive routes and others are allowed to pick their routes at random. Unfortunately the constraints of the filming schedule did not allow us to repeat these tests to try out different hybrid strategies so a decision was made to try a hybrid strategy of 25 routed drivers and 5 random drivers in the last two tests. This is clearly

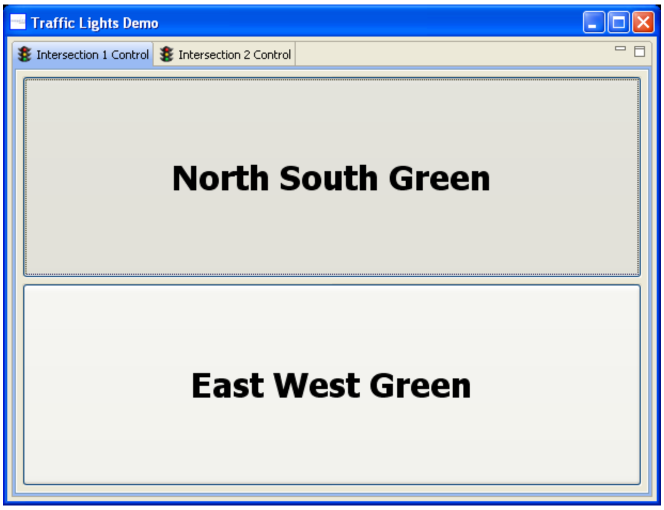

Figure 5: The remote junction control programme user interface.

not the ideal application of the scientific method but it at least gave us the opportunity collect this data and a chance of obtaining more realistic data in the last two tests. A summary of the final testing schedule used on the day is given in Table 1.

Test #	Instruction	Control method
1	Routed	Automatic
2	Random	Automatic
3	25 Routed + 5 Random	Automatic
4	25 Routed + 5 Random	Human

Table 1: Schedule for the four 15 minute test-runs of the experiment.

2.3 Sensors

The track and the vehicles were equipped with a number of sensors for monitoring the experiment. Each vehicle was equipped with a **Qstarz Bt-q1000xt** GPS logger. This logged GPS data at 1 Hz and also broadcast a bluetooth MAC address, which allowed it to be detected by six track side bluetooth sensors. There were also two high-definition cameras recording the scene on the track. The results presented in this paper are all derived from the GPS data. Results from the bluetooth sensors are reported separately in Lees-Miller et al. [2013].

3 Results

There are number of metrics that can be extracted from the GPS data to evaluate the performance of the junction. Four of these are plotted in Figure 7. These are:

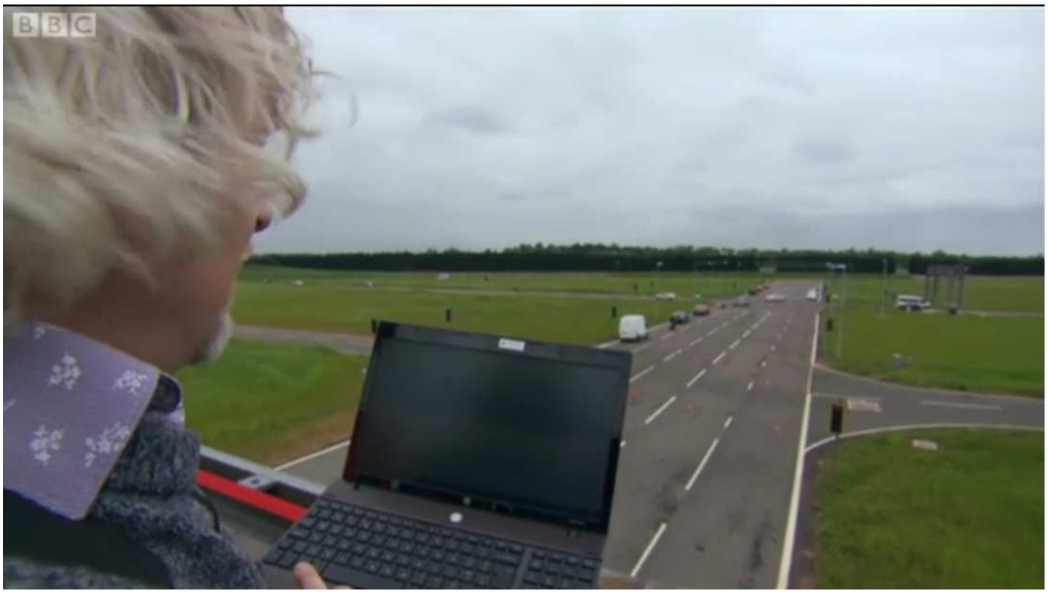

Figure 6: TV presenter Marty Jopson controlling the junction by remote control.

Distance Travelled. The closed loop nature of the test track means that the total distance travelled by vehicles during the 15 minute test run period is a proxy for junction performance. The further vehicles travel the less they are delayed during the test. The top-left plot in Figure 7 shows total distance travelled averaged across all vehicles in each test run.

Average Speed. Average speed, averaged across all vehicles is proportional to distance travelled, so the top-right plot in Figure 7 gives the same basic information as the top-left plot. This is presented to aid interpretation.

Time Stationary. Using the GPS data we can measure the total time that each vehicle spends stationary (i.e. waiting at a red light or in a queue) during the test. The bottom-left plot in Figure 7 shows total time stationary averaged across all vehicles in each test run.

Delay (Gated). When evaluating the performance of junctions in the real world and in simulations of large networks it is common to consider only a sub-region of the network surrounding the junction. To do this in our tests we implemented a virtual 'gating' approach where we consider 'trips' of vehicles to begin when they enter a region of 60 m radius around the junction and end when they leave. Travel time for these trips is measured and averaged over all trips. The delay shown in the bottom-right plot in Figure 7 is calculated from the averaged travel time by subtracting the (estimated) free-flow trip time of 15 seconds.

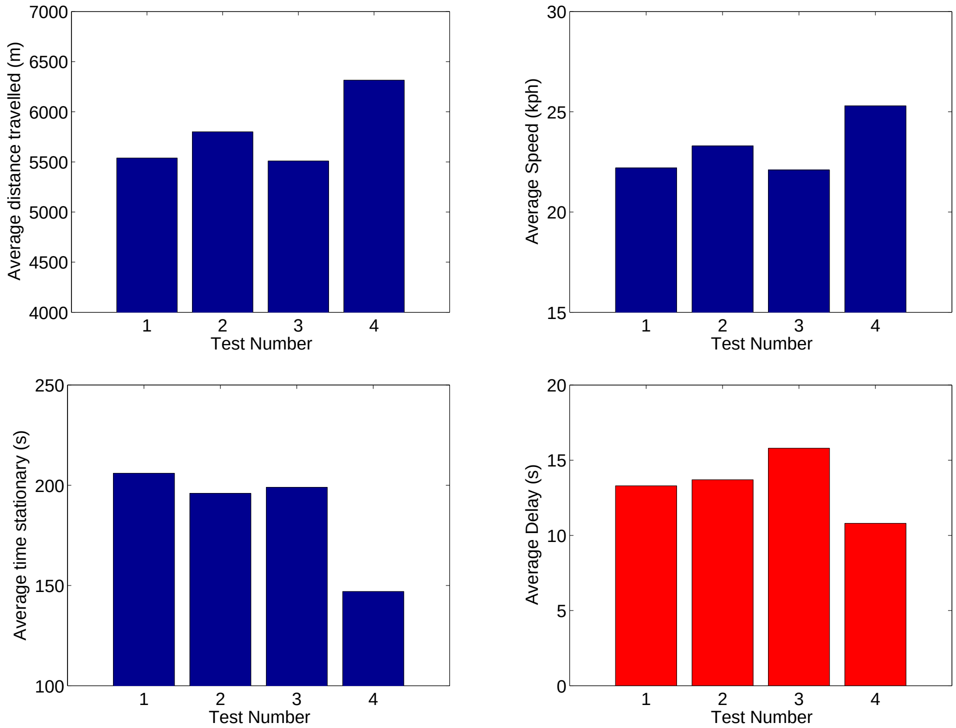

Figure 7: Four statistics for junction performance and their values in the four tests

3.1 Random vs Routed

With the exception of the delay metric, which will be discussed in Section 3.3, all metrics in Figure 7 indicate that the performance of the network is "better" during test 2 than during test 1. However as detailed in Table 1, the only difference between these two tests is the instructions given to the drivers. In Test 1 drivers were asked to follow a prescribed route, in Test 2 they were asked to pick their own route randomly.

A visual analysis of the video footage from these tests and of animations of vehicle movements generated from the GPS data suggested some possible explanations for this. In test 1 it appeared as if many of vehicles formed into platoons that were making the same turning movement at the junctions, suggesting some synchronisation effects. In test 2 it appeared that the opposite was true and that drivers preferred to make a different turning movement to the vehicle in front. It also appeared as if drivers were generally choosing routes to avoid congestion leading to a more homogeneous distribution of vehicles in test 2. It seemed reasonable that drivers in test 2 were (consciously or subconsciously) picking their route in order to minimise their personal experience of congestion and delay. If this was the case we might also expect drivers in test 2 to make less right turns, because of the increased delay associated with this turning movement.

The visual analysis of the data led us to propose three hypotheses to explain the increased performance of the network in test 2. These are:

Hypothesis 1 Vehicles in test 2 were more homogeneously distributed (less platooning) than in test 1.

Hypothesis 2 Vehicles in test 2 were less likely to make the same turning movement through the junction as the vehicle in front.

Hypothesis 3 Vehicles in test 2 made less right turning movements than vehicles in test 1.

We are able to extract evidence from the data to either support or refute these three hypotheses and this evidence is presented and discussed below.

3.1.1 Evidence for Hypothesis 1

Evidence about how homogeneous or platooned the distribution of vehicles is can be obtained using a *k-nearest neighbour* clustering analysis. At any time-stamp within the GPS data, a snapshot of vehicle positions can be taken. A clustering coefficient can be calculated using (1).

$$R = \frac{\sum_{i=1}^{i=k}(\bar{X}_i)N}{kL} \tag{1}$$

where $\bar{X}_i$ is the distance distance from a vehicle to its i^{th} nearest neighbour, averaged across all vehicles in the snapshot. N is the total number of vehicles in the snapshot and L is the length of the test track.

R will have a value close to 1 if the vehicles are evenly distributed around the the track and a value close to 0 if vehicles are tightly clustered into platoons of $k+1$ (or more) vehicles. The results presented below used a value of $k = 2$.

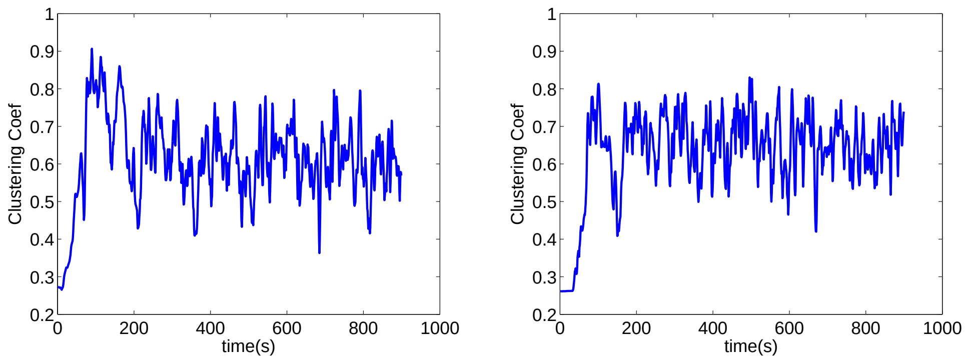

Figure 8: k-nearest neighbour clustering coefficient ($k = 2$) for test 1 (left) and test 2 (right). The clustering coefficient is calculated for every 1-second snapshot of vehicle positions and plotted over the duration of the tests.

Figure 8 shows the clustering coefficient (R) calculated for snapshots at 1 second intervals over test 1 (left-plot) and test 2 (right-plot). The clustering coefficient is low when vehicles are queueing at a red light and when vehicles are driving in close platoons. The plots indicate that vehicles are very clustered at the start of each test. This is because the vehicles were held at an *all-red* signal before the test began, so are sitting in tightly packed queues. As test 1 progresses the vehicles spread out and the clustering coefficient rises to a peak after around 1 minute. The clustering coefficient then reduces as vehicles form queues/platoons and from this point on clustering varies about a stable average. The same thing happens in test 2 except that the stable average is about the same level as the initial peak, suggesting that there is less queue/platoon formation after the vehicles have initially spread out.

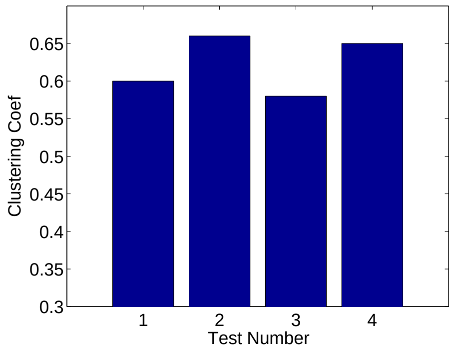

Figure 9: k-nearest neighbour clustering coefficient ($k = 2$) for for all four tests and averaged over the 900 1-second snapshots taken for each test.

The clustering coefficient averaged across all 900 snapshots for each test is shown in Figure 9. This shows that indeed the clustering is higher in test 1 than in test 2. The evidence from these data supports hypothesis 1.

Interestingly the clustering coefficient for test 3 indicates that clustering returns with the driver instructions for test 3, which makes sense because the majority of drivers in this test are following fixed routes. However platooning reduces again in test 4. Drivers in test 4 have the same instructions as in test 3, but of course the lights are now being controlled by a human, suggesting that the control strategy employed by the human is working to distribute the vehicles on the track more homogeneously. This result correlates strongly with the result in Figure 7, which indicates that the stationary time in test 4 is low, suggesting less queueing.

3.1.2 Evidence for Hypothesis 2

It is possible to extract the turning movements of vehicles in the tests from the GPS data and hence calculate the probability that a vehicle will make the same turning movement as the vehicle in front of them as they pass through the junction. This is equivalent to the fraction of all turning movements that are the same as the preceding turning movement, which is plotted in the right hand chart in Figure 10. This shows that vehicles were nearly *half* as likely to make the same turning movement as the vehicle in front during test 2. This evidence strongly supports hypothesis 2.

Note that if we assumed that A) turning movements are evenly distributed among vehicles and B) that vehicle's arrival time at the junction is uncorrelated with their turning movement;

we would expect the probability of a following vehicle making the same turning movement to be ~ 0.33. We stated above that a visual analysis of the video and animation data appeared to indicate that vehicles in test 1 were forming into synchronised platoons on the same route. We would expect this to result in a probability of the same turning movement > 0.33 in Test 1. In fact the data in Figure 10 do not show this and suggest that this view was an error in our 'human' perception.

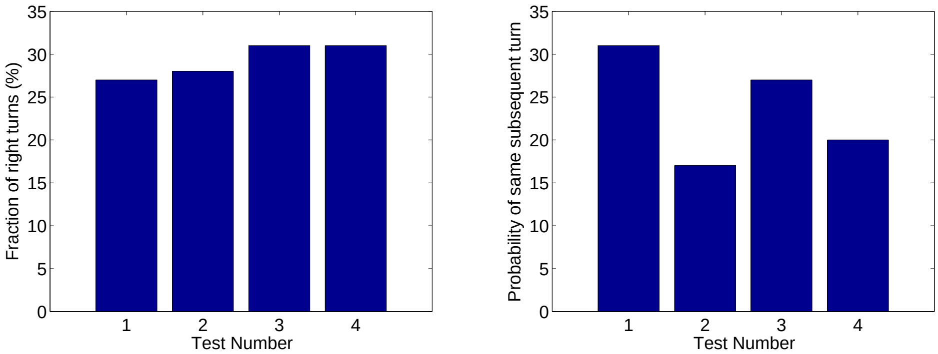

Figure 10: Analysis of vehicle's turning movements. The fraction of all turning movements that were right turns and the fraction of turning movements that were the same as the preceding turning movement.

In Test 3 the probability of equivalent subsequent turns increases as most drivers are now following a prescribed route. Interestingly in Test 4 under human control the probability of equivalent subsequent turns reduces again. This suggests that something in the control strategy of the human is biased *against* the platoon formation of vehicles on similar routes.

3.1.3 Evidence for Hypothesis 3

The number of times vehicles make a right turn can be extracted from the GPS data. The number of right turns, as a fraction of all turning movements is plotted in the left graph in Figure 10, for each of the tests. This indicates that the right turns are approximately equally common in all tests. This evidence indicates that hypothesis 3 is false. Even though drivers in Test 2 appeared to actively select routes to avoid congestion and proximity to other vehicles they did not avoid making right turns.

3.2 Human performance

All the metrics for measuring junction performance plotted in Figure 7 indicate a significant increase in performance between Test 3, where the junction was controlled by an optimized fixed time controller and Test 4, where the junction was controlled by a human. Between these two tests average speed is increased by 15%, time spent stationary is reduced by 26% and delay across the junction is reduced by 31%.

It is worth noting that the human controlling the junction in Test 4 was a BBC television presenter, with no special knowledge about traffic control and no previous experience controlling junctions.

It is tricky to extract from the data exactly what it is that the human is doing to achieve the high control performance but we can see from Figures 9 and 10 that and effect of their control is to actively discourage vehicles from forming into platoons of vehicles on similar routes. This is particularly evident in the fraction of equivalent subsequent terms, which is 10% lower than would be expected for a truly random process.

3.3 The Delay Metric

The plot for delay in Figure 7 agrees with the other performance metrics in so far as the performance of the human junction controller was better than that of the automated controller used. But on the more subtle differences in performance between Test 1, 2 and 3, the delay metric is different. The main difference between this metric and the others is that it only considers a sub section of the whole network. It might seem that, given that there are no other junctions outside of this subsection, it is reasonable to neglect these areas of the network in the analysis, but these results would indicate otherwise. It appears that the behaviour around the junction is having wider "knock-on" effects throughout the network that are affecting overall performance.

Test #	mean dist (m)	mean speed (ms^{-1})	mean time stopped (s)	mean trip time (s)	right turn fraction	repeat turn fraction
1	5539	6.17	206	28.3	0.27	0.31
2	5801	6.46	196	28.8	0.28	0.17
3	5501	6.13	199	30.8	0.31	0.27
4	6315	7.04	148	25.8	0.31	0.20

Table 2: Key statistics from the four tests extracted from the GPS data and averaged over $N = 24$ vehicles. Mean distance is the mean total distance travelled by vehicles in each 15 minute test. The mean speed is the speed averaged over all vehicles for each test. The mean time stopped is the meant time vehicles spent stationary during the test. The right turn fraction is the fraction of all turning movements that were right turns. The repeat turn fraction is fraction of turning movements that were the same as the preceding turning movement

.

4 Conclusions

Limited resources (especially time) were available for this experiment due to it being funded for a television article rather than purely research. This means that the procedure followed does not follow the ideal scientific method and the results should be interpreted in this light. The results certainly indicate that, in this case the human controller was able to control the junction significantly ‘better’ than optimised fixed time control for three performance metrics (average speed, time stationary and delay). However a full characterization of how well a human can control the traffic would require repeated tests on different junctions, with different controllers and benchmarked against better automated controllers (e.g. SCOOT and MOVA).

Results for Test 1 (where vehicles followed fixed routes) indicate that platooning of vehicles on similar routes was high relative to the other tests. But it is not clear whether this was higher than ‘normal’ and caused by feedback effects. While it was the perception of those observing the test that platooning appeared “unnaturally high”, the subsequent turn percentage in Figure 10 is approximately equal to what we could expect from a random process.

Results from Test 2 (Where vehicles instructed to follow random routes) indicate that the routes were not in fact random but that there was a strong bias *against* following vehicles taking the same route through the junction. Vehicles were also less platooned in Test 2.

Results from Test 3 (Where a hybrid strategy of 25 Routed and 5 random vehicles was used) showed similar levels of platooning as Test 1, although the subsequent turn percentage was slightly lower. The conclusion of this is that a fully routed strategy, or a hybrid strategy with low random percentage, are the preferred instructions to give to drivers in this kind of test. More work would be required to work out exactly what percentage of random drivers is the best simulation of ‘real’ conditions.

Results from Test 4 (Where a human controlled the junction) indicated good performance but it is not exactly clear why this is the case. Figures 9 and 10 indicate that in test 4 the control of the junction acted to prevent platoon formation and distribute vehicles more homogeneously around the track.

Finally small differences between the delay measured across the junction and other performance metrics that considered the network as a whole may suggest that the practise of excluding areas of the network in performance evaluation, even if they contain no junctions, can skew the results.

References

S. Box. What can traffic lights learn from us? *Traffic Engineering and Control*, 52(8), 2011.

S. Box and B. Waterson. An automated signalized junction controller that learns strategies from a human expert. *Engineering Applications of Artificial Intelligence*, 25(1):107–118, 2012a. doi: 10.1016/j.engappai.2011.09.008.

S. Box and B. Waterson. An automated signalized junction controller that learns strategies by temporal difference reinforcement learning. *Engineering Applications of Artificial Intelligence*, Online ahead of print, 2012b.

P. Hunt, R. Bretherton, D. Robertson, and M. Royal. Scoot on-line traffic signal optimisation technique. *Traffic Engineering and Control*, 23:190–192, 1982.

J. Lees-Miller, S. Box, and R. E. Wilson. Hidden markov models for vehicle tracking with bluetooth. In *13-3032 TRB Highway Traffic Monitoring Committee (ABJ35).*, 2013.

G. Vincent and J. Peirce. 'MOVA': Traffic responsive, self-optimising signal control for isolated intersections. *TRRL Research Report*, RR170, 1988.