Traffic Survey in Englishcombe Lane, Bath, September 2021

Introduction

This documents details the results of a traffic census carried out in Englishcombe Lane, Bath during September 2021 following on from a previous survey in June 2021. This updated survey takes advantage of upgraded hardware and software which is more accurate and wider ranging to enable a more accurate representation.

The information should be regarded as informative rather than definitive, but is intended to give a realistic view of traffic density and speed profiles across extended periods.

Background

Englishcombe Lane is a single carriageway urban road in Bath. It serves to connect Bloomfield Road in the East and Whiteway Road in the west, and is used as both a local connecting route and occasional commercial vehicle route into and out of Bath.

The entirety is within a 20 mph speed limit. While some sections are part of regular bus routes, the surveyed section, in front of number 93, is not. Some sections have speed humps, this section does not.

The previous study showed that this section is subject to occasional high-speed use, and heavy traffic density during School-run periods. The road serves Moorlands Schools Federation as well as routes for Beechen Cliff and to a lesser extent, Hayesfield schools.

In parallel Bath and North-East Somerset carried out their own assessment, using their equipment, in the same location.

The survey was a follow-up to the previous study, and as a comparison with the results from the BANES survey.

Equipment and Methodology

Detailed equipment design is discussed in Appendix A. In summary, the equipment uses camera and machine-learning software technology to capture and analyse images in real-time. Image analysis allows quantity, direction, speed and size of vehicle movement to be determined with suitable timestamping.

Post-processing analysis allows density and speed profiles to be presented for statistical consideration. Metrics available include vehicle movements per time period, tidal flow, vehicle speed profiling and some indication of vehicle size.

System limitations and accuracy are described in Appendix B, but in summary:

the system operates in daylight and at night with reduced accuracy,
previous system limitations on speed detection have been raised to approximately 60 mph,
the system cannot differentiate reliably between vehicles passing simultaneously in opposite directions,
system speed accuracy is now determined per event, and recorded as part of the collected data enabling accuracy filters to be applied to the analysis.

The system is not capable of, and does not try to, provide identification of vehicles or vehicle occupants. The exception is where commercial vehicles may carry some signage or branding, or where the vehicle usage is obvious such as blue-light vehicles, buses, etc.

An example of the image captured during speed analysis.

The red rectangle represents the measurement field of view, the green outline is shown as the identification of a vehicle in the image. Successive images are examined to track the progress of the detected vehicle, and the timestamp of each image is used to create a running time for the sequence.

Calibration images are taken to establish an accurate representation of pixel size versus actual size prior to use, and speeds are calculated using the speed = distance/time formula. Statistical averaging and standard deviation calculations are performed to arrive at a vehicle speed, with a margin of error that can be used to make a judgment about the quality of the measurement.

Presentation of results

Analyses are presented largely graphically, drawn from the raw database which can be made available for other analysis. The intention is to convey as simply as practicable the key results from the survey. In general results are presented as discrete interval periods, e.g averages over 60 minute intervals.

Results

The surveyed period is from Tuesday 7th September to Wednesday 15th September 2021.

This includes 6 'working' days and one weekend.

Total vehicle movements counted over this period were 24,829

.

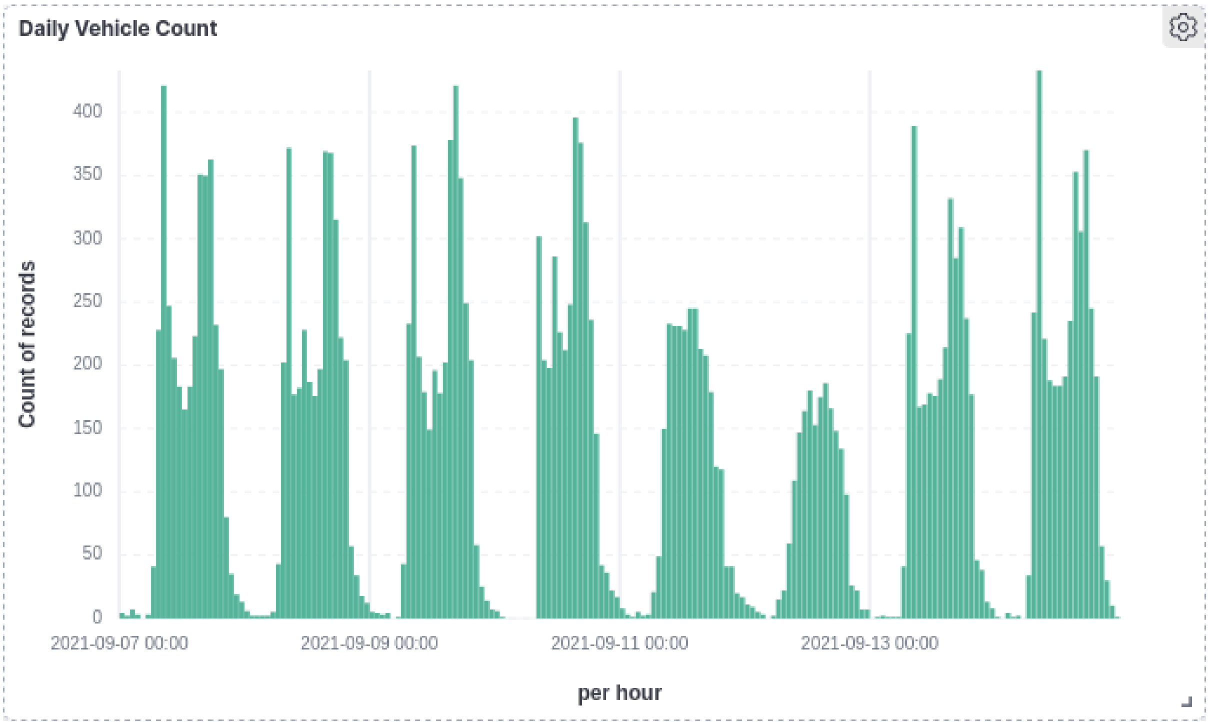

Fig. 1 - Total vehicle movements by day

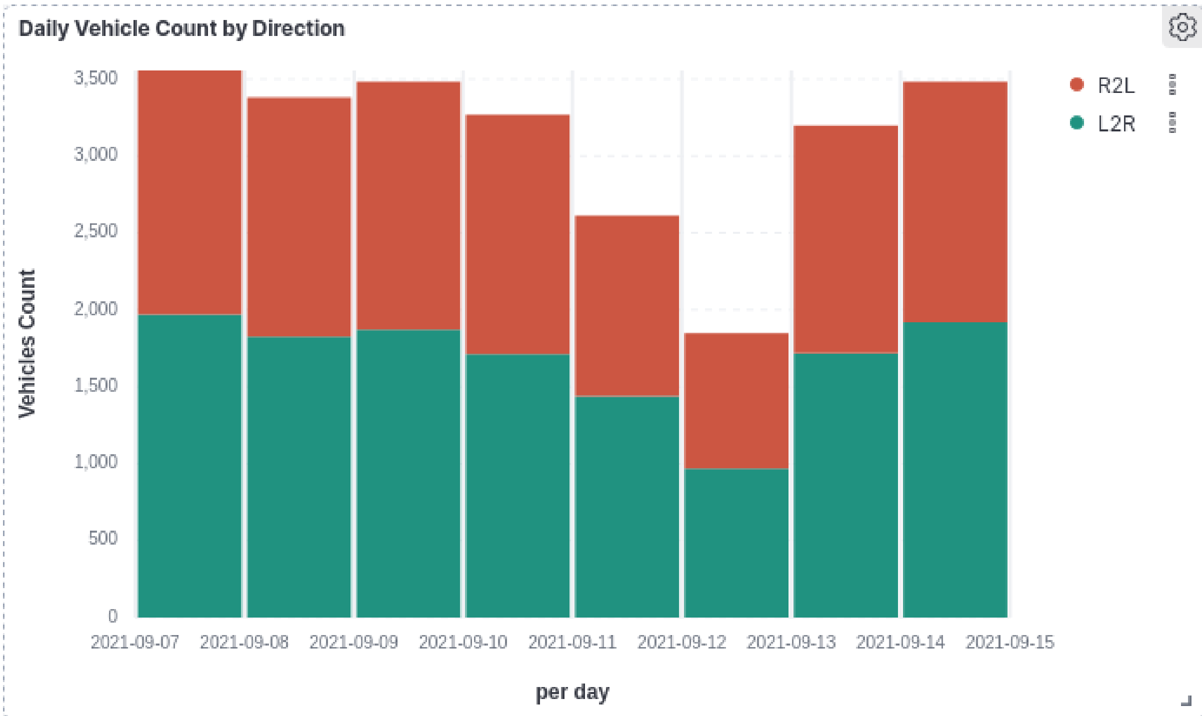

Fig. 2 Vehicle movements per day by direction

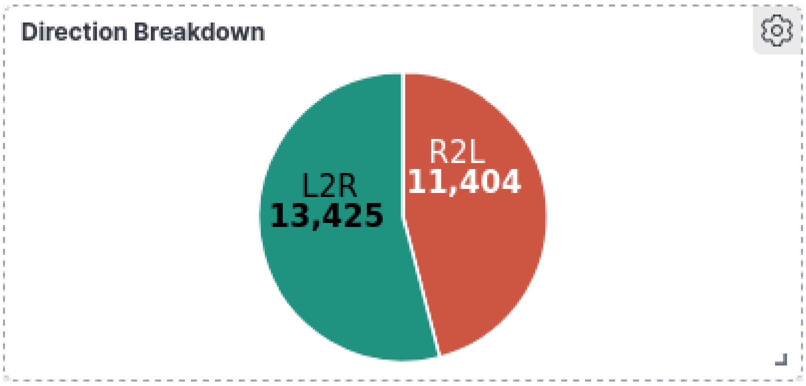

Fig. 3 –Total Vehicle movements per direction.

R2L is westbound,(46%), L2R is eastbound (54%).

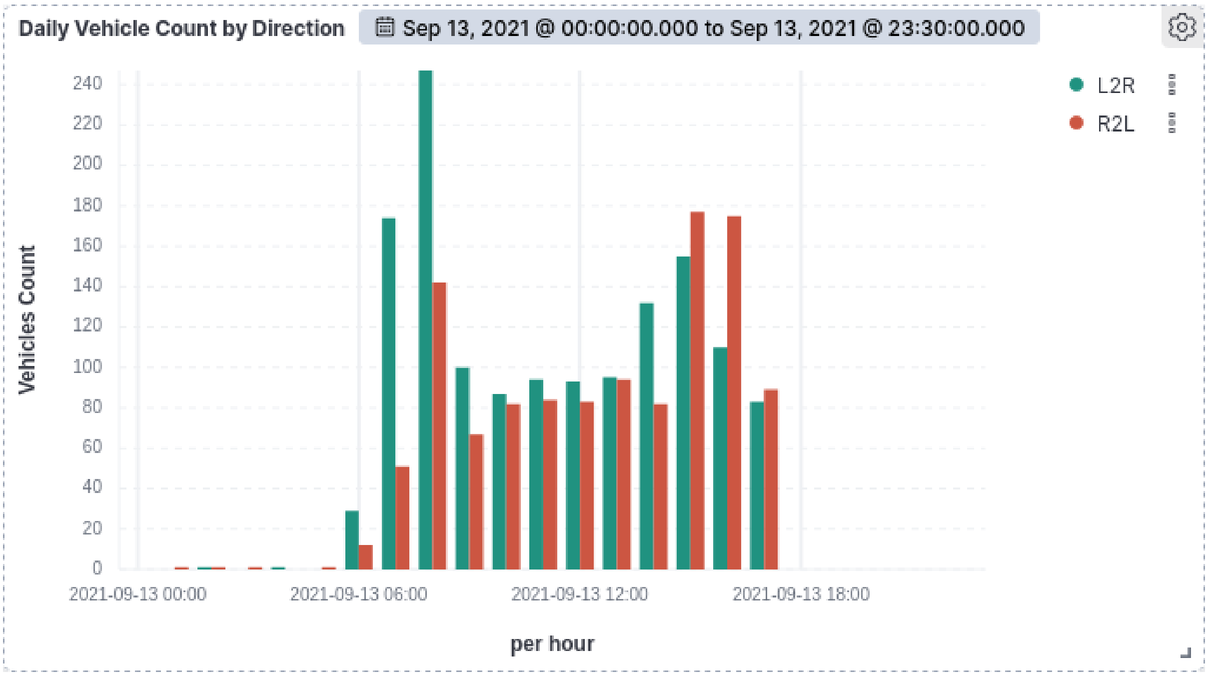

Fig 4 - Vehicle Movement for one working day by 60 minute intervals

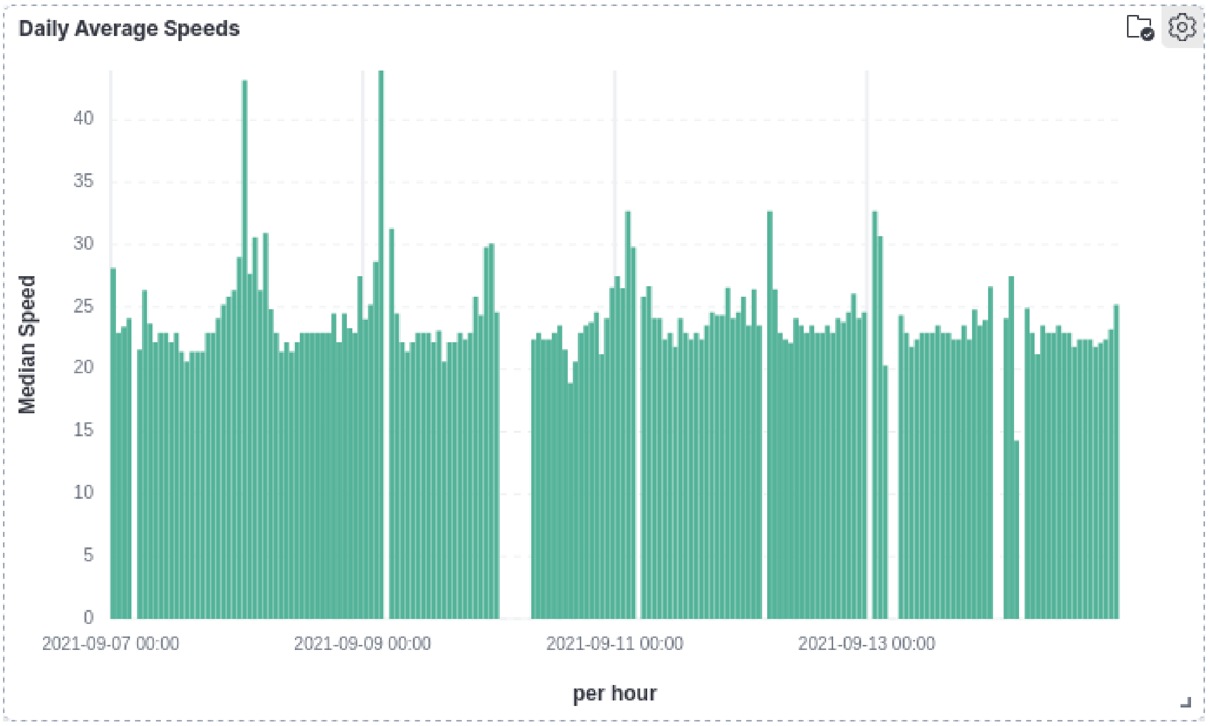

Fig 5 - Vehicle speeds average over 60 minute intervals

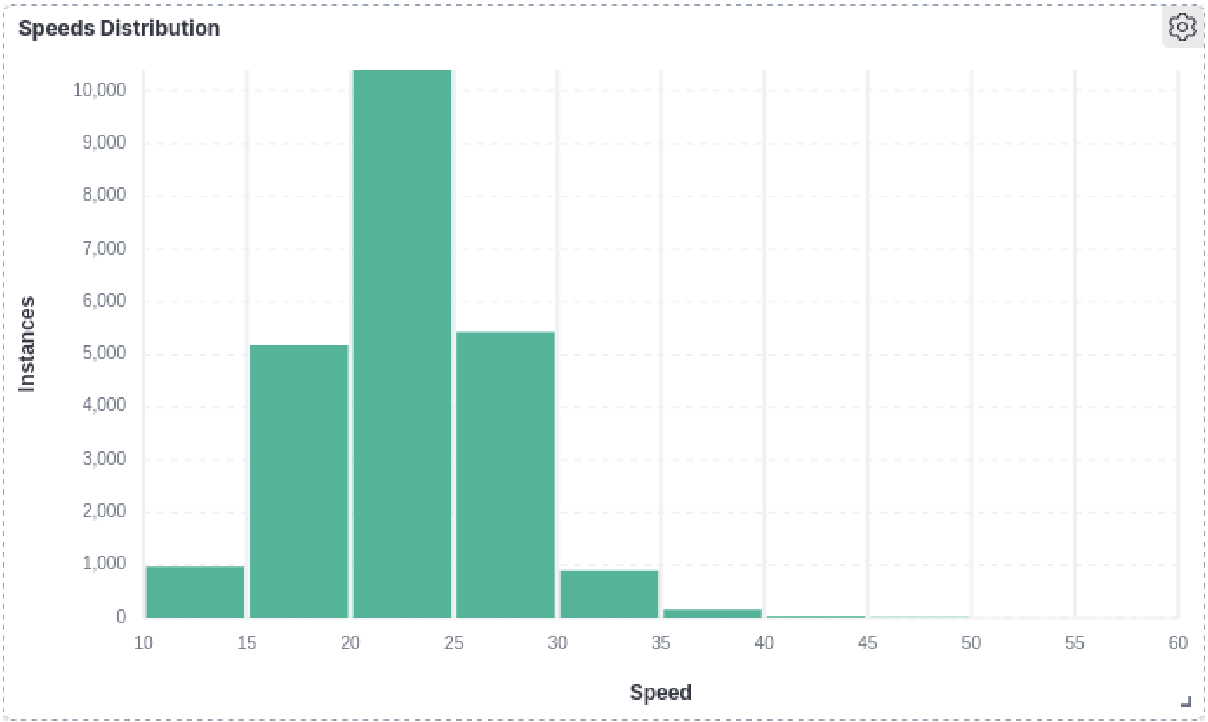

Fig 6 All speeds for the monitored period

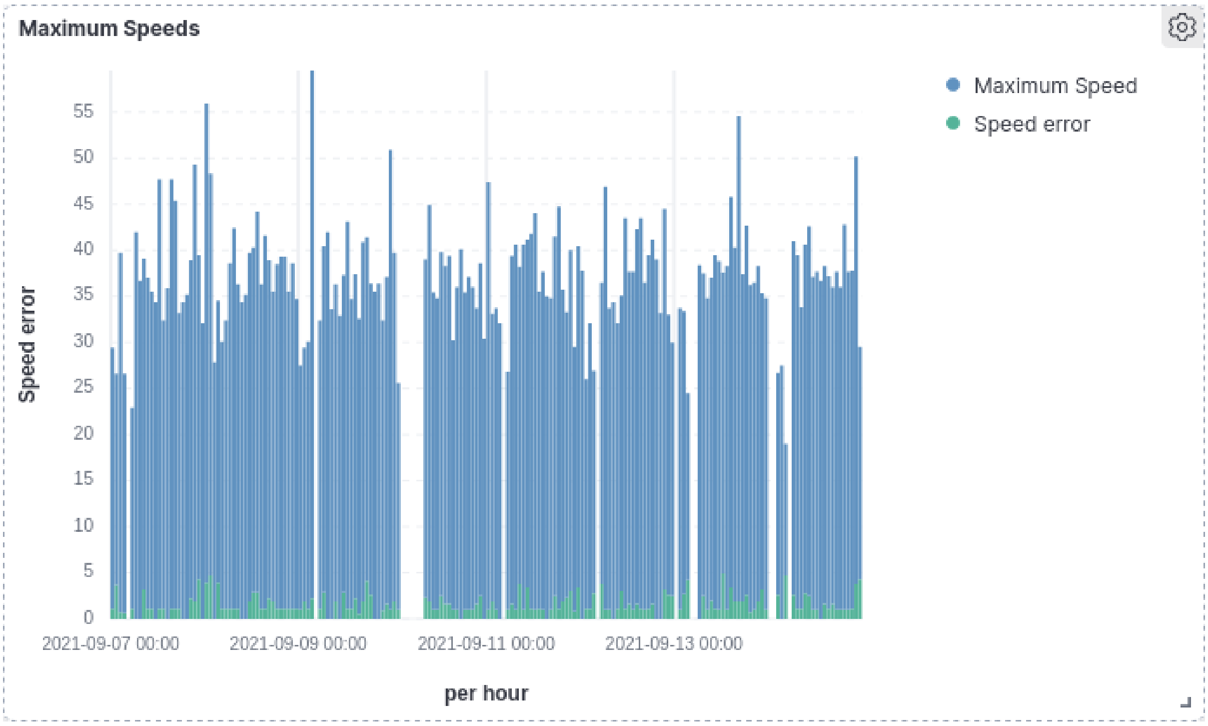

Fig 7 - Maximum vehicle speeds in 60 minute intervals

Interpretation of results

From Fig 1, weekday traffic movements are fairly evenly spread at 3000-3500 per day.

This changes at weekends, to 2000 – 2500 per day.

From Figs 2 and 3 , there is a consistent bias towards more eastbound traffic (entering the City), 54% to 46%. This is consistent across all the days of the week and is in contrast to the previous survey in June where the bias was westbound. This may be due to significant traffic diversions in operation in the City currently.

Fig 4 shows the vehicle movement across a typical work day. Most movements are around commuting/school run times, i.e. 8am to 9:30 am and 3 pm to 6pm.

Fig 5 shows the median speeds per 60 minute interval across each day showing both the higher speeds between dusk and dawn, and the consistency of median speeds above the 20 mph limit. .

In figure 7, this shows as a distribution of all median speeds across the measured period. From this it can be seen that the most likely average speed is in the 20 -25 mph range, approximately 50% of all vehicles. In addition, 20% of all speeds recorded were in the 25-30 mph range, 4% in the 30-35 mph range and 1% in the 35-40mph range. In total, 64% of all vehicle average speeds were above 20 mph. popular speed.

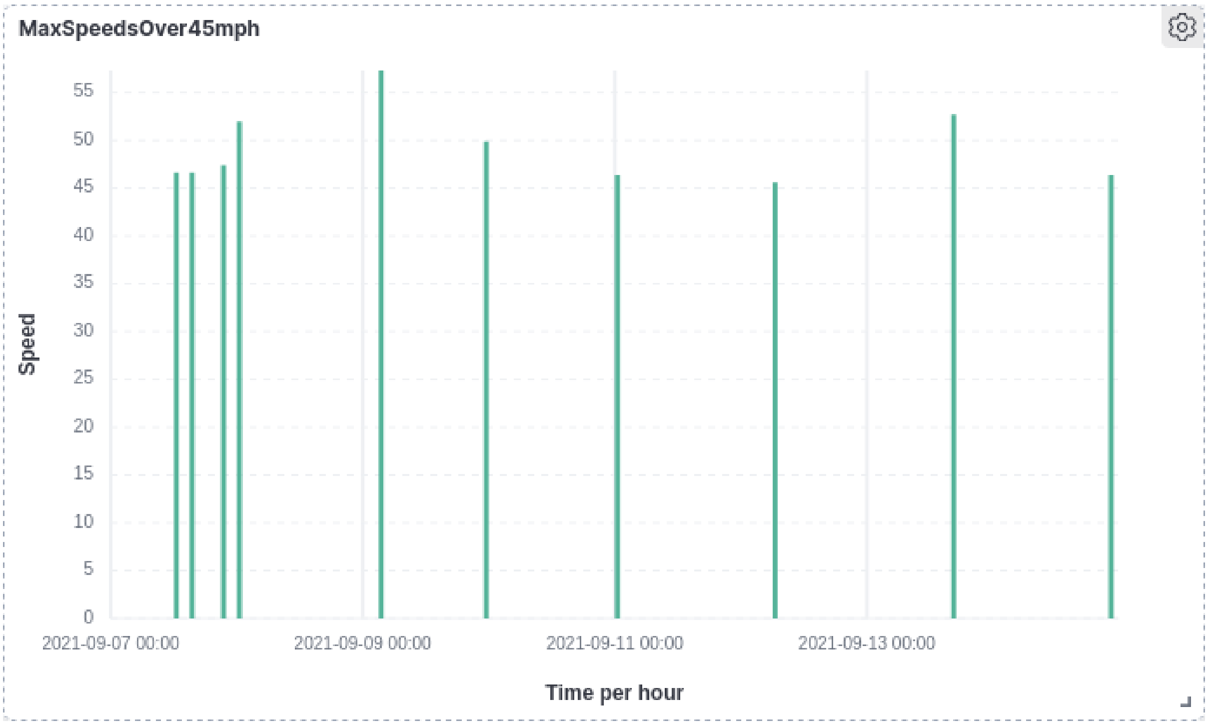

Fig 8

Over the measured period, 10 speeds were recorded above 45 mph, of which 3 were above 50 mph. The highest individual speed recorded was 57mph.

Summary and Conclusions

While this survey can only be regarded as indicative, there is a clear pattern of speed limit violation, with over 64% above the speed limit. Traffic density may also surprise some, this is a busy road especially around the time of the school run, when it may be considered there are vulnerable children using and crossing the road.

The results of the BANES survey are awaited with interest, but in general this confirms the June survey, and points to a need for further traffic calming measures.

Appendix A

Detailed equipment description

The system uses multi-core pc hardware running the Linux operating system and bespoke image capture and analysis software. It can also be run on a Raspberry PI single board computer with integrated camera.

Hardware

The desktop pc hardware is a six-core processor with 24GB RAM and GBE network interface. The hardware also includes a network IP camera with night-vision and capable of up to 30 frames per second at the chosen resolution.

The camera is mounted outside with night-vision assistance . Convenient street lighting also assists night vision.. The camera is mounted tangentially to and above the area to be surveyed, at a suitable distance to enable sufficient transit distance to be viewed.

It is necessary to have a field of view that can capture sequential video frames without interruption. Mounting at a first floor window provides a good compromise between distance and height to avoid e.g. pedestrians passing. Since the mounting is indoors, no weather-proofing or special measures are need to provide power to the device.

It should be noted that potentially any hardware platform running Linux could be used, since the system also support IP cameras of the type commonly used in domestic CCTV systems.

Software

All software used is open source and freely available. https://github.com/somerman/Vehicle-speed-detector

The operating system is Linux for speed and reliability and the software application is written in the Python programming language.

The key part of the software is the industry-standard image capture and analysis system OpenCV. This library suite is capable of sophisticated image analysis at high-speed allowing for example, vehicles to be differentiated from people.

Nevertheless, the software is process intensive and ultimately is the limiting factor in detecting high-speed vehicle movement.

The application includes calibration and environmental parameter settings which require a little time to get right and depend on the exact dimensions of the field of view. In practice, speed calibration verification is carried out by the simple trial and error process of driving at a known speed in front of the camera and adjusting parameters until the indicated speed accords with the actual speed.

Data capture consists of recording key parameters in a CSV delimited format, one record per line. Alternatively, a SQL Lite database can be invoked. Each vehicle detection is also recorded as a jpg image with annotation including detected speed and time. These can be used to verify assumptions about vehicle type, or investigate apparent anomalies.

Post-processing and Software analysis

Once a database of movement records has been captured, it is imported into the Elasticsearch engine.
Elasticsearch, together with its graphical front-end, Kibana, are tools to analyse and present timebased data in an easy-to visualise form. These tools allow aggregation, averaging, filtering and sorting of data to present histogram, bar charts, pie charts and so on. Only a selection of these are presented in this report, as images.

The raw data can be re-interpreted as required to present more detailed or more over-arching reports.

Appendix B

Limitations and Errors

The equipment used is subject to a number of error sources.

These are principally image recognition accuracy and image tracking accuracy.

Image recognition accuracy

The software-based image recognition mechanism relies on generally good raw image quality, with general form consistency, ie vehicles all look about the same. Recognition depends on pixel difference calculations using subtractive mathematics on arrays of pixel points.

Variations in lighting, especially reflections and shadows, can cause deviations in the recognised contour position, so multiple consecutive tracking positions are taken for averaging.

Filters are used to

a) remove the fixed background as far as possible, to reduce the calculation intensity
b) limit the size of image to be greater than a minimum, to prevent e.g pedestrians being falsely registered as vehicles.
c) limit the time between image processing to prevent vehicles being counted twice.

However, the proximity of vehicles closely following each other will usually result in the count being under-reported if the distances are small. This usually happens with slow-moving queued traffic.

The image processing cannot reliably distinguish between vehicles crossing in opposite directions simultaneously.

Overall, these limitations tend to an under-reporting of vehicle numbers.

Manual analysis of the speed data suggests approximately 5% of readings are inaccurate to the point of requiring exclusion from the data. The Margin of Error (MoE) is used to filter for these results. For clarity, it should be noted that this does not invalidate the counting data for these samples, just the speed data.

Image tracking accuracy

The system works reliably with individual vehicles or groups of vehicles in deciding the rate of travel. This is in part due to the setup, in general, variations in the y (vertical) position of the image are small for any particular vehicle track, since vehicles tend to stay in lane. Thus the principal change of position is in the x-axis (horizontal) portion of the image.

The system uses a high-accuracy timer to generate timestamps per image capture and these allow a reliable timebase to be established.

Calibration Process

In order to relate pixels on screen to real distances, a calculation is required. To aid this, there is a calibration mode that takes an image with an overlay 10 pixel grid. This can then be used to measure the pixel length of a target, and relate it to the actual length in millimetres. To cope with parallax errors, there is a different scaling factor for near and far lanes. Confirmation of calibration is achieved by conducting a real-time test with a vehicle travelling at a known speed, as indicated by GPS calculations, since vehicle speedometers are high-calibrated by approximately 10%..

Processing speed is ultimately limited by hardware and software, with the software portion being variable depending on the complexity of the images. Improvements to the image detection system have enabled the upper limit on speed detection to be raised. However accuracy diminishes with speeds above 45 mph, since fastmoving vehicles may have crossed the field of view before image processing is complete.

For speed calculations, a minimum of 4 consecutive frames is required which are averaged to produce the speed calculation. The speed-averaging algorithm now includes a statistical error calculation, based on the Standard Deviation of consecutive detections. This is used to create a 95% confidence level of the error distribution (in the same speed units). The error is stored with each record and is used to filter the resultant data. In general, errors in excess of +/-5 mph will result in the sample being discarded, but this is configurable during visualisations.

Overall speed reporting is most accurate between 15 and 45 mph, so outside this range should be regarded as indicative, rather than definitive.