JCT Symposium 2023

Developments in Pedestrian Metrics and their use in the Optimisation of Signalised Crossings

Calya Chetty, A
Lead, Starling Technologies Ltd

Samsen, D.
Starling Technologies Ltd

Zisch, D.
Starling Technologies Ltd

Bondi, D.
Starling Technologies Ltd

Clarke, D.
Cambridgeshire County Council

Ling, R.
Cambridgeshire County Council

Stevens, M
Cambridgeshire County Council

Pleydell, M.
PTC Ltd

Starling Technologies

Keeping people, streets and cities moving safely

Our cities are changing - safety - active travel - sustainability

Transforming the way we move around our cities is critical to reducing carbon emissions, enhancing air quality and improving national health and wellbeing. At Starling Technologies we design revolutionary AI- driven traffic management systems which make our streets and cities safer for pedestrians, cyclists and street users of all kinds, while keeping vehicle traffic flowing efficiently.   Starling Technologies street safe solutions are already in operation around the world. Our future-proof technologies enable today's active travel and make tomorrow's smart cities greener, safer, sustainable and attractive for all.

PTC

A technology consultancy providing knowledge and design services to the traffic control and ITS sectors.

Cambridgeshire County Council

Trial site hosts.

Contacts

Primary
Andrew Caleya Chetty
andrew@starlingtech.co.uk

Dr Mark Pleydell
mark.pleydell@p-t-c.co.uk

Michael Stevens
Cambridgeshire County Council

Abstract

Local Authorities have access to a lot of good quality data on the movement of vehicles through their networks and have tools available to optimise the movement of road traffic to deliver on policy objectives. When it comes to data for and interventions on behalf of pedestrians, these same authorities are less well placed.

Recent advances in detection technologies now provide not only the same functions and facilities for pedestrian detection at kerbsides and on crossings as the existing technologies, but can provide many more insights into the numbers, density, speed, and delays of the pedestrians moving along pavements and using crossings. These same technologies can also assess near-misses and jaywalking.

Having access to these additional insights opens opportunities for applying processes like MOVA and SCOOT for vehicle traffic to optimise the timings of pedestrian signals at crossings and junctions to improve the journeys of this group of previously under-served network users.

Drawing on real data from a trial site on East Rd in Cambridge kindly made available to the authors by Cambridgeshire County Council, this paper presents work on assessment of the performance of new technologies before turning to the implementation of metrics and reporting on examples of how these can be used to diagnose issues in the movement of pedestrians and then moving on to the use of these metrics in optimisation of signalised crossings.

Note: Some elements of this paper may be the subject of patent applications.

Introduction

Highway Authorities in the UK have had access to data on vehicle movements on their networks for several decades, and these data have been used to gain insights into network operation and to allow optimisation of traffic flows through tools such as SCOOT, MOVA, SPRUCE, etc. with very beneficial results, to the point where recent TTF funding was directed towards ensuring that local authorities could repair and re-instate vehicle detection to improve the realisation of the benefits of these optimisers. The situation for pedestrians has lagged far behind that for vehicles. This is largely due to the lack of comparative discipline in pedestrian movements where there are few orderly queues nor tightly defined lanes. Recent developments in technology have enabled the development of pedestrian detection that can track and characterise multiple 'actors' as they move along pavement use or misuse crossings.

This paper presents the results and conclusions for three sequential areas of activity undertaken by Starling applying these technologies and follows on from the presentation at the 2022 JCT Symposium. Using data from a site on East Rd in Cambridge that paper[1] identified that local highway authorities did not have the same level of access to data and interventions on behalf of pedestrians as they did for vehicles on their networks and outlined what could be achieved. In this paper we build on the work and findings of the previous paper and present performance data for new technologies delivering those pedestrian detection functions currently used at signalised crossings and pedestrian facilities at junctions, indicating that one detector could replace three existing devices.

We then explore the topic of metrics, looking at both pedestrians in their own right, and their interactions with other road traffic. These metrics have been applied with the trial site crossing running as a typical puffin with the aim of seeing where and when the puffin strategy might not effectively deal with the pedestrian demands.

The third part of the paper moves on to how these metrics can be used as inputs to improved decision-making for influencing the control of the crossing above the day-to-day puffin operation and presents preliminary data on interventions from the summer of 2023.

[1] https://jctsymposium.co.uk/optimisation-of-signalised-pedestrian-facilities-for-sharply-changing-demands-andrew-caleya-chetty-starling-technologies

Background

Work presented in 2022 – the Detection Performance Assessment Methodology

In September 2022, Starling presented the work that they had undertaken deploying new detection technologies to a pedestrian crossing in Cambridge on East Road outside the Anglia Ruskin University site.

A pair of Starling Detector placed on the primary signal poles diametrically opposite each other can cover not only the wait zones and on-crossing space, but can also see the pavement approaches to a crossing and monitor the vehicle traffic.

It was noted in that paper that work would be undertaken to compare the performance of the Starling Detector with the existing kerbside and on-crossing detectors with the initial objective, subject to acceptable performance and Cambridgeshire County Council's agreement, of moving control of the crossing to the outputs from the Starling detection. A parallel activity is to implement and deploy a series of metrics to not only assess the normal operation of the crossing but also to inform subsequent optimisation.

The comparative study recognised the implicit weakness of simply comparing two technologies against each other, and so the task compared both against a ground-truth derived from human observation of video and still images. The video and still images were synchronised to each other through reference to marker data events, e.g. first pedestrian into a wait zone, late night vehicle passage through the crossing, etc.

The study captured false positives (a detection when there was not a valid actor present) and false negatives (times when a valid actor was present but not detection was registered).

As with any real-world study, there was some level of uncertainty, judgement and noise present.

No attempt has been made to apply statistical methods or tests to the data.

Results from the First Phase of that Trial

During the autumn of 2022, timestamped event data was logged from both the existing kerbside, on-crossing and vehicle detection and the corresponding outputs of the Starling Detector. All these data were captured using dedicated logged equipment in the signal controller to give alignment of timings. Still images were also captured for human counting.

A selection of periods was taken covering day and night and varying weather conditions, and during those periods, the ground truth was compiled by a third-party organisation using a human observer briefed to use the definitions of TOPAS 2505 (vehicle detection), TOPAS 2506 on-crossing detection and TOPAS 2507 (kerbside detection). While the human observations are also open to a small level of interpretation, this approach did allow a better assessment of the false positives and false negatives in both the existing and the Starling Detector.

Results

The following tables and graphics provide a comparison of the detection on site with the ground truth, and not with each other.

Table 1 Kerbside Detection	**Existing**	**Starling**
Correct	1038 (86.5%)	1130 (94.2%)
False Positive	70 (5.8%)	19 (1.6%)
False Negative	92 (7.7%)	51 (4.3%)

Figure 1 Comparative Kerbside Data

Existing Kerbside Detectors

Starling Kerbside Detectors

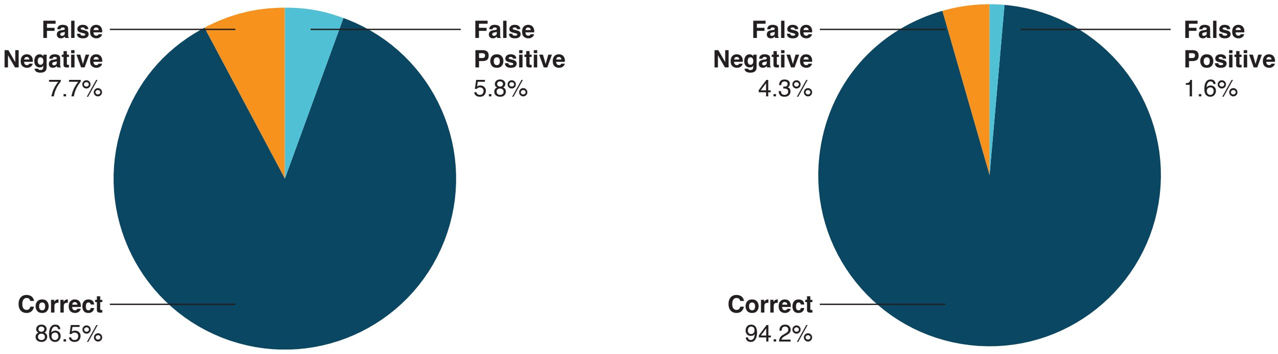

Table 2 On-Crossing Detection	**Existing**	**Starling**
Correct	997 (83.1%)	1165 (97.1%)
False Positive	173 (14.4%)	24 (2.0%)
False Negative	30 (2.5%)	11 (0.9%)

Note: The existing detection was responding to vehicles passing through the crossing as well as pedestrians as it is unable to discriminate between them. The Starling detection data excludes these vehicles.

Figure 2 Comparative On-Crossing Data

Existing On Crossing Detectors

Starling On Crossing Detectors

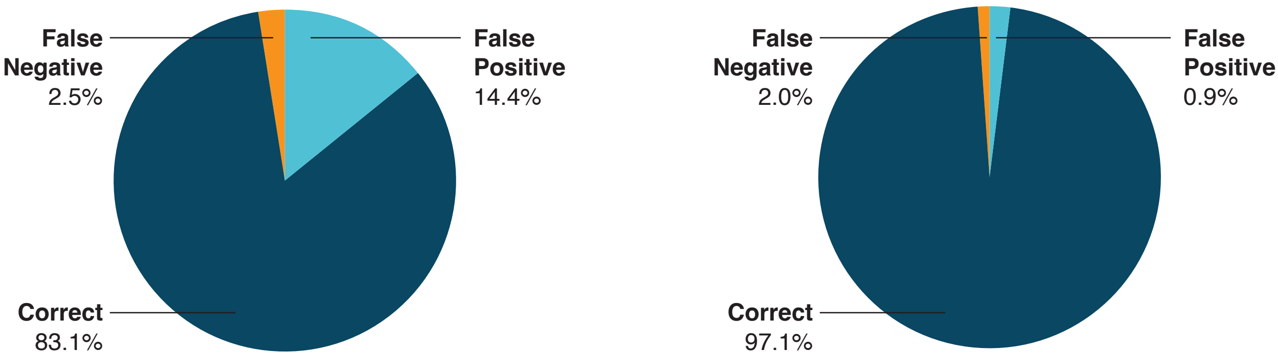

Table 3 Approaching Vehicle Detection

	Existing	Starling
Correct	916 (76.3%)	1096 (91.3%)
False Positive	206 (17.2%)	29 (2.4%)
False Negative	78 (6.5%)	75 (6.3%)

Figure 3 Comparative Vehicle Detection Data

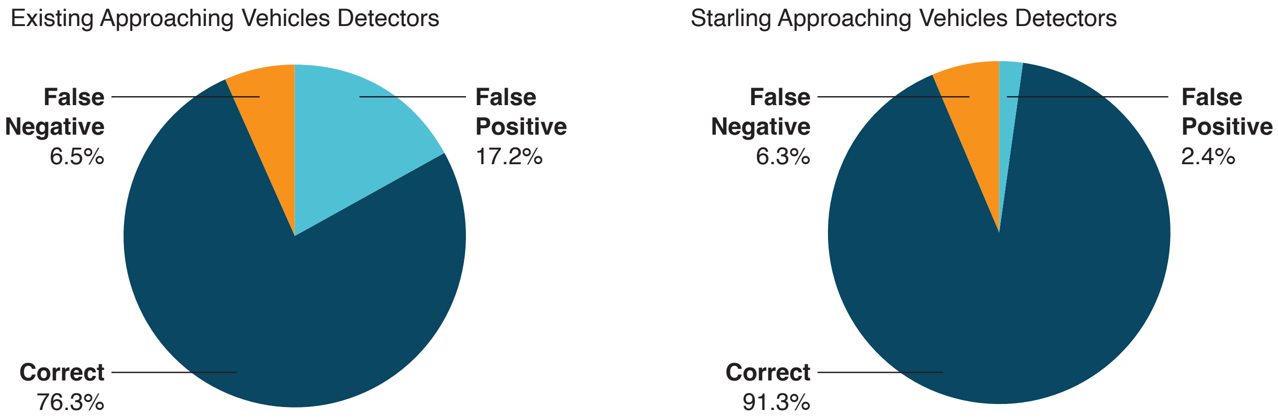

Commentary On These Results

It should be noted that no validation of the existing detection set-up and alignment was performed, as the intention of the trial was to look at typical installations. This is likely to have favoured the Starling Detector.

That being said, these results indicate generally equal or better performance of Starling Detectors. They also expose the range of real-world edge-cases that any trial throws up. However, the main observation is that, under most circumstances, for most of the time, the Starling Detector worked sufficiently well to allow it to provide the detection inputs to take over from the existing detection at the site.

The performance of the Starling detection at the site was deemed sufficient to allow the site to be migrated to control by the Starling detection. Validation of the operation is being verified by continuing to log the detection and also logging the levels of activation of the crossing.

The metric discussed in the following section also provides information on any changes in conditions for pedestrians in the before and after situations. Gathering these before and after metrics is still an on-going task.

What We Would Like To Know About Pedestrians And How To Capture That

Up to now, outside of academic studies, the detection of pedestrians and the response of traffic signal controllers to them has been quite primitive. Typically academic studies[2,3] look at addressing crowd control situations, so focus on very high-density situations and the patterns of movement that form in these crowded flows and there is less published work on modest densities.

[2] A universal function for capacity of bidirectional pedestrian streams: Filling the gaps in the literature, Feliciani C., et al, PLOS ONE I https://doi.org/10.1371/journal.pone.0208496 December 19, 2018
[3] Modelling passenger flows in public transport stations, Kırlangıçoğlu C., International Journal of Human Sciences 12.1 2015

Until recently, most detectors could only provide a single output conveying either no pedestrians or more than no pedestrians. There was no ability to say how many, or of what type. Technology now allows this information to be provided, however, the controllers are not yet able to make good use of this data, still basing pedestrian decisions largely on simple detection data. It was on this basis that the Starling detector was used to control the East Rd site.

Most traffic engineers have some insight into optimisation tools like MOVA and SCOOT and understand their detection requirements and what this detection is providing, i.e. cumulative delay, number of vehicles in a queue, saturation flow, as these all determine the demand and capacity of a junction to move vehicles. In general terms, up to now we have not had comparable information for pedestrians and so have not been able to address their needs as an equal source of demand in previous optimisation. The work by Starling goes some way into this territory, by researching and then implementing pedestrian metrics to allow future work in assessing how best to optimise crossings, and other locations for pedestrian activity. It should be noted here that the objective is not some vague promise of AI-enhanced better services. A more deterministic approach is being followed based on sound traffic engineering principles, and compatible with those systems and services that many local highway authorities already have in place. That is, the metrics will inform decisions about how we can use what we have to better effect.

How These metrics may be applied

Optimising pedestrian movement requires similar measurement capabilities, for similar reasons. We need to answer questions like how full is a wait-zone? How is that impacting on peds not using the crossing? What is the total delay of the waiting pedestrians, (which relates to their exposure to vehicle emissions)? How quickly can we clear a wait zone? And an on-crossing space? and can the receiving pavement handle the additional load? These are all questions of number, delay and capacity.

The rest of this paper looks at some of these metrics and presents preliminary data and analysis of the significance of these data.

Metrics

We have identified the following metrics as being of interest.

- Flow of pedestrians on the pavement approaches to the crossing – or any other pinch point.
- Speed and distribution of speed of pedestrians on pavements
- Number of pedestrians waiting to cross.
- Density of pedestrians waiting to cross and FRUIN value*
- Total delay of pedestrians waiting to cross.
- Crossing speed and distribution of crossing speeds of pedestrian on the crossing.
- Near misses between pedestrians and vehicles.

*The FRUIN value is a scaled measure of pedestrian comfort based on the amount of pavement space per pedestrian.

Some of these measures are of benefit to the wider active travel community and may inform changes to pedestrian facilities.

We have implemented means of configuring and collecting these metrics on the Starling Detector platform, which reports the data periodically to a clouded service from where we can access and review the data.

Examples of Pedestrian Metric Data & The Insights They Offer

Figure 4 to Figure 6 below show graphical representations of some of the metrics acquired from the East Road site. Figure 4 shows the pedestrian delay incurred by pedestrians that are static in the wait zone. The plot for the whole month has limited resolution, but there is a very strong peak for 10th June.

Figure 4 Pedestrian Delay at the Wait Zone June 2023

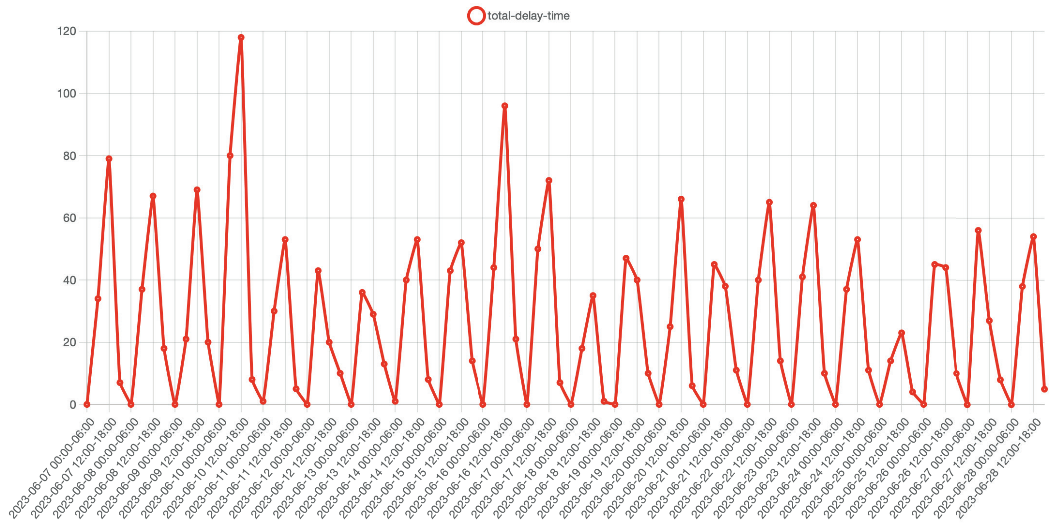

The four graphs of Figure 5 look at the 72-hour period 9th June – 11th June with more resolution. The first graph below shows pedestrian delays in the wait zone. It can be seen that the spike in Figure 4 is due to high levels of pedestrian delay throughout the day and not just attributable to a one-off incident. The corresponding FRUIN levels (measures of pedestrian comfort with higher being worse) unsurprisingly show that the wait zones were more congested. This suggests that the larger total pedestrian delay is due to more pedestrians, rather than normal levels of pedestrians waiting longer. The third graph below of Figure 5 confirms this by showing larger levels of crosser on this day than on either of the neighbouring days. The fourth graph below confirms that the vehicle flows were similar across all three days, so is unlikely to be any mechanism relating to higher levels of traffic so longer intervals between pedestrian phases are being serviced.

What is interesting about these four graphs is that we could have seen a different pattern where FRUIN levels and crosser counts stayed low and the vehicle movements were higher which would have suggested that the peds were being delayed by higher traffic loads.

Neither of these interpretations are currently accessible from conventional detection technologies, and the ability to infer pedestrian behaviour and demand opens up the possibility of giving pedestrians better levels of service at crossings.

Figure 5 Detailed View of 9th – 11th June 2023

Accumulated Crosser Delay Time

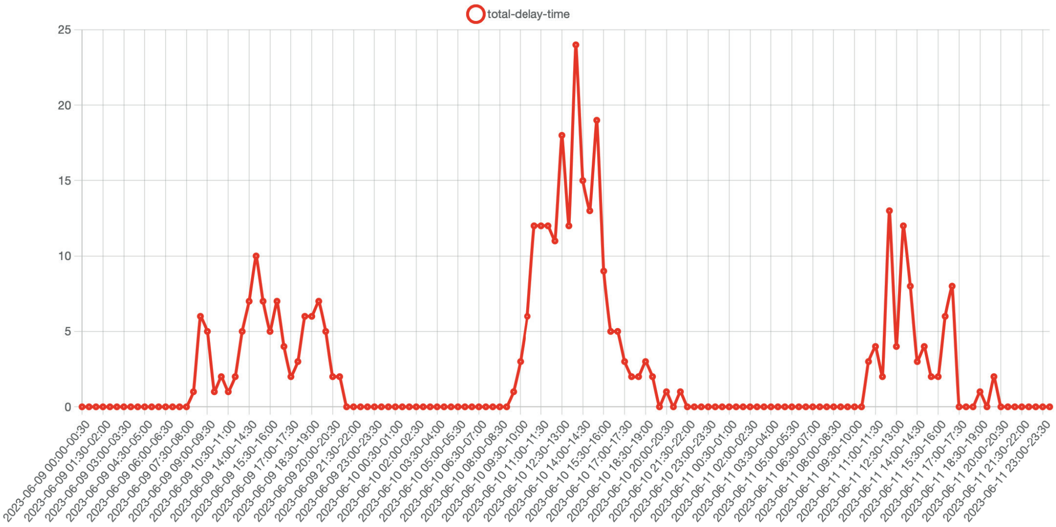

FRUIN Level at Crossing Entrance

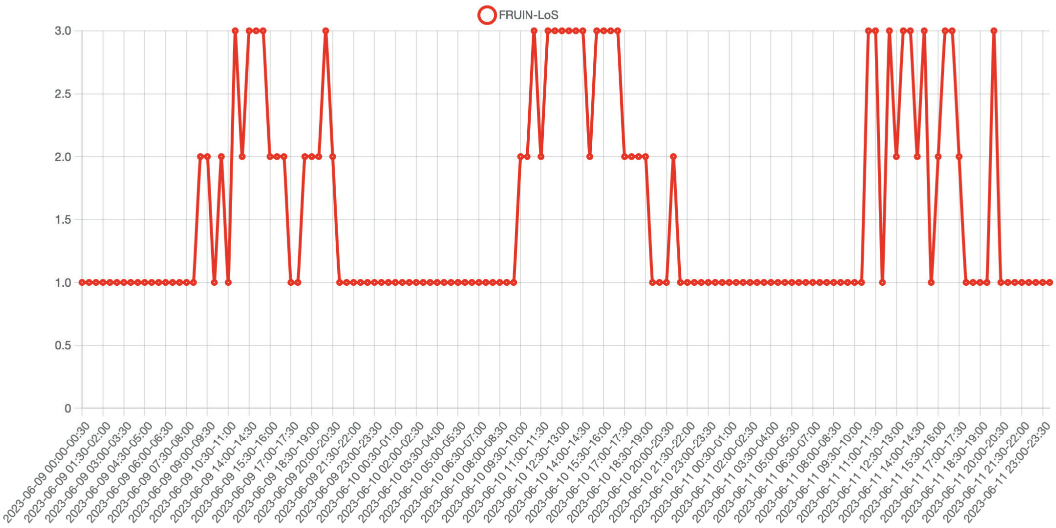

Crosser Count and Flow

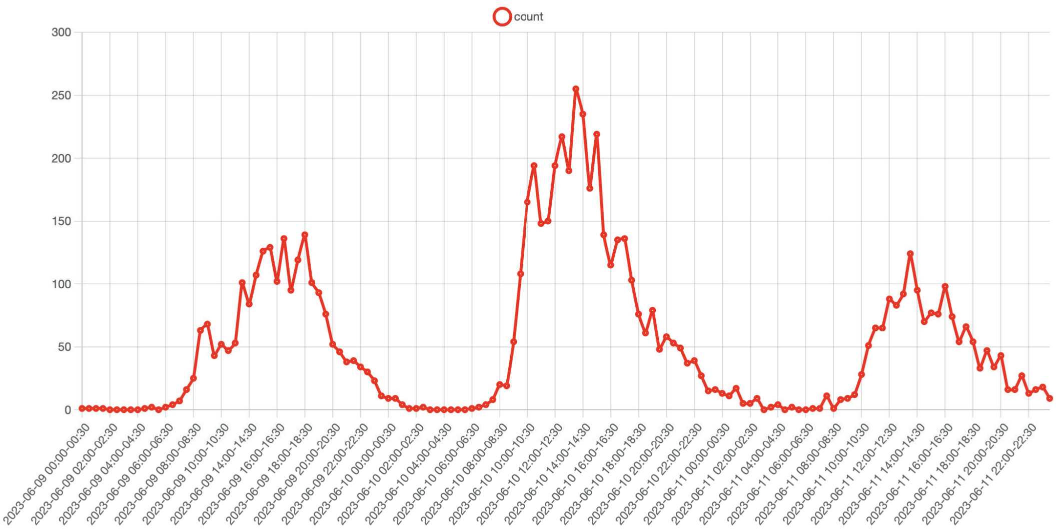

Vehicle Count
and Flow
Northbound

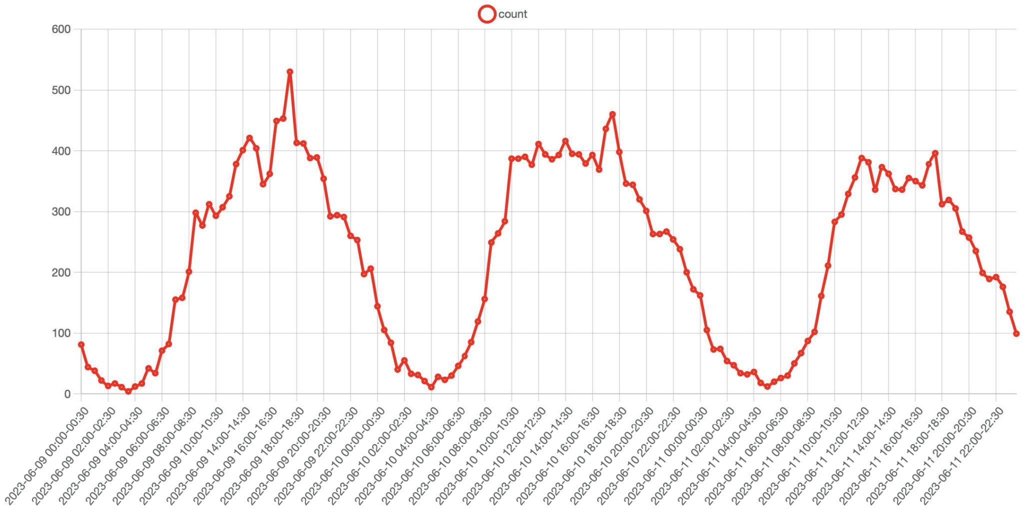

The metrics allow these insights to be extended further. Figure 6 shows the fraction of pedestrians approaching the crossing who then stop in the wait zone and cross v. the fraction that pass on along the same side of the road and do not use the crossing. The two graphs relate to the flows towards the crossing in the two opposite directions along the pavement on the same side of the road.

Figure 6 Pedestrian to Crosser Conversion Rate

East Bound

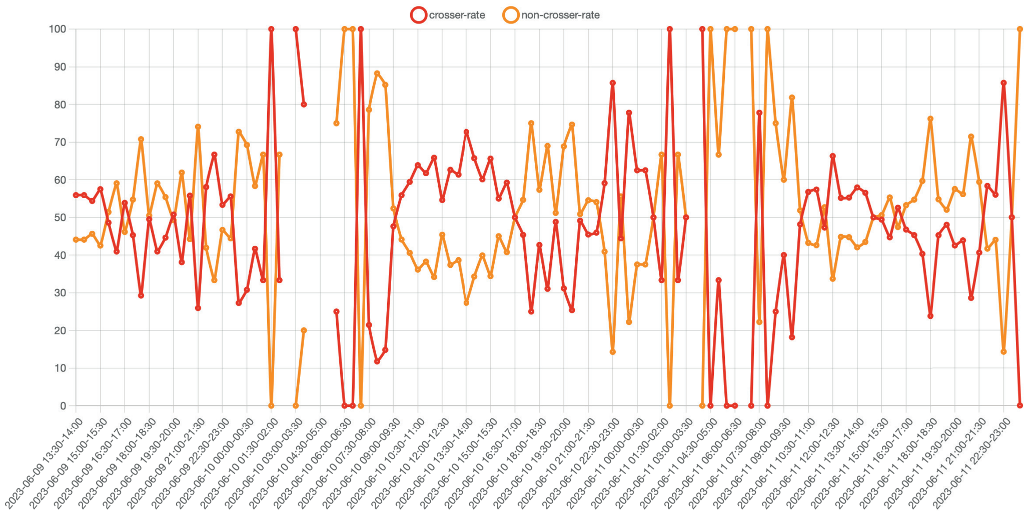

West Bound

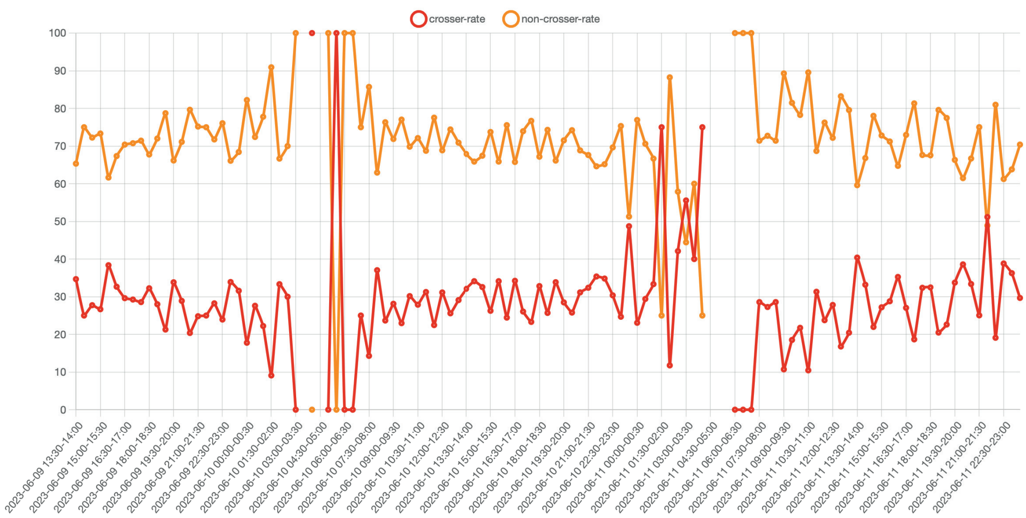

There are two immediate observations. The first is that at times of little pedestrian activity, i.e. over-night, a single pedestrian either crossing or passing by can make a massive difference to the fractional value. This explains the spiky periods in both graphs. The other observation is that on the left, there appear to be more crossers than passers-by in the mornings, but more passers-by than crossers later in the day. This would seem to show a tidal flow of pedestrians. This could be accounted for by, for example a bus stop further up the road bringing students into the city who then cross to access a college building. Of course, other interpretations are equally possible, but the important point is that we are now in a position to see the behaviour. In the other direction, on the same pavement, there is a much more stable split between crossers and passers-by. This too allows origin and destination interpretation of the pedestrian flow. These plots are, in effect, turning counts for pedestrians.

Operational Enhancements

Having demonstrated the operation of metrics, we can now look at how these metrics can be used to improve the puffin crossing operation.

The starting point for this is to note that, for the majority of the time, the puffin strategy, and the way it is set up, is handling pedestrian demands effectively. However, just occasionally the pedestrian demands, flows and accumulated delays may warrant a more incisive intervention.

As the signals are controlled by the signal controller, any enhancement has to utilise and manipulate the input and output data to and from the controller to artificially change the perception of external conditions and provoke the desired response.

By way of example, and as suggested by an expert in this field[4], on a site that is under UTC control the PV and PX bits[5] may be used to force or prevent pedestrian crossing activation. Where signal controllers are configured with these bits exposed then any external system may use them to invoke control. As these are external to the controller functionality, this approach has the benefit of not altering the settings or configuration of the signal controller allowing it to be simply and easily deployed at existing UTC sites perhaps only needing the UTC bits to be brought out to accessible IO and logic to allow take-over and hand-back by this pedestrian driven optimiser with little additional effort or cost.

This paper does not cover the specifics of the optimisation process as it is a significant subject in its own right it does consider the ways in which the application of such optimisation can be tested.

[4] Kennett, C. private communication 2023
[5] TOPAS 2523B Traffic Control Equipment Interfacing Specification p16

Verifying Optimisation

It is an inherent problem for any situation where an intervention is made, to show that the intervention has made an effective difference. While it is not so difficult to design an experimental scheme that alternates between the case with the intervention and the case without, and even to run a double-blind test when processing the results, where tests of this type are undertaken in an uncontrolled environment, i.e. the real world, other factors may also be influencing the experiments. Consider the effects of rain on the levels of pedestrian activity. It is reasonable to propose that the number of pedestrians may fall, as only those having to walk will be out in the rain, and that the average speed of those pedestrians may be higher as the less mobile may be deterred from walking more than the fully mobile. Other external factors, such as school terms/holidays, tourist season/ out of season, all could affect pedestrian behaviour. The effect is worse than just these changes to pedestrian behaviour, as these changes may alter the need for optimisation in either direction and, depending on the optimisation process, could influence the algorithm so also skew the measured effectiveness of the optimisation.

As with similar trials in different fields, the typical approach is to run small-scale tests, typically with interleaved periods of intervention and base-line operation, and depending on the level of benefit that is discovered, either alter the algorithm and/or extend the data gathering for longer periods and/or to more sites.

Results of Optimisation

What is presented below is a before and after snapshot of early data at the trial site where optimisation has been applied. Before data was gathered from 4th – 31st May (28 days) and after data from 30th June – 5th July, 9th – 12th July, and 19th to 25th July, (17 days). Data has been normalised to the average value per day for each of these periods.

Table 4 Preliminary Data from Crossing Optimisation at East Rd.

	Metric Data prior to optimisation intervention	Metric Data with optimisation
Total bidirectional crosser count (/day)	1993	2026
Total bidirectional vehicle count* (/day)	20794	20406
Accumulated pedestrian delay both wait zones (mins/day)	301	191
Max pedestrian density at crossing entrance S and N wait zones (peds/m^2)	1.01 & 0.83	1.09 & 0.94

*Comprising of ~16,000 cars and vans, with the balance being busses and trucks, motorcycles, scooters and cycles.

Observations on the Results

The data of Table 4 allow some observations to be drawn. The obvious difference between before and after is the drop in the average daily pedestrian delay from 301 minutes per day to 191 minutes per day, i.e. a 36% reduction in average pedestrian delay. The obvious challenge to that reduction is to ask the question, “Were there less pedestrians during the second period?”. The crosser count shows that there was little difference in the numbers of pedestrians using the crossing, if anything the slight increase during the intervention period would suggest that the saving of 36% is pessimistic and that crosser numbers do not account for the improvement. Similarly, there was little difference in vehicle traffic, so there was no change to the opportunities to cross. Lastly, the slight increase in pedestrian density with the intervention present might suggest that the algorithm was successfully platooning pedestrians and identifying gaps to service their needs more effectively. Currently we do not have data for the delays to vehicles during these trials so can not identify any disbenefit to the road traffic.

The optimiser used is a first-generation implementation designed to intervene only under periods of significant demand, while for most of the time, the site was running as a puffin crossing.

Conclusions

As noted in the abstract, historically, local authorities have not had good access to data on pedestrian movements through their networks. In the course of this paper we have shown that the newest round of technologies do have the capacity to substantially reduce the amount of hardware deployed at crossings to deliver the detection required for puffin operation. We have also developed, implemented, and deployed metrics for pedestrians and for their interactions with other traffic and presented examples of how we can quickly get insights into exception events.

This work led to the development of an optimisation process for crossings where, under conditions of unusual demand, strategies other than those of puffin operation could be applied and yield improvements in the experience for pedestrians using the crossing.

Returning to the opportunities that this presents to highway authorities, they now have access to equipment to reduce the amount of hardware at existing or new pedestrian crossings. They also have access to metrics and data about pedestrians moving through their network to understand where interventions may be needed and lastly, they have the means to respond to these pedestrians with far more sensitivity, improving active travel journeys, reducing exposure to pollutants and to show that this has been achieved.

Acknowledgements

Starling offer our thanks to Cambridgeshire County Council for their kind support of this trial. Specifically Dan Clarke of the Greater Cambridgeshire Partnership, Richard Ling from the Traffic Signals team and Michael Stevens their Enabling Digital Delivery Manager. Their support, insights and encouragement have been and continue to be invaluable.