

MAPPING CARBON MONOXIDE USING GPS TRACKED SENSORS

RICHARD MILTON and ANTHONY STEED

*Department of Computer Science, University College London, Gower St., London,
WC1E 6BT, United Kingdom*

{R.Milton, A.Steed}@cs.ucl.ac.uk

Abstract. In this paper we discuss a pilot study where we have mapped urban air pollution using mobile carbon monoxide (CO) sensors. Our objective is to use logs from tracked CO sensors to augment fixed sensor data in an urban environment. The concept is that by using many cheap mobile sensors in an ad-hoc fashion we will be able to facilitate new types of CO study that explore variations in CO at a fine geographic scale.

We gather data using a data-logging package that consists of a personal digital assistant (PDA), an inexpensive Global Positioning System (GPS) receiver and a CO sensor. The critical issue in data processing is treatment of the imprecise logs from the GPS. By using knowledge about the route and geometry of the buildings we are able to increase the position accuracy significantly.

Using corrected log files, we demonstrate and explore three potential scenarios of use of a mobile sensor: exploration of the situation of a fixed CO sensor, variation of CO along an urban canyon and discovery of intermittent but significant sources of CO. We conclude by discussing the results in the context of the push towards large sensor networks and mobile communications. The potential for ad-hoc mobile sensor networks may be very large.

Keywords: air pollution measurements, carbon monoxide, Global Positioning System (GPS), personal exposure, tracking mobile sensors, urban pollution

1. Introduction

Many urban environments are now monitored for carbon monoxide (CO) and other pollutants. This is because of health and quality of life issues associated with high exposure to such pollutants. Most CO monitoring is done from fixed sensors that provide continuous data feeds. However, whilst fixed sensors can tell us about temporal variation

and the range of values that might be experienced, such as the difference between a kerbside and green-field monitor, because of their geographic scarcity they do not allow us to explore variation at a scale that would allow us to see effects of particular sources, dispersal by wind or accumulation in calm areas (see Section 2 for a review of urban CO monitoring).

In this paper we describe a pilot study where we have investigated CO pollution in a city environment using a small set of tracked mobile CO sensors. The motivating concept is that small mobile sensors could be carried by the general public going about their everyday business. Mobile devices are dropping rapidly in price but increasing in capability. By its nature, data collected by non-experts in an ad-hoc manner would be less reliable and only give intermittent temporal coverage, but we expect that by aggregating many unreliable and intermittent logs files we would be able to produce a dense geographic map of variation in CO. By using tracked CO sensors we trade temporal coverage for spatial coverage. These “pollution trails” can be aggregated over time to generate maps of pollution or they can be compared with other trails that go through the same space so that temporal or transient variation can be explored. These maps and comparisons might allow us to do new types of study of CO that look at personal exposure whilst moving through a city, small-scale variation due to street canyon effects and sources and dispersal of CO at a much more detailed scale or with much less preparation and infrastructure than previously achieved.

In this paper we start by discussing further background and related work, and give an outline of the scope of the trials we have undertaken in Section 3. In Section 4 we then discuss the quality of the data that is collected. A key question is the treatment of the unreliable GPS position reports. The quality of GPS position reports varies because the satellite constellation used changes over time, but more significant is the urban-canyon effect that causes the amount of the sky that is visible to change significantly as the user moves.

In the following three sections we then outline three different types of analysis that we can do with the data. In Section 5 we do a comparison between the mobile sensor and a fixed sensor demonstrating that the two types of sensor give similar readings. In Section 6 we detail how we have built profiles along a couple of busy road sections. This is the main contribution in this paper because it suggests that we are able to produce interesting

profile data that gives the readings that might be expected given the prevailing weather and traffic conditions. In Section 7 we then outline some more novel observations, such as CO profiles whilst crossing roads and unexpected discovery of high CO levels. Finally we conclude with a discussion of the potential for this work and the field of mobile sensors.

2. Related Work

2.1. CO in Urban Environments

Independent studies of urban carbon monoxide have shown large variations caused by street layout, meteorology and traffic flow. Croxford, Penn and Hillier (1995) found peak CO concentrations in London, UK, exceeding 10ppm for short periods with nearby sensors reporting much lower values. There are several related works that aim to increase understanding of pollution variation. These include the dense sensor network of the Discovery Net e-Science project (2004), and the combined sensing and air flow modeling approach of DAPPLE (2005).

Complex numerical models exist to predict urban air quality, but these are dependent on suitable input data. The application of meteorological data to pollution monitoring was the subject of COST (European Co-operation in Science and Technology) action 715 (2004) and is still an ongoing area of research. Boddy et al. (2004), in a fixed sensor study around York, UK, use anemometers placed on opposite sides of the street to obtain accurate wind information, while DAPPLE used a number of anemometers placed along the street at different heights. The study in York is of particular interest as the fixed CO sensors used were Learian Streetboxes, which are very similar to the portable Learian sensors used in this experiment. These studies both use meteorological information generated on site, but synoptic data on this scale will not generally be available in a large-area study.

Personal exposure monitoring seeks to measure the concentrations of pollutants that people are exposed to in their daily lives. Although tracking experiments using vehicles have been very successful, as demonstrated by Weijers et al. (2004) in Northern Holland, tracking of people is only just emerging as a feasible proposition. The authors of this paper have successfully conducted a preliminary test in the Marylebone Road area of

London, UK, using sensors carried by pedestrians and cyclists to verify that expected CO features could be detected (Milton and Steed, 2005). In this paper we describe a wider scale experiment where we have been able to compare data from mobile sensors with an existing fixed location sensor. We then demonstrate that significant variations in CO level can be detected along roads. The analysis follows the accepted fixed sensor practice used by Manning et al. (1998) and Boddy et al. (2005) of plotting data by wind direction, along the street axis and perpendicular to it.

2.2. GPS Tracking

The Global Positioning System (GPS) is currently the most widely used positioning system. According to the Garmin fourth quarter financial report for 2004, 2,300,000 units were shipped that year making 10,000,000 in total. Other estimates suggest that over 100,000 GPS receivers are shipped each month, either for use in mobile phones, PDAs or in standalone units.

With so many different manufacturers producing receivers, there are a number of different options available with varying levels of accuracy and cost. We have chosen consumer-level hardware because it is important for further work that sensor packages can be fabricated cheaply. With a clear view of the sky, a consumer GPS receiver that might retail for less than \$150 has an accuracy of between 10m and 15m, although with the right satellite geometry this can be within 5m some of the time. Even at the consumer level there are a number of GPS enhancements that can improve the accuracy considerably. Many commercial GPS receivers output range and phase information when used in binary mode, which opens up the possibility of post processing the data to remove certain sources of error. The National GPS Network (see Appendix) provides many of the resources in the UK for post processing and contains other useful information about GPS systems. Space Based Augmentation Systems (SBAS) can also remove some sources of error and provide accuracies of 3m in real-time for receivers that support the feature. This includes the Wide Area Augmentation System (WAAS) that operates in the US, while the European Geostationary Navigation Overlay Service (EGNOS) covers Europe. The problem with these systems is that they require a geostationary satellite to be visible, and in an urban environment in Northern Europe, this is frequently not possible as the

satellite is not high enough in the southern sky. This problem can be alleviated using a system called SISNeT that transmits the EGNOS correction data via an Internet link. However because we log data continuously over long periods, we would need an Internet that was cheap (thus excluding GPRS) and reliable and pervasive (thus excluding Wavelan). To improve accuracy this means that we will have to retrospectively correct the data logs using knowledge of building footprints and carrier behavior (see Section 4.2).

3. Trial Logistics

3.1. Equipment

The equipment used for the data collection consists of an HP Jornada PDA running our own data logging software, a Garmin 15L GPS receiver with external aerial and an ICOM CO sensor manufactured by Learian Environmental. The ICOM CO devices use an electrochemical CO sensor and are accurate to 0.1ppm. References to further information on these devices are in the appendix. All the equipment is housed in a rucksack with the sensor open to the air in a mesh pocket on one side (see Figure 1). Other than switching everything on and starting the data logging software, no further interaction is necessary until it needs to be switched off.



Figure 1: The sensor package in a small rucksack. The ICOM sensor is in the mesh pocket on the right and the GPS aerial is visible on the top of the sack. The PDA and GPS unit are stowed in the outside pocket, leaving the main rucksack compartment empty.

The CO level and GPS position are stored once every second and written to log files on the PDA. Once the logging software is stopped, these files are uploaded to a server for analysis.

3.2. Server and Tools Support

The CO and GPS data are stored in a Postgres database and uploaded or viewed using web-based services. This allows data to be analysed in the field where a wireless Internet connection is available. More complex analysis tools requiring direct access to the database have also been developed and these have been used to produce the composite charts in section 6. Various forms of 3D visualisation have also been used during the data analysis process.

3.3. Geographic Area and Time-Scale

As discussed in Section 2.1 our previous experiments have concentrated on a single road, or route around a small area. In the current study, the emphasis is on proving the potential of more ad-hoc data collection over longer periods of time. Thus, the area is much larger (1km x 1.5km, see Figure 2) and the period of data collection is 2.5 months (November 2004 – February 2005). This particular area of London was chosen because it was also being studied by the Vivacity2020 project. That project placed four fixed pollution sensors in the local environment and conducted other related studies around the area in that period.

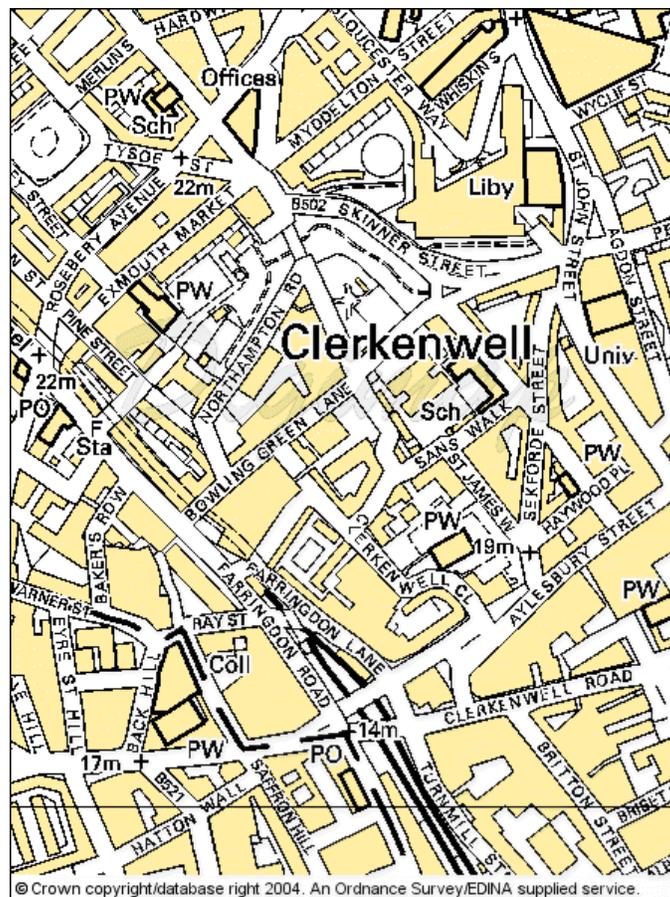


Figure 2: The Clerkenwell area with Exmouth Market, Farringdon Road and Clerkenwell Road shown.

As mentioned in Section 1, the concept motivating these trials is that the data would be collected by a passive data collector who visited an area occasionally whilst going about their everyday business. However, at this stage there were only a few devices available, so it was necessary to concentrate the effort on a few smaller areas to collect enough data for analysis. The data collectors were instructed to walk on the kerb-side whenever

possible to further reduce ambiguity in the dataset. In the area chosen, we hypothesized that most pollution and variation in pollution would be found on Clerkenwell Road and Farringdon Road because they have the highest amount of traffic. We also explored Exmouth Market, which is a pedestrianised area, bordering Farringdon Road. Other roads around the area were covered in less detail as the data collectors explored the area.

4. Initial Data Processing

In total, data was sampled for 94 separate paths between 17 November 2004 and 4 February 2005. This resulted in a total of 328,601 data points, of which 183,830 had valid GPS positions. The first author of this paper collected all of the data presented, apart from two days in November when a second person was available.

Three or more GPS satellites could be seen, and thus a GPS fix received, for 55.9% of the time. This is slightly higher than the 42% in our previous experiments on Marylebone road, see Milton and Steed (2005), and is comparable to figures achieved with a consumer GPS unit in a study of GPS availability in London carried out by Ochieng et al. (2003).

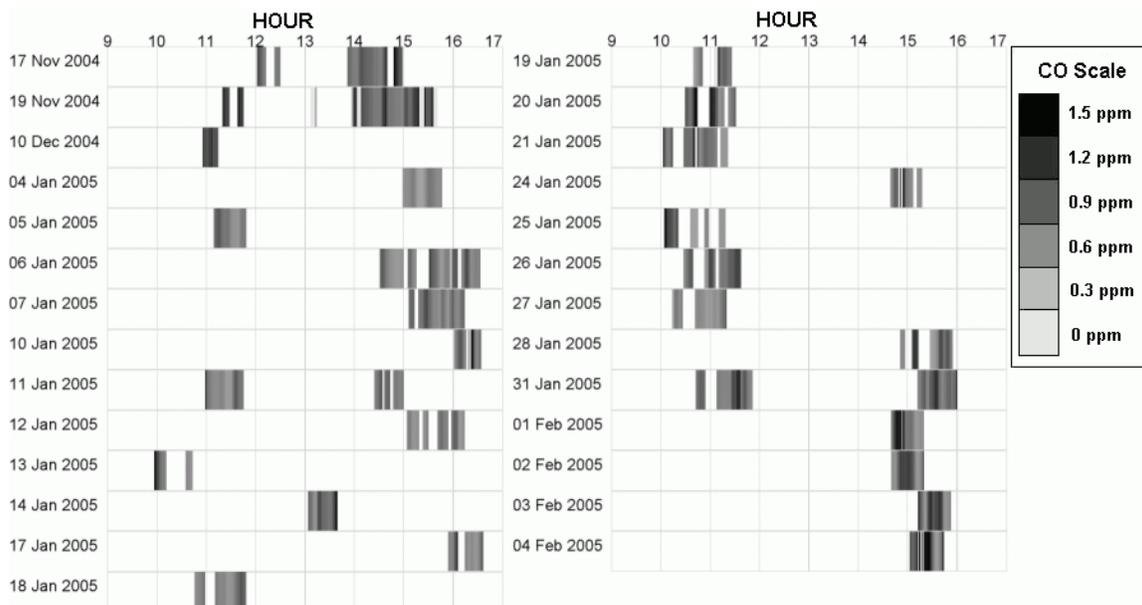


Figure 3: All data collected in the Clerkenwell area between 17 November 2004 and 4 February 2005.



Figure 4: the complete set of data collected. This is a plot of mean CO values on a 10m grid, after GPS correction (see following sections).

4.1. GPS Accuracy

The GARMIN GPS receivers used for the trial report an estimate of the horizontal position error in metres, in addition to the usual horizontal dilution of precision (HDOP), or standard deviation of position errors. Using the horizontal error, the percentage of the time that the GPS error was within certain limits can be calculated. Figure 5 shows the cumulative probability of obtaining a GPS fix with an accuracy better than the radius in metres indicated on the X axis. Using this data, it can be seen that the GPS positions are within 15 metres 64% of the time, or better than 45 metres 95% of the time. The accuracy of the GPS devices used for this experiment should be within 15 metres 95% of the time, but in an urban environment this is significantly reduced due to a limited view of the sky forcing the GPS to use non-optimal satellites for the position fix.

Occasionally, multipath reflections from buildings can cause a much larger error than this estimate based on satellite geometry would suggest. In this situation, the reflected signal from the satellite will be attenuated and show a longer path to the receiver than the direct line of sight, displacing the location fix by many metres.

One final piece of information reported by the GPS is the track angle. This is a compass heading reported by the GPS giving the direction of travel. The angle reported is not particularly accurate, but is good enough to highlight 90 degree changes in direction

when crossing roads. A track angle can also be calculated manually from successive GPS points, but the track angles reported by the GPS will also work while the receiver is stationary.

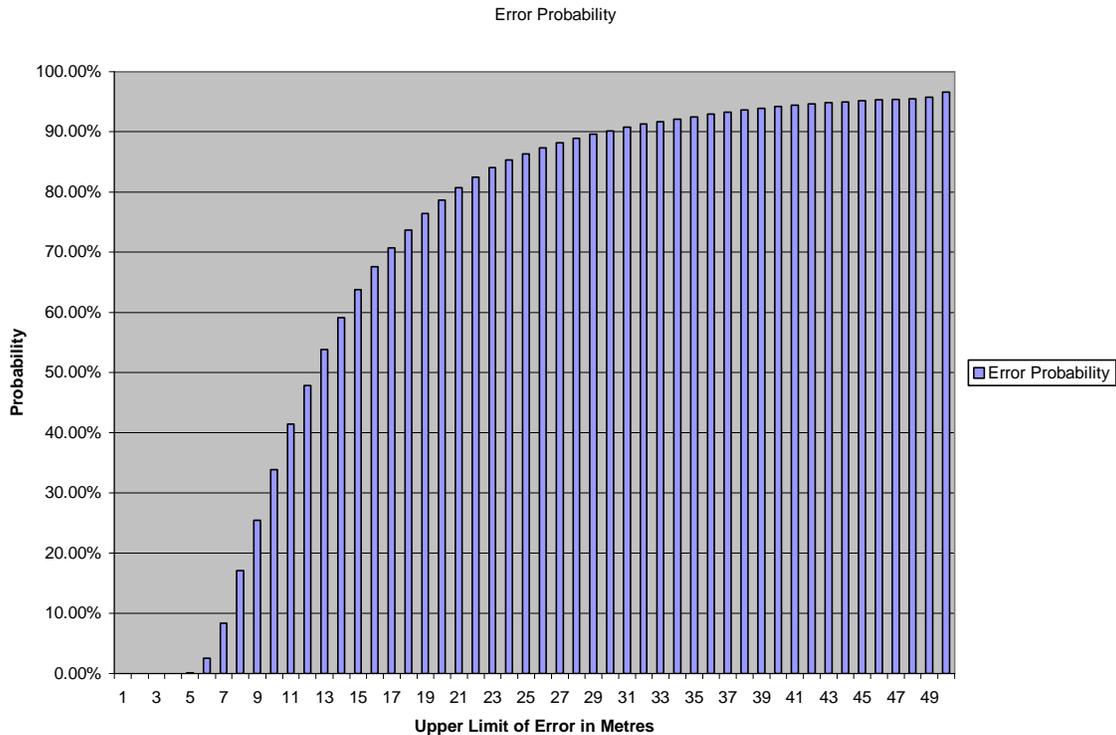


Figure 5: The cumulative error of GPS accuracy from all log files. This data can only be plotted for times when a 2D or 3D position fix was possible.

4.2. GPS Correction

The GPS accuracy discussed in the previous section indicates that, on its own, GPS is not sufficient to position a user on a particular side of the street. No matter what tracking technology is used, there will always be some ambiguity, but street-side is very important for the scenarios we envisage, since CO levels can differ greatly due to traffic flow, building layout, wind speed and direction. Using a single sensor for the pilot study, differences between sides of a street will be hard to measure unless we can guarantee that conditions do not change in the time it takes to cross the road.

Figure 6 shows an extract from a section of a log along Clerkenwell Road. The section investigated (between labels A and D) contains three bus stops and a petrol station. It is a

short stretch of road approximately 15m wide from building to building. The sensor travels from west to east on the north side, east to west on the south side and then repeats the route, giving two traces for each side of the road. It is not immediately apparent which side is which from the figure.

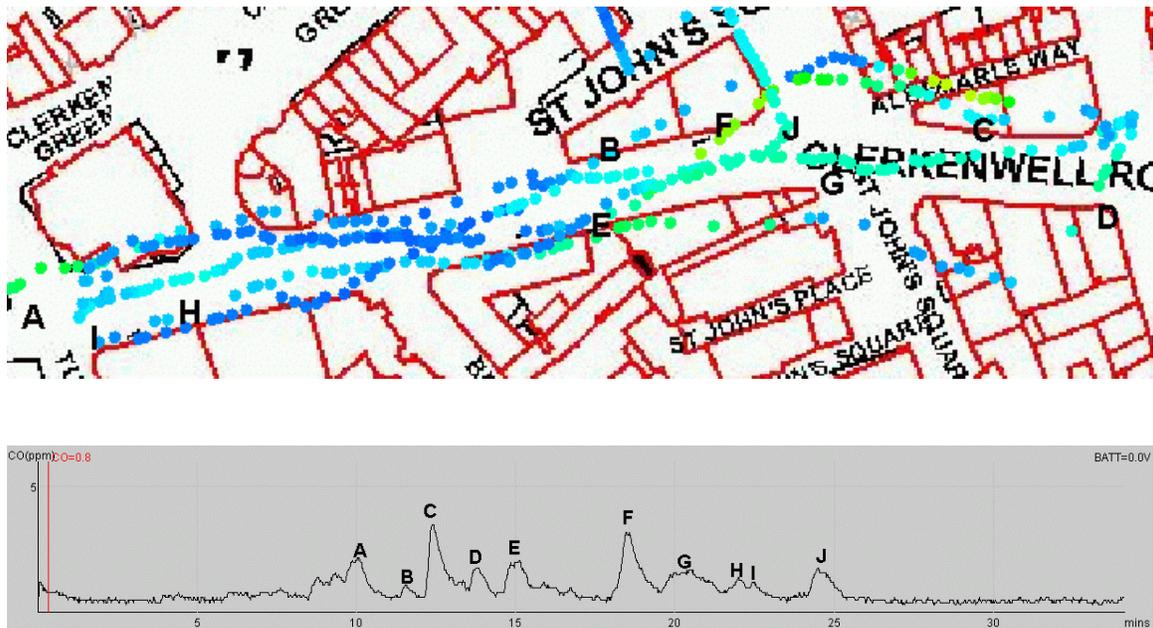


Figure 6: The raw GPS and CO traces along Clerkenwell Road on 4 February 2005 at 15:05, sampled west to east on the north side, east to west on the south side and then repeated to give two traces for each side of the road. North is upwards. Labels A-J are discussed in the main text.

Given we know the expected behavior of the carrier, we can note some characteristics of these behaviors in the GPS log. Direction of travel is an obvious one, and this would allow us to segregate north and south in this particular example given we know the route. A further observation is that points where the road was crossed can be identified from where the carrier of the GPS waits by the side of a road before crossing, or from the track angle. This opens up the possibility of automatically extracting crossing profiles from the collected data to investigate how CO varies across the width of the road.

When analyzing the data in more detail, it is apparent that there are distinct CO peaks sampled along the road where the levels are rising or falling between successive samples. Points B and E and also F, G and J in Figure 6 show examples of this. The GPS device

reports the standard deviation of the position errors calculated from the satellite geometry at each sample point, and it is known that the sensor was always carried along the pavement by the kerb, so, by correcting the raw GPS positions to the closest kerb-side point, the positions can be corrected to within a reasonable degree of error.

In order to perform the correction, the Ordnance Survey data for the area is used to find the closest pavement edge to the reported GPS position. If no edge is found within a 20m radius, then the GPS point is not moved. Side road crossings are detected from the change in track angle and points in the road can be interpolated from the raw data and previous modified point. Figure 7 below shows the results of applying this algorithm.



Figure 7: Clerkenwell Road 4 February 2005 at 15:05, GPS corrected to the pavement with detection of road crossings based on a change in track angle. This does not use any additional route knowledge.

In general though, correcting to the nearest kerb point can only be applied in conjunction with a threshold less than half the road width. If the corrected data in Figure 7 is animated it is apparent that a significant percentage of points have been placed on the wrong side of the road as the position reported by GPS has crossed the road centre line. This example has been deliberately chosen as a test case as the correct side of the road can be worked out from the direction of walking: left to right on the north side and right to left on the south side. In this case, to expose the CO peaks on opposite sides of the road, the only available option is to use route knowledge.

In summary, the raw GPS data from a current consumer GPS shows errors less than 15m for 65% of the time. Thus, automatically placing a point in the correct position on

the pavement, or even just placing it building side or kerb side, is unrealistic. However, by using additional map knowledge and knowledge about carrier behavior, these errors can be reduced in a semi-automatic editing process. This allows various types of analysis on the CO data to be performed. We discuss the general problem of semi-automatic log file correction in another paper by Milton and Steed (2005).

5. Comparison with Fixed Sensors

5.1. CO Values in the Vicinity of a Fixed Sensor

During the period the mobile sensors were being used in Clerkenwell, there was also a fixed location CO sensor in the area as part of the Vivacity2020 project. This fixed sensor, sited in Exmouth Market, outputs a 15 minute mean, four times an hour. A comparison was made by taking the average of all mobile sensor data collected in a rectangle 43m by 22m immediately in front of the fixed sensor. This area was chosen to cover most of the street in front of the sensor and allow for small GPS errors. There were 26 separate comparisons possible during January and February, but one unavoidable problem with the results was that the mean CO levels throughout the period were very low. The mobile sensors are accurate to 0.1ppm, while the fixed sensor mean for the whole period is 0.3ppm. Despite this, and other sources of error, there is a good correlation between the fixed and mobile sensors. The overall RMS error of the difference between the fixed and mobile sensors is 0.17ppm, while the 95% confidence level is 0.06ppm. This gives a 95% probability that the fixed and mobile difference will be +/-0.06ppm from the mean, very close to the accuracy level of the sensors.

The Exmouth Market sensor is mounted 3m above the ground in a pedestrianised area, so the background CO level should dominate the readings with few traffic effects. The data produced by the fixed and mobile sensors are not directly compatible, leading to a number of sources of error. Situations where the CO level is slowly rising or falling throughout the 15 minute period will cause the mobile sensor's short period mean to be too high or low. In this particular experiment, this did not cause much of problem due to the generally low values, so no attempt to correct for it has been made.

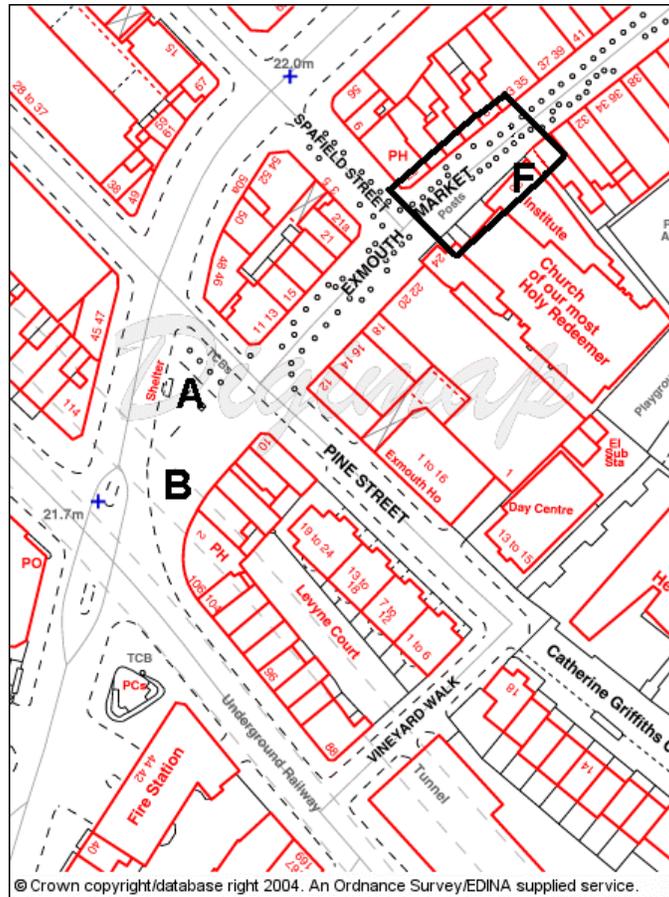


Figure 8: Exmouth Market with the fixed sensor position labelled 'F' and the catchment area outlined in black. Two nearby CO peaks are labelled 'A' and 'B'. The Distance between 'A' and 'F' is approximately 80 metres.

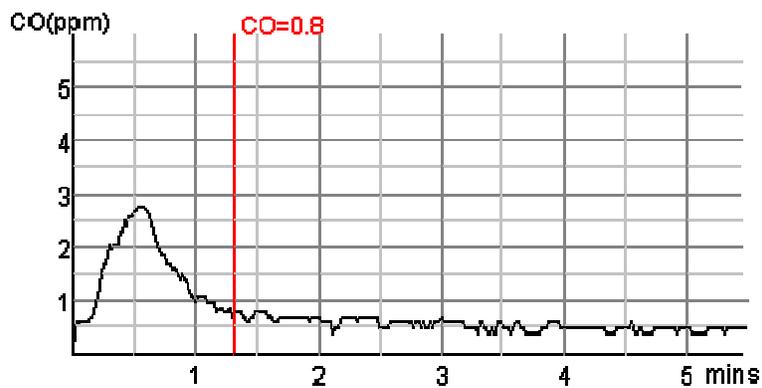


Figure 9: CO Peak 'A', 20 January 2005.

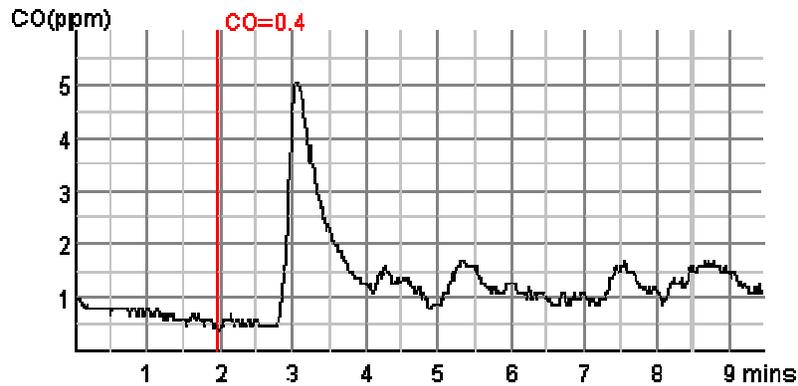


Figure 10: CO Peak 'B', 1 February 2005.

The secondary aim of comparing fixed and mobile sensor data was to identify how CO levels varied around the location of the sensor. Data for 20 January and 1 February are particularly good examples of where the mobile sensor sampled large CO peaks (2.8ppm and 5.0ppm respectively) at a major junction about 80m from the fixed sensor. On both occasions the minimum recorded mobile sensor reading is closer to the fixed value than the mean, due to the reduction in CO as the fixed sensor is approached.

The vertical CO lines on the graphs in Figure 9 and Figure 10 indicate the location of the fixed sensor. Figure 9 shows a gradual drop in CO from an area where food was being cooked on portable gas cookers near the junction to the fixed sensor location. Figure 10 shows a high peak closer to the junction itself and a definite CO gradient from 0.8ppm to 0.4ppm between the junction and the fixed sensor.

In general, the CO trend from the fixed sensor shows how the background levels are varying with time, while the variance of the mobile samples around the area show how they vary with distance. These two facts can be used to improve the comparison between the fixed and mobile sensors.

Given the limitations inherent in comparing the fixed and mobile data, together with the low CO levels during the winter, the result of 0.007 ppm mean +/-0.06 ppm with 95% confidence is surprisingly good and is within the accuracy of the sensors.

5.2. CO Levels at Fixed Times of the Day

Given that there is a good match between mobile readings and static readings, the next logical step is to plot how the CO levels vary throughout the day. However, here we hit a

problem of data density. Because the device is moving, there is good spatial coverage of the area, but very little data for specific locations e.g. within 20 metres of a fixed sensor. Both daily running means and scatter graph plots were attempted using the mobile data by selecting points around the Exmouth Market fixed sensor. In this case, the mobile data shows an unrepresentative picture when compared to the fixed sensor of the Vivacity2020 project, due to the poor temporal coverage of the mobile sensor. For this type of data to be plotted using mobile sensors, much more data is required to give a uniform spread of samples throughout each day of the sampling period.

6. Composite Profiles Along Roads

In the following results, both automatic GPS correction and manual correction based on knowledge of the route have been combined to identify features at a finer scale than is possible with the raw GPS data. Profiles along the length of roads have been plotted to show the variation in CO levels. This data could be used to decide the best locations for fixed sensors, or to calculate how critical the positioning of a fixed sensor is. The method of plotting is to average all the values within a radius that spans the road, so a GPS correction algorithm that places samples on the correct side of the road is not necessary in this context. The general CO trend along a road can be calculated from the raw data to an acceptable level of accuracy, while simple linear interpolation of missing GPS points can increase the accuracy further. More complex correction schemes are not necessary as long as the data is bounded by the confines of the road.

6.1. Clerkenwell Road

Using wind data for the closest synoptic station and time, composite profiles of days with similar wind directions can be produced. Clerkenwell Road runs along bearing 255 degrees, 15 degrees rotated from west to east. With a 25 degree window either side, the two parallel and two perpendicular wind groupings are shown in Table 1.

| Along Road | Number of Days | Mean Wind Speed | Bearing |
|-------------------|-----------------------|------------------------|----------------|
|-------------------|-----------------------|------------------------|----------------|

| | | | |
|----------------------|---|----------------------|------------|
| West to East | 9 | 5.78ms^{-1} | 230 to 280 |
| East to West | 1 | 5.0ms^{-1} | 050 to 100 |
| Perpendicular | | | |
| North to South | 5 | 3.4ms^{-1} | 320 to 010 |
| South to North | 0 | | 140 to 190 |

Table 1: Clerkenwell Road CO data grouped by time and wind direction relative to the road.

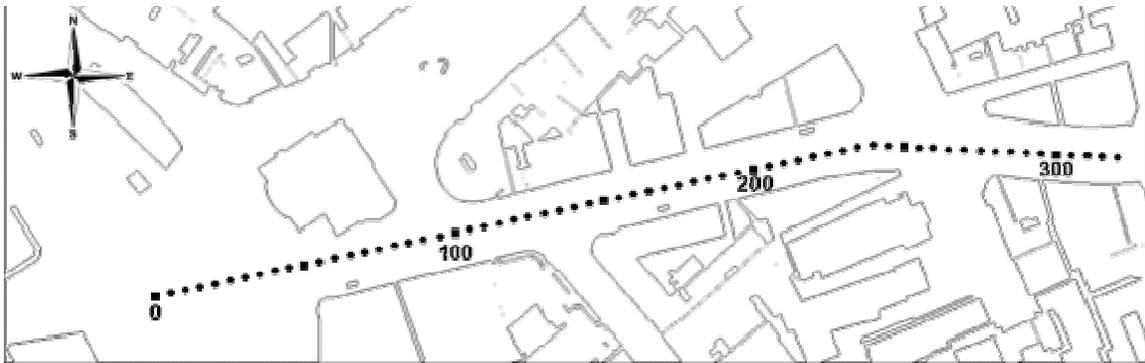


Figure 11: Clerkenwell Road showing the distances in metres corresponding to the following Figure 12 and 13.

CO Mean and Variance Along Clerkenwell Road

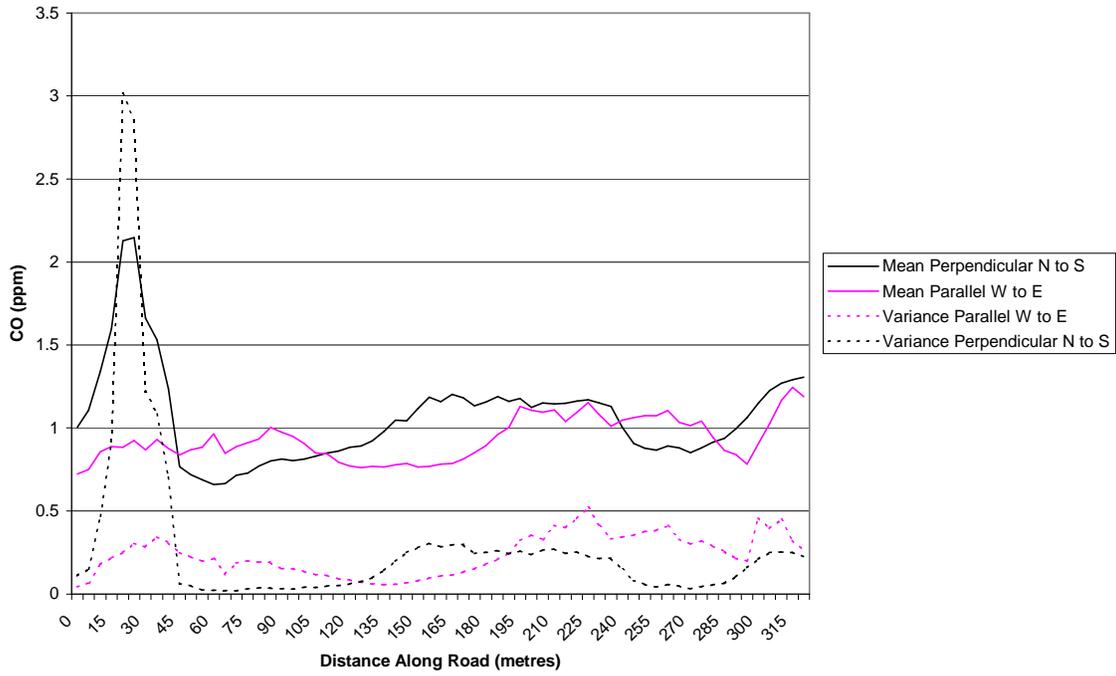


Figure 12: Comparing profiles for W/E (parallel) winds and N/S (perpendicular) winds.

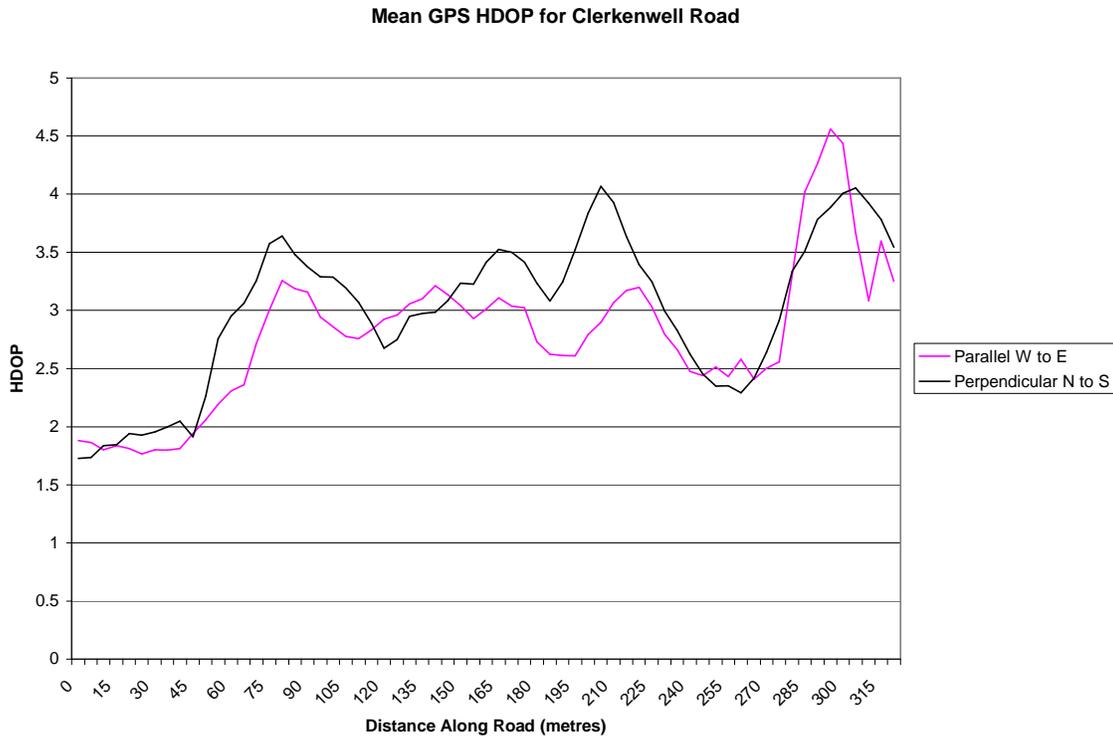


Figure 13: Horizontal Dilution of Precision (standard deviation of position errors) reported by the GPS receiver for the position data used in Figure 12.

Figure 12 shows similarly shaped CO traces for both parallel and perpendicular winds, except for the peak of 1.8ppm at the west end of the road on the perpendicular wind profile. The only other difference is that the parallel trace is shifted approximately 45 metres west. It is interesting to see that the variance around the 1.8ppm peak at 30 metres is significantly higher than anywhere else along the road, while there are also large areas of very low variance.

Figure 13 shows the ‘Horizontal Dilution of Precision’ (HDOP) reported by the GPS for the same data as used in Figure 12. The reported HDOP value is the standard deviation of the position errors as calculated by the GPS, based on the satellite geometry used for the position fix. The graph in Figure 13 is plotted by taking the mean of all the HDOP values within the same positions along the road profile as the previous Figure 12, giving an estimate of the GPS accuracy along the road. In order to obtain a rough estimate of the position accuracy in metres, the HDOP value should be multiplied by the receiver’s basic accuracy, in this case 11m.

From Figure 13, it can be seen that the GPS accuracy in Figure 12 varied between 22m and 49.5m along the length of the road. The most noticeable feature of Figure 13 is the similarity in the shape of the two curves. This is an important result as it shows that there was no bias in the CO data of Figure 12 caused by differences in GPS accuracy. Given that each is a plot of data from different days and so no data is shared between the two curves, it suggests that building layout along the road is affecting the satellites visible to the GPS. HDOP is a function of the position of the satellites, which you would not normally expect to change so quickly in such a small area, unless the view of the sky is occluded and the GPS is forced to use non-optimal satellites for the fix. This is supported by the fact that the two lowest points of the graph are at 0 metres and 255 metres. From 0 to 45 metres there are no surrounding buildings as can be seen from Figure 11. At 255 metres there is a large gap in the buildings at St John's Square on both sides of the road. The HDOP graph matches the position of this gap between the buildings exactly and is the same width at 45 metres. This sudden widening of the road causes an interesting CO gradient, while the junctions at either end also produce CO peaks where traffic is idling. This means that we may be fortunate in that where we expect interesting pollution events, we might expect to get better GPS availability and accuracy.

Taking the GPS accuracy into account, it is possible that the CO peaks on the two traces in Figure 12 should line up as they are approximately 45 metres out of phase but show the same trough, peak, trough, peak pattern from 110 metres onwards. It is also possible that this 45 metre offset could be caused by the difference in wind direction, but additional information is needed to prove this. The high CO variance and mean from 0 to 45 metres when the wind is northerly is probably due to the wind blowing straight down Farringdon Road and across the very open junction with Clerkenwell Road. Traffic build-up here is common due to the road layout and this area was shown to be particularly bad during the whole of the collection period. Areas where the CO variance is high need to be treated with particular care, as previous studies have shown that high CO variance and areas of low GPS accuracy are often linked. In this case, the GPS HDOP value is only 2.0, so the high CO variance is almost certainly a real CO characteristic rather than an effect of badly positioned samples. One other feature worth noting in Figure 12 is the final trough on both traces at the 260 metre (northerly wind) and 300 metre (easterly wind) positions. This corresponds to the point where the road opens out on both sides into

St. John’s Square. When the wind is blowing along the road from the east, this trough is very narrow, but when the wind is from the north there is a drop in CO the same width as the gap in the buildings. This is also matched by extremely low variance. On many occasions during data collection, CO gradients around this corner were measured.

6.2. Farringdon Road

Farringdon Road runs along bearing 305, approximately northwest to southeast. Allowing a 25 degree window either side, the two parallel and two perpendicular wind groupings are shown in Table 2.

| Along Road | Number of Days | Mean Wind Speed | Bearing |
|---------------------------|-----------------------|------------------------|----------------|
| Northwest to Southeast | 7 | 5.57ms ⁻¹ | 280 to 330 |
| Southeast to Northwest | 0 | | 100 to 150 |
| Perpendicular | | | |
| Northeast to Southwest | 7 | 4.43ms ⁻¹ | 010 to 060 |
| Southwest to Northeast | 6 | 5.17ms ⁻¹ | 190 to 240 |

Table 2: Farringdon Road CO data grouped by time and wind direction relative to the road.

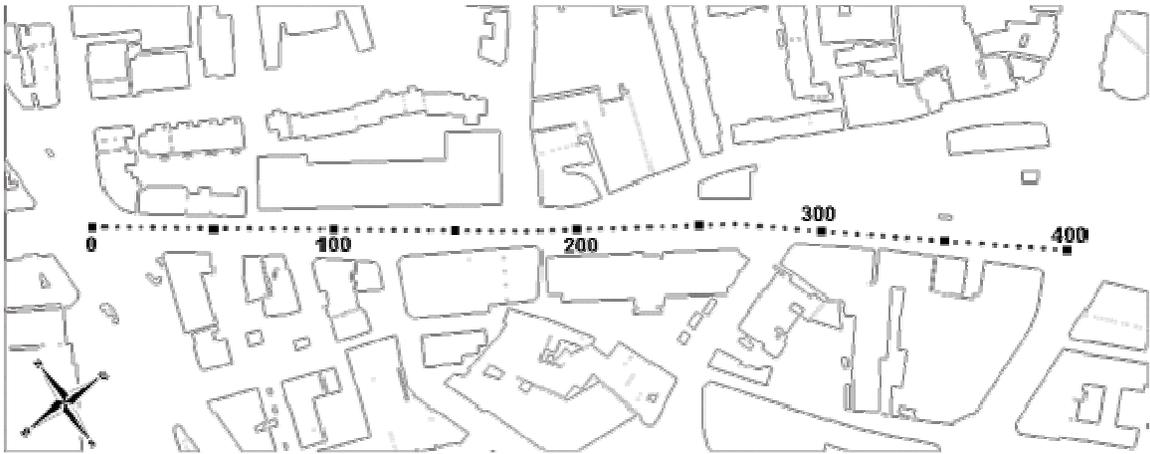


Figure 14: Farringdon Road showing the distances in metres corresponding to the following Figure 15 to Figure 17. The large squares are 50 metres apart, while the small squares are 5 metres apart.

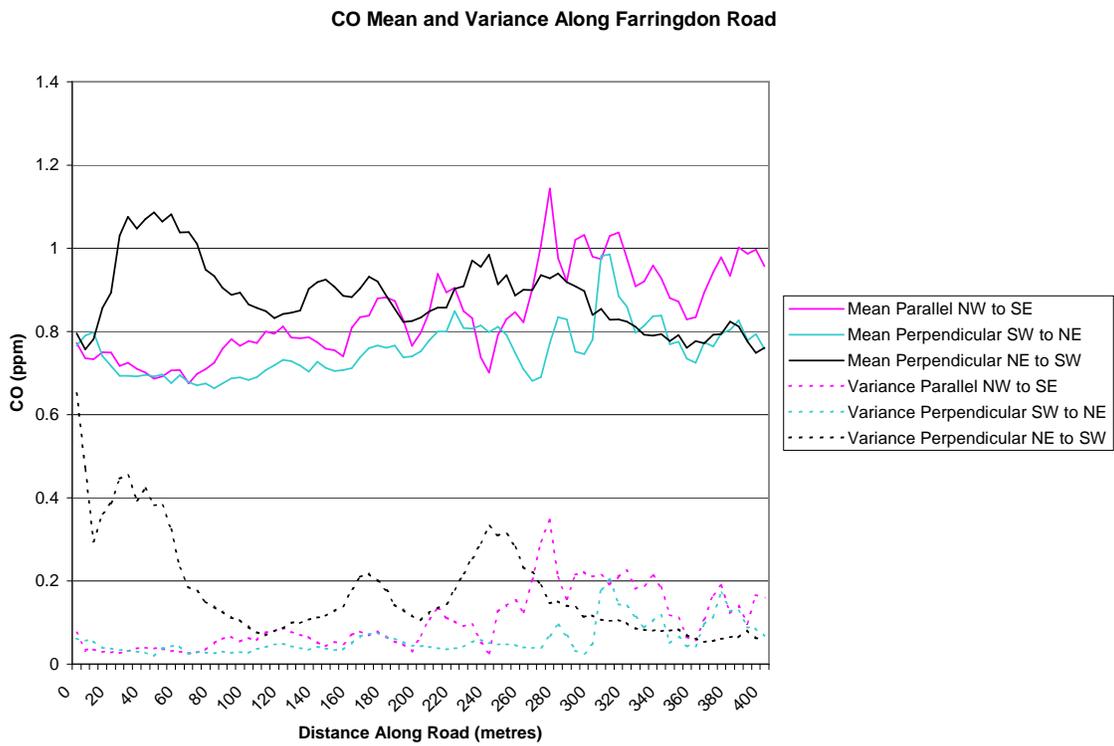


Figure 15: Comparing CO profiles for NW/SE (parallel) winds, SW/NE (perpendicular) winds and NE/SW (perpendicular) winds along Farringdon Road.

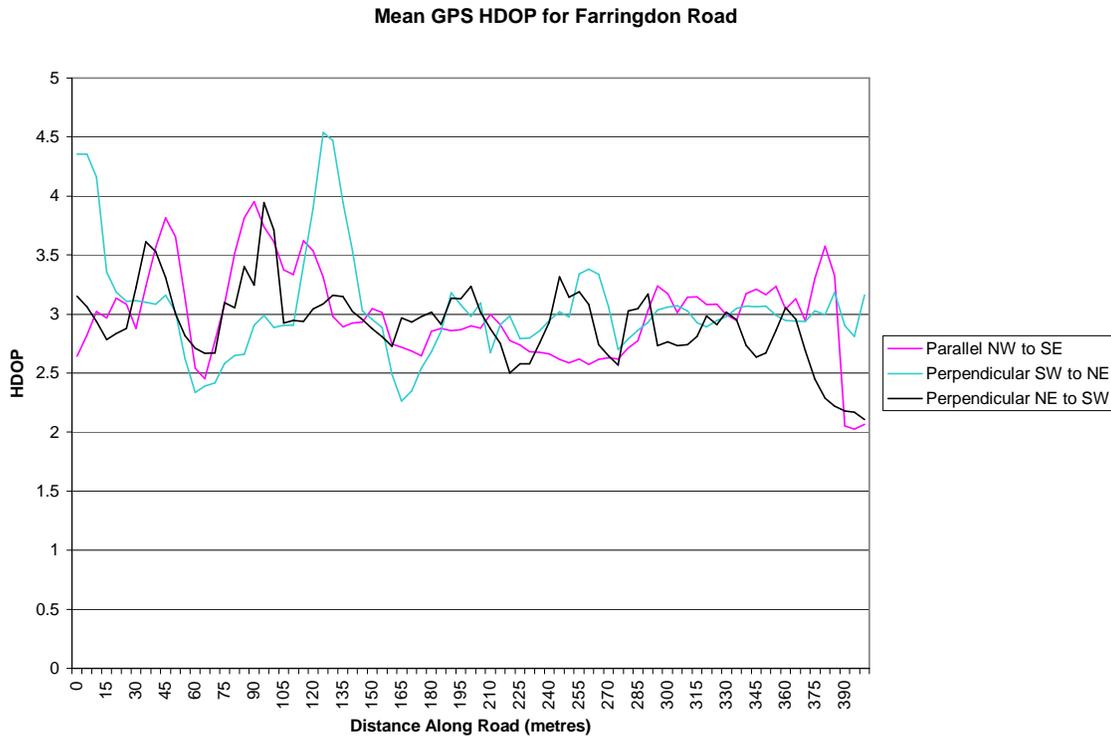


Figure 16: Horizontal Dilution of Precision (standard deviation of position errors) reported by the GPS receiver for the position data used in Figure 15.

Figure 15 shows a strong similarity between the parallel profile and the perpendicular profile in the SW to NE direction with a small shift of less than 30 metres. In this case, a second perpendicular profile in the opposite direction (NE to SW) has been plotted, which does not follow the same trend as the other two. This profile has a distinct peak with a width of 45 metres at the north end of the road (20 to 65 metres), finishing just before the car park. As the adjacent Rosebury Avenue joins Farringdon Road from the northeast, this might explain the peak for this wind direction. Towards the 400 metre end, where the junction with Clerkenwell Road is, the two perpendicular profiles become very similar while the parallel profile finishes with a much higher mean level. This is interesting as it corresponds to the 2.2ppm peak on the Clerkenwell Road North to South profile in Figure 12 earlier. With the wind in similar directions on both roads, there are significant peaks around the same junction.

The GPS accuracy plotted in Figure 16 shows the error to be around 30 metres for most of the length of the road. There is generally good similarity between the GPS data used

for the three profiles in Figure 15, so it is unlikely that any consistent bias has been introduced into the CO data by the GPS. Looking at the CO data in Figure 15 together with the GPS accuracy of Figure 16, a few points do look as though they might have been displaced along the road. In particular, 240 metres on the NW to SE trace and 275 metres on the SW to NE trace look similar and happen to coincide with significant differences in GPS accuracy between these two data sets in Figure 16. The two peaks on the same trace at the 0 and 130 metre positions also suggest that the position might be displaced in the three plotted profiles at these locations.

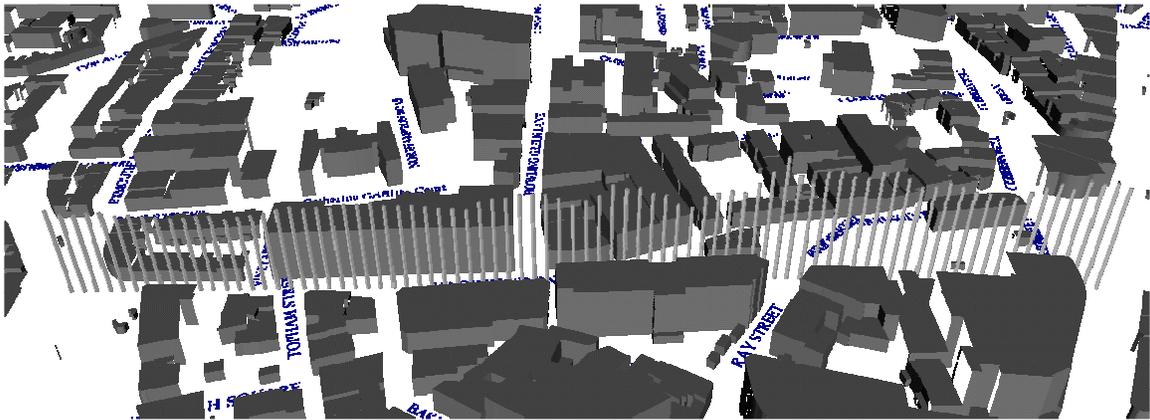


Figure 17: 3D plot of Farringdon Road parallel profile.

6.3 Implications

The link between GPS accuracy, building layout and CO levels has some interesting implications. At junctions where CO tends to build up due to idling traffic, the unobstructed view of the sky often leads to significantly improved GPS accuracy. Even where the GPS at junctions is bad, finding the point where the carrier waits to cross, or where the track angle changes by 90 degrees can often position the crossing accurately. In dense urban canyons, where buildings often block the GPS signal, CO levels are likely to be higher where traffic is stopped, or when re-circulation occurs. In this case it is probably sufficient to know the correct side of the street and an approximate position.

7. Other Observations

7.1. Profiles Across Roads

In addition to sampling data along the length of roads, a significant amount of data was collected while crossing from one side of the road to the other. This turns out to be relatively easy to detect manually if the approximate crossing points on a route are known, as it does not rely on the actual reported GPS position. Points of zero velocity while waiting to cross and 90 degree changes in direction, either from the GPS location track or the reported track angle, are indicative of road crossings. This enables a crossing point to be positioned accurately even when the GPS position is 50m out.

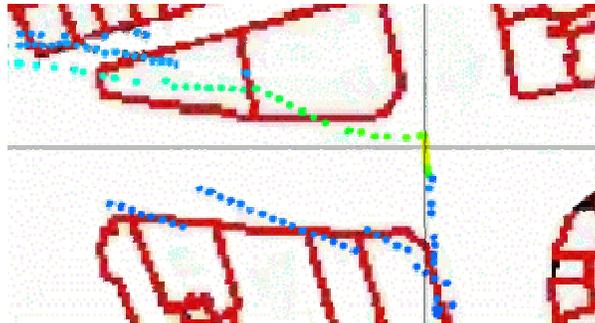


Figure 18: GPS Trace crossing Clerkenwell Road on 31 January 2005, 11:33.

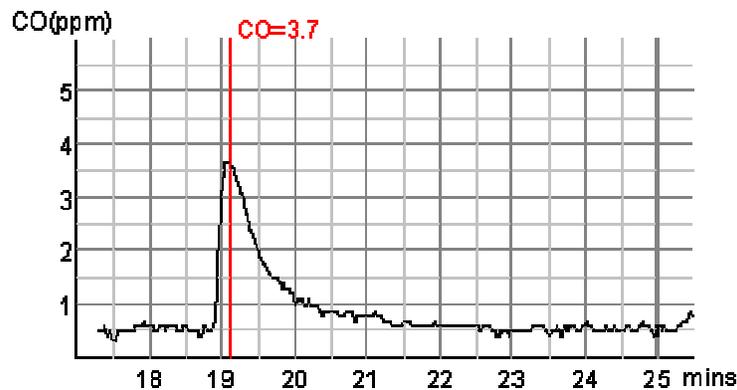


Figure 19: CO Profile while crossing Clerkenwell Road on 31 January 2005, 11:33.

Figure 18 and Figure 19 show Clerkenwell Road being crossed at the St. John Street end. In this case, the GPS accurately identifies the south carriageway as 0.5ppm and the north carriageway as 3.7ppm. This is probably due to the traffic waiting at the lights on the north side while the traffic on the south side is free-flowing. The same crossing was made five minutes later and is labeled at 25 minutes in the diagram, but without any noticeable peak.

On numerous occasions, CO peaks were sampled multiple times at the same location, or on opposite sides of the road. Figure 6 and Figure 20 show examples of CO peaks.

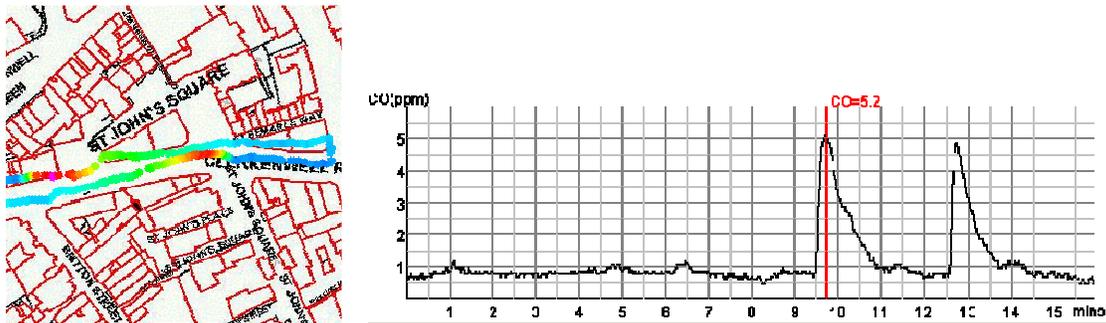


Figure 20: 24 January 2005 15:14 showing identical peaks on opposite sides of the road.

Figure 20 shows a peak of 5.2ppm detected when walking past the petrol station on the north side of Clerkenwell road. Heading towards the junction at the east end, the levels drop. After crossing to the south side and walking west, an almost identical peak is detected 3 minutes later on the opposite side of the road.



Figure 21: GPS accuracy around the double CO peak on 24 January 2005.

Figure 21 shows the GPS errors for this trace as a radius around the reported point. The initial peak outside the petrol station shows errors of 3m to 5m, while the second trace on the other side of the road is around 12m. This is not enough to line the peaks up across the width of the road and suggests that it is either a large area of CO, or one that has drifted in the three minutes between samples. Alternatively, it is possible that there were two independent sources of CO that produced a peak of the same magnitude.

Without any additional information, it is impossible to know what caused these two peaks. No accurate wind measurements were made in the street, but the prevailing wind was reported as coming from 020 degrees at 6 ms^{-1} . This stronger than average wind could be responsible for causing the CO peak to move east along the road and additionally, the second peak is positioned right on the corner leading south into St John's Place. Data showing what happened next would be useful to see whether the peak continues round the corner southwards. It is also significant that the two measured peaks are in an enclosed street canyon. The CO levels are elevated for the length of the road between Clerkenwell Green (west end) and St John's Square (east end) where the road becomes more open.

This set of data is interesting as it suggests that we may be able to track moving peaks using the mobile sensor. Although in this pilot trial we have explored only passive data collection, active collectors might be able to track such peaks.

7.2. Unexpected Sources

During the course of the data collection, there were a number of unexpected CO sources discovered. Table 3 lists some of them with a probable explanation for each.

| <i>CO Level</i> | <i>Location</i> | <i>Possible Explanation</i> |
|-------------------------------|------------------|---|
| 14.9 ppm | Malet Street | Caused by a motorbike starting up close by. Unfortunately it is not known how long this persisted for, but as this location contains a motorbike parking bay, it might be an interesting location to monitor. |
| 4.5 ppm | Exmouth Market | Occasional high peaks at the end of Exmouth Market possibly caused by food being cooked on gas stoves in the open air and by portable generators. |
| 2.4 ppm | Exmouth Market | Lorries making deliveries in this pedestrianised area show distinct tracks when the CO trace is plotted. |
| 3.4 ppm | Indoors | Inside a coffee shop and inside a house. On both occasions the level outside was less than 1.0 ppm. |
| 2.2ppm to 0.9ppm in 30 metres | St. John Square | The CO gradient around St John Square on Clerkenwell road was not unexpected, but is shown particularly well by the tracked sensors. |
| 3.9 ppm | Malet Street | A portable generator on the corner throughout November and December was seen to produce CO peaks at the times when it was running. |
| 2.3 ppm | Guildford Street | Peak observed outside a hospital over a period of months where building work was going on. |

Table 3: Unusually high readings and possible explanations

8. Conclusions

In this paper we investigated the use of GPS tracked sensors to map CO pollution in an urban environment. Specifically, we were interested in whether GPS could be used to place CO samples accurately enough so that expected environmental features could be observed. From the GPS data obtained over three months, it was found that for 95% of the time, the samples could only be positioned to within 50m, with GPS only available for around 50% of the time. Some GPS correction schemes and the use of additional knowledge were investigated in an attempt to improve this accuracy, resulting in data at a much higher resolution for selected features. Road crossings in particular are easy to detect, so peaks at junctions can be positioned with 5m accuracy. Separation and comparison of data on opposite sides of the road is also possible using similar techniques, while differences in the CO profiles along roads due to wind direction can be built up over time.

The mobile sensors were successfully compared with existing fixed sensors and showed how large CO peaks could exist less than 100m from a fixed sensor and not be detected by it. Using the mobile sensors, we were able to take CO readings in areas not normally accessible, for example inside a coffee shop, on the tube or on buses. Some use was made of the equipment's real-time display to locate CO peaks as they were forming and to track them, but the low CO levels during the winter made this almost impossible.

Using a single mobile sensor in an urban area has already provided much useful data that would not have been possible using fixed sensors alone. If the number of sensors can be increased into the hundreds, then high-resolution maps of the changing CO levels in an area of the city can be built. This could provide a valuable resource in the understanding of how pollution dissipates at street level by providing data for selected areas at a finer resolution than the current sensor network.

Acknowledgements

This work was supported by the UK projects Advanced Grid Interfaces for Environmental e-science in the Lab and in the Field (EPSRC Grant GR/R81985/01) and EQUATOR Interdisciplinary Research Collaboration (EPSRC Grant GR/N15986/01).

We would also like to thank Ben Croxford and Gemma Moore and the Vivacity2020 project for providing the fixed sensor data used in section 5 and for their assistance throughout this project.

The vector data used was supplied by the UK Ordnance Survey.

Appendix

Further information on the GPS device used for the experiment can be found on the GARMIN website at: <http://www.garmin.com/products/gps15h/>.

The CO sensors used for the experiment are manufactured by ‘Learian Environmental Systems’ and use the ‘Citicel A3CO’ semiconductor device. Further information can be found on the Learian Website at: <http://www.learian.com/pdfs/icom.pdf>. The datasheet for the A3CO semiconductor sensor can be found at: <http://www.citytech.com/PDF-Datasheets/a3co.pdf>.

The ‘National GPS Network’ is the main source of GPS information in the UK, enabling the download of differential GPS information from fixed sites for post-processing, as well as containing technical information on the GPS system itself. For a good introduction to the different options available, see: http://www.gps.gov.uk/additionalInfo/whatisGPS_01.asp.

References

- Aquilina, Noel, Micallef, Alfred: 2003, ‘Evaluation of the Operational Street Pollution Model Using Data From European Cities’, *Environmental Monitoring and Assessment*. Volume 95, Numbers 1-3, pp. 75-96.
- Boddy, J.W.D., Smalley, R.J., Dixon, N.S., Tate, J.E., Tomlin, A.S.: 2005, ‘The Spatial Variability in Concentrations of a Traffic-Related Pollutant in Two Street Canyons in York, UK—Part I: The influence of background winds’, *Atmospheric Environment*. Volume 39(17), pp. 3147-3161.
- COST Technical Committee on the Environment: 2004, 9th International Conference on Harmonization within Atmospheric Dispersion Modelling for Regulatory Purposes. COST Action 715.
- Croxford, B., Penn, A., Hillier, B.: 1995, ‘Spatial distribution of urban pollution: civilizing urban traffic’. *5th Symposium on Highway and Urban Pollution. Copenhagen, Denmark. Science of The Total Environment*, Volumes 189-190. 28 October 1996, pp. 3-9.

DAPPLE (Dispersion of Air Pollutants and their Penetration into the Local Environment): 2005. URL: <http://www.dapple.org.uk/>.

GARMIN fourth quarter financial report for 2004. Available online: <http://www.garmin.com/aboutGarmin/invRelations/releases/O42004EarningsPressRelease.pdf>.

Ghanem, M. Guo, Y., Hassard, J., Osmond, M., and Richards M.: 2004, 'Discovery Net Sensor Grid for Air Pollution Monitoring'. In Proceedings of the 3rd UK e-Science All-hands Conference AHM 2004, pp. 106-113. Nottingham UK. ISBN 1-904425-21-6. Available online: http://www.discovery-on-the.net/papers_documentation.html#published.

Manning, A. J., Nicholson, K. J., Middleton, D. R., and Rafferty, S. C.: 1998, 'Field Study of Wind and Traffic to Test a Street Canyon Pollution Model', *Environmental Monitoring and Assessment*. Volume 60, Number 3, pp. 283-313.

Milton, R. and Steed, A.: 2005, 'Correcting GPS Readings from a Tracked Mobile Sensor', Thomas Strang, Claudia Linnhoff-Popien, in: *Location- and Context-Awareness: First International Workshop, LoCA 2005*, Springer-Verlag GmbH, May 12-15, 2005, Oberpfaffenhofen (near Munich), Germany, Lecture Notes in Computer Science Volume 3479/2005, pp. 83-94. Springer-Verlag GmbH. Available online: <http://loca2005.context-aware.org/>

Ochieng, W. Y., Polak, J., Noland, R., Park, J.-Y., Zhao, L., Briggs, D., Gulliver, J., Crookell, A., Evans, R., Walker, M., Randolph, W.: 2003, 'Integration of GPS and dead reckoning for real-time vehicle performance and emissions monitoring', *GPS Solutions, Springer-Verlag GmbH*. Volume 6, Number 4. March 2003, pp. 229 – 241.

Steed A., Spinello, S., Croxford, B., Greenhalgh, C.: 2003, 'e-Science in the Streets: Urban Pollution Monitoring'. Proceedings of the UK e-Science AHM Conference, Nottingham, UK. Available online: <http://www.nesc.ac.uk/events/ahm2003/AHMCD/pdf/009.pdf>

Vivacity2020 Project. URL: <http://www.bartlett.ucl.ac.uk/research/space/vivacity-overview.htm>

Weijers, E.P., Khlystov, A.Y., Kos, G. P. A., Erisman, J.W.: 2004, 'Variability of particulate matter concentrations along roads and motorways determined by a moving measurement unit'. *Atmospheric Environment*. Volume 39(19), pp. 2993-3002.