Work supported by the
Swedish Transport and
Communications Research Board

Traffic subflow estimation and bootstrap analysis from filtered counts

Urban Hjorth
Department of Mathematical Statistics
Chalmers University of Technology, 41296 Gothenburg, Sweden

1 Background

The fluctuations of traffic flow counts in one position can often be regarded as time independent or at least uncorrelated random variables and the dependency between such counts at different places gives information about the car flow between the places. This is explicitly or implicitly the basis of much work on dynamic traffic flow estimation. For a background of such estimation we refer to Cremer and Keller (1987), Nihan and Davis (1987), Bell (1991), Davis and Nihan (1991), Cascetta et all (1993), Yu and Davis (1994), Hjorth (1999) and further references given there. However, the use of such estimates in more complicated traffic structures is not always easy. The presence of traffic lights violates the time independence of measured fluctuations near the lights since the traffic streams will then be either highly periodical if the lights are so, or at least autocorrelated if the lights for example are traffic dependent. If there are several lights without a synchronised periodicity we have to consider unsynchronised readings and dependency between both such readings at one place and at different places.

We also have slow variations in the traffic volumes of the whole network due to rush hours and more quiet periods in between, which is largely a deterministic effect, and on top of that other more randomly occurring slow variations affecting measured values in a larger area but not necessarily due to car movements between the measurement places. We think of such variations with time periods between about 15 minutes up to whole days. Such slow variation causes positive covariance or positive regression coefficients between readings at different places for other reasons than vehicles travelling between the places. However, under normal circumstances the fast fluctuations will not be correlated for such reasons. These high frequency variations act as random deviations from the low frequency levels and, when the traffic counts and the different traffic lights are all unsynchronised, the natural cause for dependency between this high frequency part of the counts at different places is vehicles travelling between the places. A possible source of unwanted dependency at higher frequencies are traffic lights with matching periods at two measurement places. We will exclude this situation from our analysis but if this case appears we can either take the light period as our sampling period or identify the traffic light frequency and filter it out. Operationally we can see if an estimation method suffers from the problems above by studying estimated parameters of car movements or other measures of dependency for impossible time differences such as simultaneous counts at two places where it is impossible to count the same vehicle.

Both the estimates of vehicle movements and the uncertainty analysis of such estimates must consider the influence of all these possible effects. In Hjorth (1999), an analytical analysis of uncertainty was given for the situation with uncorrelated fluctuations at the origin. Here we will use a different analysis in order to incorporate all other sources of variation.

2 Data case

As our example we use a measurement program at the Industrigatan in Linköping during four weeks (41, 43, 44, 45) of 1997. The traffic system is shown in Figure 1 where four measurement places are also indicated as circled numbers. Industrigatan has a structure with many crossings and traffic lights and a few roundabout. It is a link from north east to south west which was once the natural road for long distance traffic going through, but is now more like a main link for traffic to industries and commercial areas. Only traffic flow in one direction, from north east to south west was registered so our streams are going from measurement places with lower numbers to those with higher numbers except that no traffic moves from place 3 to place 4 since both are outgoing streams. In each position the passing times of individual vehicles were registered. In the processing of the data we have turned them into counts of the number of passing vehicles in rather short time intervals. The counting intervals are not synchronised with traffic lights and there is no common fix periodicity in the lights, so in our modelling we will regard the lights as having random phase compared to one another and to the counting period.

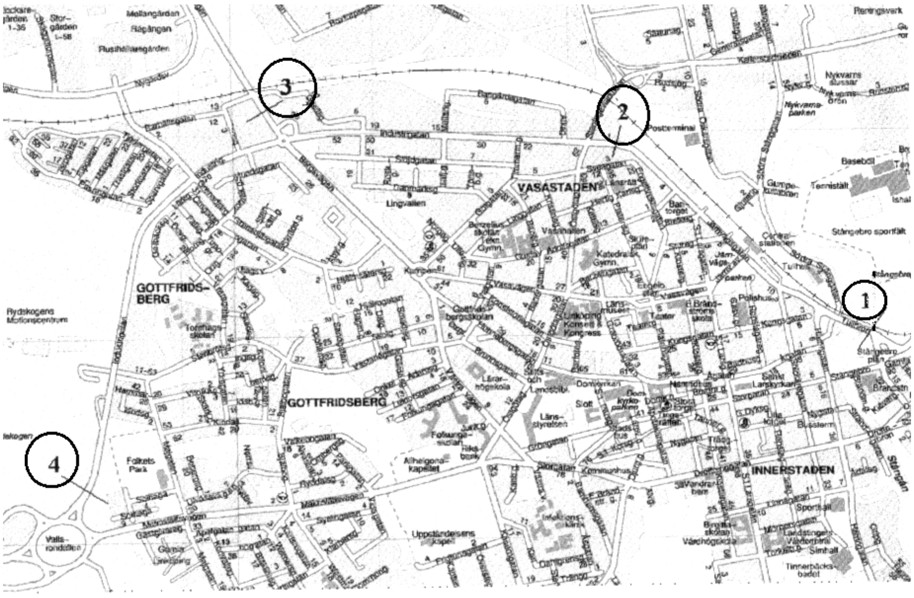

Figure 1: Map of Industrigatan

The measured traffic is seldom congested but since both the route choices, the travelling times, the traffic volumes and the traffic light administration may vary with the hour of the day, we have sectioned our analyses to consider such effects, and have estimated traffic parameters and filters for the different time periods.

3 The model

We will use a stochastic process approach to the series of traffic counts. We may in principle think of these counts as functions of a more basic model where the traffic is regarded as a random (point) process in space and time. Our series of counts becomes a discrete time process which is as multivariate as the number of places where counts are made, but we will only need the two-variate projections of this process since we are comparing the counts at two places at a time. Our model is non-stationary due to the varying traffic volume over the day which is a combination of a deterministic and a stochastic variation both regarded as slowly varying. Using the language of signal processing we want to filter out this low frequency part of our process in order to use the remaining high frequency part as information about parameters describing vehicle movements. This filtering also removes much of the non-stationarity and we are left with a series of high frequency variations which can be treated as approximately stationary over well defined time periods. This supposes that also the traffic flow patterns are stationary over the same periods.

Let $\tilde{X}_i(t),\ t = 1, 2, \ldots,$ denote traffic counts in position i aggregated to some rather short time interval δ and with t enumerating the intervals. Let a sub-stream of the flow past i go on to another position j and use the counts $\tilde{X}_i(.)$ and $\tilde{X}_j(.)$ together in order to study this sub-stream. Define $\tilde{X}_{ij}(t-u, t)$ as the (unobserved) number of vehicles moving from i to j and counted at $t-u$ and t respectively. Also introduce $\mathbf{X}_i(t)$ as a notation for the counts at i up to and including time t and $\mathbf{X}_i$ for all the counts at i without any time restriction.

Consider a time section where the distribution of route selection and travelling speed is approximately stable. Let $p_{ij}(t-u, t)$ denote the probability that a randomly selected vehicle counted at the position i at time $t-u$ will also be counted at j at time t. By a stable route and speed distribution we mean that $p_{ij}(t-u, t)$ are stationary i.e. independent of t during the analysed time period so that each vehicle can be associated with the same probabilities, otherwise our method estimates the average distribution. Let

$$p_{ij}(u) = p_{ij}(t-u, t)$$

denote this stationary value. When there are traffic lights, this assumption includes the assumption about a random phase between the lights and the counting intervals.

We have

$$\tilde{X}_j(t) = \sum_u \tilde{X}_{ij}(t-u, t) + \tilde{Z}_j(t).$$

where $\tilde{Z}_j(t)$ denotes the counted vehicles at j coming from other sources than i. Conditioning on the flows at i we have $E[\tilde{X}_{ij}(t-u, t)|\mathbf{X}_i(t)] = \tilde{X}_i(t-u)p_{ij}(u)$ and we may write

$$\tilde{X}_j(t) = \sum_u \tilde{X}_i(t-u)p_{ij}(u) + \tilde{\epsilon}_j(t) + \tilde{Z}_j(t), \tag{1}$$

where $\tilde{\epsilon}_j$ is the deviation of i, j-traffic from the conditional expectation. Our interest is to make inference about the probabilities $p_{ij}(u)$. However, we can not use this equation as it stands for the estimation. When data from different weeks and a few hours per day are analysed together we can expect slow variations in the traffic volumes which may affect all the traffic flows in similar ways as we discussed in the background section. This may create strong dependency between for example X_i and Z_j creating false large values of the probabilities p if not recognised. For the more rapid variations around the average level, the high frequency components, we expect no such correlation between traffic from different sources (except for congested situations when the sub-flows compete or at synchronised traffic lights which we have excluded). Thus our strategy is to filter out the low frequency components and then estimate $p_{ij}(u)$ from the correlated high frequency part of the data.

4 The kernel regression filter

There are many versions of high pass filters but we will use one which relates to the well-known statistical method of local regression. We first estimate a local mean $m_i(t)$ of $\tilde{X}_i(t)$ by a kernel regression where we for each t minimise

$$Q(\beta, t) = \sum_{u=-u_0}^{u_0} k^2(u)(\tilde{X}_i(t-u) - \beta_0 - u\beta_1 - u^2\beta_2)^2 \tag{2}$$

with respect to $\beta = (\beta_0, \beta_1, \beta_2)'$ and take $\hat{m}_i(t) = \hat{\beta}_0$.

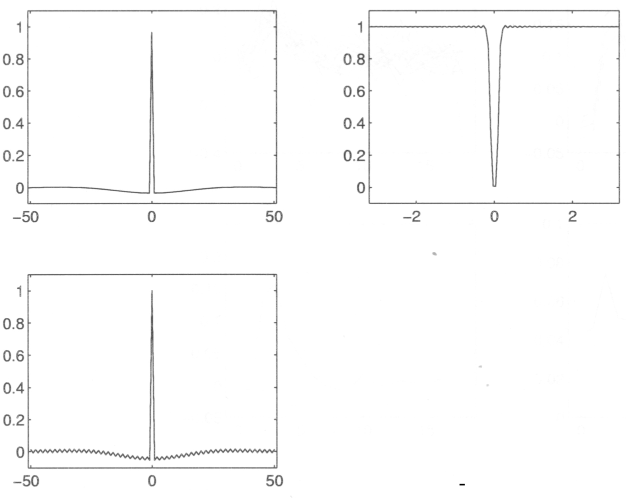

Figure 2: Filter $h(u)$ removing low frequency variation (left) and spectral weight function for the filter (right), and the autocorrelation function of filtered white noise (below).

The kernel is chosen as a truncated Gaussian $k(u) = exp(-u^2/d^2), -u_0 \le u \le u_0$. Defining the vectors $u = (-u_0, \ldots, u_0)'$, $k = k(-u_0), \ldots, k(u_0))'$ and denoting element-wise multiplication as $.*$ we find the solution $\hat{\beta} = (X'X)^{-1}X'(k.*Y)$, where $X = (k \quad k.*u \quad k.*u.*u)$ and $Y = (X_i(t-u_0), \ldots, X_i(t+u_0))'$. Since only $\hat{\beta}_0$ is used as an estimate of $m(t)$ we only need row one and may see the result as $\sum_u \tilde{h}(u)\tilde{X}_i(t-u)$. Subtracting this estimated level from the observation we get the filtered series

$$X_i(t) = \sum_u h(u)\tilde{X}_i(t-u), \tag{3}$$

where $h(0) = 1 - \tilde{h}(0)$ and $h(u) = -\tilde{h}(u),\ u \neq 0$. The filter weights are plotted in Figure 2 for $u_0 = 50$ and $d = u_0/\sqrt{2}$. The corresponding frequency domain function is also given in the figure and illustrates the filtration of low frequencies. Since we sometimes use this filter on uncorrelated data, we also show the auto-correlation function of the output in that situation. The negative parts explain a tendency towards negative estimated covariances around estimated peaks in the cross covariance function which may lead to the same tendency for (unrestricted) estimated probabilities unless all the probabilities are estimated together as we will suggest here. Since no traffic light period was synchronised with our data we were satisfied with the low pass filtering only, otherwise further filtering might be necessary.

5 Relations between filtered series

We apply the filter (3) to all the measured series. This will change the autocorrelation of the series but will not change the probabilities describing the relation between counts at different places as is seen from the following computation. From model (1) the filtered process at j can be written

$$\begin{aligned} X_j(t) &= \sum_u h(u)(\sum_v p_{ij}(v)\tilde{X}_i(t-u-v) + \tilde{\epsilon}_j(t-u) + \tilde{Z}_j(t-u)) \\ &= \sum_v p_{ij}(v)\sum_u h(u)\tilde{X}_i(t-u-v) + \sum_u h(u)\tilde{Z}_j(t-u) + \sum_u h(u)\tilde{\epsilon}_j(t-u) \\ &= \sum p_{ij}(v)X_i(t-v) + Z_j(t) + \epsilon_j(t). \end{aligned} \tag{4}$$

We have here removed the $\tilde{}$-symbol for components filtered by the kernel regression filter, this means also for $Z_j(t)$ and $\epsilon_j(t)$. Under the assumption that the high frequency variations for Z and X_i are uncorrelated, the $p_{ij}(.)$ can be found from the auto and cross covariances, or equivalently by multivariate regression of $X_j(t)$ on the $X_i(t-v)$. We also see from the relation that the probabilities $p_{ij}(v), v = 0, 1, \ldots$ define a filter relating the counts at i as input to the part $X_{ij}(.)$ of the counts at j as output. The filter is linear and causal for each analysed section of data, but this does not exclude non-linearities when regimes of different traffic conditions are compared since we may then estimate significantly different such filters. An overall linear such filter would require the same route selection probabilities and travelling time distribution to be always valid and this is seldom reasonable in true traffic.

6 Estimation

The covariance structure of the filtered counts will be estimated from several pieces of data collected during the same hours at different days. This follows the standard procedures of analysing stationary processes in the time domain via their covariance functions. We have between 12 and 20 useful working days for the different combinations of places and hours. The lower number is for combinations involving place 1 and at least 16 days are available when the other places are combined. As before we consider traffic from i to j. Let the index d be an enumeration of the useful days. Estimate first the cross covariances on every piece of data

$$r_{ij,d}(v) = \sum_t X_{i,d}(t) X_{j,d}(t-v)/T_d, \tag{5}$$

where the sum is over $t = v+1, \ldots, T_d$ and let the auto-covariances be given by the same expression with $j = i$. Since the T_d are approximately equal we take the averages

$$r_{ij}(v) = \sum_{d=1}^{D} r_{ij,d}(v)/D, \tag{6}$$

as our estimated covariance functions for the combined data. The auto-covariances at position 2 and cross covariances between 2 and 3 are given in Figure 3 and the periodic nature in some of the plots shows how the traffic light becomes important during some periods but not during other. In Figure 4 we show covariances estimated on individual days together with bootstrap results for the average covariances as described later.

Define a maximal lag v_0 and the covariance matrix C with elements $c_{kl} = r_{ii}(|l-k|)$ for $0 \leq k, l \leq v_0$ and the vector r_{ij} of cross covariances between place i and j for the same lags. Then

$$\hat{p} = C^{-1} r_{ij} \tag{7}$$

gives the unrestricted vector of estimated probabilities $p_{ij}(v)$ for the corresponding lags. We may also compute the solution restricted to positive probabilities. One such solution is given by minimising $p'Cp - 2p'r_{ij}$ under this restriction. We made a faster solution by forcing the most negative estimate to 0 and recalculating $\hat{p}$ iteratively until all remaining probabilities were positive. This solution saves time in the bootstrap calculations and will often lead to the same result.

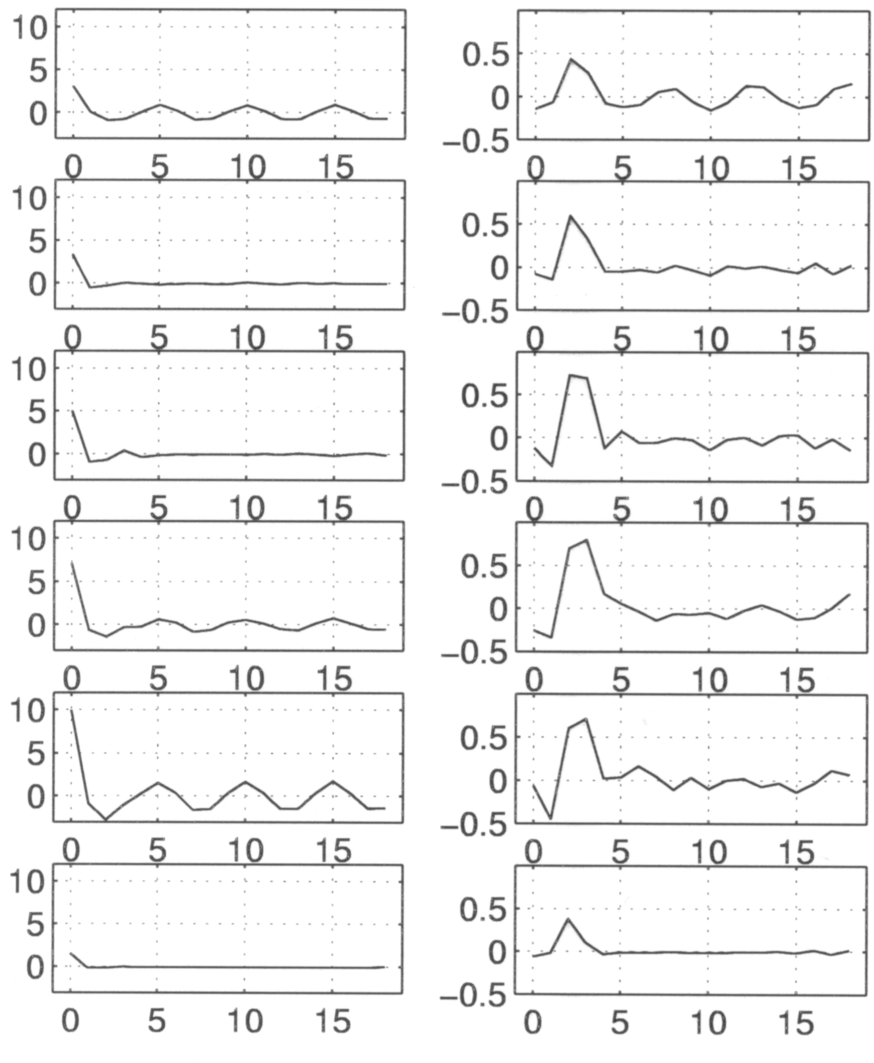

Figure 3: Estimated auto-covariance functions of high-pass filtered counts at place 2 (left) and cross covariances between place 2 and 3 (right) for time periods 1 (top) to 6.

Unrestricted estimates of $p_{ij}(u)$ are shown as a central line in the bootstrap results of Figure 5 and the corresponding restricted estimates are given in Figure 6 together with plotted bootstrap results described in the next section. Notice that the estimates are low for impossible time differences which indicates that we were successful in eliminating other sources for the covariances than the car flow between the places.

7 Bootstrap analysis of estimation uncertainty

Every working day gives about the same information and can be considered independent. Our estimation and subtraction of mean levels by the kernel filter are also made on each day independently of the others. Thus, in the covariance estimates we are averaging over as many independent daily estimates as there are useful days. This makes it very natural to re-sample at a high level and not go down into the series of counts and re-sample there by some of the methods (block, residual or spectral re-sampling) that have been studied for processes. Readers interested in such process re-sampling may consult for example Davison and Hinkley (1997), Efron and Tibshirani (1993), Hjorth (1994), Kunch (1989), Nordgaard (1996) where bootstrap methods are discussed.

If the days $d = 1, 2, \ldots, D$ are available for the covariance estimates and for the estimates of the probabilities p etc., then in every bootstrap run we draw D days at random and independently from the set $d = 1, \ldots, D$. This produces a vector $k^* = (k_1^*, \ldots, k_D^*)'$ of nonnegative integers summing to D and describing how many times each original day is drawn in the bootstrap sample. Our bootstrap covariance estimates are then given as

$$r_{ij}^*(v) = \sum_{d=1}^{D} k_d^* r_{ij,d}(v)/D. \tag{8}$$

Such bootstrapped covariances are illustrated in Figure 4.

From these covariances the probabilities $\hat{p}_{ij}^*(v)$ are estimated just as in the original estimates and by simulating this we get an illustration of the uncertainties for the estimates. We have made 499 such bootstrap simulations of each case. (This number splits the probability distribution into 500 pieces with the same expected probability 0.002.) In Figure 5 we give one example of such bootstrap results for unrestricted estimates and in Figure 6 the corresponding bootstrap results are given for positively restricted estimates. In both illustrations we draw a central line connecting the original estimates and we also draw two outer lines embracing 90% of the bootstrap distribution. In the plots we use

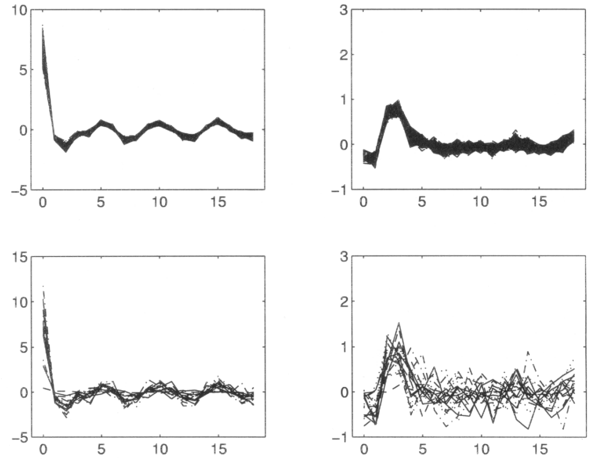

Figure 4: Re-sampled autocovariances at place 2, upper left, and cross covariances between place 2 and 3, upper right. Corresponding daily estimated covariance functions on which re-samples are based (below). Afternoon, time period 4.

the trick of randomly disturbed time values in order to show the data without too much overlap.

In the next section we show several cases of estimates of this type without plotting the individual bootstrap data. Since some of the estimates have skewed distributions and especially the restricted estimates of low probabilities will have an atom at 0 the exact interpretation of the 90% bootstrap interval is difficult to make. We do not state this as a 90% confidence interval although Efron (1993) has done so by one of his methods in other skewed situations. However, the bootstrap gives us the order of uncertainty and for the more symmetrical unrestricted estimates the confidence interval interpretation seems adequate.

8 Results

Each day has been divided into 6 time periods corresponding approximately to the morning rush, before lunch, lunch rush, afternoon, afternoon rush and

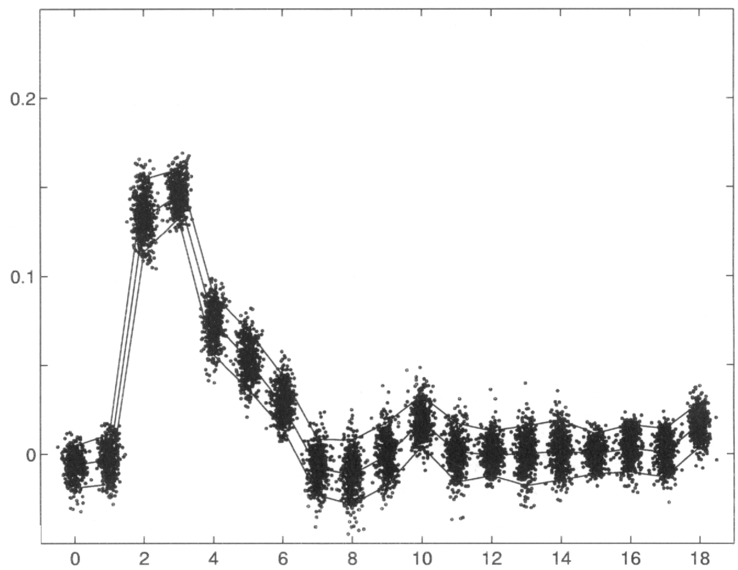

Figure 5: Bootstrapped unrestricted probability estimates with central line showing original estimates and with (pointwise) 90% of the bootstrap results inside outer lines. Traffic from place 2 to 3, time period 4.

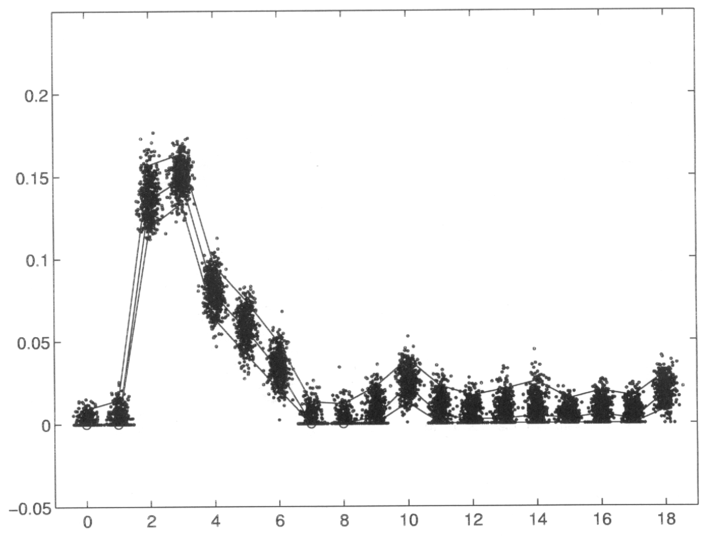

Figure 6: Bootstrapped positively restricted probability estimates, same case as Fig. 5.

evening. More exactly we use

Period 1: 05.30 - 08.30,
Period 2: 08.30 - 11.00,
Period 3: 11.00 - 13.30,
Period 4: 13.30 - 16.00,
Period 5: 16.00 - 18.30,
Period 6: 18.30 - 24.00.

Next, we have studied the traffic from place 1 to the places 2, 3 and 4, and the traffic from 2 to 3 and 4. This gives 5 combinations of origin and destination. (The combination 3 to 4 is not useful since the measured flows at 3 and 4 are both leaving the area in different directions.) We have also limited ourselves to working days mainly because this gives a good basis for the bootstrap analysis but also because this is from some aspects the most interesting days for a link dominated by industry and commercial traffic.

The setting of parameters was the following. As time resolution we selected 20 seconds. This was our compromise between noisiness and resolution of travelling times. The kernel filter (3) was used with $u_0 = 50$ which means that it uses 17 minutes before and after for fitting and subtracting a second degree polynomial but with a strong down-weighting of the outer ends. The kernel filter was run on unbroken data sequences for whole weeks before the data were divided into days and time periods. Thus we have minimised the effects of boundaries, but at the beginning and the end of the weeks a truncation of the filter was made so that as much as possible of the same weights was used. For each studied origin and destination and for each time period we have estimated the probabilities and bootstrapped the estimates as described above. A set of such results for unrestricted estimates are plotted in Figure 7 in a similar way as Figure 5 above with boundaries around the original estimates covering 90% of the simulated bootstrap distribution, but the individual points are not given there because of the small scale.

The estimates related to traffic from place 2 are very clear with a well defined peak as can be seen in Figure 7 for unrestricted estimates and in Figure 8 for positively restricted estimates. The estimates related to traffic from place 1 are more or less dominated by noise. It seems as if the probabilities are very small and that the vast majority of the vehicles has taken other routes than past positions 2,3,4. Nevertheless, for the unconstrained estimate the sum of estimated probabilities from position 1 to 2 (the total route selection probabilities $\sum_u \hat{p}_{12}(u)$) shows a positive result which is supported by another 499 bootstrap simulations. In Table 1 we give these estimates together with the standard deviations from the bootstrap. No significant results were found for the traffic from 1 to 3 or 4 and they are therefore left out.

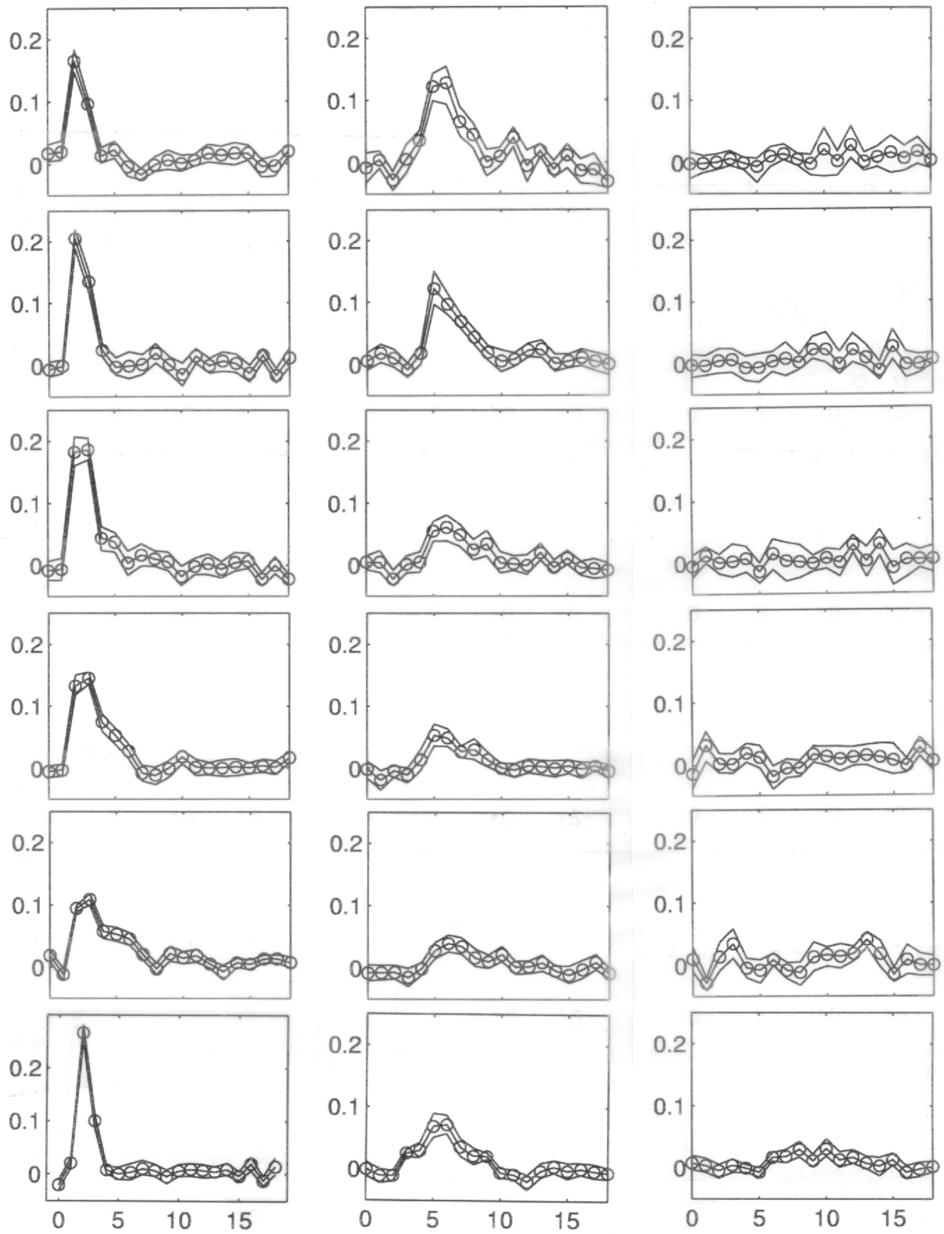

Figure 7: Unrestricted estimated travelling probabilities. Limits covering the central 90% of 499 bootstrap simulations. Traffic from place 2 to 3 left, from 2 to 4 middle and from 1 to 2 right.. Time periods 1 to 6 from above.

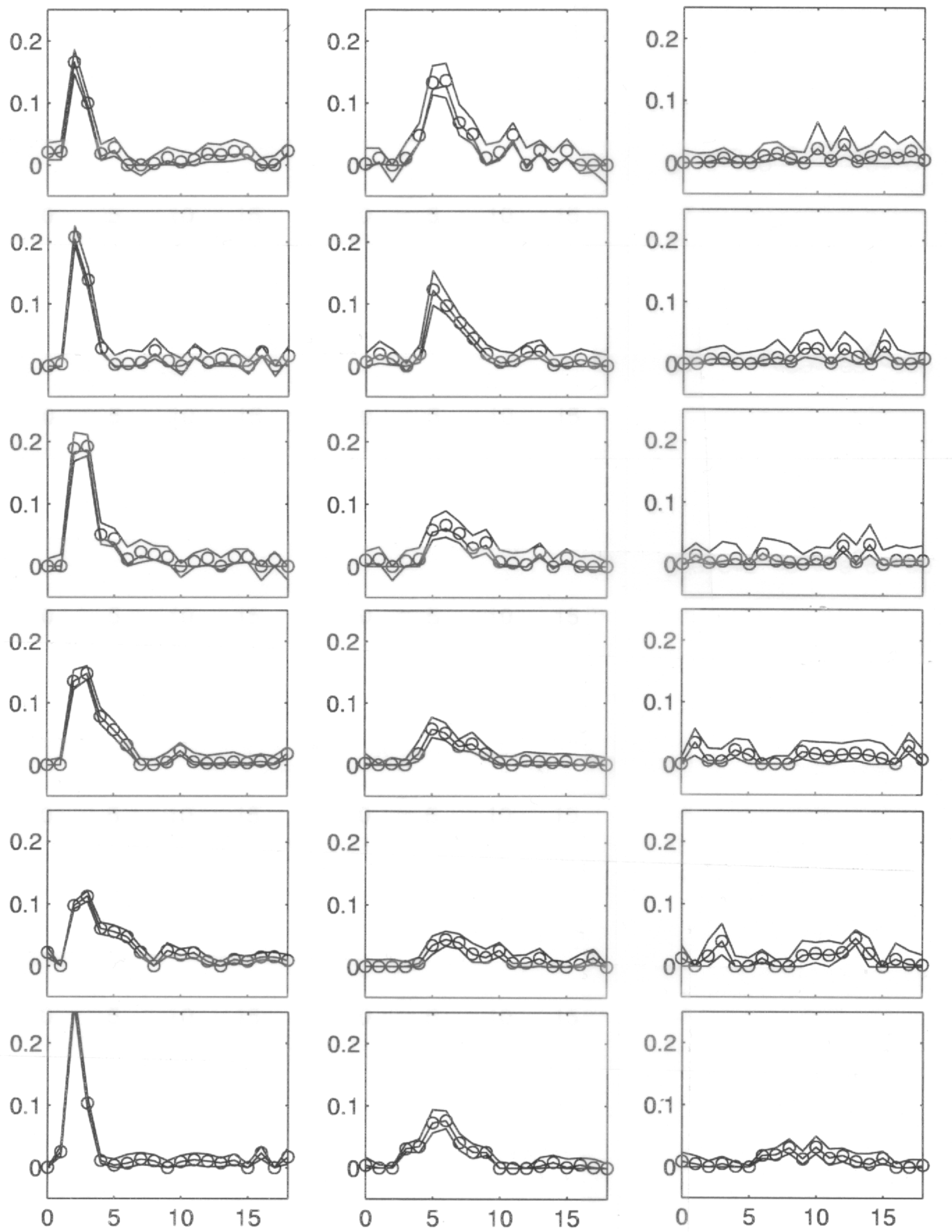

Figure 8: Positively restricted estimated travelling probabilities, as Figure 7.

The constrained estimates are for obvious reasons biased upwards for small probabilities and their sum has also a large positive bias compared to unconstrained estimates (which we believe are nearly unbiased), so there is no, or at least not enough, observed compensation for this in the larger estimated probabilities. A bootstrap study of the true bias is difficult in our high level bootstrapping because there is no obvious true value of the parameter for the empirical distribution.

Table 1. Estimated total route selection probabilities from position 1 to 2, $\sum_u \hat{p}_{12}(u)$, and standard deviations from 499 bootstrap simulations in parenthesis. Unrestricted estimates.

Time period	Estimate	(std)
05.30 – 08.30	.14	(.10)
08.30 – 11.00	.10	(.08)
11.00 – 13.30	.09	(.05)
13.30 – 16.00	.14	(.06)
16.00 – 18.30	.12	(.05)
18.30 – 24.00	.15	(.06)

9 Classical uncertainty analysis

It is possible to produce a set of iid estimates of $p_{ij}(u)$ if we skip the averaging of the daily covariance estimates and instead use the estimates of type (7) on covariances (5) estimated from each day by itself. Since we have 20 useful days of data between points 2 and 3 this gives 20 estimated $\hat{p}_{23}$-vectors plotted in Figure 9. Averaging these estimates gives another unique estimated $\hat{p}$-vector. We expect the averaging of covariance estimates (the bootstrapped unrestricted version) to be more stable than the averaging of 20 $\hat{p}$-vectors, since the inversion of the matrix C could be crucial, but since the data are rich, both versions turn out to be useful here. Individual 90% confidence intervals were produced by standard methods for the normal distribution, i.e. the estimated average $\hat{p} \pm as/\sqrt{D}$ where a is the appropriate value from the Student's distribution with $D-1$ degrees of freedom (1.73 if $D = 20$) and $s = s(u)$ is the estimated standard deviation for the estimates at time u. The results of this computation for time period 4 and places 2,3 are shown in Figure 9 and the confidence intervals given in one of the sub-plots are only marginally wider than the corresponding bootstrap intervals, so the two methods give in fact very similar results.

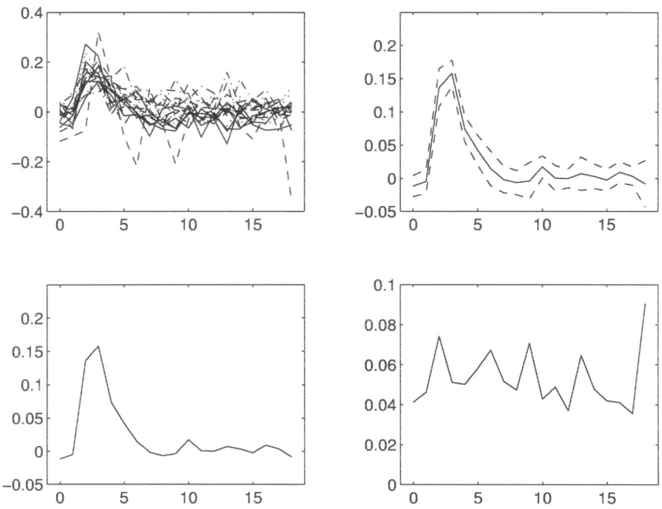

Figure 9: Daily estimated $\hat{p}$ and the average (left), 90% confidence intervals and estimated standard deviations of the daily estimates (right). Traffic from 2 to 3 in time period 4. "Classical" analysis.

10 Discussion

In the presence of traffic lights the process of traffic counts is like an unsynchronised random process with a strong periodicity and considering the phase as random it can be modelled as a stationary process during limited periods of time. This is of course an approximation since the traffic intensity is actually varying all day, but what really matters is that the route selection probabilities and the travelling time distribution will stay constant. Since the dependency and the structure of the data is fairly complex the uncertainty analysis by classical methods using individual counts would be difficult and based on approximations. However, the statistical similarity between different traffic days gives us a sample of highly multivariate data sets, one for each day, which are both equally distributed and independent by natural assumptions. Thus the high level bootstrap analysis comes as a natural tool. Of course this is also in a sense an approximate method, but it does not use any simplified assumptions about the inner structure of the data from different traffic days, only the iid properties of the whole daily data sets. As an example of another and different application of bootstrap methods to dependent processes in traffic we refer to the block re-sampling used in Bergendorff (1999). As it turned out, a classical high level analysis leads to almost the same result as our bootstrap method. We could not be sure of this in advance, since the estimate is non-linear and uses matrix inversion and the uncertainties of estimated covariances seem larger than just small disturbances. In retrospect, however, this is a positive sign for both methods.

The modelling of counts and filters acting on the counts has many similarities with classical signal processing in other areas such as tele communication. Not only the high pass filter used as a tool in the estimation but also all filters described by the probabilities $p_{ij}(u)$ relating upstream data with subflows downstream are of interest. Identifying filters for different traffic elements is one way to describe the system with much appeal to communication engineers and may lead to some new useful understanding. This work is just one step in such a direction.

Acknowledgement This work depends heavily on cooperation and exchange of data with the Swedish Road and Traffic Research Institute in Linköping. I thank in particular Mats Wiklund and Hans-Åke Cedersund for their time and interest. I also thank the Swedish Transport and Communications Research Board (KFB) for their financial support.

References

Bell, M.G.H. (1991). The real time estimation of origin-destination flows in the presence of platoon dispersion. *Transportation Research* 25B, 115-125.

Yu, P. and Davis, G.A. (1994). Estimating freeway origin-destination pattern using automation traffic counts. *Transportation Research Record* 1457, 127-133.

Bergendorff, P. (1999). Extracting information on traffic patterns from classified traffic counts. Linköping Studies in Science and Technology, Thesis 762, Mathematical statistics, Linköping University, 58183 Linköping, Sweden.

Cascetta, E., Inaudi, D. and Marquis, G. (1993). Dynamic estimators of origin-destination matrices using traffic counts. *Transportation Science* 27(4), 363-373.

Cremer, M. and Keller, H. (1987). A new class of dynamic methods for the identification of origin-destination flows. *Transportation Research* 21B, 117-132.

Davison, A.C. and Hinkley, D.V. (1997). Bootstrap Methods and their Applications. Cambridge University Press.

Davis, G.A. and Nihan, N.L. (1991). Stochastic process approach to the estimation of origin-destination parameters from time series of traffic counts. *Transportation Research Record* 1328, 36-42.

Efron, B. and Tibshirani, R.J. (1993). *An Introduction to the Bootstrap.* Chapman and Hall.

Hjorth, U. (1994). *Computer Intensive Statistical Methods – Validation, Model Selection, Bootstrap.* Chapman and Hall.

Hjorth, U. (1999). The inherent precision of regression estimated route probabilities. Transp. Res. B 33, 593–607.

Hjorth, U. (1999 b). Flow estimation and accident modelling. Studies in applied probability and statistics 1999:5, Dep. of Mathematical statistics, Chalmers/GU. 41296 Göteborg, Sweden.

Kunch, H.R. (1989). The jackknife and bootstrap for general stationary observations. Annals of Statistics 23, 1217 - 1241.

Nihan, N.L. and Davies, G.A. (1987). Recursive estimation of origin-destination matrices from input/output counts. *Transportation Research* 21B, 149-163.

Nordgaard, A. (1996). Computer-intensive methods for dependent observations. Dissertation, Mathematical statistics, Linköping University, Sweden.