Estimating Vehicle Speed from Traffic Count and Occupancy Data

Martin L. Hazelton
University of Western Australia

Abstract: Automatic vehicle detectors are now common on road systems across the world. Many of these detectors are based on single inductive loops, from which data on traffic volumes (i.e. vehicle counts) and occupancy (i.e. proportion of time during which the loop is occupied) are available for 20 or 30 second observational periods. However, for the purposes of traffic management it is frequently useful to have data on (mean) vehicle speeds, but this is not directly available from single loop detectors. While detector occupancy is related in a simple fashion to vehicle speed and length, the latter variable is not measured on the vehicles that pass.

In this paper a new method for speed estimation from traffic count and occupancy data is proposed. By assuming a simple random walk model for successive vehicle speeds an MCMC approach to speed estimation can be applied, in which missing vehicle lengths are sampled from an exogenous data set. Unlike earlier estimation methods, measurement error in occupancy data is explicitly modelled. The proposed methodology is applied to traffic flow data from Interstate 5 near Seattle, during a weekday morning. The efficacy of the estimation scheme is examined by comparing the estimates with independently collected vehicle speed data. The results are encouraging.

Key words: Bayesian inference, inductance loop, Metropolis-Hastings algorithm, measurement error, missing data.

1. Introduction

Road traffic management is becoming increasingly reliant on the availability of real-time traffic flow data. In Melbourne, for example, SCATS (Sydney Co-ordinated Adaptive Traffic System) makes use of such data to optimize signals over the road network. Similar schemes operate in many other cities throughout the world. Data collection is typically done by inductive loop vehicle detectors embedded in (or lying on) roadways. A single detector loop provides information on traffic volumes (i.e. vehicle counts over an observational period) and occupancy (the proportion of an observational period during which the loop senses

the presence of a vehicle). However, an increasing number of ITS (intelligent transport system) initiatives aimed at congestion relief require estimates of (average) vehicle speeds, often as a precursor to estimating travel times. A prime example from the United States is SWIFT (Seattle Wide Area Information for Travelers). Accurate estimation of vehicle speeds is therefore a significant problem for traffic engineers. Furthermore, given the widespread use (and relatively low cost) of single loop inductive detectors, methods for obtaining speeds from the count-occupancy data produced by these detectors are of particular interest. See Persuad and Hall (1989) and Dailey (1992, 1999) for related comments.

Data from loop detectors is (typically) not available at an individual vehicle level, but rather aggregated over pre-set time intervals (often of 20 or 30 second durations). Given information on vehicle lengths, one can hope to compute an estimated speed based on the total length of vehicle passing the loop, and the time it takes to do so (calculated from the occupancy). While single loop detectors do not provide vehicle length data, large exogenous data sets of vehicle lengths are available. Most estimation procedures proposed to date have been based on an application of first order method of moments at each time interval. The recent work of Dailey (1999) improved on this relatively crude methodology by incorporating a second order correction, and applying a Kalman filter to smooth mean speeds from consecutive intervals.

In this paper we take a novel approach to the problem of speed estimation, concentrating on modelling at the level of individual vehicles. This avoids bias due to aggregation which is inherent in earlier techniques. We also allow for measurement error in the recorded occupancies, a feature of the data that is widely recognised (see Coifman, 1999, for example) but has been ignored in published work on speed estimation. Modelling at a disaggregate level means that we are confronted with a great deal of missing data – namely the speeds and lengths of each individual vehicle. Markov chain Monte Carlo (MCMC) methods provide powerful tools for estimation in the presence of missing data. See Diebolt and Ip (1996), for example. We use these techniques to obtain Bayesian estimates of the mean vehicle speed over each time interval.

The paper is structured as follows. In the next section we describe our basic model. Details of our MCMC algorithm are given in section 3. Issues regarding the block structure of this algorithm and its relation to the mixing rate (see Gamerman, 1997) are discussed. In section 4 our methodology is used to estimate speeds from loop traffic count-occupancy data collected between 4:00am and 9:30am on a weekday from Interstate 5 in Seattle. This data is available over the World Wide Web from `http://www.its.washington.edu/tdad`, thanks to the Traffic Data Acquisition and Distribution (TDAD) project managed by the Intelligent Traffic Systems group at the University of Washington. The speeds es-

timated from this analysis are compared with independent observations obtained from a speed trap positioned close to the detector loop under consideration. This comparison produces encouraging results.

2. Modelling

An inductive loop vehicle detector incorporates an insulated electric wire through which an alternating current is driven to set up an electromagnetic field. Any (metallic) vehicle which passes through this field will decrease the inductance of the loop, and cause the detector to register the vehicle's presence. While the shape and size of the loops can vary, many are squares of side approximately 6 ft. (and we shall assume this shape henceforth). A loop detector registers that it is occupied from the moment that the front end of a vehicle enters its region of sensitivity on one side of the square to the time when the rear of the vehicle leaves that region on the other side. The region of sensitivity may extend a little beyond the physical boundaries of the loop. The 6 ft. loops used in the Seattle area have a sensitivity of about 8 ft. according to documentation from the Washington State Department of Transport, although this range varies a little from detector to detector. See Kell *et al.* (1990) for further details on this type of vehicle detector.

Consider a vehicle of length λ ft. entering a loop with sensitivity range of $\lambda_0 = 8$ ft. The time, x, (in seconds) for which the vehicle occupies the loop is also the time required for the vehicle to travel a distance $l = \lambda + \lambda_0$, the *effective vehicle length*. Hence the speed of the vehicle (in feet per second) is given by $s = lx^{-1}$. Now, data is collected at an aggregate level over a number of consecutive intervals each of length δ (20 seconds for the Seattle data) so we shall extend the notation so that s_{ij} and l_{ij} are the speed and effective length respectively of the jth vehicle during the ith interval (where $i = 1, \ldots, m$) . Let n_i denote the vehicle count during the ith interval, and let x_i be the length of time in that period during which the loop was occupied (i.e. the sum of the individual x's from each vehicle). (Note that traffic engineers typically refer to occupancy as a proportion, $x_i \delta^{-1}$.) Then the individual vehicle effective lengths and speeds are related to the (exact) occupancies through

$$x_i = \sum_{j=1}^{n_i} \frac{l_{ij}}{s_{ij}}. \tag{2.1}$$

The observed data are pairs (n_i, y_i) $(i = 1, \ldots, m)$, where y_i is the occupancy for the ith interval as recorded by the vehicle detector. In practice x_i and y_i will not be identical due to measurement error. (This measurement error is due to a number of factors, including variations in the profile of metal distribution across

the bodies of different vehicles.) For a modern, well calibrated inductive vehicle loop, the magnitude of this error will typically be around 5%; see Coifman (1999) and the references therein, for example. Denoting the measured occupancy in the ith interval by y_i, we model the measurement error in a multiplicative fashion by

$$y_i = x_i(1 + z_i) \tag{2.2}$$

where $z_1, \ldots, z_m$ are independent and normally distributed, each with mean zero and variance σ_z^2.

Almost all methods in the literature for estimating the mean speed $\bar{s}_i$ during the ith interval have used a first order method of moments approach. See Kurkjian *et al.* (1980) and Leutzbach (1988), for example. Following this methodology provides an estimator $\check{s}_i = \hat{\mu}_l y_i^{-1}$, where $\hat{\mu}_l$ is an estimator of mean effective vehicle length, obtained from some exogenous data set. A disadvantage of $\check{s}_i$ is that it is biased, a result of interchange from harmonic to arithmetic mean in the derivation of this estimator from equation (2.1) (with exact and observed occupancies interchanged). A second drawback is that $\check{s}_i$ makes no use of the temporal dependence that one would expect to exist in the sequence $\bar{s}_1, \bar{s}_2, \ldots, \bar{s}_m$. We take account of this characteristic of the traffic flow processes by assuming that the individual vehicle speeds form a random walk. Specifically,

$$s_{ij} = s_{i(j-1)} + \epsilon_{ij} \qquad (i = 1, \ldots, m \quad j = 1, \ldots, n_i \quad i, j > 1) \tag{2.3}$$

where $\{\epsilon_{ij}\}$ are independent $N(0, \sigma_\epsilon^2)$ random variables, and we use the convention $s_{(i+1)0} = s_{in_i}$. This model was chosen in light of the lack of stationarity in traffic speeds during time periods containing a congested 'peak hour' (as is the case with our Seattle data).

Having developed a basic model for vehicle speeds we now turn our attention to the missing effective vehicle lengths in equation (2.1). We assume speed-length independence so that *a priori* $\{l_{ij}\}$ is a random sample from a distribution f_l. The distribution f_l can be estimated from an exogenous data set of vehicle lengths, $\lambda_1, \ldots, \lambda_\nu$ (to which λ_0 must be added to obtain effective vehicle lengths). In some applications a random sample of vehicle lengths for the road in question will be available. This is the case for Seattle Interstate 5, where we have such a data set of size $\nu = 17528$. On other occasions length data from traffic on a comparable type of road will have to suffice. The distribution of the Seattle length data is displayed in Figure 1 using a kernel density estimate (see Wand and Jones, 1995). We actually plot the density of the natural logarithms of the effective vehicle lengths, since important structure is most easily seen on the log scale. The three clear modes in the distribution suggest that the distribution of the lengths is a mixture with three components. These components may be naturally interpreted (in order of increasing vehicle length) as corresponding to cars, vans and small lorries, and large lorries and road trains respectively.
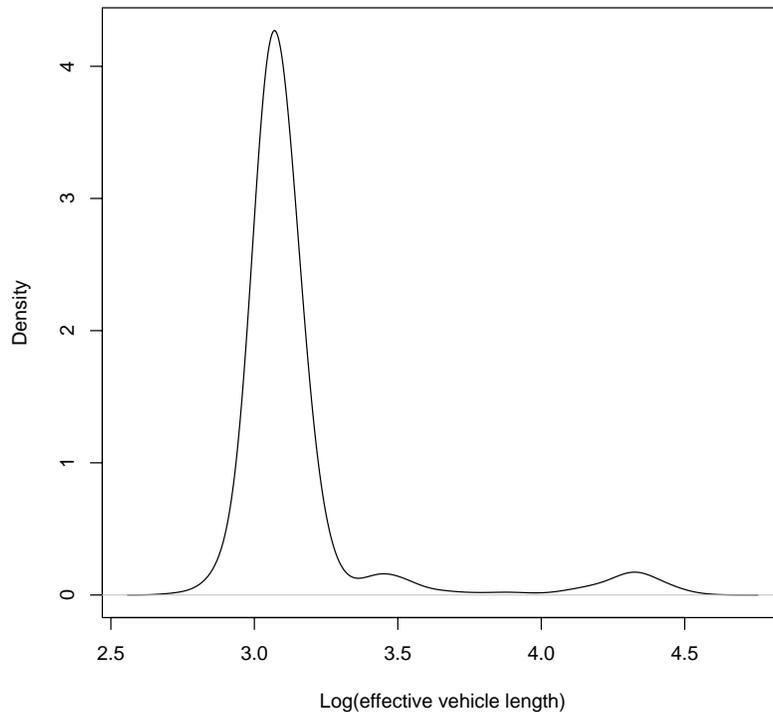
Figure 1. Kernel density estimate of natural logarithms of effective vehicle lengths (in feet) from Seattle Interstate 5. Sample size $\nu = 17528$, bandwidth $h = 0.050$.

To complete the specification of our model (within a Bayesian framework) it remains to select prior distributions for s_{11} (the speed of the first vehicle observed), σ_ϵ^2 and σ_z^2. For the Seattle data we employed a uniform prior on the interval $(0, 150)$ feet per second for s_{11}, although a diffuse normal prior would have been equally appropriate. Gamma priors were employed for the precisions $\tau_\epsilon = \sigma_\epsilon^{-2}$ and $\tau_z = \sigma_z^{-2}$ throughout our work. For τ_ϵ we considered both vague Gamma$(0.001, 0.001)$ as well as informative priors such as Gamma$(0.11, 1)$, which gives most weight to values of σ_ϵ close to 3 feet per second (i.e., about 2 miles per hour) reflecting a belief that speeds of consecutive vehicles are unlikely to differ by much more than 5 miles per hour. For τ_z we considered only informative priors. The observed data provide very little information about σ_z, so prior information plays an important role with this parameter. Our preferred prior for τ_z was Gamma$(400, 1)$ which gives most weight to values of σ_z close to 0.05, in line with the '5% measurement error' mentioned above. We also looked at priors

corresponding to values of σ_z close to 0.01 and 0.1. Sensitivity of speed estimates to changes in the priors is discussed in section 4.

3. MCMC Inference

In this section we discuss the exploration of the posterior distribution by Markov chain Monte Carlo methods (see Gamerman, 1997, for example). Derivation of the full set of conditional posterior distributions in closed form does not seem possible. We therefore use a Metropolis-Hastings algorithm for sampling from the posterior distribution. (See Chib and Greenberg, 1995, for an overview of the Metropolis-Hasting methodology.) This approach produces a chain of simulations $\{(\boldsymbol{s}^{(t)}, \boldsymbol{l}^{(t)}, \boldsymbol{z}^{(t)}, \sigma_\epsilon^{(t)}, \sigma_z^{(t)}) : t = 1, 2, \ldots\}$ (where $\boldsymbol{s}$ and $\boldsymbol{l}$ denote the vectors of vehicle speeds and corresponding effective lengths ordered chronologically, and $\boldsymbol{z}$ the vector of multiplicative measurement errors). The basic structure of our algorithm is as follows.

1. Initialize: $t = 1$, $\boldsymbol{s}^{(t)} = \boldsymbol{s}_0$, $\boldsymbol{l}^{(t)} = \boldsymbol{l}_0$, $\boldsymbol{z}^{(t)} = \boldsymbol{z}_0$, $\sigma_\epsilon^{(t)} = \sigma_{\epsilon 0}$ and $\sigma_z^{(t)} = \sigma_{z0}$.

2. For i in $\{1, 2, \ldots, m\}$:

   (a) Generate vector of candidate vehicle speeds in interval i:

   $$\boldsymbol{s}_i^\dagger \sim [\boldsymbol{s}_i \,|\, \boldsymbol{s}_{i-}^{(t+1)}, \boldsymbol{s}_{i+}^{(t)}]$$

   Here $\boldsymbol{s}_{i-}$ is the vector of vehicle speeds in intervals before i, and $\boldsymbol{s}_{i+}$ is the vector of vehicle speeds in intervals after i. Also $[a\,|\,b]$ denotes the conditional distribution of a given b.

   (b) Generate vector of corresponding candidate effective lengths for all vehicles in interval i:

   $$l_{ij}^\dagger \sim \hat{f}_l \qquad j = 1, \ldots, n_i$$

   where $\hat{f}_l$ represents the empirical distribution of effective vehicle lengths (obtained with the use of exogenous data).

   (c) Define

   $$z_i^\dagger = y_i \left\{ \sum_{j=1}^{n_i} \frac{l_{ij}^\dagger}{s_{ij}^\dagger} \right\}^{-1} - 1$$

   following equations (2.1) and (2.2). Writing $\boldsymbol{l}_i^\dagger$ for the vector of candidate lengths in interval i, accept $(\boldsymbol{s}_i^\dagger, \boldsymbol{l}_i^\dagger, z_i^\dagger)$ with probability

   $$\beta = \min \left(1, \exp \left[\frac{(z_i^{(t)})^2 - (z_i^\dagger)^2}{2\sigma_z^2} \right] \right).$$

(d) If candidate accepted then $s_i^{(t+1)} = s_i^\dagger$, $l_i^{(t+1)} = l_i^\dagger$ and $z_i^{(t+1)} = z_i^\dagger$. Otherwise $s_i^{(t+1)} = s_i^{(t)}$, $l_i^{(t+1)} = l_i^{(t)}$ and $z_i^{(t+1)} = z_i^{(t)}$.

3. Generate

$$\tau_\epsilon^{(t+1)} \sim [\tau_\epsilon | \, s^{(t+1)}, \, l^{(t+1)}, \, z^{(t+1)}]$$

and let $\sigma_\epsilon^{(t+1)} = \sqrt{\tau_\epsilon^{-(t+1)}}$. Similarly, generate

$$\tau_z^{(t+1)} \sim [\tau_z | \, s^{(t+1)}, \, l^{(t+1)}, \, z^{(t+1)}]$$

and let $\sigma_z^{(t+1)} = \sqrt{\tau_z^{-(t+1)}}$

4. $t \leftarrow t + 1$. Go to step 2.

Remark 1: At initialization, all elements of l_0 were set equal to the mean of the (estimated) effective length distribution. The vehicle speeds for the ith interval were initialized at $\breve{s}_i$, the first order method of moments estimator described above. Since these initial speeds and lengths satisfy (2.1) with $y_i = x_i$, so all elements of z_0 were set to zero. The parameter σ_ϵ was initially set at $\sigma_{\epsilon 0} = 3$. The parameter σ_z was initialized at the mean of its marginal prior distribution.

Remark 2: Computation of the conditional distributions at stage 2(a) of the algorithm is straightforward because of the random walk model for vehicle speeds. Details are provided in the appendix, where we also describe the conditional distribution of τ_ϵ and τ_z.

Remark 3: The algorithm is structured so that parameters from each interval are blocked together for updating. It is possible to update each single (scalar) parameter in turn, or to update parameters for each single vehicle in turn. However, while these approaches would provide a higher acceptance rate than obtained using the algorithm above, the resulting Markov chains would mix very poorly. This is due to the fact that the length distribution is essentially a mixture with three components. Transitions of a given vehicle length between components (and in particular, from the 'road train' component to the other components) are highly unlikely because they give rise to substantial changes in x_i and hence relatively large (and thus improbable) values for z_i. By defining blocks in terms of speeds and lengths over an entire interval, our algorithm mixes at a reasonable rate without the acceptance rate becoming too small.

In applying this algorithm to the Seattle data the chain was run for 100000 iterations. Simulations from the first 20000 iterations were discarded as the burn-in period. Convergence of the chain after this burn-in was confirmed by employing
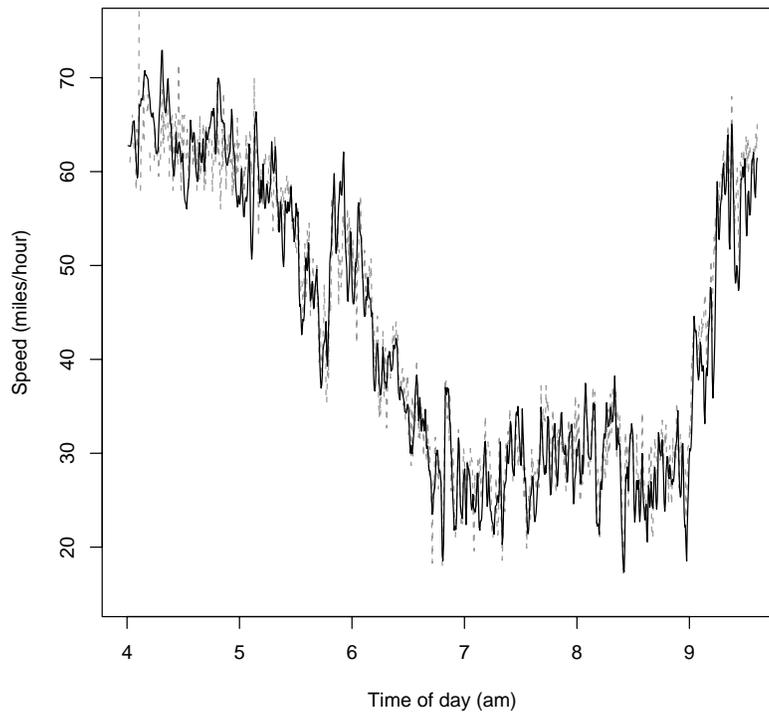
Figure 2. Estimates of interval mean speeds for traffic on Interstate 5, near Seattle. The grey dashed line depicts mean speeds obtained from a speed trap, while the solid black line depicts mean speeds estimated from count-occupancy data from a single vehicle detector.

Geweke's (1992) methodology, and (utilizing results obtained from a parallel simulation) by Gelman and Rubin's (1992) method. The acceptance rate (at stage 2c of the algorithm) was about 7% in equilibrium. Naturally a higher rate would have been preferable, but it appears to be impossible to obtain such an improvement without a substantial slowing of the mixing rate of the algorithm (cf. Remark 3 above). A thinning interval of 10 was applied to the simulation output for the sake of parsimony in computer storage (a matter of some importance since 1000 mean vehicle speeds had to be stored at each monitored iteration). All computation was done using the statistics language R (Gentleman and Ihaka, 1996) running on a 1900MHz PC operating under Linux.
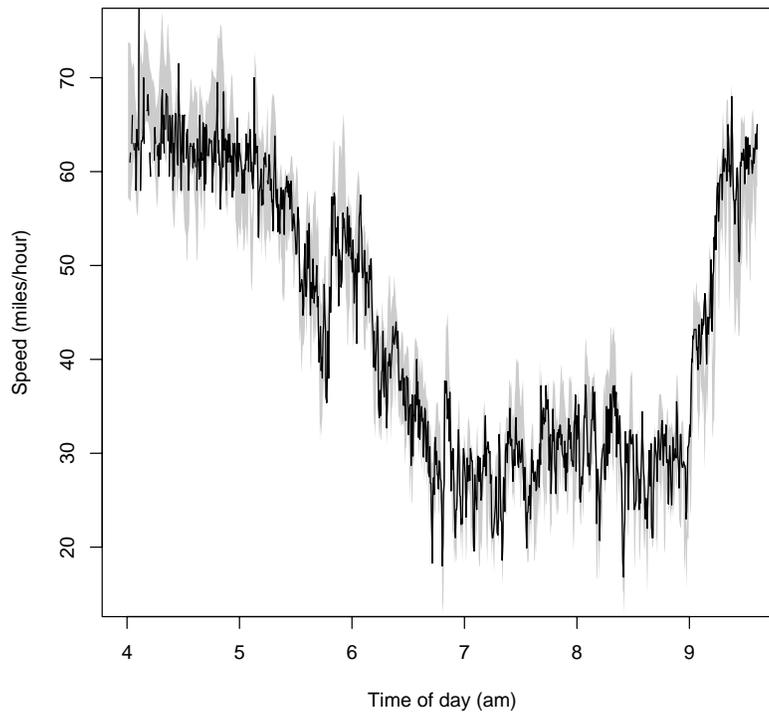
Figure 3. The envelope of a 95% credible intervals for interval mean vehicle speed is shaded grey. Speed trap data is plotted as a solid black line.

4. Discussion of Results

In this section we discuss the results from an analysis of the Seattle data. Data were observed over 1000 twenty second intervals (from 4:00am until just after 9:30am). The quantities of principal interest are the mean vehicle speeds for each interval, $\{\bar{s}_i\}$. The posterior means for these speeds constitute natural point estimates, and are plotted (after conversion to miles per hour) in Figure 2. A useful comparison for our estimates is provided by a set of independent observations on interval mean speeds obtained from a speed trap located close to the inductance detector under study on Interstate 5. We shall refer to the speed trap data as *measured speeds* (although it is important to recognise that speed traps are by no means entirely error-free). These measured speeds (in miles per hour) are also plotted in Figure 2.

The results in Figure 2 indicate that our methodology is doing a very reasonable job of reproducing the measured speeds. The root mean squared difference
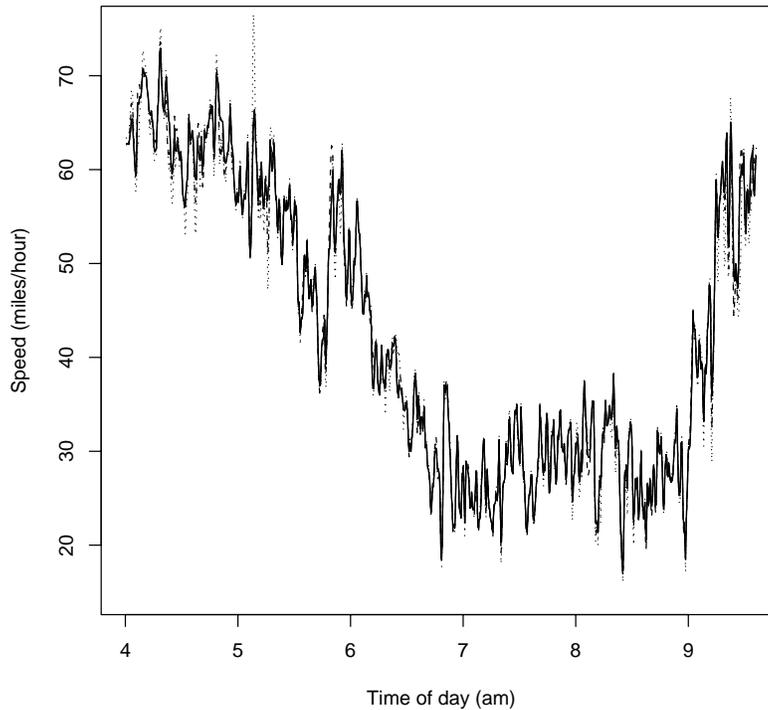
Figure 4. Speed estimates using the following priors for τ_Z: Gamma$(400, 1)$ (solid black line), Gamma$(100, 1)$ (dashed black line), and Gamma$(10000, 1)$ (dotted black line).

between our estimates and the speed trap data is 4.3 miles per hour. This is a very substantial improvement on the figure of 10.2 miles per hour which we obtained using the first order method of moments approach described in section 2. Not surprisingly, the first order method of moments results can be improved by smoothing the estimates $\{\breve{s}_i\}$ using locally weighted regression. However, even when the smoothing parameter is chosen optimally to minimise the root mean squared distance from the measured speeds (an unrealistically good choice for practical purposes), the resulting speed estimates are still worse than our MCMC estimates.

So far we have concentrated on point estimates of speeds. We can obtain (pointwise) 95% credible intervals from the quantiles of the sampled values of $\{\bar{s}_i\}$ in a straightforward manner. The envelope of these credible intervals is shaded in Figure 3 (since this is easier to visualize than the pointwise limits themselves which tend to be obscured due to the rapid fluctuations in speed estimates). The

plot of the measured speeds is superimposed for comparative purposes. Note the speed trap data lie within the credible intervals for the vast majority of the time.

The sensitivity of the results to choice of prior was investigated. Not surprisingly, the precise forms of vague priors for s_{11} and τ_ϵ were of no practical significance in terms of the final speed estimates. This was also the case when a moderately informative Gamma$(0.11, 1)$ prior was assigned to τ_ϵ. In all cases the posterior mean for σ_ϵ was between 1.4 and 2.0 feet per second (i.e. between 1.1 and 1.4 miles per hour). The choice of prior for τ_z was considered a far more serious, since the data provide almost no information about this parameter. We consider three choices of prior – Gamma$(400, 1)$ (corresponding to approximately 5% measurement error); Gamma$(100, 1)$ (corresponding to approximately 10% measurement error); and Gamma$(10000, 10)$ (corresponding to approximately 1% measurement error). The resulting speed estimates are displayed in Figure 4. It is difficult to distinguish many differences in speed estimates using these different priors, suggesting that the exact choice of prior for τ_z is not critical.

We conclude with some comments on a possible refinement of our methodology. In developing our estimation methodology we have implicitly assumed that the distribution of vehicle lengths remains constant over time. Nonetheless, on some roads one might expect a small but significant change in this distribution through the day, with (for example) large freight carrying vehicles having a higher relative frequency during the night than during the morning peak hour. Such temporal variation in the length distribution will have an effect on the speed estimates, since (intuitively speaking) a long vehicle must be travelling at a higher speed than a short vehicle in order to give the same occupancy. It follows that any inadequacies in using a constant length distribution may well be visible in terms of a time varying bias in the speed estimates. While our results do not provide clear evidence of such behaviour in the Seattle Interstate 5 data, the need for a time dependent length distribution will be dependent on the road system under study, and possibly on the time of year. When temporal variation in the length distribution is required, a possible approach is to model the log-length data using a normal mixture model in which the mixture proportions depend on time. (Visual inspection of Figure 1 suggests that the use of a normal mixture is not unreasonable.) This could be achieved using a 'constructive definition' (or 'stick breaking' representation) of the mixture proportions (see Walker *et al.*, 1999), and modelling each the logit of each relative probability by employing cubic splines. We believe that this approach may open up some interesting avenues for further work on speed estimation.

Acknowledgment

The author acknowledges the helpful comments of an anonymous referee.

Apendix: Conditional Distributions

Sampling from the conditional distribution

$$[s_i \,|\, s_{i-}^{(t+1)}, s_{i+}^{(t)}]$$

is rendered straightforward by the Markov property of the random walk of (2.3). There are three distinct cases that need to be covered.

(i) When $i = 1$ the conditional density is given by

$$f(s_1 \,|\, s_{1+}) = f(s_{1n_1} \,|\, s_{21}) \prod_{j=1}^{n_1-1} f(s_{1j} \,|\, s_{1(j+1)})$$

$$\propto \phi_\sigma(s_{1n_1} - s_{21}) \prod_{j=1}^{n_1-1} \phi_\sigma(s_{1j} - s_{1(j+1)}) I(0 < s_{11} < 150)$$

where $\phi_\sigma(x)$ is a $N(0, \sigma^2)$ density evaluated at x, and $I(A)$ is the indicator of the event A.

(ii) When $i = m$ we have (in similar fashion)

$$f(s_m \,|\, s_{m-}) = \phi_\sigma(s_{m1} - s_{(m-1)(n_{m-1})}) \prod_{j=2}^{n_m} \phi_\sigma(s_{mj} - s_{m(j-1)}).$$

(iii) When $1 < i < m$,

$$f(s_i \,|\, s_{i-}, s_{i+}) = \prod_{j=1}^{n_i} \phi_{\sigma_{ij}}(s_{ij} - \mu_{ij})$$

where

$$\mu_{ij} = \frac{(n_i - j + 1)s_{i(j-1)} + s_{(i+1)1}}{n_i - j + 2},$$

and

$$\sigma_{ij}^2 = \frac{(n_i - j + 1)\sigma^2}{n_i - j + 2}.$$

Finally, note that

$$[\tau_\epsilon \mid \boldsymbol{s}, \boldsymbol{l}, \boldsymbol{z}] = [\tau_\epsilon \mid \boldsymbol{s}] = \text{Gamma} \left(2^{-1}n + 0.001, \; 2^{-1} \sum_{i=1}^{m} \sum_{j=1}^{n_i} (s_{ij} - s_{i(j-1)})^2 + 0.001 \right)$$

where we define $s_{10} = s_{11}$ and a Gamma$(0.001, 0.001)$ prior is assigned to τ_ϵ; and

$$[\tau_z \mid \boldsymbol{s}, \boldsymbol{l}, \boldsymbol{z}] = [\tau_z \mid \boldsymbol{z}] = \text{Gamma} \left(2^{-1}n + 400, \; 2^{-1} \sum_{i=1}^{m} z_i^2 + 1 \right)$$

when a Gamma$(400, 1)$ is employed for τ_z.

References

Chib, S. and Greenburg, E. (1995). Understanding the Metropolis-Hastings algorithms. *The American Statistician*, **49**, 327-335.

Coifman, B. (1999). Using dual loop speed traps to identify detector errors. *Transportation Research Record no. 1683*, 47-58.

Dailey, D. (1992). Travel time estimation using cross-correlation techniques. *Transportation Research*, **27B**, 97-107.

Dailey, D. (1999). A statistical algorithm for estimating speed from single loop volume and occupancy measurements. *Transportation Research*, **33B**, 313-322.

Diebolt, J. and Ip, E. (1996). Stochastic EM: method and application. In *Markov Chain Monte Carlo in Practice* (edited by W. Gilks, S. Richardson and D. Spiegelhalter), 259-274. Chapman and Hall, London.

Gamerman, D. (1997). *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*, Chapman and Hall.

Gelman, A. and Rubin, D. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science* **7**, 457-472.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to calculation posterior moments. In *Bayesian Statistics 4* (edited by J. Bernardo, J. Berger, A. Dawid and A. Smith), 169-193. Oxford University Press.

Ihaka, R. and Gentleman, R. (1996). A language for data analysis and graphics. *Journal of Computational and Graphical Statistics* **5**, 299-314.

Kell, J., Fullerton, I. and Mills, M. (1990). *Traffic Detector Handbook (2nd edition)*. Federal Highways Administration.

Kurkjian, A., Gershwin, S., Houpt, P., Willsky, A. and Chow, E. (1980). Estimation of roadway traffic density on freeways using presence detector data. *Transportation Science* **14**, 232-261.

Leutzbach, W. (1988). *Introduction to the Theory of Traffic Flow.* Springer.

Persuad, B. and Hall, F. (1989). Catastrophe theory and patterns in 30-second freeway data – implications for incident detection. *Transportation Research* **23A**, 103-113.

Walker, S., Damien, P., Laud, P. and Smith, A. (1999). Bayesian nonparametric inference. *Journal of the Royal Statistical Society, Series B* **61**, 485-528.

Wand, M. and Jones, M. (1995). *Kernel Smoothing.* Chapman and Hall.

Martin L. Hazelton
Department of Mathematics and Statistics
University of Western Australia
35 Stirling Highway
Crawley WA 6009
Australia
martin@maths.uwa.edu.au