UC Berkeley

Earlier Faculty Research

Title

A generalized and efficient algorithm for estimating transit route ODs from passenger counts

Permalink

<https://escholarship.org/uc/item/17m7k4vm>

Author

Li, Yuwei

Publication Date

2007

A generalized and efficient algorithm for estimating transit route ODs from passenger counts

Yuwei Li, Michael J. Cassidy *

Institute of Transportation Studies, Department of Civil and Environmental Engineering, University of California, Berkeley, CA 94720, USA

Received 19 October 2005; received in revised form 1 December 2005; accepted 5 April 2006

Abstract

An algorithm is presented for transit passenger OD estimation. The algorithm, like some other existing OD estimation techniques, generates estimates based upon passenger boarding and alighting counts at each stop along the route. The algorithm is distinct in that it does not only estimate an OD matrix for the vehicle trip from which the boarding and alighting counts were taken. Rather, it further estimates the passenger alighting probabilities at every stop on the route and these are more apt to remain fixed across transit trips. Therefore, when coupled with projected boarding counts, the alighting probabilities better characterize OD patterns on the route. These probabilities, moreover, are estimated in such way as to reflect the passengers' latent tendencies to travel to and from "major activity centers" where trip-making is induced. The algorithm is therefore a more general-use method than is the estimation technique proposed by Tsygalnitsky. Since the algorithm does not require a seed matrix, and since the number of iterations required for generating estimates is specified *a priori*, the algorithm is easier to apply and more computationally efficient than the balancing method of OD estimation.

© 2006 Elsevier Ltd. All rights reserved.

Keywords: OD estimation; Transit boarding and alighting counts; Trip purpose

1. Introduction

Need frequently arises to estimate the origin–destination (OD) pattern of a transit route's ridership. Large-scale direct sampling of passenger ODs is costly, especially since such sampling would be required for each period marked by distinct demand patterns. This kind of sampling, moreover, is prone to biases and errors (Furth and Navick, 1992).

In contrast, accurate passenger boarding and alighting counts can generally be obtained at all stops along a route. (Perhaps the only notable exception occurs when transit vehicles have insufficient capacity to carry

* Corresponding author. Tel.: +1 510 642 7702; fax: +1 510 642 1246.

E-mail address: cassidy@ce.berkeley.edu (M.J. Cassidy).

demands, such that residual passenger queues form at stops.) Hence the advantage of OD estimation methods that require boarding and alighting counts in conjunction with only limited, and readily obtainable, added information.

Consider the OD matrix for a transit vehicle's trip along a route, with rows representing origin stops and columns representing destination stops. The numeric sum of a row or a column is a stop's boarding or alighting count, respectively, measured during the vehicle's trip. A given set of such sums can correspond to many OD matrices. The task is to select from among these many matrices the one that is most appropriate.

A common estimation approach is the so-called balancing method (Ben-Akiva et al., 1985; Bregman, 1967; Lamond and Stewart, 1981). As its starting point, the method relies on a seed matrix, i.e., a matrix that supposedly incorporates knowledge about passenger travel patterns. It may be an old, outdated OD matrix or one developed from an on-board survey of a small number of passengers. The objective is to pick the OD matrix that matches the boarding and alighting counts at each stop and that provides minimum information over the seed matrix, i.e., the objective is to pick the OD matrix with maximum likelihood.

When this objective is feasible, the balancing method is proven to converge to the optimal solution (Lamond and Stewart, 1981). It does so by (i) multiplying each element in a matrix row by a constant, such that the row's sum matches the stop's actual boarding count; (ii) repeating this for each matrix row in sequence; (iii) repeating the previous two steps for the matrix columns so that sums match alighting counts; and (iv) iterating in this fashion until reaching convergence.

The balancing method is subject to the so-called "problem of non-structural zeros" (Ben-Akiva et al., 1985): if an element in the seed matrix is zero, that element retains a zero value in every iteration. If, moreover, the seed matrix contains a row (column) of zeros but the boarding (alighting) count for that stop is non-zero, then a feasible solution does not exist and the method will not converge.

Tsygalnitsky (1977) proposed a simpler, non-iterative algorithm. Furth and Navick (1992) have shown this algorithm to be equivalent to the balancing method under a special case: namely, when the seed matrix reflects equal probabilities for all OD combinations. Tsygalnitsky treats onboard passengers as a mixed fluid: boarding passengers are added to this fluid and alighting passengers are drawn from it with equal probability (Simon and Furth, 1985). This assumption is often unrealistic, particularly for routes that include major stops that serve activity centers, such as commuter train stations, large businesses, etc.

We propose an OD estimation algorithm that is well suited to routes that serve stops near these activity centers (referred to here as "major" stops) along with stops in other zones, such as residential areas ("minor" stops). Notably, the designation of "major" and "minor" is not a distinction that is necessarily based on the passenger demand for a stop. Rather, the distinction is intended to differentiate stops that serve activity centers (where travel may be induced) from stops that do not.

One can imagine, for example, that residential stops might serve high boarding counts during weekday mornings as passengers commute to their workplaces and schools. During the evenings, these same stops might have large alighting movements. The stops may be designated as minor ones nonetheless.

By distinguishing a stop based on the zone it serves (and not its demand), the proposed algorithm is different from Tsygalnitsky's method. Since both methods generate estimates from boarding and alighting counts, both automatically distinguish stops with high demand from other stops. But only our proposed algorithm can, for example, systematically assign different (presumably lower) likelihoods to passenger trips that both start and end at what we call minor stops.

The example just cited (that passengers are disinclined to travel to and from minor stops) is justified in light of certain empirical evidence, like what is shown in Table 1. It displays survey results of passengers who use AC Transit, an agency with a fleet of more than 650 buses serving two counties in California's San Francisco Bay Area. Review of the table reveals that approximately 85% of passenger trips involve travelling for purposes of work, school, shopping, recreation and medical appointments. As these are activities that commonly occur in zones served by major stops, we can reasonably assume that most or all of these trips involve the use of at least one major stop.

The remaining 15% of trips, which we will call miscellaneous trips, may include those that both start and end at minor stops; e.g. trips whereby passengers travel from home to visit friends or family. Of course, even some of these trips will involve use of a major stop if a residence is near such a stop or perhaps if a trip requires a transfer.

Table 1
AC transit passenger trip purpose survey result, 1999

Trip purpose	Frequency	Percent (%)
Work	946	43.7
School	479	22.1
Shopping	278	12.9
Recreation/social activity	88	4.1
Personal business	141	6.5
Medical appointment	56	2.6
Other	175	8.1
Total	2163	100

In any event, we can say that in some (perhaps many) instances, the trip purposes of passengers who board at a route's minor stops will be similar to the purposes of the overall population of passengers using that route. This occurs on routes in which (i) miscellaneous trips are relatively rare, as in Table 1; and (ii) a large proportion of trips originate at minor stops, as will often be the case when the route carries commuters who travel to work by boarding at minor stops near their homes.

It follows that for such a route with trip purposes distributed as in Table 1, no more than about 15% of the passengers who boarded at some minor stop will alight at minor stops as well. In contrast, Tsygalnitsky's method can over-predict this percentage. (A hypothetical example is presented in an Appendix to this paper.)

By the same token, passenger inclination to travel both to and from major stops might be very different from what is assumed in Tsygalnitsky. This latter inclination may depend upon certain features of the route, as we will illustrate with real data later in the manuscript. The point is that by capturing inclinations of this kind, our algorithm can be applied more generally than Tsygalnitsky's method.

Compared to the balancing method, our algorithm is more computationally efficient since convergence is not an issue, and the number of iterations required is small and specified *a priori*.

Very importantly, our proposed algorithm is distinct from the above two existing techniques in that it does not only estimate an OD matrix for the transit trip from which the boarding and alighting counts were measured. It further estimates the passenger alighting probabilities at every stop. These probabilities characterize passenger OD patterns on the route; they are likely to remain (approximately) constant across transit trips made under similar conditions. Thus when coupled with the (distributions of the) projected boarding counts, the alighting probabilities are better suited for predicting ODs in future transit trips.

The algorithm is described in the following section. Description of the rational metric used to quantify the fitness of estimates is also included there, as is brief discussion on how the transit planner might distinguish a route's major and minor stops. We also explain how the algorithm can be calibrated so as to always furnish estimates that are at least as good as those from Tsygalnitsky's method, even when stops have not been suitably ("correctly") assigned designations of major or minor.

Outcomes from a real-world application of the algorithm are presented in Section 3. There we illustrate that by capturing passenger trip-making tendencies, including passenger disinclination to make short trips, the algorithm furnishes better-fit estimates of alighting probabilities.

Concluding remarks are offered in Section 4. In Appendix A, we present justification for a key aspect of the algorithm's iterative calibration process (described in due course). Finally, in Appendix B we highlight some of the algorithm's computational details by applying it to a simple, hypothetical example.

2. The algorithm

Consider a route and time when transit passengers tend to take trips that involve boarding or alighting at major stops. At a moment immediately before the transit vehicle reaches stop s , take a to be a randomly-selected onboard passenger who boarded at a major stop. The probability that this passenger alights at stop s is denoted p_{as} . Take b to be an onboard passenger who boarded at a minor stop. The probability that she alights at s is p_{bs} . At stop s ,

$$\frac{p_{as}}{p_{bs}} = \frac{1 - \alpha}{\alpha}, \quad (1)$$

where $\alpha \in (0, 1)$. We will assume that $\alpha = \alpha_a$ for all major stops; otherwise $\alpha = \alpha_b$. (Tsygalnitsky's estimation method is suitable for those special cases where $\alpha = 0.5$ for all stops, such that distinctions between a route's major and minor stops do not exist.)

Passengers alighting at stop s are drawn from those onboard following the rule specified in the following theorem.

Theorem. Suppose that immediately upstream of stop s , there are N_a onboard passengers who had boarded at major stops and N_b who had boarded at minor ones. Suppose too that among the n passengers who actually alight at this stop, a total of n_a boarded at major stops. If (1) holds, the expectation of n_a is given by

$$E[n_a] = \frac{(1 - \alpha)N_a}{(1 - \alpha)N_a + \alpha N_b} n. \quad (2)$$

Proof. We randomly select an onboard passenger immediately upstream of stop s and define the following events:

- off: the passenger alights at s ,
- a or b : the passenger boarded at a major or minor stop, respectively.

Eq. (1) means that

$$\frac{p(\text{off}|a)}{p(\text{off}|b)} = \frac{1 - \alpha}{\alpha}. \quad (3)$$

Since

$$p(a|\text{off}) = \frac{p(\text{off}|a) \cdot p(a)}{p(\text{off})}, \quad (4)$$

and

$$1 - p(a|\text{off}) = p(b|\text{off}) = \frac{p(\text{off}|b) \cdot p(b)}{p(\text{off})}, \quad (5)$$

we have

$$\frac{p(a|\text{off})}{1 - p(a|\text{off})} = \frac{1 - \alpha}{\alpha} \frac{p(a)}{p(b)}. \quad (6)$$

Thus,

$$p(a|\text{off}) = \frac{(1 - \alpha)p(a)}{\alpha p(b) + (1 - \alpha)p(a)}. \quad (7)$$

Since $p(a) = \frac{N_a}{N_a + N_b}$, $p(b) = \frac{N_b}{N_a + N_b}$, and $E[n_a] = np(a|\text{off})$, we have Eq. (2). $\square$

If $n - N_b \leq E[n_a] \leq N_a$, our estimate of n_a is given by $\hat{n}_a = E[n_a]$. However, Eq. (2) yields $E[n_a] > N_a$ when $\frac{\alpha}{1 - \alpha} < \frac{n - N_b}{N_b}$. In this latter case, we set $\hat{n}_a = N_a$. Similarly, when $\frac{1 - \alpha}{\alpha} < \frac{n - N_b}{N_a}$, (2) gives $n - E[n_a] > N_b$. In this case, we set $\hat{n}_a = n - N_b$.

The $\hat{n}_a$ for a stop is drawn proportionally from each of the major origins upstream, such that the origin contributing most (least) to N_a gives the greatest (smallest) contribution to $\hat{n}_a$. The $\hat{n}_b = n - \hat{n}_a$ is drawn in like fashion from all minor origins upstream.

Having just described how the algorithm draws alighting passengers for any given α , we now present the process for choosing α_a and α_b to obtain the OD estimates. The boarding and alighting counts needed for the process might be collected from multiple vehicle trips made during a certain period (e.g. the morning peak)

on either a single day or across multiple days.¹ Suppose we use data from a set of M vehicle trips to estimate OD matrices for specified ranges of α_a and α_b , and data from a set of J trips to choose the α_a and α_b that give the best-fit OD estimates. (The two sets of trips may be different or can even be the same.) Finally, suppose that the route has I designated stops. The iterative procedure is as follows:

Step 1: OD Matrix estimation for initial choice of α_a , α_b . For each of the M vehicle trips, we obtain the OD matrix by drawing alighting passengers at each stop from the mix onboard using the logic described above. At each major stop, the $\hat{n}_a$ (and thus the $\hat{n}_b = n - \hat{n}_a$) are calculated for some initial selection of α_a . The selected value for α_b is used to calculate $\hat{n}_a$ and $\hat{n}_b$ at each minor stop. The OD matrices for all M trips are then averaged (element by element) to obtain the aggregate OD matrix. The fitness of this aggregate matrix is assessed *ex post* using the following steps:

Step 2: Conversion of aggregate OD matrix to alighting probabilities. An alighting probability matrix, P , is generated by dividing each element in the aggregate OD matrix by its row total.² For any row total of zero in the OD matrix, one might fill the corresponding row in P with alighting probabilities estimated for some neighboring stop. If the set of J trips is a subset of the M trips, one can simply fill the corresponding row in P with zeros, since no boarding occurs at the stop during any of the J trips.

Step 3: Prediction of alighting counts for each trip. For each of the J trips, the vector of alighting counts is predicted using the boarding counts for that trip and the matrix P . The alighting count, A_s , for each stop, s , is predicted as $A_s = \sum_{i=1}^I B_i P_{is}$, where B_i is the actual boarding count at stop i and P_{is} is the probability that passengers boarding at i alight at s (as computed in Step 2).

Step 4: Fitness assessment of alighting probabilities. To estimate the fitness of the alighting probability matrix, one might compute the (Euclidean) distance between the vector of predicted alighting counts and the vector of measured alightings for each vehicle trip. Such a metric has little physical meaning, however. Instead therefore, the algorithm assesses fitness based upon the vehicle's average load, i.e., the average number of passengers onboard weighted by the route distance (the transit vehicle's revenue kilometers). A vehicle trip's average load is determined from what we call its "characteristic plot", an example of which is shown in Fig. 1. The top curve in that figure displays the cumulative boarding count vs location along the route (as measured from some "origin"). The lower curve shows the cumulative alightings.

The vertical separation between the curves corresponding to any point along the route is the number of onboard passengers (the load) there. This interpretation does not require passengers to board and alight in FIFO fashion. Thus, average load is the area between the curves divided by the route distance.

For each vehicle trip j , the fitness of the alighting probability matrix is determined by comparing the actual average load (obtained from real boarding and alighting counts) with the estimated average (obtained by replacing the cumulative alighting curve of the actual counts with one constructed from the alighting counts predicted in Step 3). To measure overall fitness of the alighting probability matrix across all J trips, we denote the observed average load for trip j as x^j and the estimated average load for that trip as x_{α_a, α_b}^j . We then use

$$D = \sqrt{\frac{1}{J} \sum_{j=1}^J (x_{\alpha_a, \alpha_b}^j - x^j)^2} \quad (8)$$

as the overall fitness measure of the OD matrix corresponding to given α_a and α_b .

Step 5: Iteration. The α_a and α_b are jointly selected through iteration to produce the alighting probability matrix that corresponds to the lowest value for the fitness measure D .

¹ The transit vehicle should have sufficient capacity to accommodate the passenger demands and the usual estimation issues of sample size, unbiased sampling, etc. apply.

² We show in Appendix A that these alighting probabilities are the maximum likelihood estimates if the "true" OD data were (somehow obtained and) used for each of the M trips.

Fig. 1. An example characteristic plot.

Evidence that this reliance on the smallest D is rational is provided in the following section. There we will use synthetic data to show that the true alighting probability matrix yields the smallest D .

As a notable aside, Steps 2–4 above can also serve to quantify the fitness of alighting probabilities when the OD matrix for each vehicle trip is estimated via some methods other than Step 1 of our iterative process.

As regards to the practical matter of distinguishing between a route's major and minor stops, an analyst might designate as "major" only those stops that serve obvious (large) activity centers. Designations can even be assigned through trial-and-error calibration when a small number of stops are in question.

Ultimately, the task of assigning designations to stops requires judgement and this may be viewed as a limitation of the algorithm. We hasten to add, however, that the calibration process can be made to generate OD estimates that are always at least as good as those from Tsygalnitsky's method – even if many or all stops along a route are assigned unsuitable designations. One need only include $\alpha_a = \alpha_b = 0.5$ in the calibration so that alighting passengers are drawn from those onboard with equal probabilities and designations of "major" and "minor" thus become irrelevant.

Assigning stop designations that are generally suitable, on the other hand, adds information that can improve OD estimates. The benefit of this added information is illustrated in the following section.

3. A case study

Here we provide a "macro" level illustration of the algorithm's application. The focus is on the calibration of α_a and α_b for a real bus route (served by AC Transit). The route, shown schematically in Fig. 2, is 26 km in length and serves 58 stops in total. We designate 10 of these as major stops, since each provides easy access to one of four activity centers, the BART (train) station, the college and two shopping malls shown in the figure. Boarding and alighting counts used for estimating the aggregate OD matrix ("Step 1" in the iterative process) as well as for computing the overall fitness of the alighting probability matrix ("Step 4") came from six bus trips made during a 3-h-long morning peak. (Fig. 1 is actually the characteristic plot obtained from the six-trip averages of boarding and alighting counts at each stop along this route.)

Like Tsygalnitsky's method, our algorithm can accommodate an assumption of "intervening opportunities" (Ben-Akiva et al., 1985): a minimum riding distance can be specified so as to reduce or eliminate short trips that in reality are typically made by some other means, such as walking. For any specified minimum trip length L , onboard passengers who have travelled distances greater than L have priority in alighting at a given stop. The N_a and N_b are each partitioned into two groups – those that have travelled distances greater than L and those that have not. Passengers in the latter group alight (in FIFO fashion) only when those with priority have all alighted pursuant to the rule established in Section 2.

Fig. 2. A bus route serving four activity centers.

In this application, L varied from 0 to 4.8 km in increments of 0.4 km. The α_a and α_b each varied from 0.1 to 0.9 in increments of 0.1. As such, the calibration compares 1053 scenarios, each giving rise to an alighting probability matrix.

The minimum value of D is 0.397 and was obtained for $L = 3.2$ km, $\alpha_a = 0.1$ and $\alpha_b = 0.4$ (or 0.3), as shown in Table 2. Had α_a and α_b been calibrated in finer increments, D would be minimum with $\alpha_a \approx 0$ and $\alpha_b \approx 0.35$. An $\alpha_b \approx 0.35$ indicates that at a minor stop, a passenger who had earlier boarded at a major stop is about twice as likely to alight than one who had boarded at a minor one. An $\alpha_a \approx 0$ indicates that at major stops, near-absolute alighting priority should be given to those who had boarded at major stops. These parameter values suggest that, on this route, a large proportion of passengers who boarded at the upstream-most major stop (near the BART station) alighted at major stops.

Of further interest, the choice of $\alpha_a = \alpha_b = 0.5$ (as per Tsygalnitsky's method) generated a minimum value of $D = 0.464$ when $L = 3.2$ km. This too is shown in Table 2. The reader will note that the OD matrix estimated by our algorithm exhibited a better fit than did the estimates from Tsygalnitsky's method.

To illustrate that the calibration process can produce better estimates of alighting probabilities, we simulated 1000 bus trips along the route shown in Fig. 2. For each simulated trip, boarding counts were randomly generated from the distributions inferred from the real data and passenger destinations were randomly assigned using the alighting probability matrix that coincides with $\alpha_a = 0.1$ and $\alpha_b = 0.4$ (consistent with the calibration summarized in Table 2). This matrix thus constituted the "true" alighting probabilities for our synthetic experiment. Each trip's resulting average passenger load was taken as an "observed" value.

We used the same alighting probability matrix to obtain predicted passenger loads. These loads came by calculating alighting counts in the manner previously described in Section 2 (Step 3 of the iterative calibration process). The resulting fitness measure, D , was 0.627.

Values of D were then recalculated for a range of predicted passenger loads. These latter loads came by using the alighting probability matrices for the range of α_a and α_b previously used for the calibration summarized in Table 2. All of these matrices were, of course, inferior to the "true" one that coincided with $\alpha_a = 0.1$

Table 2

Fitness measure of estimated alighting probabilities for different α 's, with $L = 3.2$ km

D	$\alpha_a = 0.1$	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
$\alpha_b = 0.1$	0.413	0.417	0.418	0.419	0.431	0.450	0.453	0.457	0.461
0.2	0.405	0.408	0.409	0.411	0.427	0.447	0.451	0.455	0.460
0.3	0.397	0.401	0.400	0.405	0.425	0.447	0.452	0.458	0.465
0.4	0.397	0.401	0.403	0.413	0.438	0.462	0.470	0.479	0.488
0.5	0.407	0.412	0.419	0.435	0.464	0.490	0.500	0.511	0.523
0.6	0.423	0.434	0.446	0.468	0.499	0.527	0.539	0.553	0.566
0.7	0.451	0.469	0.485	0.512	0.543	0.571	0.586	0.601	0.617
0.8	0.492	0.513	0.535	0.563	0.592	0.618	0.634	0.650	0.665
0.9	0.539	0.563	0.589	0.615	0.641	0.665	0.682	0.700	0.718

and $\alpha_b = 0.4$. And notably, all the latter D were larger than 0.627, the value calculated when predicted loads were obtained from the “true” matrix. (When, for example, $\alpha_a = \alpha_b = 0.5$ as per Tsygalnitsky, the value of D was 0.643.)

This exercise shows empirically that the best-fit alighting probability matrix (in this case, the “true” matrix) is the one that generates the lowest prediction error in load (the smallest D). The passenger OD patterns on the route follow from these alighting probabilities, i.e., the distribution of passenger flow between an OD pair is a mixture of binomial distributions.

4. Conclusions

The algorithm proposed here has several advantages. First, as we have noted, it is computationally efficient. The algorithm does not require a seed matrix; convergence is not an issue and the number of the iterations required by the algorithm is specified *a priori*. Second, by considering passenger trip-making tendencies (primarily through the selection of α_a and α_b), the algorithm is a more general-use method than the one proposed by Tsygalnitsky.

The algorithm further generates best-fit estimates of alighting probabilities and we have noted that from these probabilities the passenger ODs for a route can be obtained. The distribution of the passenger flow between each OD pair on a route can be calculated from these probabilities and the projected distribution of boarding counts at each stop. Mean values of the OD distributions can serve as the scaler elements of a matrix (though information is obviously lost with such a transformation).

In Appendix A, we furnish justification for the algorithm’s use of OD matrices to obtain alighting probabilities. In Appendix B, we present further details related to application of the algorithm.

Acknowledgements

This work has been funded by the California PATH Program. The authors gratefully acknowledge personnel at AC Transit, particularly Tony Divito, Robin Little and Howard Der for providing their ongoing assistance. We also thank the anonymous referee for supplying helpful comments via the review process.

Appendix A

We demonstrate here that maximum-likelihood estimates of alighting probabilities can be obtained from the “true” OD matrices for multiple transit trips along a route. This demonstration justifies our use of an aggregate OD matrix to estimate alighting probabilities.

Suppose we have OD matrices for M trips along a route with I stops. We denote T_{is}^m as the OD flow from stop i to stop s on the m th trip. The underlying alighting probability matrix,³ (P_{is}), is unknown but is the same for all M trips, and passengers’ alighting behaviors are mutually independent.

From the (T_{is}^m) we obtain boardings, $B_i^m = \sum_s T_{is}^m$, and alightings, $A_s^m = \sum_i T_{is}^m$. Given the (B_i^m) and (P_{is}), the relative frequency in which (T_{is}^m) arises should be

$$f_1 = \prod_{m=1}^M \prod_{i=1}^I \frac{B_i^m! \prod_{s=1}^I P_{is}^{T_{is}^m}}{\prod_{s=1}^I T_{is}^m!}. \quad (9)$$

Because (T_{is}^m) and (B_i^m) are data, maximizing f_1 is equivalent to maximizing

$$f_2 = \prod_{m=1}^M \prod_{i=1}^I \prod_{s=1}^I P_{is}^{T_{is}^m}, \quad (10)$$

³ We use parenthesis () to denote a matrix.

which is equivalent to maximizing

$$f_3 = \sum_{m=1}^M \sum_{i=1}^I \sum_{s=1}^I T_{is}^m \ln P_{is}. \quad (11)$$

The set of constraints are

$$\sum_{s=1}^I P_{is} = 1 \quad (12)$$

for $i = 1, \dots, I$.

Applying the method of Lagrangian Multipliers, we obtain the maximum likelihood estimate of P_{is} as

$$P_{is}^* = \frac{\sum_{m=1}^M T_{is}^m}{\sum_{s=1}^I \sum_{m=1}^M T_{is}^m}. \quad (13)$$

Thus, we can obtain (P_{is}^*) by dividing the elements of the aggregate OD matrix by the corresponding row total, as per Step 2 of our iterative calibration process.

Appendix B

This appendix illustrates more of the algorithm's details by applying it to a very simple, hypothetical transit route.

The route in Fig. 3 has four stops. Those numbered 1 and 4 are designated as major stops (we suppose they serve business districts). Although on some trips we will assign relatively large demands to stops 2 and 3, they are designated as minor ones – passengers are presumably disinclined both to start and end trips at these two stops.

The hypothetical sets of boarding and alighting counts represent those from two (hypothetical) vehicle trips. They are presented in Table 3.

Let $L = 0$ so that no short trips are excluded, and $\alpha_b = 0.25$. (For this example, the value of α_a does not matter because all onboard passengers alight when the vehicle arrives at stop 4.) For trip 1, values of N_a and N_b are first calculated at stop 2. These are used by the algorithm along with the alighting count n , to calculate $\hat{n}_a$ and $\hat{n}_b$. These alighting passengers are drawn from the upstream origin (stop 1). The process is repeated for stops 3 and 4; alighting passengers at each are drawn from the origin stops and the numbers from each origin that remain onboard are determined as well. Key estimates for each step are shown in Table 4. The resulting OD matrix is shown in Table 5.

Fig. 3. A route with four stops.

Table 3
Boarding and alighting at stops

Stop number	1	2	3	4
Stop type	Major	Minor	Minor	Major
Trip 1 boarding	2	6	0	0
Trip 1 alighting	0	0	2	6
Trip 2 boarding	6	2	0	0
Trip 2 alighting	0	0	6	2

Table 4
OD matrix estimation process for trip 1

	N_a	N_b	n	$\widehat{n}_a$	$\widehat{n}_b$
<i>Current stop: 2</i>					
$\alpha_b = 0.25$	2	0	0	0	0
Origin stop number	1	2	3	4	
Stop type	Major	Minor	Minor	Major	
Number alighting	0				
Number remaining	2				
New boarding count		6			
<i>Current stop: 3</i>					
$\alpha_b = 0.25$	2	6	2	1	1
Origin stop number	1	2	3	4	
Number alighting	1	1			
Number remaining	1	5			
New boarding count			0		
<i>Current stop: 4</i>					
Origin stop number	1	5	6	1	5
Number alighting	1	2	3	4	
Number remaining	1	5	0		
New boarding count	0	0	0	0	

Table 5
OD matrix for trip 1, with $\alpha_b = 0.25$

Stop number	1	2	3	4
1	0	0	1	1
2	0	0	1	5
3	0	0	0	0
4	0	0	0	0

The OD matrix for trip 2 is estimated in the same fashion. The result is shown in Table 6. By averaging the matrices for both trips, we obtain the aggregate OD matrix shown in Table 7.

The alighting probability matrix is shown in Table 8. The sums for its third and fourth row elements are zero. This does not present any problem for the algorithm; it merely indicates that there is no boarding at stops 3 and 4.

Table 6
OD matrix for trip 2, with $\alpha_b = 0.25$

Stop number	1	2	3	4
1	0	0	5.4	0.6
2	0	0	0.6	1.4
3	0	0	0	0
4	0	0	0	0

Table 7
Aggregate OD matrix, with $\alpha_b = 0.25$

Stop number	1	2	3	4
1	0	0	3.2	0.8
2	0	0	0.8	3.2
3	0	0	0	0
4	0	0	0	0

Table 8
Alighting probability matrix for $\alpha_b = 0.25$

Stop number	1	2	3	4
1	0	0	$\frac{4}{5}$	$\frac{1}{5}$
2	0	0	$\frac{1}{5}$	$\frac{4}{5}$
3	0	0	0	0
4	0	0	0	0

Table 9
Predicted alighting counts for two trips

	Stop number	1	2	3	4
Trip 1	Actual boarding	2	6	0	0
	Actual alighting	0	0	2	6
	Predicted alighting	0	0	2.8	5.2
Trip 2	Actual boarding	6	2	0	0
	Actual alighting	0	0	6	2
	Predicted alighting	0	0	5.2	2.8

Table 10
Boarding and alighting at stops (alternate example)

Stop number	1	2	3	4
Stop type	Major	Minor	Minor	Major
Trip 1 Boarding	2	6	0	0
Trip 1 Alighting	0	0	6	2
Trip 2 Boarding	6	2	0	0
Trip 2 Alighting	0	0	2	6

Alighting counts are predicted for each individual trip using that trip's boarding counts and the alighting probability matrix. The results for trips 1 and 2 are shown in Table 9.

The actual average load for both trips 1 and 2 is 5.33, while the predicted average loads are 5.07 and 5.60, respectively. The fitness measure D for $\alpha_b = 0.25$ is computed to be 0.27.

The reader can verify that $D = 0$ when α_b is a small value approaching zero. As such, the algorithm predicts that virtually none of the passengers boarding at minor stop 2 also alight at minor stop 3. In sharp contrast, Tsygalnitsky's algorithm (with $\alpha_b = 0.5$) produces a $D = 0.50$ and predicts that for trip 2, 75% of those boarding at stop 2 alight at stop 3.

Admittedly, $D = 0$ for an $\alpha_b \approx 0$ is a consequence of the inputs used in this example; i.e., the boarding and alighting counts for the two trips shown in Table 3. Had we used instead the inputs shown in Table 10, our algorithm would predict $\alpha_b \approx 1$, and all passengers who boarded at minor stop 2 alight at minor stop 3. This result indicates that the boarding and alighting counts in Table 10 cannot be reconciled with a distribution of trip purposes similar to that in Table 1. For example, for trip 1 in Table 10, at least four passengers who alight at stop 3 must have boarded at stop 2. Thus, we should conclude that the distribution of passenger trip purposes in this new example is different from that in Table 1. Essentially, α_b is a proxy for trip purposes when the distribution of these purposes is unknown.

References

- Ben-Akiva, M., Macke, P., Hsu, P., 1985. Alternative methods to estimate route level trip tables and expand on-board surveys. *Transportation Research Record* 1037, 1–11.
- Bregman, L., 1967. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics* 7, 200–217.

- Furth, P.G., Navick, D.S., 1992. Bus route *o-d* matrix generation: Relationship between biproportional and recursive methods. *Transportation Research Record* 1338, 14–21.
- Lamond, B., Stewart, N., 1981. Bregman's balancing method. *Transportation Research Part B* 15 (4), 239–248.
- Simon, J., Furth, P., 1985. Generating a bus route *o-d* matrix from on-off data. *ASCE Journal of Transportation Engineering* 111 (6), 583–593.
- Tsygalnitsky, S., 1977. Simplified methods for transportation planning. M.S. Thesis, Department of Civil Engineering, Massachusetts Institute of Technology, Cambridge.