COMPARING THE QUALITY OF OD MATRICES: IN TIME AND BETWEEN DATA SOURCES

Tim Pollard, Nick Taylor, Tom van Vuren
Mott MacDonald

Abstract

Essential but expensive to carry out, roadside interviews are important building blocks for travel demand models. As a rule of thumb, they cost approx. 10 euros per trip record, and capture around 10% of the total passing traffic – expansion factors of 10 are not uncommon. Many (highway) assignment models rely heavily on these types of survey.

The UK Department for Transport requires data in models to be not older than 6 years. In a period of austerity, it is not unreasonable to query whether older roadside interviews will suffice for a model application, or at least to try and use older RSIs, appropriately weighted, in OD matrix estimation. Alternatively, other sources for observed travel patterns are considerably cheaper, but somehow the profession demands that their value is proven by comparison with traditional roadside interviews. An interesting paper by Potter et al (2011) raises concerns about the dependence on roadside interviews in modern transport models.

Recent literature (Djukic et al, 2013) suggests a comparator between OD matrices, based on a similarity index developed for comparing images at pixel level.

In our paper we discuss the following:
-The complexities in comparing OD matrices using standard statistical tests,
-The application of the MSSIM test as described by Djukic et al

We conclude that the MSSIM test is a valuable addition to the transport modeller's toolkit for comparing matrices, in time, between sources, and pre and post matrix estimation from counts. However, further work is required in the operational application of the MSSIM test, which we hope to present in a future paper.

1 INTRODUCTION

Origin-destination, OD or trip matrices are crucial elements of travel demand models, reflecting the travel patterns in the area of study. Often they are the most expensive, and certainly the most complex element in the implementation of a transport model, but it is difficult to determine the quality of either the matrix components or the overall matrix. There are many reasons why there is value in understanding the quality of an OD matrix:

- Avoiding unnecessary surveys where existing data may suffice
- Comparing matrices derived from cheaper (BIG) data sources with those derived by traditional means, to confirm their acceptability
- Comparing the quality of synthesized matrices with observed ones, to strengthen our belief in the underlying model

Comparing the difference between prior and estimated matrix, when refining OD matrices by using counts

Essential but expensive to carry out, roadside interviews are important building blocks for trip matrices. As a rule of thumb, they cost approximately 10 euros per trip record, and capture around 10% of the total passing traffic – expansion factors of 10 are not uncommon. Many (highway) assignment models rely heavily on these types of survey.

The UK Department for Transport guidance in DfT (2012) requires data in models to be no more than 6 years old and we expect similar rules apply elsewhere in Europe. In a period of austerity, it is not unreasonable to query whether older roadside interviews or any other data containing travel pattern information will suffice for a model application, or at least to try and use older data sources, appropriately weighted, in OD matrix estimation. An interesting paper by Potter et al (2011) raises concerns about the dependence on roadside interviews in modern transport models.

Alternatively, other (passively collected) sources for observed travel patterns may be considered. We have extensive experience in the development of OD matrices from GPS data (Van Vuren and Carey, 2011). A recent paper by Calabrese et al (2013) explored mobile phone data. These data sources are substantially cheaper, but until better established the profession demands that their value is proven by comparison with traditional sources, including roadside interviews or other observation samples.

Finally, when using matrix estimation techniques from counts, it is a requirement to prove that the estimated matrix does not differ too much from the original, prior matrix. The UK DfT's WebTAG guidance states maximum allowable differences in the form of slopes, intercepts and R^2values of a regression line between before and after OD cells and trip ends. But are those appropriate comparators for spatial patterns? Are correlations ignored? And is an R^2 of 0.99 a real necessity or just a convenient number?

In our paper we explore alternative ways of assessing the quality of OD matrices, and we intend to identify not only more appropriate metrics for comparison, but also appropriate values for these metrics that provide more evidence that the quality of matrices is actually sufficient. In this exploration we make use of a new comparison measure, the MSSIM, as presented by Djukic et al (2013).

The rest of our paper is structured as follows:

- In section 2 we describe other published work in the area;
- In section 3 we discuss in more detail the MSSIM and how we have implemented this in practice;
- In section 4 we describe the results of our comparisons, and investigate at what values the MSSIM might indicate matrix comparisons of sufficient quality;
- In section 5 we draw conclusions, explore deficiencies in our work to date and how future work may improve on this.

2. PREVIOUS WORK

There is remarkably little published material on the comparison of trip matrices. Perhaps this is not surprising as they are in general not observed, but constructed from different datasets. If they are based on observations, these datasets may have known statistical properties (e.g. determined by sample rates) and these are taken into consideration when constructing matrices. Other datasets, however, such as synthesized matrices from models, have unknown statistical properties. Anyway, even if such a comparison were possible, the results will not be able to show if they are statistically similar enough.

As a result, comparisons of matrices and determination of the acceptability or not of their similarity is done on the basis of their practical application, after assignment. If assigned flows are similar, the matrices are considered to be similar as well. The main problem of this approach is that it smooths over differences between matrices and as a result can give a false impression of similarity. Figure 1 below illustrates how two structurally different matrices lead to exactly the same link flows.

In Figure 1 the black numbers indicate observed flows in-between the 4 junctions/zones of 10, 15 and 12 vehicles. The red numbers show two possible matrix solutions, one with many short trips, and the second with fewer trips, but over longer distances.

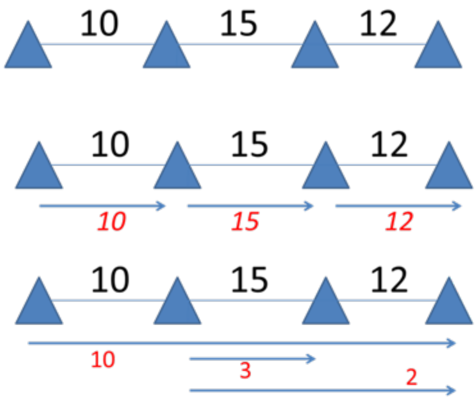

Figure 1: structurally different matrices lead to exactly the same link flows

Sometimes, a further comparison of the trip length distribution in a matrix will be made. This would certainly have detected the structural difference between the matrices in Figure 1, but still there is no clear indicator that will identify an acceptable level of similarity.

Djukic et al (2013) describe the Mean Structural SIMilarity (MSSIM) index in their paper and make a convincing case why this is potentially a more appropriate comparator of matrices. They refer to an earlier, original paper by Zhou et al (2004), who describe the index against a background of comparing images. From a structural point of view, if we liken an image to an OD matrix, pixels equate to individual OD cells; this is attractive as neighbouring pixels / cells are expected to have a degree of correlation just as in images, and the MSSIM index is designed to capture this.

Whereas Djukic et al describe tests that illustrate the value of the MSSIM index in comparing estimated matrices, they do not go as far as exploring which values of the index might indicate a certain level of acceptability. In this paper we discuss tests that we have carried out during the re-estimation of the PRISM model (see van Vuren et al (2004)) for the Greater Birmingham region in the UK, based on new surveys carried out between 2009 and 2011, and also incorporating Big Data from GPS satnav sources.

The work reported in this paper does not yet lead to firm conclusions. To a large extent, only when trying to pursue this line of enquiry did we start recognising practical issues to consider when using the MSSIM index as a matrix comparator. However, our initial findings are encouraging in that they appear to add information over and above the traditional statistical comparisons on the matrix itself before and after assignment.

3. THE MSSIM IN A TRAFFIC CONTEXT

The MSSIM as described by Zhou et al (2004) is a method for comparing two greyscale images. Briefly, the MSSIM computes statistics on groups of pixels and then takes the average (mean), rather than computing statistics based on all the pixels in the image together (such as an R^2test).

Before we consider how or if this may be applicable to transport applications we shall first look at the details of how it is calculated.

Imagine two pictures A and B (of the same size in pixels). By converting each pixel to a number we can consider a collection of numbers arranged in a matrix instead of a rectangular picture. Consider these two picture-matrices side by side, and then look at an NxN square window in the top left of each (in the example Fig 2 below we are using a 4x4 square window). For each of the windows (a and b) calculate the mean and the variance of each, and then the covariance of the two windows together. In the notation below let μ_a and $\mu_b$refer to the mean of numbers in window a of picture-matrix A and the mean of numbers in window b of picture-matrix B. Similarly, σ_a^2, σ_b^2 refer to the variance in windows a and b, and σ_{ab} to the covariance.

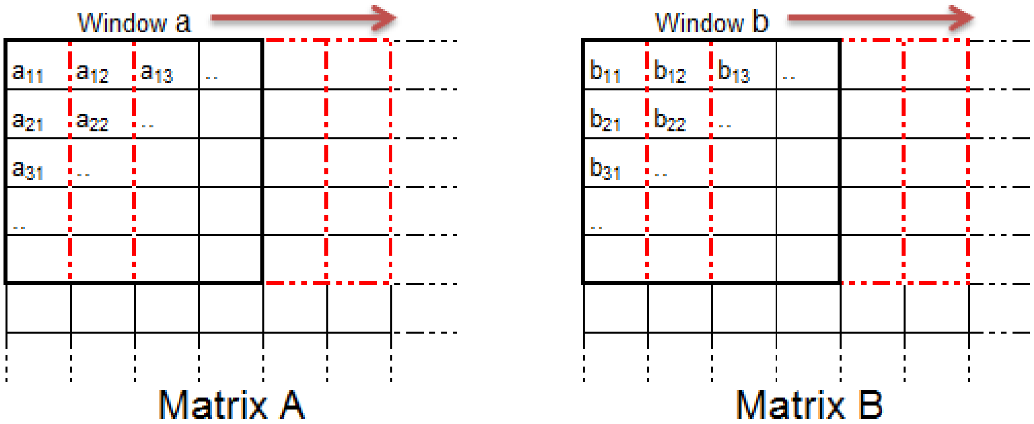

Figure 2: Comparison of two pictures using MSSIM

These numbers are then combined in equation [Fig 3] to form a single Structural SIMilarity Index (SSIM). The equation also has constants C_1 and C_2; these will be explored later in this paper.

$$SSIM(a,b) = \frac{(2\mu_a\mu_b + C_1)(2\sigma_{ab} + C_2)}{(\mu_a^2 + \mu_b^2 + C_1)(\sigma_a^2 + \sigma_b^2 + C_2)}$$

Figure 3: SSIM between two windows a and b.

After calculating the SSIM, the windows a, b, are then shifted one element to the right, and another SSIM is calculated, followed by another shift to produce another SSIM. Once it has moved all the way to the right, it starts again from the left but moves one element down.

Finally, once all possible SSIMs have been calculated from all possible windows within the picture-matrices, the mean of these SSIMs is taken to produce the MSSIM.

$$MSSIM = \frac{1}{N}\sum_{i=1}^{N} SSIM(a_i, b_i)$$

Figure 4: MSSIM between two matrices.
N is the number of windows the SSIM is calculated on

The MSSIM is bounded (between -1 and +1) and is designed to give a value of 1 for identical pictures, and 0 (or close to 0) for two random pictures.

By grouping pixels, this technique assumes that if two pictures are similar then on average the windows will be similar in magnitude of brightness (comparing means), have a similar range of brightness (comparing the variances), and that the arrangement of light-to-dark pixels is similar in both (compared with the covariance). The paper by Zhou et al (2004) goes into detail showing how this technique is combing three comparisons (on mean, variance and covariance) and that the MSSIM provides additional information when comparing pictures making it more useful than standard image comparison tests.

When applying this technique (or any other matrix comparison technique from a different field) to transport OD matrices, one must consider how this relates to the physical meaning of the matrix, and hence if it is still a useful comparison tool.

For example, consider two OD matrices A and B, and associated 6x6 windows a and b. Each window represents movements from 6 origin zones to 6 destination zones (Fig 4).

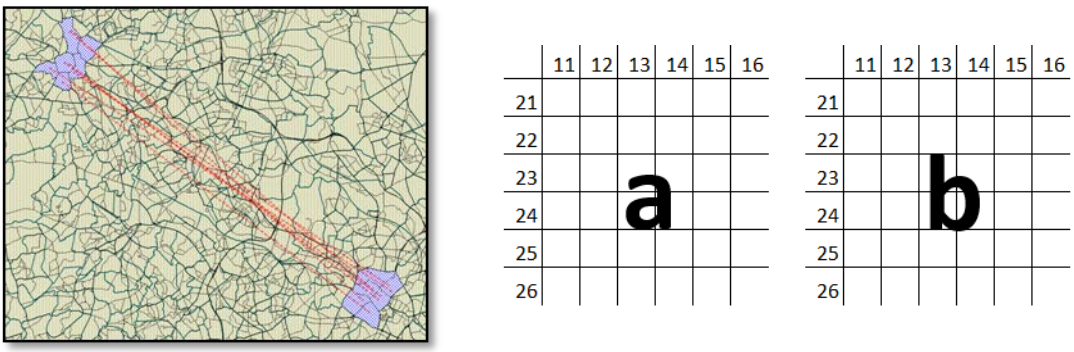

Figure 5: OD matrices A and B
with associated 6x6 windows a and b.

If the two matrices describe similar transport patterns then, as with pixels, it is reasonable to expect similarities on windows of OD patterns. Within each pair of windows a, b on two similar matrices we would expect:

- the total volume and hence mean volume of traffic to be similar in each;
- the variance of those travel patterns to be similar (e.g. very small variance if both are city centres with high demand, or perhaps a high variance if one of

the zones has a very low population);
- the pattern of travel demand to be similar.

Since the three components of the SSIM measure (mean, variance and covariance comparisons) compare aspects of transport patterns that we would expect to be similar in similar OD matrices, it becomes clear that the MSSIM index should be as valuable in comparing OD matrices as with images.

3.1 Testing MSSIM on OD Matrices

In this paper we calculate the MSSIM between matrices, and compare these results with other statistical tests. Before applying the MSSIM to a matrix we need to consider both the values of constants C_1, C_2 and the size of the windows being used. The default values suggested by Zhou et al (2004) of constants C_1 and C_2 taking values 6.5… and 58.5…[1] and using a window of 10x10 pixels did not seem applicable to our context.

The constants in the SSIM equation (fig 3) stabilise the result in the case where either the mean or variance is close to zero. We had chosen to use a matrix from the PRISM model (see van Vuren et al (2004)) as the basis for comparison, however the mean cell value in the matrix is close to 10^{-6}. Clearly for such a matrix the constants will dominate the SSIM equation for most windows, resulting in MSSIM values of near 1 regardless of the matrices you compare with. This was indeed confirmed by testing.

In trying to calculate suitable constants for MSSIM in this context, we decided to consider the relationship between the constants and typical values in their original application – namely greyscale images where each pixel ranges between 256 possible values of grey Zhou et al (2004). If we assume pixel values are randomly[2] chosen between 0 and 255 then the mean squared is ~10^4 and variance is ~10^3. In the SSIM equation (fig 3) the constant C_1 stabilises the effects of the mean squared, and the default value for C_1 is 4 orders of magnitude less than this expected mean from a greyscale image. Likewise the variance is stabilised by C_2 which is 2 orders of magnitude less than the expected variance from the greyscale. We used these crude relationships to guide us in choosing more appropriate constants for the equation for an OD matrix context.

Taking median of values in the PRISM validated matrix we found a mean squared of ~10^{-6}, and variance ~10^0. Hence we adjusted our constants to values of C_1 $=10^{-10}$, C_2 $= 10^{-2}$.

When considering the size of the window, a detailed look at the size of the prism zones was undertaken, and a typical consecutive collection of 6 zones numbers produced a pattern of 6 adjacent zones approximately 3km across in total. Further work is needed in the future to compare the impact to MSSIM values of different

[1] Zhou et al (2004) calculates c = k*L^2 with k equalling 0.01 and 0.03 for C_1 and C_2.
[2] A test was carried out on 100 such samples, from which median values of mean and variance were used.

window sizes.

4. RESULTS OF COMPARISONS

The method and results of the investigation can been split into two parts. Firstly we looked at the possibility of using the MSSIM index to see how altering the structure of a matrix affects the flows on an assignment. Secondly we looked at comparing the MSSIM index with another measure for assessing the difference between matrices; the R^2 measure. Our overall aim was to see how the MSSIM index behaves with structurally different matrices and how it can be used in a practical situation alongside other measures.

In this paper we focus on one particular matrix, the 2011 Employer-Business matrix from a large and detailed transport model named PRISM. PRISM, is a multi-model, multi-purpose transport model based on the West Midlands area in the UK described in more detail in van Vuren et al (2004).

We have chosen the Employer-Business matrix in particular as it has been calibrated and validated to a high standard using nearly 1500 traffic counts and it is reasonable large containing nearly 50,000 trips in total. The matrix was the result of a merge of a Road Side Interview (RSI) matrix constructed from surveys undertaken from 2009 to 2011; a GPS matrix built from various GPS devices including satellite navigation devices; and a Synthetic matrix created from a demand model applied to the network using 2011 land use data. We shall use this matrix as the 'ground truth' matrix for our comparisons.

To enable testing, we artificially created a set of matrices that are structurally different from the ground truth matrix. We added different types of ‘noise’ to the Employer Business matrix by adding varying amounts of the source matrices. We then scaled the resulting matrices so that their totals were equal to the total of the original matrix.

Four types of matrix have been created:

Type 1. $B_1 = A + \alpha\mathrm{Random}$

Type 2. $B_2 = A + \alpha\mathrm{RSI}$

Type 3. $B_3 = A + \alpha\mathrm{RSI} + \beta\mathrm{GPS}$

Type 4. $B_4 = A + \alpha\mathrm{RSI} + \beta\mathrm{GPS} + \gamma\mathrm{Random}$

In the list above, A denotes the Employer-Business matrix; α, β and γ are random numbers from -1 to 1; Random denotes a matrix of random values from 0 to 1 (inclusive); and B is the resulting altered matrix[3]. After this it was important to scale each matrix so that its total size was equal to that of the ground truth. This is to remove bias from the test, which would otherwise be created as some of the matrices created would be larger than others due to the method in which they were made. It was felt such a removal of bias was necessary as in any real-world matrix build

[3] It should be noted that we capped our matrix values below at 0, and so the resultant B matrices will contain some 0 values. Also the random matrix used is different for each new test matrix produced.

application (before comparing two matrices for a merge) they would be scaled in some way to result in similar volumes of traffic.

The resulting matrices were then compared to Matrix A (the original Employer-Business Matrix) using the MSSIM index.

The first series of tests compares the MSSIM index between matrices with the assigned flows of these same matrices. To compare assigned flows of different matrices we have made use of the UK's WebTAG validation criteria and the GEH measure defined in DfT (2012). This states that given a modelled link flow M and an observed traffic count C, the GEH is as follows:

$$\mathrm{GEH} = \sqrt{\frac{2(\mathrm{M}-\mathrm{C})^2}{\mathrm{M}+\mathrm{C}}}$$

Figure 6: Definition of GEH from WebTAG guidance

DfT (2012) recommends aiming to achieve a GEH of less than 5 for 85% of links. In place of the C values in the equation above we have used the assigned flows from the original Employer-Business Matrix and for the M values we have used the flows from assigning a test matrix. This is because we are only interested in drawing comparisons with flows from the ground truth matrix. During the network assignment, as well as the Employer-Business matrix being assigned there are other matrices (i.e. other Car, LGV, and HGV) that are also assigned. These fixed matrices are assigned so that the resulting route choice and flows for Employer-Business would be as realistic as possible. These other matrices are fixed, in that they do not change between assignments of the test matrices. Assigning these extra matrices has some implications to the results. Although this was taken account of in our test for GEH (only comparing the resulting flow from the Employer-Business matrix, rather than total flow), this may still have stabilised the GEH results, as the capacity on links, and route choices available will have been affected by the consistent assignment of these fixed matrices.

The results of the first comparisons can be seen below in figure 7.

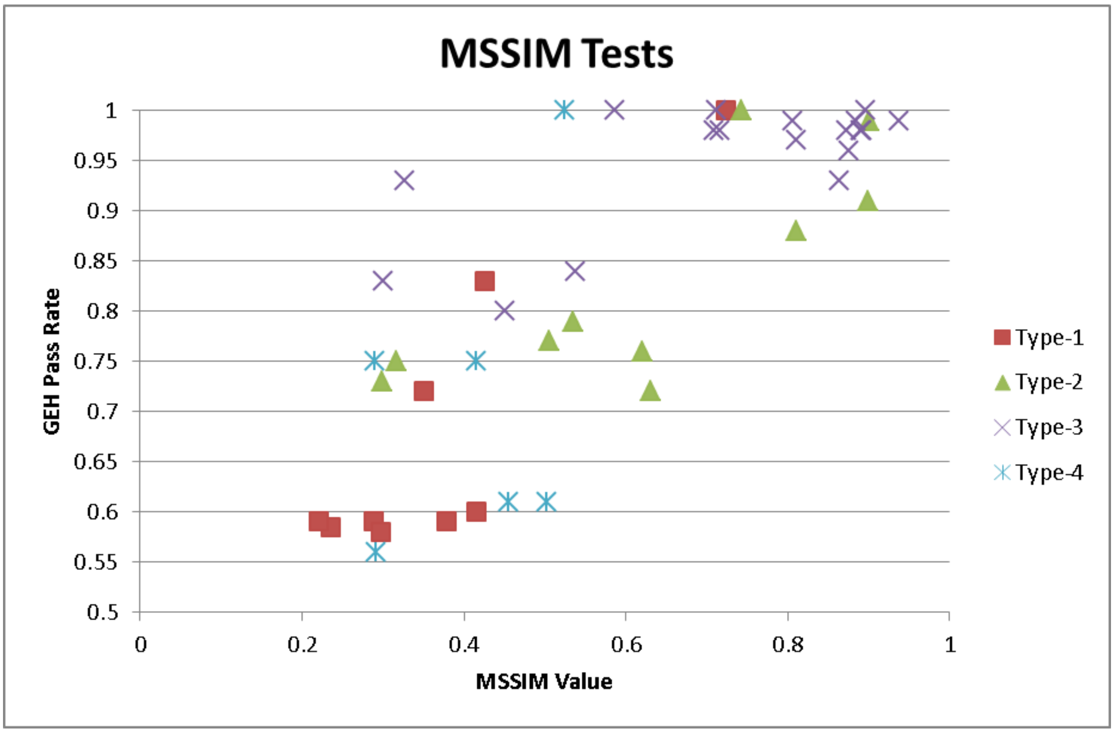

Figure 7: Comparison of GEH pass rate and MSSIM index

Interestingly in our test we did not manage to create a matrix with MSSIM less than 0.2 or with GEH pass rate less than 50% compared to the original (even with 100% random entries) – the latter is due to the stability introduced by the fixed matrices in assignment remaining constant and thus affecting the GEH pass rate (as mentioned earlier). There is a weak correlation between the MSSIM values and the pass rates, but not as strong in the way we had hoped.

It is still possible to infer something from these results.

The worst results for GEH pass rates and MSSIM values came from type 1 matrices, which contain a large random component and tend cluster to the bottom left of the graph.

The type 2 Matrices seem to have a lower GEH per MSSIM compared with the type 3 matrices. Remembering that the type 2 matrices resulted from multiples of the RSI matrix being taken away or added while type 3 matrices have multiples of RSI and GPS matrices taken away or added we can begin to understand this result. Perhaps as the RSI matrix have low cell coverage (i.e. small percentage of non-zero values) then they can make a large change to flows through specific links at the interview sites used to make the RSI matrix, affecting GEH, without making many changes to the overall matrix, and hence MSSIM. On the other hand with type 3 matrices, the GPS matrix has a very large coverage, and so can affect many more OD pairs, affecting the MSSIM, without necessarily affecting the flow as much.

Results for the type 4 matrices, which like type 1 also include a large random component, cluster towards the left of the graph.

Generally lower MSSIM values match to low GEH pass rates and high MSSIM values match to high GEH pass rates. In particular something seems to happen around an MSSIM value of 0.65.

We can split the results roughly into two groups (as shown in figure 8). To the right of 0.65, matrices generally have a GEH pass rate of 85% or above (85% being the normal criteria for passing) and to the left of 0.65 most of the results are below a pass rate of 85%. As shown in the graph we had obtained some results with a low MSSIM value and a high GEH pass rate, which underlines the earlier point in the introduction that assignment of matrices alone is not enough to determine similarity.

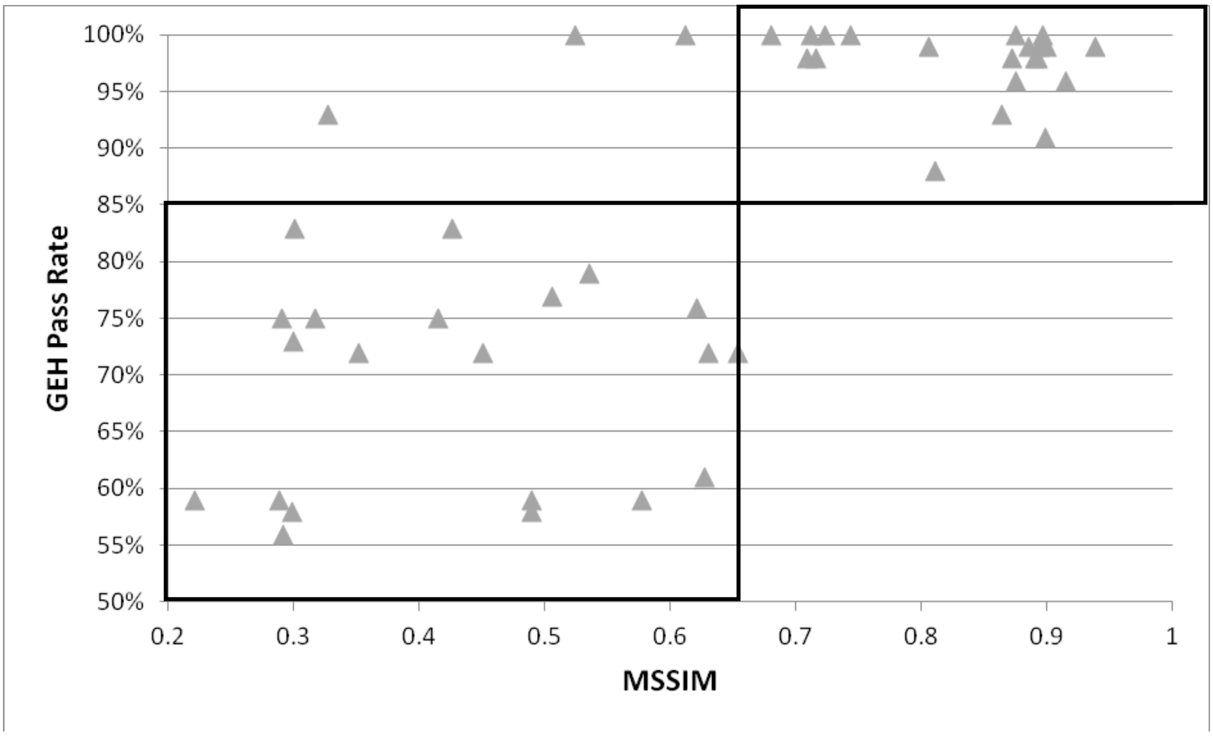

Figure 8: Comparison of GEH pass rate and MSSIM index with splits

To try and get a better understanding of what is happening we decided to carry out a second series of tests on the matrices using the R^2 measure. DfT (2012) section 8.3.13 describes this measure as being recommended before and after matrix estimation with traffic counts. A selection of 8 matrices has been randomly chosen to complete this test. 4 matrices have been chosen with a MSSIM value of less than 0.65 and 4 matrices have been chosen with an MSSIM value of greater than 0.65. These matrices are depicted in Figure 9.

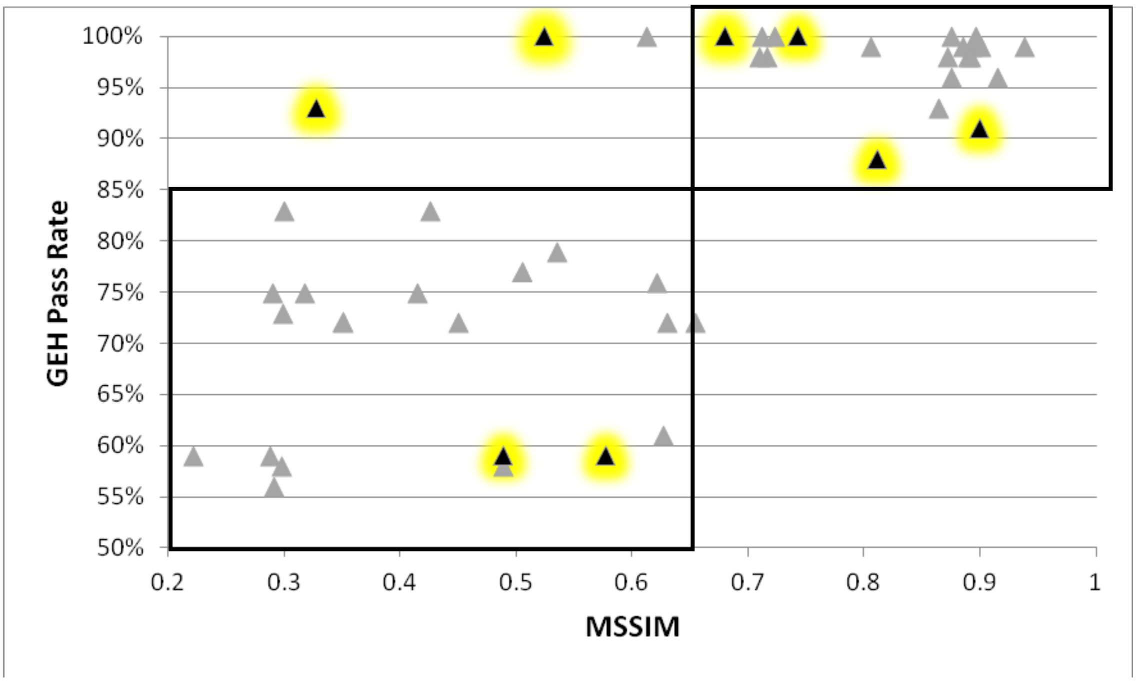

Figure 9: Matrices chosen for R^2 test

The following graph (Fig 10) shows a comparison between the calculated MSSIM values and the related R^2 values. The graph has been ordered from left to right with decreasing MSSIM values.

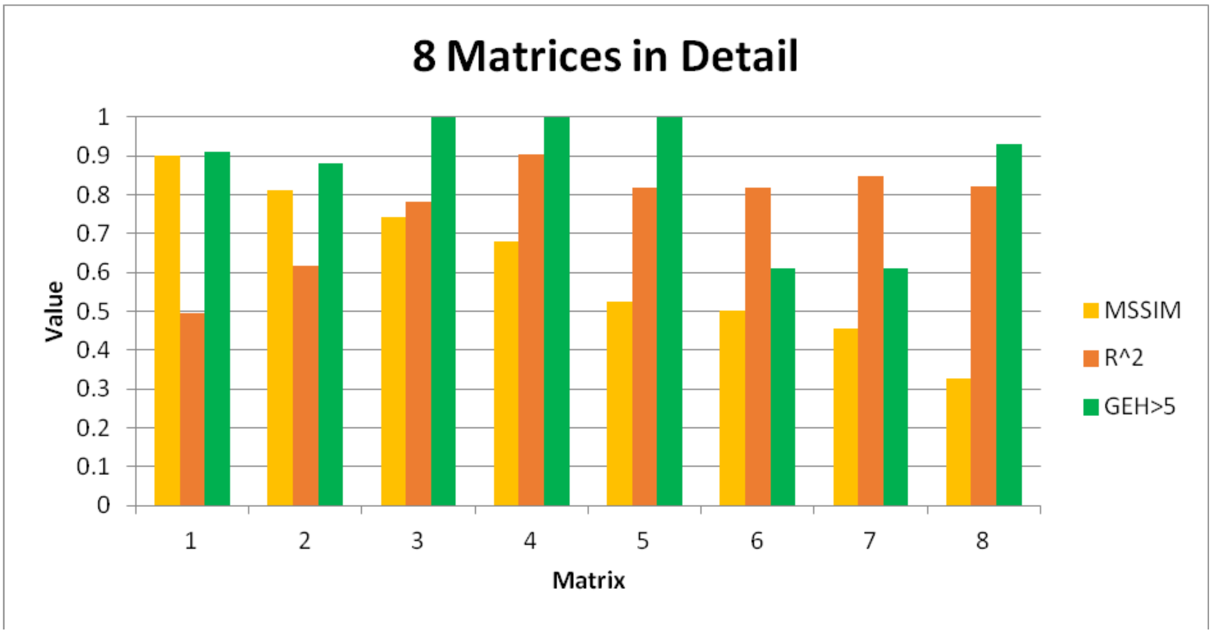

Figure 10: Correlation between calculated MSSIM and R^2.

Again there seems to be no obvious correlation between the values, however as the MSSIM decreases we seem to get an increase in R^2 which stabilises for the lower MSSIM values. This indicates that the R^2 measure does not pick up the same level of change in matrices as the MSSIM index. The R^2 measure itself relates to individual cell values and not the structural aspects of how cells relate to one another. The results also illustrate that the R^2 measure is not a reliable approximation of GEH pass rate.

To try and understand differences further, a final series of tests was done on 2 types of random matrices. In the figure below 4 matrices (left) are constructed by looping through every cell of the original Employer-Business matrix and multiplying by a random number between 0 and 1 exclusive (a different random number for each cell), a further 4 (right) were created purely from random numbers between 0 and 1 exclusive. As before these test matrices were then scaled to have the same total as the original Employer-Business matrix. The tests GEH, R^2 and MSSIM were then carried out, with results displayed below in figure 11.

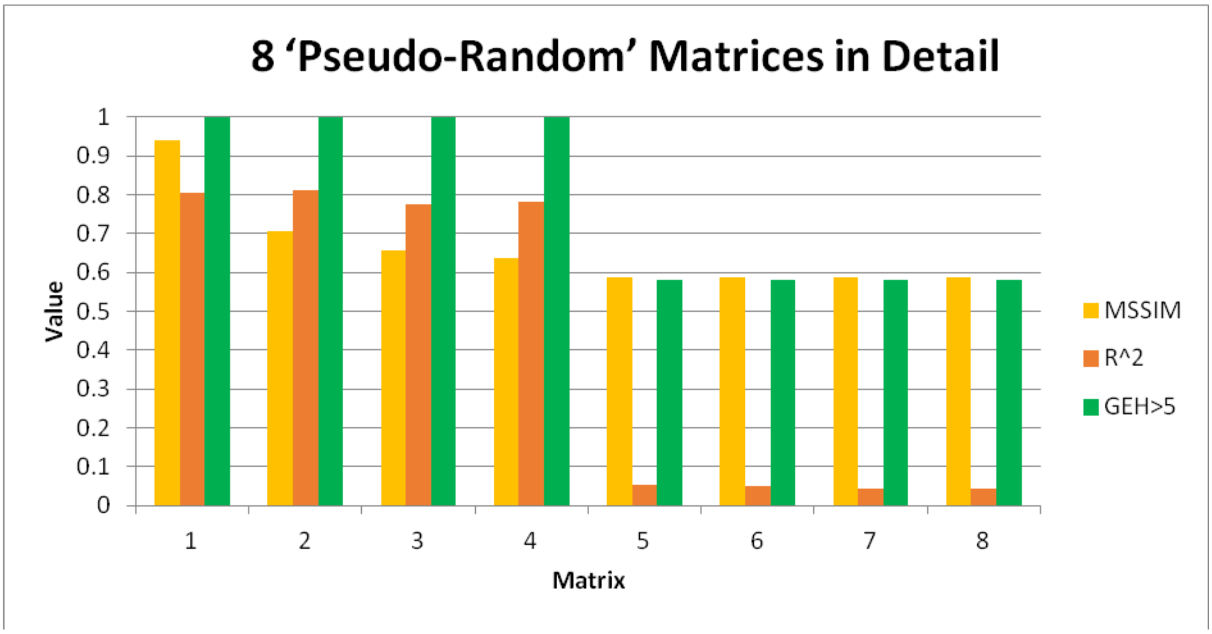

Figure 11: MSSIM and R^2 comparison for B_1 and B_4 matrices.

The 4 purely random matrices consistently have an MSSIM around 0.59, which is higher than some of the values we had before. This is likely because these random matrices were constructed to have no 0 values unlike the previous test matrices. R^2 and GEH seem to give consistent values within both sets of matrices, however the MSSIM has more variation.

There is a clear barrier between the two sets of matrices in the results, with those of a purely random nature scoring much lower on comparisons. These results seem to back up the indication we were getting that GEH was being stablised around 50%, and that something key happened with MSSIM along the 0.6 value. Looking back to results in figure 7, type 1 and type 4 matrices (those that had a strong purely random element but including some zeros) were consistently to the left of 0.65 and often with low GEH. This would indicate that our method for calculting stabilising constants needs to be refined even further so that purely random matrices compared with Employer-Business have an MSSIM closer to 0. Nevertheless looking at the matrices in the right of figure 7, combined with the data from figures 10 and 11 it seems clear that the MSSIM is giving us new information about the structure of these matrices that the standard measures do not reveal.

It is clear that further investigation needs to be carried out, however we believe we have shown that there is more information that can be learnt about matrix similarity from the MSSIM index.

5. CONCLUSIONS AND FURTHER WORK

There are many reasons why in practical transport modelling we want to compare OD matrices and determine if they are similar enough:

- Avoiding unnecessary surveys where existing data may suffice
- Comparing matrices derived from cheaper (BIG) data sources with those derived by traditional means, to confirm their acceptability
- Comparing the quality of synthesized matrices with observed ones, to strengthen our belief in the underlying model
- Comparing the difference between prior and estimated matrix, when refining OD matrices by using counts

And because there are so many reasons why we need to compare matrices we feel more work is required to ensure we understand how the MSSIM indicator can be applied best in practice.

We have used the MSSIM indicator to compare artificially constructed matrices with a ground truth matrix, and compared those results with results from other comparison techniques. Based on our tests, we believe that the MSSIM indicator should be refined further when used on transport matrices, as they are different from the pixel matrices for which the indicator was originally developed.

We believe that the care needs to be taken in calculating constants used within each step of the MSSIM calculation paying attention to known relationships within the field of image comparisons (as we have done) but also with the result of comparisons with a random matrix.

While Djukic et al (2013) pointed out that zoning systems needed to be consistent in both matrices being compared; we feel also the zone numbering needs to be refined so that consecutive numbers always result in adjacent zones. For various modelling reasons this isn’t always the case but before using an MSSIM measure in practice this should be reviewed.

Another possible avenue for further work is to consider weighting some windows (i.e. some area to area movements) as being more or less important, to produce a weighted-Mean-SSIM. This situation may occur when the model the matrices are being built for is more interested in movements into and out of eg a city-centre, and less concerned with surrounding traffic.

The MSSIM index may not be the only comparator that captures the spatial basis of OD matrices; we have come across others, for example the method discussed by

Turner et al (1989). There appears to be a richer literature out there than has been recognised by the transport modelling profession.

Concluding, we believe these results are encouraging enough to continue our investigations, perhaps using different models. We have identified a number of further strands of investigation and testing which should help to provide solid guidance on the application of MSSIM in an OD context that would add value over and above currently used comparators.

6. Acknowledgements

This paper is based on a number of tests with the PRISM strategic transport model for the West Midlands region in the UK. None of the test results reflects an actual model application for either policy or infrastructure investment. We are grateful for the PRISM Management Group's permission to use PRISM for these investigations. The results and their interpretation are solely the responsibility of the authors and cannot be attributed to their employers or the model owners.

7. References

Calabrese, F; Diao, M; Di Lorenzo, G., Ferreira Jr, J; and Ratti, C (2013) **Understanding individual mobility patterns from urban sensing data: A mobile phone trace example**, Transportation Research Part C: Emerging Technologies, Volume 26, Pages 301-313, http://www.sciencedirect.com/science/article/pii/S0968090X12001192, January 2013.

DfT (2012) WebTAG **3.19 Highway Assignment Modelling**, www.dft.gov.uk/webtag, August 2012.

Djukic, T; Hoogendoorn, SP; and Lint, JWC van (2013) **Reliability assessment of dynamic OD estimation methods based on structural similarity index**, presented at TRB, January 2013

Potter, A; and Birks, S (2011) **Are Road Side Interviews (RSIs) still what we think they are?**, presented at 9th Transport Practitioners Meeting, Liverpool, July 2011.

Turner, MG; O'Neill, RV; Gardner, RH; and Milne, BT (1989), **Effects of changing spatial scale on the analysis of landscape pattern**, Landscape Ecology vol. 3 nos. 3/4 pp 153-162

van Vuren, T; and Carey, C (2011) **Building Practical Origin-Destination (OD/Trip) Matrices from Automatically Collected GPS data**, European Transport Conference http://abstracts.aetransport.org/paper/index/id/3737/confid/17 Glasgow, October 2011.

Van Vuren, T; Gordon, A; Daly, A; Fox, J; and Rohr, C (2004) **PRISM: Modelling 21st Century Transport Policies in the West Midlands Region**, Proceedings of European Transport Conference, Strasbourg, Seminar on Applied Transport Methods, October 2004.

Zhou, W; Bovik, AC; Sheikh, HR; and Simoncelli, EP (2004**) Image quality assessment: from error visibility to structural similarity**, Image Processing, IEEE Transactions on, 13(4): p. 600-612, April 2004.