

Learning to walk: Modeling transportation mode choice distribution through neural networks

EPB: Urban Analytics and City Science

2021, Vol. 48(1) 186–199

© The Author(s) 2019

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/2399808319862571

journals.sagepub.com/home/epb

Gideon DPA Aschwanden ,
Jasper S Wijnands , Jason Thompson,
Kerry A Nice , Haifeng Zhao  and
Mark Stevenson

University of Melbourne, Australia

Abstract

Transportation mode distribution has a large implication on the resilience, economic output, social cost of cities and the health of urban residents. Recent advances in artificial intelligence and the availability of remote sensing data have opened up opportunities for bottom-up modeling techniques that allow understanding of how subtle differences in the urban fabric can impact transportation mode share distribution. This project presents a novel neural network-based modeling technique capable of predicting transportation mode distribution. Trained with millions of images labeled with information from a georeferenced transportation survey, the resulting model is able to infer transportation mode share with high accuracy ($R^2 = 0.58$) from satellite images alone. Additionally, this method can disaggregate data in areas where only aggregated information is available and infer transportation mode share in areas without underlying information. This work demonstrates a new and objective method to evaluate the impact of the urban fabric on transportation mode share. The methodology is robust and can be adapted for cases around the world as well as deployed to evaluate the impact of new developments on the transportation mode choice.

Keywords

Transportation modeling, neural networks, machine learning, transportation mode share, walkability

Corresponding author:

Gideon DPA Aschwanden, University of Melbourne, Building 133, Masson, Melbourne, VIC 3010, Australia.

Email: gideon.aschwanden@gmail.com

Introduction

Population increases in past decades have led to new, car dependent development on the urban fringe, resulting in increased travel times (Vandersmissen et al., 2003). The cost of increased travel times can be calculated as either lost time or wasted energy. They can be accrued by each driver directly or impact the society as a whole, namely through higher costs of production, lower productivity, and increased pollution and greenhouse gas emissions (Urban Transportation Task Force, 2012). Time spent traveling is negatively associated with life satisfaction (Hilbrecht et al., 2014), reduced productivity and increased absenteeism (Van Ommeren and Gutiérrez-i-Puigarnau, 2011), and decreased self-reported health (Oliveira et al., 2015). Transportation mode choice, however, can play a critical role in mitigating these costs. For example, commuting by public transport increases physiological energy expenditure and leads to weight loss without exposing people to additional risks such as air pollution, particularly particulate matter (Cepeda et al., 2017; Morabia et al., 2010).

Understanding the factors influencing transportation mode choice is essential for identifying and implementing potential solutions to improve public health and economic output. Increased density and land use diversity can have an impact on transportation patterns and lead to overall health benefits, in particular, through reducing non-communicable diseases such as diabetes, cardiovascular diseases, and respiratory disease (Stevenson et al., 2016). Land use planning decisions have a direct influence on demographics and the socio-economic makeup of an area and, therefore, an indirect impact on temporal distribution, volume, and makeup of traffic (Stewart, 1948). Stewart (1948) also showed that the effect of land use change decreases with distance to a main street, while Aschwanden et al. (2012) showed that increased entropy in land use leads to a reduction in commuting distances and emissions. Santos et al. (2013) investigated factors influencing transportation modal split in European cities and found that density and population size do not have an influence on motorized mode share. However, a negative association was found between public transport subsidies as well as the existing presence of light rail on motorized mode share.

A large body of research has investigated the relationship between the built environment and pedestrian volume, most notably spatial aggregated parametric modeling methods. For example, Cervero and Kockelman (1997) identified the parameters of density, diversity, and design as key indicators for walking; Frank and Engelke (2001) showed the health implications; and Cerin et al. (2009) highlighted the impact of greenery and socio-economic status on walking behavior. As this evidence shows, the parameters that might influence transport volume and mode choices are manifold. Drawing causal inference between urban design factors and outcomes is, therefore, challenging.

Common methods used to model transport patterns use parameters described above and compare them with transportation data collected by government or other agencies. Data are often collected through household surveys, automated traffic recorders (ATRs), and short-term traffic counts (STTCs). ATRs are induction-based counters permanently installed in the pavement and combined with STTCs to estimate average daily traffic, with a relative high error of ATR averaging 24.6% (Gadda et al., 2007). McCord et al. (2003) showed that a combination of aerial photographs and satellite images from highways and ground-based estimates can reduce the error of estimates but can be expensive to implement. High error and expensive collection methods highlight the complexity of the problem and the need for better methods and techniques to estimate traffic volume and mode share.

Neural networks (NNs) can be used in image recognition to detect objects such as roads (Mnih and Hinton, 2010) or vehicles from satellite images (Chen et al., 2014). The use of NNs in transportation forecasting has been proven to yield better insights than traditional

statistical models in volatile cases such as time series, traffic speed (Vlahogianni and Karlaftis, 2013), short-term passenger flow (Wei and Chen, 2012), and vehicle traffic on urban highways (Kumar et al., 2013).

Convolutional NNs are able to incorporate many features without making them a priori explicit (Schmidhuber, 2015). This project, therefore, uses a supervised learning methodology that deploys a NN in combination with labeled satellite images to train a model that estimates transportation mode share. The labeled satellite images contain parameters that would be available through GIS land use plans (e.g., street width, land use, and location of bus shelters) as well as information not available in traditional maps (e.g., greenery, actual building coverage, and building sizes and material). The output is a novel method for incorporating the effect of multiple manifest and latent parameters on transport mode use without limiting of input parameters.

Methods

A NN is trained with satellite images accessed from Google Maps (Google, 2017a) that are labeled with transportation mode distributions from a geo-located trip data set. This section will first introduce the NN methodology, then indicate the details of the training procedure, and lastly explain the validation procedures.

Data

The study area is the state of Victoria, Australia. Victoria contains 6.3 million people located across major cities and agricultural areas and is spread over more than 227,000 km². This paper uses two data sets: georeferenced trips from the Victorian Integrated Survey of Travel and Activity (VISTA) (Victorian State Government and Department of Economic Development, Job, Transport and Resources, 2013) and satellite images from Google Maps.

VISTA data

Trip data were sourced from the Victorian State Government's survey of household travel activity. VISTA contains a random sample of 5780 households asked to complete a travel diary for a specific day in 2013. A total of 14,520 people contributed to the data set. Each household was asked to provide information about trips conducted by all members of the household, the purpose, time, and choice of transportation mode, leading to 63,365 individual trips (Victorian State Government and Department of Economic Development, Job, Transport and Resources, 2013). Geographically, the majority of households in the survey (4130 or 71.5%) are located in the metropolitan area of Melbourne, a representative sample, as 75.1% of Victorians live in metropolitan Melbourne, removing the need to weight trips spatially.

Each trip's information was associated with the origin and destination Statistical Area level 1 (SA1). Trips (origin or destination) are not always associated with the residential location of the respondent. SA1 contain at least 200 residences and can have up to 800 persons residing permanently in it (for details, see <http://abs.gov.au>). With an average of 400 permanent residences, the area of the SA1s range between 4000 m² and 102 km². A sufficient number of trips are required to estimate the transportation mode distribution. With outliers created in areas with a low number of trips, only SA1s with more than 10 trips (origins or destinations) were taken into consideration, leading to a subset of 2177 SA1s out of 13,339 across Victoria (see Figure 1).

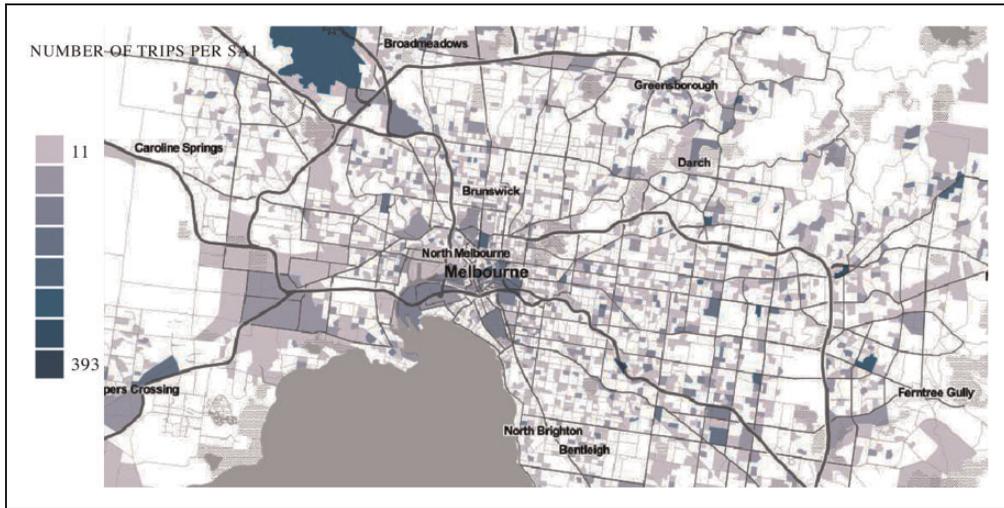


Figure 1. Map of Greater Melbourne, Australia, indicating the number of trips per SA1 (white < 10 trips) with the large area of 30.3 km² in the north containing the airport, an outlier of 393 trips.

Environmental data

To capture the environmental characteristics of each SA1, satellite images were sampled. Satellite images capture not only the coarse configuration of the urban fabric, such as street width and building typology, but also minute differences such as tree coverage and differences in roof tiles. Both indicate differences in socio-economic makeup of the area, which has a direct impact on the availability of different modes and the decision to use them.

To have a data set that is both large enough for training and evenly distributed across all SA1s regardless of its extent, a random set of 1000 satellite images were downloaded from Google (see <https://cloud.google.com/maps-platform/>) for each SA1. Each image is $\sim 400\text{ m} \times 400\text{ m}$. A distance of 400 m is used in several cities as a “rule of thumb” for public transport network distances (Daniel and Mulley, 2013). The downloaded images were 320×320 pixels. This $\sim 1\text{ m}^2$ per pixel strikes a balance between capturing the details while including the wider makeup of an area. To accommodate for inconsistencies due to different capturing times, shadows, etc., preprocessing steps were deployed for each image to increase the robustness of the model (i.e., mirror, adjusting hue, saturation, and contrast). See the “Training Procedure” section for details.

Neural networks

NNs are built from individual neurons that combine multiple inputs $\alpha_x^{(y)}$ linearly by multiplying them with adjustable weights $w_{x,y}$ and creates an output parameter $\alpha_1^{(y)}$ limited by a bounded activation function. The individual neurons are combined in a network architecture that consists of three basic parts: input, hidden, and output layers connected by weights that are adjusted during the learning phase.

Two fundamental paradigms in machine learning exist: supervised and unsupervised learning. Unsupervised learning is deployed in cases where no classification is available and used to extract features. In this study, where the mode share of the SA1 and the respective satellite images are known, supervised learning was deployed. Supervised learning

Table 1. Transportation mode choice distribution and number of trips in the VISTA data set.

Mode of transportation	Number of trips	Share in Victoria (%)
Bicycle	1053	1.66
Motorcycle	148	0.23
Private vehicle	48,115	75.93
Walking	9866	15.57
Bus	903	1.45
Tram	3135	4.95
Other	145	0.23

uses the input variables (i.e., satellite images) and compares the modeled output with the labels (transportation mode share) to adjust the internal weights ($w_{x,y}$) of the neurons via backpropagation.

Three different computational toolkits were explored: DIGITS (NVIDIA, 2016), CNTK (Microsoft, 2017), and TensorFlow (Google, 2017b). All of them require a similar workflow that includes the labeling and classifying of the input images into discrete groups. Since the transport demand of each SA1 is not discrete but follows a probabilistic distribution for each mode choice, the NN's output layer needs to be adjusted (i.e., labels were a probability distribution rather than a binary category). TensorFlow was the platform used for the final model since it allowed for the required customization.

This study uses an adapted network architecture based on a successful design for image recognition (Schmidhuber, 2015) that includes max-pooling in intermediate layers and a SoftMax layer output layer, namely Inception V2 (Szegedy et al., 2016). The Inception V2 and most image recognition NNs are designed to identify a single category from a set of input parameters. Transportation mode choices cannot be classified into discrete states but follow a probability distribution. Therefore, to use this network architecture, the distribution of mode share labels is incorporated using the following adjustments (i.e., vehicle = 0.7593, walking = 0.1557, tram = 0.0495, bicycle = 0.0166, bus = 0.0145, motorbike = 0.0023, and other = 0.0023; see Table 1).

Probabilistic input and output labels have implications on how they are fed into the model and the error calculation for the training process. This study uses cross entropy (CE) to calculate the error between the labeled vector p and the inference vector q . CE uses the unweighted mean absolute error for each transportation mode calculated as

$$CE_{(p,q)} = -\sum_{x \in \chi} p(x) \log q(x)$$

Training Procedure

After the data preparation, labeling the individual images with the transportation mode shares at their location, the images are used to train the NN. To improve the model and its training time, this study uses a pre-trained network with Inception V2 architecture that has been used for image recognition. The network architecture was adjusted to incorporate the probabilistic output layer indicating the transportation mode share and uses an error calculation with CE in both the training and validation process.

The training data set, containing 2,177,000 images, is split randomly into a training and a validation data set in a 3:1 ratio. The individual images have a resolution of 320×320 pixels. To fit them into the network architecture, they are randomly cropped to 256×256 pixels. Additionally, random preprocessing steps are performed each time that an image is loaded during the training, improving the robustness of the model to accommodate color and azimuth variations in the satellite images. The preprocessing steps include random flipping to accommodate for different shadow directions, randomly varying the brightness within one-eighth of the total range, and randomly adjusting the color saturation and contrast by $\pm 50\%$ from the base image.

Training was performed using batches of 64 images, where the accuracy of the model was evaluated by accumulating the error between labels and the classification of the model. This ADAM algorithm (Kingma and Ba, 2014) uses the sign (+/-), the scale of the gradient, to update the weight every 40 epochs and two momentum functions to overcome local minima. The momentum functions have an exponential decay value of 0.9 and 0.999, respectively. A weight decay value of 0.00002, leading to 0 value for some weights over time, makes the network sparser and therefore more robust. The learning rate, following a polynomial decay function (factor 0.999), started at 0.01 with a minimal value of 0.001.

Three stages of validation were performed. The first compared the inferred mode shares with the observations from the VISTA data set during the training of the model every 600 epochs. The training reached a CE of 0.5927. The second validation step compared the modeled mode shares from a set of new images with a benchmark model, and the last step compared the modeled mode shares with the mode share at their location.

Validation

To validate the NN's modeling capability, the inferred values are compared to the values from the VISTA data. Two sets of satellite images were extracted in two grids across the state of Victoria at a resolution of 0.011 degrees (~ 1 km East–West and ~ 1.25 km North–South) and 0.05 degrees (~ 400 m East–West and ~ 500 m North–South).

Validation against a benchmark model

The trained model is compared to a benchmark model. The benchmark model uses the average of all mode shares found in the data set and infers the same ratio to all areas. To measure the accuracy of both models, the CE is calculated between the VISTA data set observations and the predictions of both models. Only predictions within SAIs with more than 10 trips were considered, removing outliers generated by low numbers of observations. The error of the trained model (CE error = 0.593) was 14.1% lower than the benchmark model (CE error = 0.676).

Validation against local distribution mode share

Using the correlation of the modeled mode share and the VISTA data in each SA1 to evaluate the predictive capabilities of the model for each mode share shows that the model has skills in predicting vehicle ($R^2 = 0.587$) and walking ($R^2 = 0.411$) mode share, but it is not suitable to infer the other modes accurately (e.g., bicycle $R^2 = 0.244$ and bus $R^2 = 0.031$) (see Figure 2). Low skill in some modes results from a method that tries to reduce the error. Since the underlying data have very low variance for some of the mode shares, predicting the same value does not lead to the same high error. In short, since many mode shares present a stable value (close to 0), the optimization cost, indicated by a high

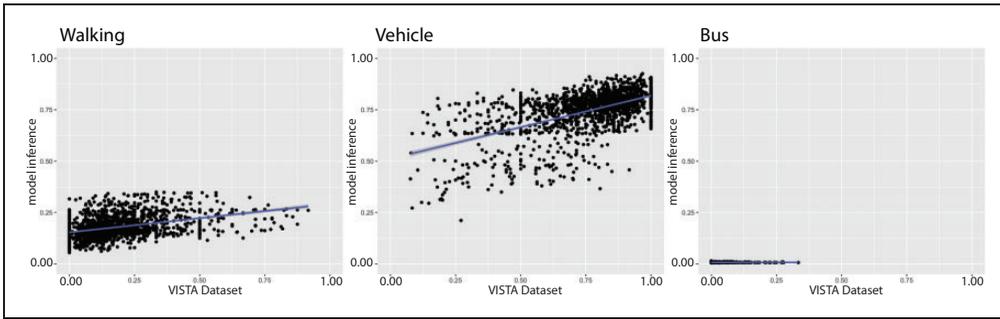


Figure 2. Comparisons of the values from the VISTA data set and the inferred values from the model by SA1.

error, is greater where values of the variance are higher and explains why the model is better in predicting vehicle and walking trips (see Figure 3).

Inference to places without underlying data

The data set containing gridded imagery of two different resolutions was also used for inference. The results inferred from these are overlaid on the map showing the transportation mode choices across Victoria (see Figure 4, left) and metro Melbourne (see Figure 4, right). The NN identified a high proportion of walking trips in high-density urban areas as well as in nature reserves and forests.

The model's inference capability was also deployed to disaggregate distributions within areas where only an average/aggregated number is available despite the fact that only aggregated information on SA1 level was used to train the model (for details, see <http://abs.gov.au>). This highlighted the model's ability to extract characteristics from satellite images at a small scale. Figure 5 shows both the VISTA data (left), where each SA1 is colored to show the share of walking trips, as well as the grid-based prediction of the model (right). This shows that the methodology can be applied to problems where a smaller aggregation level is required than the data provided. On average, the inferred mode shares are in line with those of the full SA1.

By focusing on the extreme values for different modes, the extreme characteristics of the urban fabric are amplified. To do so, we looked at two maps of clusters where either the predicted grid points share of motorized trips are >95% or self-propelled (walking + cycling) transportation modes are >30% (see Figure 6). Areas with a high density of high car use in Greater Melbourne are, with a few exceptions, clustered just outside the ring road. Areas with a high share of self-propelled trips are clustered around the CBD as well as the inner north and south eastern suburbs.

Looking at the urban fabric where the highest predicted shares occur for the main transportation modes, conclusions can be drawn regarding whether building structures and land uses of the urban configurations are conducive to walking, bus, driving, and train (see Figure 7). Walking mainly occurs in the CBD and its adjacent suburbs in the north and south. These areas have high to medium population density and a narrow street grid with trees. Areas where bus trips are common are in the middle suburbs close to major transportation corridors. Train prevalence is incoherent geographically but is detected in areas where a mix of large and small building footprints is prevalent. The share of trips conducted

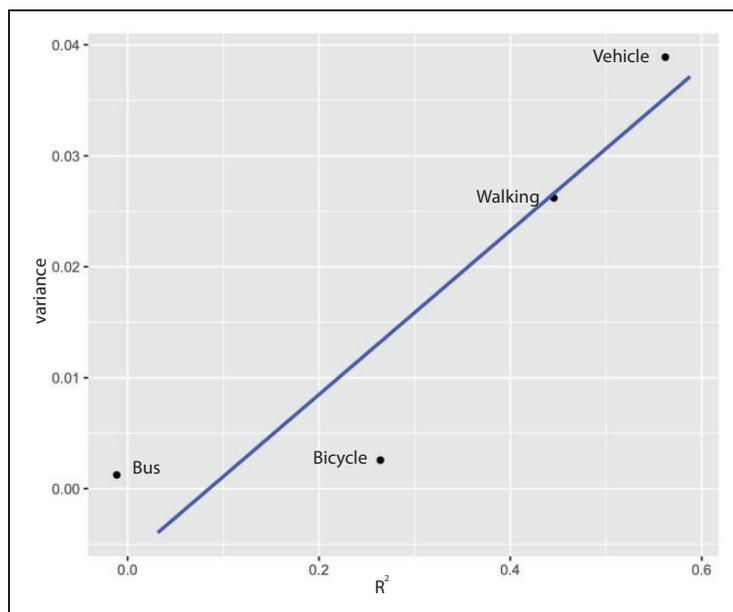


Figure 3. Comparison of the variance in the data set and the accuracy of the model measured in R^2 for different transportation mode shares with a fitted line (blue) through them.

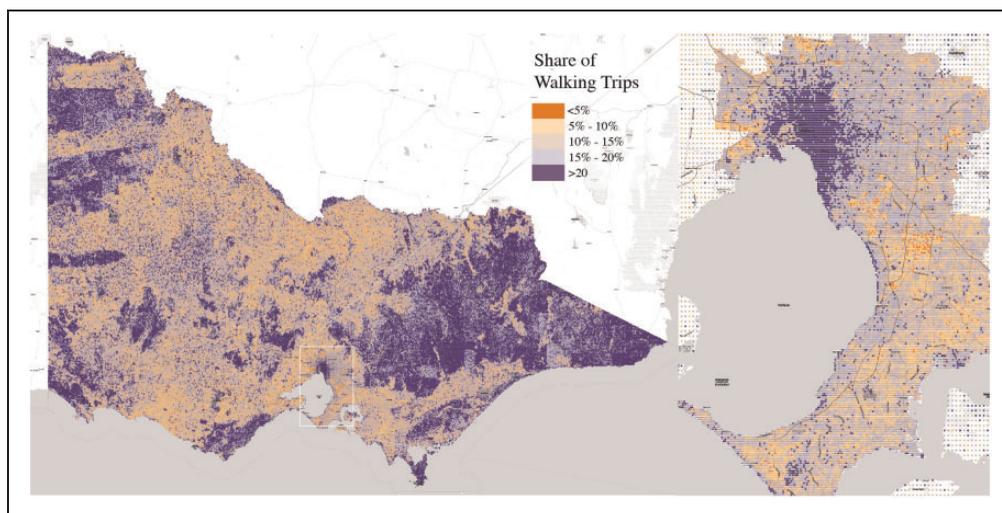


Figure 4. Map of Victoria (left) and Metropolitan Melbourne (right), showing modeled share of walking trips from red (min = 1.2%) to purple (max = 35.3%).

with private vehicles is particularly high in areas where mainly green space can be found (agricultural land use or nature reserve). It is difficult to find associations for the other modes of transportation since the number of observations are low and their share in the overall transportation is negligible (bicycle, motorcycle, and other).

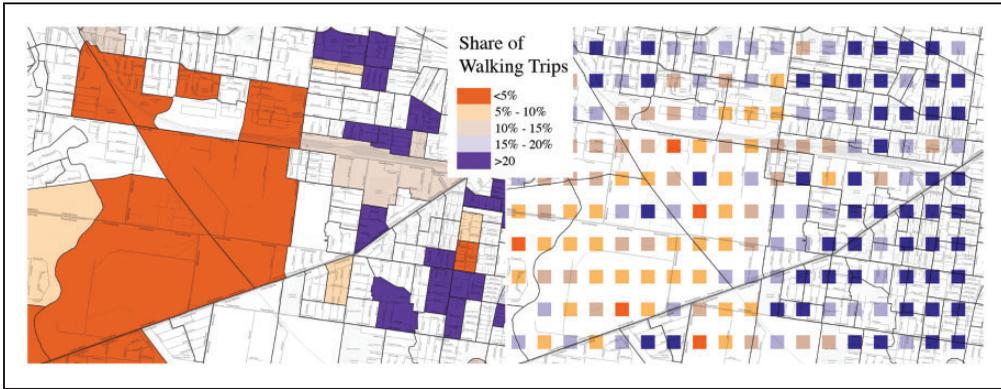


Figure 5. Share of walking trips in each SAI with more than 10 trips (left) and modeled values (right).

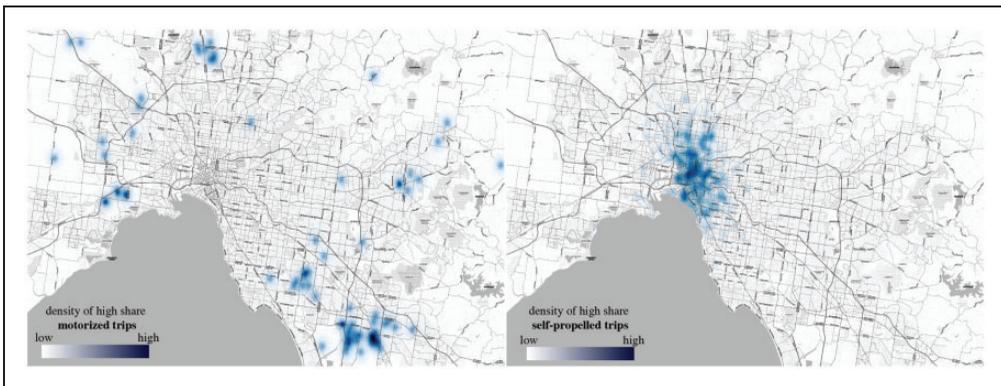


Figure 6. Density of inferred areas with a high share of motorized trips (left) and self-propelled trips (right).

Discussion

This study assessed the impact of the urban fabric on transportation mode share. We have demonstrated a parsimonious model that is able to infer transportation mode shares at an unprecedented scale of granularity. The model infers the distribution of different modes from satellite images where it extracts a wide range of input parameters. It has also shown that the model can be deployed beyond the training and validation data set to provide finer detail for aggregated data sets and to areas where no observation data exist to infer walking trips.

Even though the VISTA transportation survey is limited and is distributed across a small area, the method and the resulting model have been shown to produce transportation mode choice predictions that have high validity indicated by a low error term and the high R^2 . Looking at the satellite images with the highest and lowest values for the individual transportation mode share indicates the presence of distinct urban configurations in areas where a mode share is more or less prevalent. These predictions are consistent with what might be expected given a detailed analysis.



Figure 7. Satellite images from Google on 17 March 2017 with the highest and lowest prediction value for each transportation mode across the state.

The presented method can predict, at a fine-grained scale, the likely transport mode share of existing and new urban developments. Planners and designers now have a tool that allows them to anticipate the impact of their intervention during the design phase without laborious preparation, just by adjusting existing satellite images with images of the new design. The method is highly efficient in evaluating changes with respect to mode share to the built environment and negates the requirement for detailed and costly surveys implemented after neighborhoods have already been established, when the implications of a non-walkable neighborhood are too late or too expensive to change.

The underlying data used in this study are geographically specific, but the method in combination with local transportation data sets can be applied globally. Satellite images are available for all locations, and transportation data sets such as VISTA exist in many countries: e.g. HITS (household interview travel survey) in Singapore; NTS (national travel survey) in England; Mobility and Transportation Micro census in Switzerland; and NHTS (National Household Travel Survey) in the United States. A limitation of our data set is the use of aggregates of trips for a day in SAIs and does not consider the types of trips since splitting it into smaller temporal or topical groups would reduce the data down to a size that is insufficient for drawing inferences.

Since this project only uses satellite images, it is susceptible to all the limitations that they contain. Individual images are collected once at different times over a large area. Their pixel resolution is fixed, and they capture only the visible upper layer. The methodology can, therefore, be augmented by including additional layers of information, including land use, temporal changes and access to other modes of transportation such as public transport or amenities.

Conclusion and future work

This paper showed that a NN was able to identify the urban patterns that are more conducive for walking or other modes of transportation. To this end, two data sets were

combined: satellite images and a transportation survey. Satellite images were labeled with the transportation mode shares of the corresponding SA1 and used as a training data set. To generate this model, the NN was trained with distributions rather than discrete groups of labels. The resulting model was validated first with individual images linked to a distribution as well as calculating the distribution error across a known area. The results show that satellite images, despite their limited information density, can provide accurate estimates of transportation mode choices.

Extending the work by comparing cities in different countries and their transportation mode share highlights differences in preferences or could identify common drivers of transportation mode share decisions worldwide. A related project deploying NNs to compare cities larger than 300,000 people using NNs uses the confusion as a measure of similarity. Another project looking at similar areas of a city found that depending on the training data (satellite, maps, or Google Street View) used to train the NN, different characteristics will be significant (Nice et al., 2018). These studies show that the combination of NNs and global imagery data yields new insights not directly through the model but by creative interpretation and uses of them.

Acknowledgements

This project was made possible thanks to accessible hardware of the THUD (Transportation Health and Urban Design) Research Hub at the University of Melbourne. The data sets, including raw data and the labeled images, are available on request to the corresponding author. We only presented the final results and model generated with TensorFlow.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Gideon DPA Aschwanden  <https://orcid.org/0000-0003-0315-7778>

Jasper S Wijnands  <https://orcid.org/0000-0002-5832-4134>

Kerry A Nice  <https://orcid.org/0000-0001-6102-1292>

Haifeng Zhao  <https://orcid.org/0000-0002-3051-6417>

References

- Aschwanden GD, Wullschleger T, Müller H, et al. (2012) Agent based evaluation of dynamic city models: A combination of human decision processes and an emission model for transportation based on acceleration and instantaneous speed. *Automation in Construction* 22: 81–89.
- Cepeda M, Schoufour J, Freak-Poli R, et al. (2017) Levels of ambient air pollution according to mode of transport: A systematic review. *The Lancet Public Health* 2: e23–e34.
- Cerin E, Leslie E and Owen N (2009) Explaining socio-economic status differences in walking for transport: An ecological analysis of individual, social and environmental factors. *Social Science Medicine* 68: 1013–1020.
- Cervero R and Kockelman K (1997) Travel demand and the 3Ds: Density, diversity, and design. *Transportation Research Part D: Transport and Environment* 2: 199–219.

- Chen X, Xiang S, Liu C-L, et al. (2014) Vehicle detection in satellite images by hybrid deep convolutional neural networks. *IEEE Geoscience and Remote Sensing Letters* 11: 1797–1801.
- Daniels R and Mulley C (2013) Explaining walking distance to public transport: The dominance of public transport supply. *Journal of Transport and Land Use* 6: 5–20.
- Frank LD and Engelke PO (2001) The built environment and human activity patterns: Exploring the impacts of urban form on public health. *Journal of Planning Literature* 16: 202–218.
- Gadda S, Magoon A and Kockelman KM (2007) Quantifying the uncertainty in annual average daily traffic (AADT) count estimates. In: *11th World Conference on Transport Research Society*, Berkeley, CA, 24–28 June.
- Google (2017a) Google maps. Available at: <https://www.google.com/maps> (accessed 17 March 2017).
- Google (2017b) *TensorFlow: An Open-Source Software Library for Machine Intelligence*. California: Google.
- Hilbrecht M, Smale B and Mock SE (2014) Highway to health? Commute time and well-being among Canadian adults. *World Leisure Journal* 56: 151–163.
- Kingma DP and Ba J (2014) Adam: A method for stochastic optimization. *arXiv Preprint arXiv:1412.6980*.
- Kumar K, Parida M and Katiyar V (2013) Short term traffic flow prediction for a non urban highway using artificial neural network. *Procedia – Social and Behavioral Sciences* 104: 755–764.
- McCord M, Yang Y, Jiang Z, et al. (2003) Estimating annual average daily traffic from satellite imagery and air photos: Empirical results. *Transportation Research Record: Journal of the Transportation Research Board* 1855: 136–142.
- Microsoft (2017) *Microsoft Cognitive Toolkit (CNTK): An Open Source Deep-Learning Toolkit*. Redmond: Microsoft.
- Mnih V and Hinton GE (2010) Learning to detect roads in high-resolution aerial images. In: Daniilidis K, Maragos P and Paragios N (eds) *European Conference on Computer Vision*. Berlin: Springer, pp.210–223.
- Morabia A, Mirer FE, Amstislavski TM, et al. (2010) Potential health impact of switching from car to public transportation when commuting to work. *American Journal of Public Health* 100: 2388–2391.
- Nice K, Thompson J, Wijnands J, et al. (2018) The Paris end of town? Urban typology through machine learning. In: *Symposium on New Horizons in Human Dynamics Research*. Available at: <http://hdl.handle.net/11343/212314> (accessed 27 June 2019).
- Nvidia. 2016. DIGITS: Interactive deep learning GPU training system. Available at: <https://developer.nvidia.com/digits> (accessed 27 June 2019).
- Oliveira R, Moura K, Viana J, et al. (2015) Commute duration and health: Empirical evidence from Brazil. *Transportation Research Part A: Policy and Practice* 80: 62–75.
- Santos G, Maoh H, Potoglou D, et al. (2013) Factors influencing modal split of commuting journeys in medium-size European cities. *Journal of Transport Geography* 30: 127–137.
- Schmidhuber J (2015) Deep learning in neural networks: An overview. *Neural Networks* 61: 85–117.
- Stevenson M, Thompson J, De Sá, TH, et al. (2016) Land use, transport, and population health: Estimating the health benefits of compact cities. *The Lancet* 388: 2925–2935.
- Stewart JQ (1948) Demographic gravitation: Evidence and applications. *Sociometry* 11: 31–58.
- Szegedy C, Vanhoucke V, Ioffe S, et al. (2016) Rethinking the inception architecture for computer vision. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Piscataway, NJ: IEEE, pp. 2818–2826.
- Urban Transportation Task Force (2012) *The High Cost of Congestion in Canadian Cities*. Canada: Council of Ministers Transportation and Highway Safety.
- Van Ommeren JN and Gutiérrez-i-Puigarnau E (2011) Are workers with a long commute less productive? An empirical analysis of absenteeism. *Regional Science and Urban Economics* 41: 1–8.
- Vandersmissen M-H, Villeneuve P and Thériault M (2003) Analyzing changes in urban form and commuting time. *The Professional Geographer* 55: 446–463.
- Victorian State Government and Department of Economic Development, Job, Transport and Resources (2013) Victorian Integrated Survey of Travel and Activity (VISTA). Available at: <http://economicdevelopment.vic.gov.au/transport/research-and-data/vista> (accessed 27 June 2019).

Vlahogianni E and Karlaftis M (2013) Testing and comparing neural network and statistical approaches for predicting transportation time series. *Transportation Research Record: Journal of the Transportation Research Board* 2399: 9–22.

Wei Y and Chen M-C (2012) Forecasting the short-term metro passenger flow with empirical mode decomposition and neural networks. *Transportation Research Part C: Emerging Technologies* 21: 148–162.

Gideon DPA Aschwanden is a Lecturer of Urban Analytics at the University of Melbourne and his research focuses on big data and neural networks to evaluate the urban fabric for health, transportation, and economic opportunities. He teaches in urban design and urban planning and property and leads research projects on open data, quantitative policy evaluation, and blockchain development. Prior to this, he taught and researched at Princeton University on digital fabrication methods and building systems as part of the CHAOS lab (Cooling and Heating Architecturally Optimized Systems) in the Andlinger Centre for Energy and the Environment. He gained his MSc in Architecture and PhD in Science from the ETH Zurich and is a founding member of the Future Cities Laboratory in Singapore.

Jasper S Wijnands has a background in applied mathematics, with extensive experience in quantitative analyses and mathematical methods throughout his education, research, and consultancy career. His research has led to publications in various domains, including transport, climate science, and the financial services sector. His PhD in mathematics (University of Melbourne) focused on machine learning algorithms and other mathematical techniques to improve the accuracy of tropical cyclone forecasts in Australia and the South Pacific Ocean. Jasper's research currently focusses on artificial intelligence, especially deep learning using neural networks, in relation to transport, health, and urban design. Prior to this, Jasper was a manager and consultant in EY's Advisory practices (financial risk management) in Melbourne, New York, Zurich, and Amsterdam.

Jason Thompson work is focused on the translation of research into practice across the areas of transportation, heavy-vehicle safety, public health, post-injury rehabilitation, and health system design. Thompson has expertise in methods of computational social science for the modeling of urban, social, organizational, safety, rehabilitation, and active transport systems and has pioneered the use of agent-based models in areas of traditional health systems and transportation research. In 2017, Thompson was awarded an Australian Research Council Discovery Early Career Research Fellowship (DECRA). This project is focused on the effect that introduction of autonomous vehicles will have on Australia's \$5billion personal injury insurance scheme market.

Kerry A Nice Building on a background in software engineering and urban climates, his work uses modeling and artificial intelligence to study urban environments. His PhD at Monash University focused on the creation and use of an urban micro-climate model (VTUF-3D) to assess the positive human thermal comfort impacts in urban areas of increased urban vegetation and water sensitive design (WSUD) infrastructure. Kerry's research currently focuses on the investigation of urban factors impacting the accessibility of active transport; assessing the impacts of urban vegetation on transport, health, and micro-climates; and using artificial intelligence, especially deep learning using neural networks, to assess the influence of urban characteristics on urban environments and ultimately on the people who live there.

Haifeng Zhao has a background in Civil Engineering, Road and Railway Engineering, and Geomatics, with extensive experience in software development, optimisation, agent-based modelling and simulation. His PhD in Geomatics at the University of Melbourne investigated the role of spatial information in evacuation planning and addressed the uncertainty of planning with outdated information using context-awareness and peer-to-peer data sharing. Haifeng's research currently focuses on agent-based simulation, machine learning and spatial analysis to investigate the effects of innovative transportation modes on the health of residents of cities and to improve the design of road infrastructures using artificial intelligence.

Mark Stevenson is an epidemiologist and a Professor of Urban Transport and Public Health at the University of Melbourne. His appointment is across the Melbourne Schools of Design, Engineering, and Population and Global Health. He is a National Health and Medical Research Council (Australia) Research Fellow, an Honorary Professor in the Peking University Health Science Centre, China, and an advisor for injury to the Director General of the World Health Organization. Stevenson has a PhD (Distinction) from the University of Western Australia and a Master's degree in Public Health from Curtin University and became a Fellow of the Australasian College of Road Safety in 2008. He has published over 220 peer-reviewed articles, books, book chapters, and technical reports and procured more than \$31 million in competitive research funding including funding from the NHMRC, ARC, and the US National Institutes of Health. Stevenson is the director of the Transport, Health and Urban Design research hub comprising a cross-disciplinary research team exploring how the effects of urban form and transportation influence the health of residents in cities.