# The 'Paris-end' of town? Urban typology through machine learning.

Kerry A. Nice,
Jason Thompson, Jasper S. Wijnands,
Gideon D.P.A. Aschwanden, Mark Stevenson

Transport, Health, and Urban Design Hub, Faculty of Architecture, Building and Planning, University of Melbourne

AAG 2018, New Orleans, 12 April 2018

# The 'Paris-end' of town? Urban typology through machine learning.

Kerry A. Nice[1*], Jason Thompson[1], Jasper S. Wijnands[1], Gideon D.P.A. Aschwanden[1], Mark Stevenson[1],

**1** Transport, Health, and Urban Design Hub, Faculty of Architecture, Building, and Planning, University of Melbourne, Victoria 3010, Australia

* kerry.nice@unimelb.edu.au

## Abstract

The confluence of recent advances in availability of geospatial information, computing power, and artificial intelligence offers new opportunities to understand how and where our cities differ and also, how they are alike. Departing from a traditional 'top-down' analysis of urban design features, this project analyses millions of images of urban form (consisting of street view, satellite imagery, and street maps). A (novel) neural network-based framework is trained with imagery from the largest 1692 cities in the world and the resulting trained models are used to compare within-city locations from Melbourne and Sydney to determine the closest connections between these areas and their international comparators. This work demonstrates a new, consistent, and objective method for understanding the relationship between cities around the world, and the health, transport, and environmental consequences of their design. The results show specific advantages and disadvantages using each type of imagery, and we draw conclusions about the best use of each for specific analytic goals. Finally, and perhaps most importantly, this research also answers the age-old question, "Is there really a 'Paris-end' of your city?".

# What makes Paris look like Paris?


**Figure 5:** *Books on Paris architecture are expressly written to give the reader a sample of the architectural elements that are specifically Parisian. We consulted one such volume [Loyer, 1988] and found that a number of their illustrative examples (left) were automatically discovered by our method (right).*



**Figure 9:** *Geographically-informative visual elements at the scale of city neighborhoods. Here we show a few discovered elements particular to three of the central districts of Paris: Louvre/Opera, the Marais, and the Latin Quarter/Luxembourg.*

Doersch et al. (2012)

# Convolutional Neural Networks

Algorithms for understanding and categorising images

Can be trained to understand edges, colours and patterns unique to individual categories of images

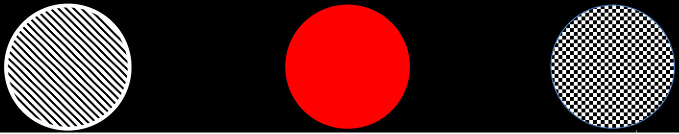

The CNN is made up of millions and millions of individual 'neurons' that each fire in response to various combinations of RGB channel data contained within the image, itself

Over time, weights associated with each neuron are 'adjusted' in response to the correct or incorrect categorisation of images until the model 'converges'

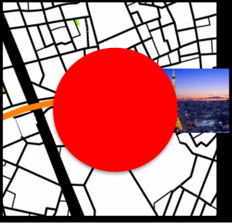

# Neural network training

Using 1.692 million images
1692 cities Cities with > 300k residents
Random Sample of 1000 locations from each city
75% Training, 25% Validation
Selection area scaled to reflect population size p^.85
Training takes a number of weeks
Watercooled GPUs

Each model was trained until convergence for a total of 150 epochs, using the Microsoft Cognitive Toolkit (CNTK) (Yu et al., 2015).

# Which is Melbourne? Or Paris?

The most representative images of Melbourne (left) and Paris (right) according to neural network predictions.

# Training data for 3 neural networks

Four sample Google Maps (GM) neural network training data images for Paris, France. Neural network was trained with 1.665 million images from 1665 cities. Reached a validation accuracy of 73.2%

Sample Google Maps Satellite (GS) neural network training data images for Adelaide, Australia and Beijing, China. Neural network trained with 1.688 million images from 1668 cities. Reached a validation accuracy of 99.4%

# Training data for 3 neural networks


Sample Google Street View (GSV) neural network training data image from Sydney, Australia (Google Maps, 2017b) (A) and the processed segmented version (B). Sample Baidu Street View (BSV) neural network training data image from Beijing, China (Baidu, 2017) (C) and the processed segmented version (D). Neural network was trained with 1.074 million images from 1074 cities. Reached a validation accuracy of 43.1%

# Sampling urban imagery

The sampling area for each city was chosen as a circular area aligned to the city's centre, where the radius $r$ (km) of the sampling area was determined based on the population size $p$ according to Barthelemy (2016)

$$r = \sqrt{\frac{28.27}{\pi}\left(\frac{p}{300{,}000}\right)^{0.85}} \tag{1}$$

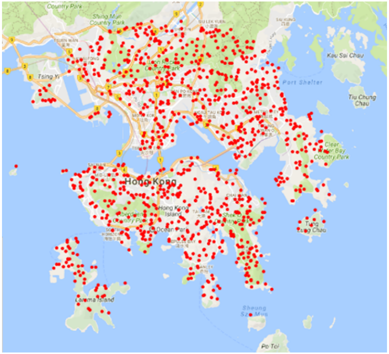


Randomly generated locations to sample urban form in Hong Kong. 1000 images sampled from each city. Large water-bodies (e.g., oceans but not coastlines) were removed from the sampling area.

# Evaluation locations for Melbourne and Sydney

Melbourne (23,027) and Sydney (24,596) evaluation locations at 400m resolution. No training was performed using imagery from Melbourne or Sydney, so the evaluation forced the neural networks to pick the most similar city to each evaluated location.

# Top 20 cities like Melbourne/Sydney using Google Maps imagery

| Predicted city | Melbourne evaluation Matches | Melbourne evaluation % matching | Sydney evaluation Matches | Sydney evaluation % matching |
|---|---|---|---|---|
| Brisbane, Australia | 4062 | 17.6 | 2922 | 11.9 |
| Beer Sheva, Israel | 1671 | 7.3 | 3479 | 14.1 |
| Canberra, Australia | 1074 | 4.7 | 1770 | 7.2 |
| Sunshine Coast, Australia | 960 | 4.2 | 456 | 1.9 |
| McAllen, United States of America | 794 | 3.5 | 140 | 0.6 |
| Valparaíso, Chile | 431 | 1.9 | 383 | 1.6 |
| Haifa, Israel | 372 | 1.6 | 723 | 2.9 |
| Gold Coast, Australia | 361 | 1.6 | 265 | 1.1 |
| Newcastle and Lake Macquarie, Australia | 347 | 1.5 | 1214 | 4.9 |
| Toronto, Canada | 325 | 1.4 | 209 | 0.9 |
| Kitwe, Zambia | 303 | 1.3 | - | - |
| Pretoria, South Africa | 303 | 1.3 | - | - |
| Auckland, New Zealand | 285 | 1.2 | 293 | 1.2 |
| Johannesburg, South Africa | 239 | 1.0 | - | - |
| Perth, Australia | 216 | 0.9 | 111 | 0.5 |
| Philadelphia, United States of America | 185 | 0.8 | - | - |
| Washington, D.C., United States of America | 174 | 0.8 | - | - |
| Port Elizabeth, South Africa | 162 | 0.7 | - | - |
| Virginia Beach, United States of America | 143 | 0.6 | 235 | 1.0 |
| Jerusalem, Israel | 142 | 0.6 | 393 | 1.6 |
| Tasikmalaya, Indonesia | - | - | 238 | 1.0 |
| London, United Kingdom | - | - | 199 | 0.8 |
| Adelaide, Australia | - | - | 152 | 0.6 |
| Denton-Lewisville, United States of America | - | - | 143 | 0.6 |
| Hong Kong, China, Hong Kong SAR | - | - | 121 | 0.5 |
| Osaka, Japan | - | - | 113 | 0.5 |

# Plotting color scheme

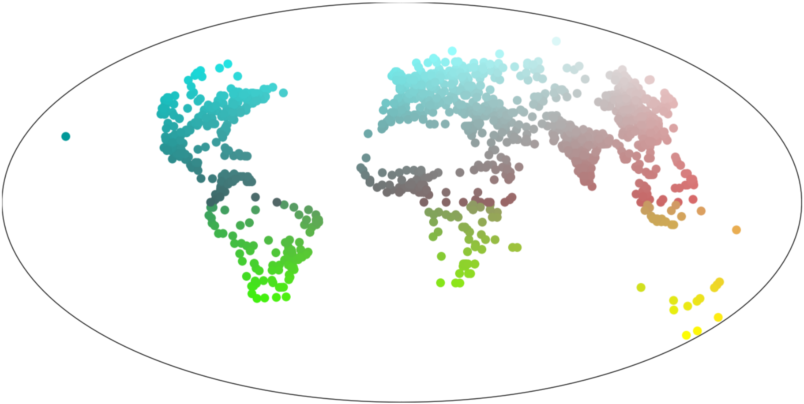

Latitude/longitude based color scheme for plotting cities-like for Melbourne and Sydney evaluations.

# Are Melbourne/Sydney like Paris using Google Maps imagery?

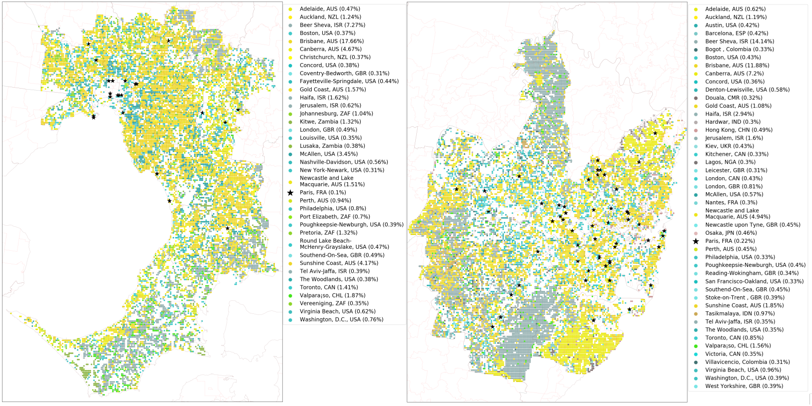


Predicted similar cities using the GM neural network. Top predicted cities plotted using color scheme for Melbourne evaluation locations and Sydney. Predicted Paris locations marked with black stars.

# Melbourne CBD GM neural network evaluation locations

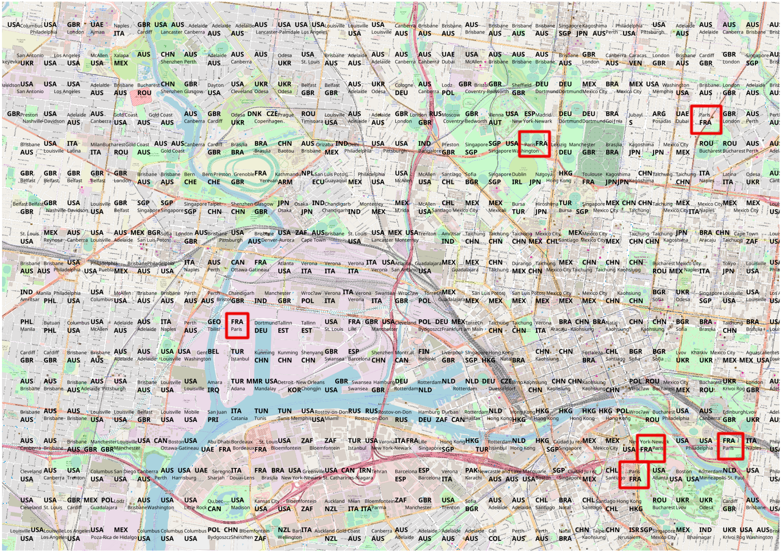

Detail of Melbourne CBD, with predictions of Paris highlighted in red squares.

# GM model Paris-like locations

Gallery of 'Paris-like' locations in Melbourne using the GM neural network.

Gallery of 'Paris-like' locations in Sydney using the GM neural network.

# Are Melbourne/Sydney like Paris using Google Satellite imagery?

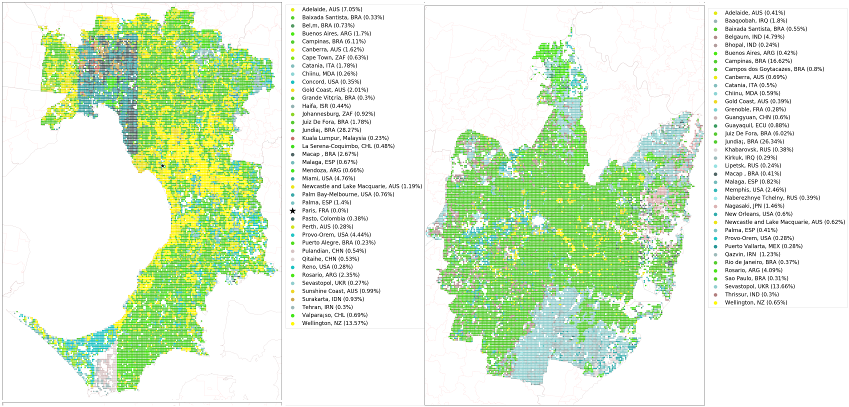


Predicted similar cities using the GS neural network. Top predicted cities plotted using color scheme for Melbourne evaluation locations and Sydney. Predicted Paris locations marked with black stars.

# GS model predicted similar cities

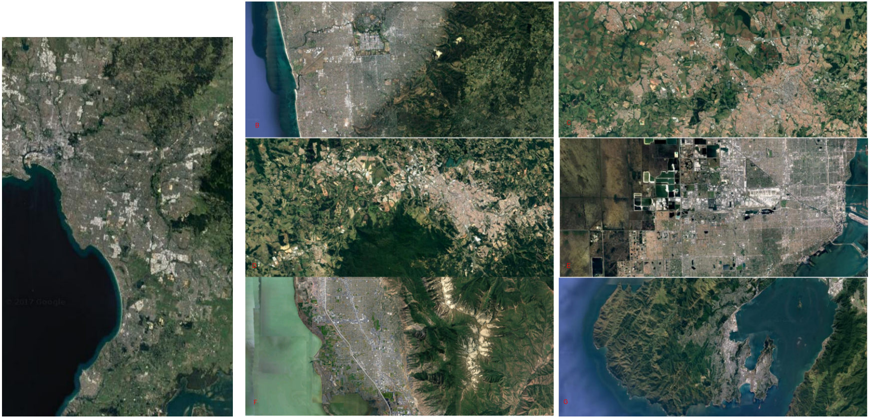

Satellite imagery of Melbourne, Australia (A), Adelaide, Australia (B), Campinas, Brazil (C), Jundiaí, Brazil (D), Miami, USA (E), Provo, USA (F), and Wellington, NZ (G) Google Maps (2017a).

# Are Melbourne/Sydney like Paris using Google/Baidu Street View imagery?

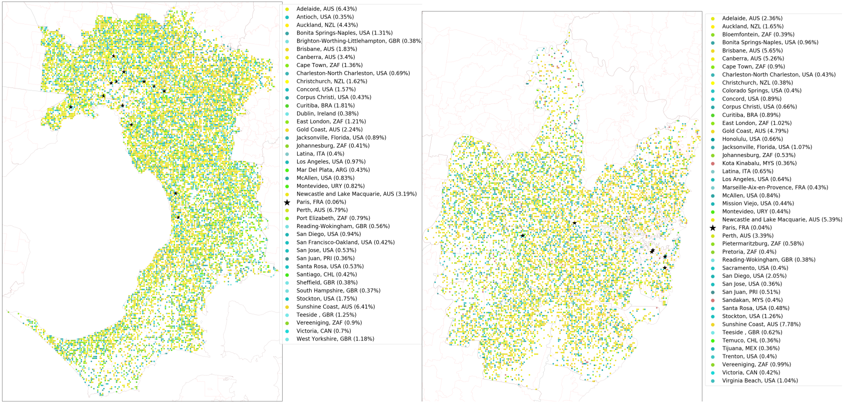


Predicted similar cities using the GSV-BSV neural network. Top predicted cities plotted using color scheme for Melbourne evaluation locations and Sydney. Predicted Paris locations marked with black stars.

# GSV-BSV model Paris-like locations

Gallery of ‘Paris-like‘ locations in Melbourne/Sydney using the GSV-BSV neural network.

# Is there a 'Paris-end' of Melbourne or Sydney?

No

# Related work - City clustering based on confusion

Most algorithms of this type are interested in accuracy - i.e., what percentage of images can the algorithm correctly classify?

80% +

Instead, we were interested in what the network got wrong – when it got confused

Each instance of confusion created a link between two cities, generating city 'clusters'

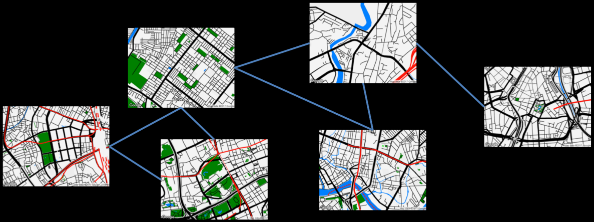

Road networks, public transport, open space, density, water, topography

# Clustering of similar cities - Examine health outcomes

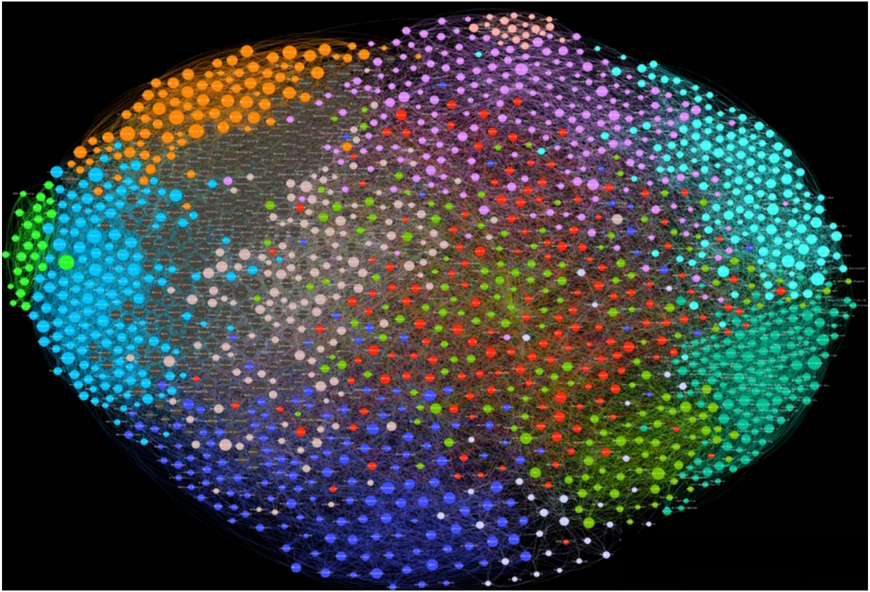

Thompson et al. (2018)

# Summary and insights

- We have created a method for objectively and consistently defining city typologies and clusters.
- Method uses easily accessible and globally consistent imagery data.
- Can explore how urban design can impact public health and other urban issues.
- Neural networks have no preconceived notions of what they will find and use all of the data. Cherry-picking is right out.
- Sometimes the neural networks can produce mysterious or surprising results, requiring cross-disciplinary interpretation (and perhaps some reverse engineering).
- Different types of imagery can emphasise different urban characteristics.
- Sadly, neither Melbourne or Sydney can claim to have a 'Paris-end' of town.
- The trained networks are available. Contact me if you want to try it on your own town.

# Bibliography


Baidu (2017), Baidu Street View API, Available from http://api.map.baidu.com/. (accessed 15 June 2017).

Barthelemy, M. (2016), *The Structure and Dynamics of Cities: Urban Data Analysis and Theoretical Modeling*. Cambridge University Press.

Doersch, C., Singh, S., Gupta, A., Sivic, J. and Efros, A. (2012), What Makes Paris Look like Paris? *ACM Transactions on Graphics, Association for Computing Machinery*, 31(4).

Google Maps (2017a), Google Static Maps API, Available from https://developers.google.com/maps/documentation/static-maps. (accessed 15 June 2017).

Google Maps (2017b), Google Street View API, Available from https://developers.google.com/maps/documentation/streetview/. (accessed 15 June 2017).

Thompson, J., Stevenson, M., Wijnands, J., Nice, K., Aschwanden, G., Silver, J. and Nieuwenhuijsen, M. (2018), Linking derived city typologies and health; A new global perspective. *In-preparation*.

Yu, D., Eversole, A., Seltzer, M.L., Yao, K., Huang, Z., Guenter, B., Kuchaiev, O., Zhang, Y., Seide, F., Wang, H., Droppo, J., Zweig, G., Rossbach, C., Currey, J., Gao, J., May, A., Peng, B., Stolcke, A. and Slaney, M. (2015), An Introduction to Computational Networks and the Computational Network Toolkit. Microsoft Technical Report MSR-TR-2014–112. Technical report.

**Transport, Health, and Urban Design (THUD) Research Group**

Dr Kerry Nice (Software Engineer, Urban Climate)
https://mothlight.github.io/
@mothlight

Dr Jason Thompson (Clinical Psychology)
Dr Jasper Wijnands (Mathematician)
Dr Gideon Aschwanden (Architect, Urban Analytics)
Professor Mark Stevenson (Epidemiologist)
Dr Haifeng Zhao (Complex Systems Modelling)

THE UNIVERSITY OF MELBOURNE | msd Melbourne School of Design