Aggregation of Dots

Martijn Tennekes

Keywords: visualization, spatial data, compositional data

1. INTRODUCTION

The dot map is a powerful visualization method to explore the spatial distribution of data. Each data point is represented by a dot on a map. A famous example is the dot map by John Snow in which he showed that the locations of cholera deaths were correlated with the location of an infected water pump (Figure 1, left).

Historically, the dots were coloured black. When colouring the dots according to a category in the data, dot maps can not only be used to analyse the spatial distribution of data, but also its composition. This is illustrated by the Racial Dot Map [1], where the ethnic background of the US population is shown in an interactive map (Figure 1, right).

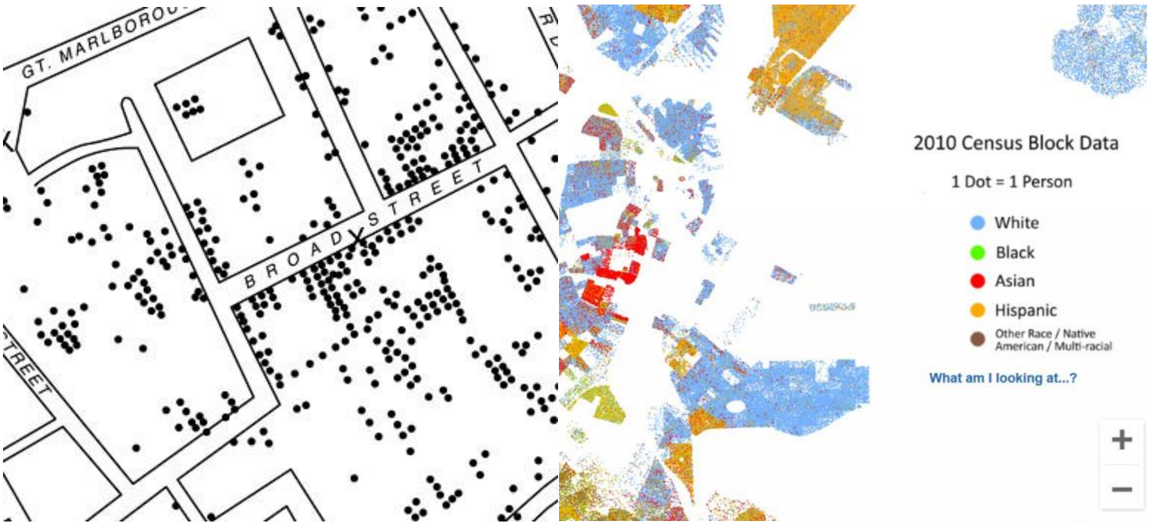

Figure 1. The classic dot map by John Snow (left) and the Racial Dot Map (right)

The popularity of interactive maps introduces a new methodological challenge, namely how to aggregate the dots. At the highest resolution, each individual data element can be represented as a dot. However, at lower resolutions, so when zoomed out, there are often not enough pixels available to plot all dots. Therefore, dots have to be aggregated somehow.

Our scope is unordered categorical data, such as ethnic background, gender, and marital status. Note that categorical data often have an explicit or implicit ordering, such as age class, level of education, and level of health. These data can also be visualized using a sequential or diverging colour scheme. Aggregating such data can be done by visualizing averages. However, in many such applications, it is interesting to observe the distribution of classes, for instance the composition of young, middle aged and elderly people in a certain neighbourhood.

We propose two different approaches. In the first approach we use blended colours. Although the dots are not distinguishable anymore, the blended colours could give a sense of the immensity of the data. The key aspect of this method is to how to pick the most appropriate colours such that no (blended) colour dominates the others. The second approach is to aggregate the dots explicitly to super dots. Here, we focus on two questions: how to aggregate the dots to super dots, and how to assess the quality of such aggregations?

2. BLENDED COLOURS

The analogy of using blended colours is to look at a pointillism style painting: nearby, the dots are distinguishable, but at a certain distance, only a gradual blending of colours can be perceived. Another analogy of a dot map using blended colours is a beach of sand, while composed of grain of sand, it looks like a yellowish shape on aerial pictures.

The Racial Dot Map [1] (Figure 1, right) also uses blended colours. The individual dots are rendered semi transparently. The main downside is that the order of rendering influences the colour of the pixel. Also, five categories (and thus five colours) are used, which makes it even harder to read the composition from the dot map. Although the Racial Dot Map is an effective visualisation, the used method may not a suitable in general.

2.1. Method

For the method we propose, we do not blend colours like we blend paint or light beams. Instead, we propose a method to pick the 'blended' colours from the Hue Luminance Chrome colour space (Figure 2). This colour space has nice perceptual properties [2], for instance, by varying the hue parameter, the luminance (brightness) and chroma (saturation) remain perceptually approximately constant.

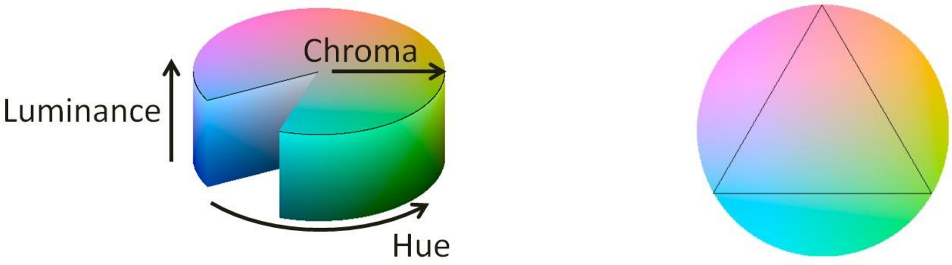

Figure 2. HCL colour space

The density is encoded with luminance; the darker the colour, the denser the dots. A dark pixel may therefore be interpreted as a stacking of dots, which is inevitable at lower zoom levels.

The composition is encoded with hue and chroma. Each category is assigned a hue value. If multiple categories are present, the pixel hue is a mixture of these hue values. The chroma indicates the dominance of a category; if only one category is present, the chroma is maximal, whereas if all categories are equally present, the chroma is zero, which results in grey. Although this method works for any number of categories, the best results are achieved with three categories: each category is assigned a colour at a corner of the triangle depicted in the Figure 2 (right). A mixture of categories results in a position in the triangle where the distance to the corners reflects the relative presence of the corresponding category.

2.2. Application

The application of a dot map using blended colours has been published as a CBS beta product [4, https://research.cbs.nl/colordotmap/EN]. At the highest resolution, each dot represents a person, where the colour represents the ethnic origin. This dot map is created with publicly available data on neighbourhood level, where the positions of the dots are uniformly distributed per neighbourhood [5].

A couple of improvements and experiments have been done regarding the level of neighbourhoods, the legend, and the disclaimer popups. These dot maps have not been published online yet. A screenshot of an improved version is provided in Figure 3. The right hand side image is a screenshot where individual dots can be distinguished. In the left hand side screenshot, blended colours are used.

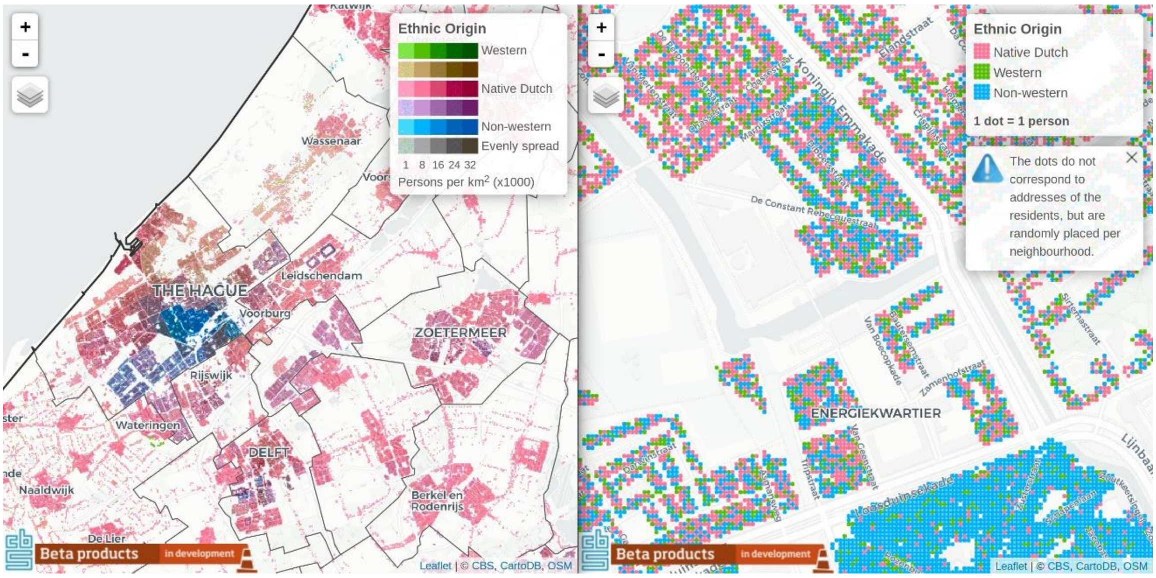

Figure 3. Dot map of the Dutch population by ethnic origin for two zoom levels.

2.3. User study

A user study has been done to assess the quality of the dot map [7]. Respondents were asked questions while interactively using the dot map. Audio has been recorded and transcribed, and eye moments has been captured with eye-tracking equipment (Figure 4, left).

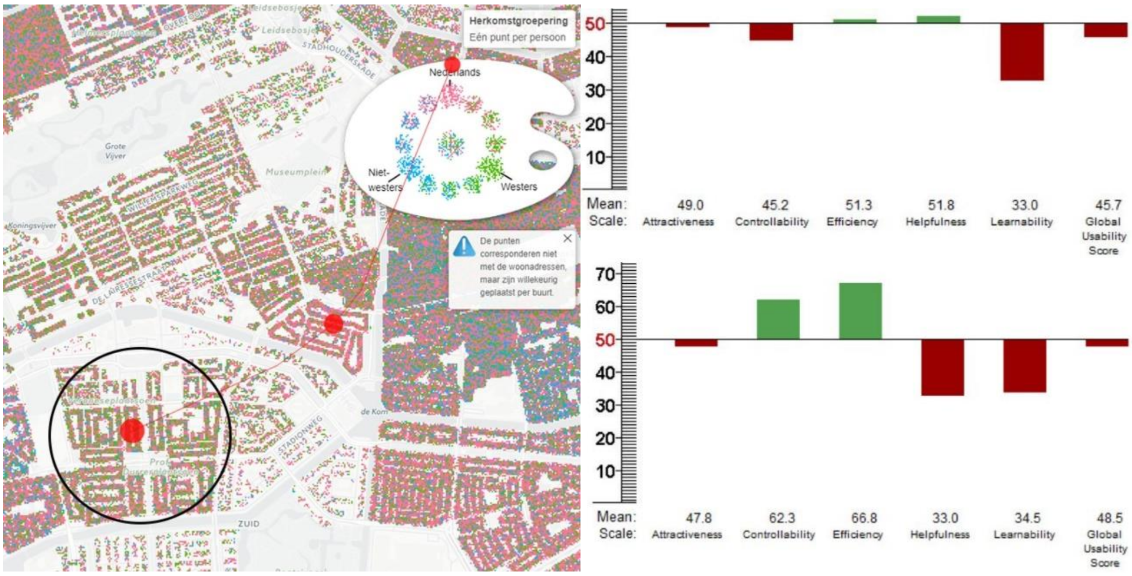

Figure 4. Left: screenshot during an eye-tracking session. The circle was drawn for specific questions about this area, and the path of red dots shows eye movement. Right: quality scores for the WAMMI questionnaire for two versions. The quality indicators shown from left to right are: attractiveness, controllability, efficiency, helpfulness, learnability, and global usability score.

The main results were:

- In general, the respondents were able to determine the composition correctly, while they had more trouble with density. Sometimes they experienced inconsistencies regarding the composition when zooming in.
- Respondents found the colour legend difficult to interpret.
- None of the respondents understood the fact that the dots were uniformly distributed per neighbourhood. Moreover, none of them read the disclaimer shown in Figure 3 (right). Instead, they thought that data on person level was shown.

The visualisation has also been tested with WAMMI, a website analysis service (www.wammi.com). The results for two versions are shown in Figure 4 (right). In general, the visualisation scores a little lower than the average visualisation tested using this service.

3. Super Dots

Another approach to aggregate the dots, is by creating 'super' dots which represent the small dots. For interactive dot maps, where the dots are placed on a grid, this means that 2 by 2 (or 4 by 4) dots are replaced by one large dot when zooming out. Since interactive maps (including Google Maps, Bing Maps, and Open Street Maps) scale by a factor of 2 by 2 at each zoom level, the size of the grid should preferable also be a factor of 2.

It is a methodological challenge how to determine the location and colour of these super dots. We would like to address two questions: which algorithm can be used to aggregate the dots, and how to access the quality of the result? We start with the last question, because an answer to that may provide useful insights in the development of algorithms.

3.1. Quality assessment

A small example of a dot map with four categories is shown in Figure 4, where at the top left, the original dot map is shown and at the top middle, the aggregated dot map. Conceptually, we thread white as a separate category, although for the dot map aggregation, white is less important than the other categories.

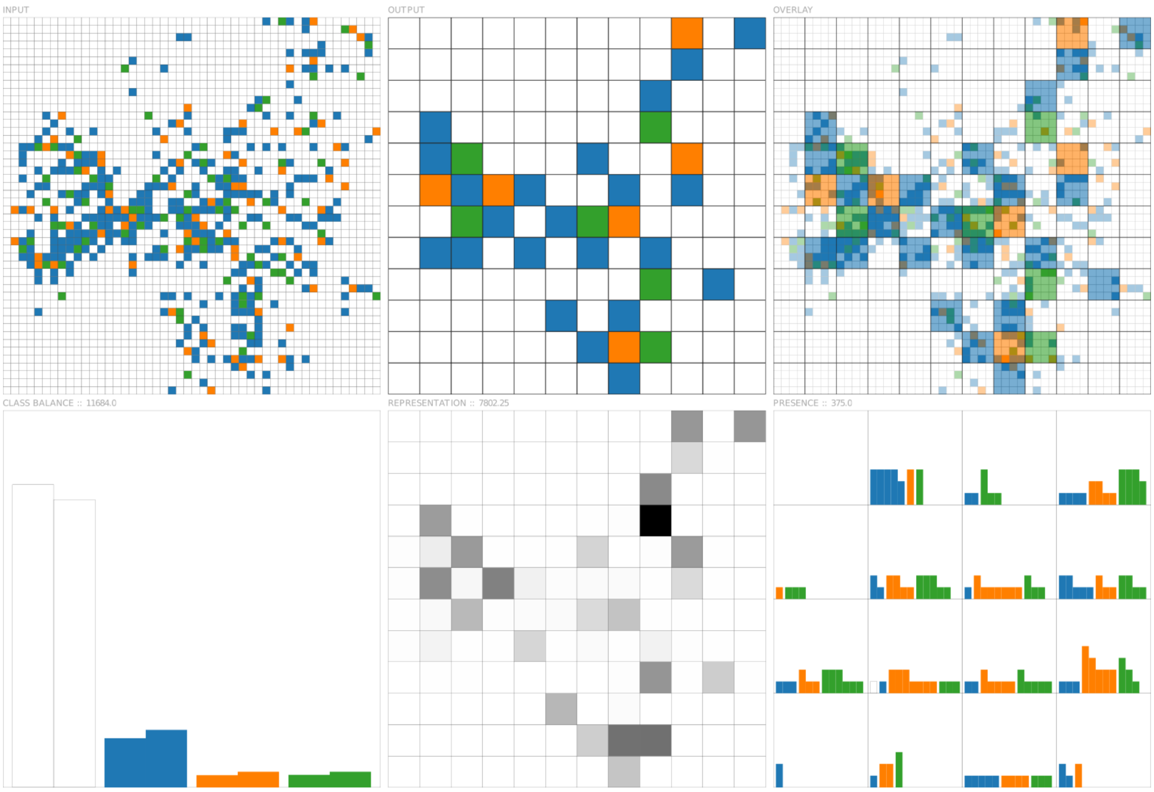

Figure 5. Dot map aggregation. Top row from left to right: original dot map, aggregated dot map, overlay. Bottom row from left to right: class balance, representation, presence.

We propose the following three dimensions to access the aggregation quality:

- **Colour balance**. The total number of super dots per class should represent the number of small dots. In the example shown in Figure 5, the number of super dots per category should ideally be equal to the number of small dots divided by 16 (since we apply a 4x4 grid cell aggregation). The colour balance is shown at the bottom left chart of Figure 5.
- **Representation**. How well do the super dots represent the small dots? Each small dot is represented once, and each super dot can represent at most 16 small dots. A heat map for representation is shown in the bottom middle chart in Figure 5. The darker the colour, the farther the represented small dots are away from the super dot.
- **Presence**. How well are the small dots represented by the super dots? For each small dot, the distance to the nearest super dot is measured. The bar chart at the bottom right of Figure 5 shows the total distances per class.

These quality dimensions are described in detail in [7].

3.2. Algorithms

The development of algorithms is still work in progress. We can already propose two algorithms, which still in development stage.

The first algorithm explicitly uses the quality concepts described in the previous paragraph:

1. Start with an empty map, and pick the class which the largest class imbalance. Since white is often less important than the other categories, a lower priority may be assigned to it.
2. For this class, find an empty spot for a super dot which achieves the best representation.

Another algorithm uses two-dimensional kernel densities:

1. For each class, estimate two-dimensional kernel densities.
2. Stack the kernel densities, and find the threshold value such that the number of grid cells (in the aggregated dot map) for which the stacked density is higher than this value, is equal to the desired total number of super dots.
3. For each grid cell for which the stacked density is higher than the threshold value, place a super dot, where the class is sampled from the classes with probabilities proportional to the density values.
4. Calculate the number of dots per class in the aggregated dot map. When the class balance is above a certain threshold, go to step 2 where the kernel densities are stacked with weights that are proportional to the class balance per class. Do this, until a certain number of iterations has been reached.

Before applying an aggregation algorithm, it is important to determine at which level(s) the small dots are aggregated and what the aggregation grid size is (i.e., a super dot may represent 2x2, 4x4, 8x8, or 16x16 small dots). It is hard to define general guidelines for that, since it depends on the application and the content of the data. Note that the aggregation grid size changes how many data units are represented by each dot. For instance, when a small dot represents 1 persons, a super dot represents 4, 16, 64, or 256 persons.

3.3. Application

The aggregation to super dots has been applied to a dataset with the number of children that follow primary education per neighbourhood, categorized by the distance between home and school into three categories: 0 to 500 meters, 500 to 1000 meters and 1000 meters or more. A screenshot of this dot map, which has not been published yet, is shown in Figure 6.

The dots are aggregated with a 2x2 grid. So each super dot represents 4 small dots. In this application, one dot represents 10 children when zoomed in (Figure 6, right), and 40 children when zoomed out (Figure 6, left). For the aggregation, the second aggregation algorithm, which uses two-dimensional kernel density estimators, has been applied.

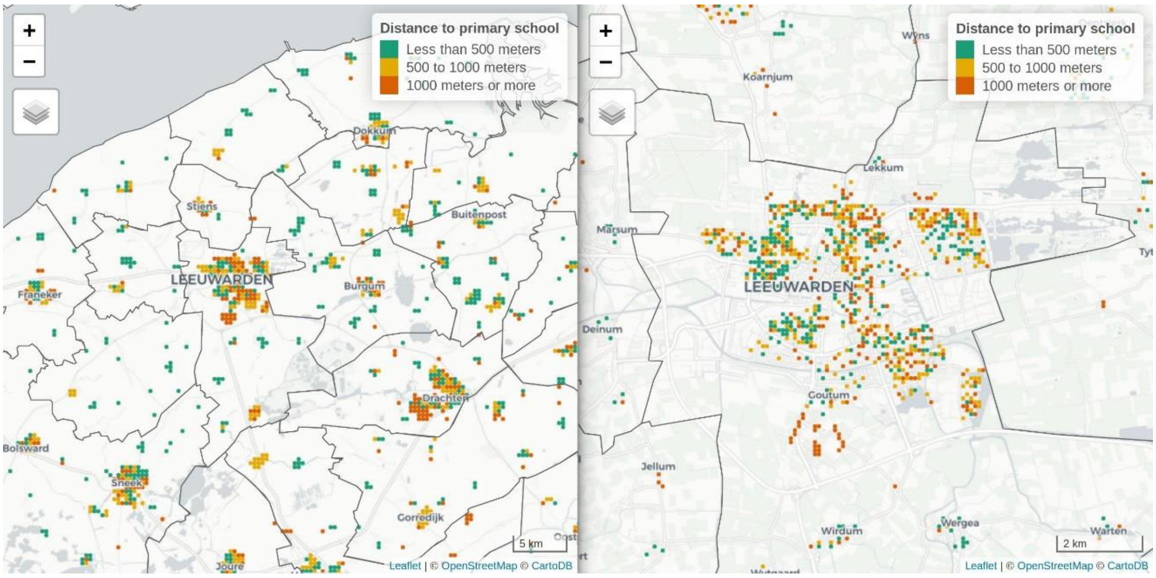

Figure 6. Dot map of children where the colour indicates how far they live from school. Each dot in the left hand side map represents 40 children, and in the right hand side map 10 children.

4. Conclusion and Discussion

The dot map is a powerful visualisation method to show the spatial distribution and composition of data. In principle, it is easy to understand in the sense that the location of a dot represents the location of a certain unit. The map by John Snow (Figure 1, left) is easy digestible once the application is clear.

The dot map can be made interactive [1, 4], which makes visualisation of big spatial data possible. The main problem is that number of units in the data easily exceeds the number of pixels when zoomed out. Therefore, the dots need to be aggregated for these zoom levels. This is far from trivial, as shown in this paper.

We have proposed two solutions: blended colours and the aggregation to super dots. Dots maps using blended colours will look more impressive; by zooming in and out, the immensity of the data can be experienced, similar to zooming in and out with an aerial camera directed at a full football stadium.

Super dots, on the other hand, are easier to read. Regarding density, it is easier to count super dots in a specific area than to read the luminance of a colour and use the legend to interpret it. Regarding composition, it is easier to perceive the composition of super dots than to read the hue and chroma of a colour and interpret it with the legend.

Both methods will work best with two or three categories. This is especially true for the blended colour method, since it is impossible to convert a blended colour back to more than three basic colours. Moreover, blended colours do not work for colour blind people. However, colours palettes with up to 8 colours are available that are colour blindness proof [8, 9]. Therefore, the methods of super dots can be used by colour blind people.

A downside with using super dots, is that it requires some effort to find and tweak a suitable algorithm for the application at hand. Different applications may require different parameter settings. For further research, we recommend to continue the development of aggregation algorithms and to provide general guidelines how to choose and use them.

Acknowledgements

I would like to thank Edwin de Jonge (CBS), Wouter Meulemans (TU Eindhoven), Chantal Melser (CBS), and François Engelen (Zuyd Hogeschool) for a pleasant and fruitful collaboration.

References

[1] D. Cable, The Racial Dot Map (2013). Online publication.
URL: http://www.coopercenter.org/demographics/Racial-Dot-Map

[2] A. Zeileis, K. Hornik, P. Murrell, Escaping RGBland: Selecting colours for statistical graphics, Comput. Stat. Data Anal. (2009) 53-9, 3259–3270.

[3] M. Tennekes, E. de Jonge, Colouring Interactive Compositional Dot Maps (2016) Poster paper presented at the 18th EG/VGTC Conference on Visualisation (EuroVis 2016), Groningen.

[4] Statistics Netherlands, CBS experimenting with dot maps (2017). Online publication https://www.cbs.nl/en-gb/our-services/innovation/project/cbs-experimenting-with-dot-maps

[5] M. Tennekes, E. de Jonge, The compositional dot map: a visualization of spatial data (2017) Paper presented at the 2017 New Techniques and Technologies for Statistics (NTTS) conference, Brussels.

[6] F. Engelen. Dot Map Evaluation (2018). Report Hogeschool Zuyd, Maastricht.

[7] W. Meulemans, M. Tennekes, Assessing Dot-Map Aggregations (2018). Poster paper to be presented at the VIS2018 conference, Berlin.

[8] C.A. Brewer, G.W. Hatchard, M.A. Harrower, ColorBrewer in Print: A Catalog of Color Schemes for Maps, Cartography and Geographic Information Science (2003), 30(1): 5-32.

[9] M. Okabe, K. Ito, Color Universal Design (CUD) - How to make figures and presentations that are friendly to Colorblind people (2002). URL= http://jfly.iam.u-tokyo.ac.jp/color/