Assessing dot-map aggregations

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
• The final author version and the galley proof are versions of the publication after peer review.
• The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

Download date: 08. Feb. 2024

assigned one of the data classes in a set $\mathscr{C}$. For a class $c \in \mathscr{C}$, we use $F(X,c)$ to denote the cells in matrix X that have class c. We use $\mathscr{C}$ also as a function, mapping an entry to its assigned class in $\mathscr{C}$.

2 MEASURING AGGREGATION QUALITY

We consider three dimensions to quantify aggregation quality: class balance, representation and presence.

Class balance To capture the overall composition, the relative number of dots of each color should remain the same. We can quantify deviation in various ways. We use a simple sum of squares on the differences, expressed in the formula below. We use $k^2 F(A,c)$ as each entry in A has the size of k^2 dots in the input map.

$$\sum_{c \in \mathscr{C}} \Big(|F(M,c)| - k^2 |F(A,c)| \Big)^2$$

Representation Each aggregated dot a in A represents up to k^2 input dots in M of the same class. Each smaller dot can be represented only once. That is, we assume that each aggregated dot has a set S_a containing up to k^2 dots of the same class, such that all sets S_a are pairwise disjoint. To capture spacial aspects, we want an appropriate representation: that is, a should be a good summary of those k^2 smaller dots in S_a. The further away the smaller dots are from a, the worse a represents them. We propose to quantify this via the sum of squared Euclidean distances, to penalize large distances:

$$\sum_{a \in A} \Big((k^2 - |S_a|)D^2 + \sum_{s \in S_a} \|a - s\|^2 \Big)$$

where $\|a - s\|$ indicates the Euclidean distance between the center of a and of s and D is a constant, penalizing aggregated dots that represent fewer that k^2 input dots. We can find the optimal sets S_a via a minimum-cost flow computation. The parameter D then ensures that the distance between a and a dot in S_a is at most D.

Presence Complementing representation, we ideally want that each input dot is present in the aggregated map in the form of a nearby aggregated dot. This particularly considers the effect of local minorities, which are not necessarily captured by the other dimensions. We measure the presence of a dot as its squared Euclidean distance to the nearest aggregated dot of the same class.

$$\sum_{m \in M} \min_{a \in F(A, \mathscr{C}(m))} \|m - a\|^2$$

We explicitly choose to not use a matching between small and large dots, to keep this dimension independent from representation. Specifically, a large dot may be the closest for more than k^2 small dots.

3 A VISUAL-ANALYTICS APPROACH

To support our reasoning and validate our measures, we developed a simple visual-analytics tool that will allow us to view, compute, analyze, and interact with dot maps and their aggregations.

Algorithm Starting from an empty map A, our approach iteratively picks the class c that minimizes $F(A,c)/F(M,c)$, i.e., that has the largest relative class imbalance. For this class c, it then finds the empty entry $a \in A$ that achieves the best representation if a is assigned class c. It sets the class of a and S_a accordingly[1] and removes S_a from M such that they cannot be used in next iterations.

[1]These computed S_a are not optimal, but can be recomputed afterwards.

Figure 1: Main view of our tool to explore dot map aggregation. Top row shows the dot maps, whereas the bottom row shows analysis results.

The visualization tool The main view of our tool is shown in Fig. 1. At the top are three views showing input, output and their overlay. Hovering dots allows linking and exploring the maps according to the sets S_a used to compute representation.

At the bottom are views that allow visual assessment of the three dimensions of aggregation quality. Each such view is an interactive small multiples, showing analogous graphics in an array of frames. The detail level (i.e., the number of small frames) can be configured, starting from a frame for each output dot to a single frame for the entire map. Each view is also annotated with the overall score in that dimension. From left to right, the views are used as follows:

A bar chart to show class balance (for the entire map): the left bar of a color indicates the fraction of dots in the input map of that class, the right bar shows the fraction of dots in the output map.

A heat map to show representation (each output dot individually): darker dots show worse representation. Values are scaled such that the worst representation is black.

A chart to show presence (frames match the underlying 3×3 output submatrices): the 12 input dots with the worst score are each shown as a bar, its height indicating the presence score for that dot.

4 DISCUSSION

A first observation is the special role of blank cells. We seem to intuitively focus on the colored dots that represent data, whereas the blanks are easily overlooked. While this is certainly desirable to some degree, it does cause a seemingly odd effect: in a region with many blanks, but also some colored dots, the best choice for class balance and representation (which are mostly driving our heuristic approach) will be to keep the blank as aggregated dot; for presence, however, it would be beneficial to inject more color into such places. We used this also in our algorithm, which can be configured to give a structurally lower priority to the blank class – this also is the reason

why the other three classes are overrepresented in the result of Fig. 1 (bottom left).

Generally, our three dimensions cannot always be optimized simultaneously as they measure conflicting aspects. Beyond validating our measures, we need to understand which dimensions are more important and to what degree.

Algorithms that optimize for different dimensions are necessary to further explore these trade-offs. Eventually, it would be worthwhile to design algorithms that use parameters that reflect these trade-offs. The question is how (if at all) we can efficaciously combine the dimensions into a single judgment on aggregation quality. Realistically, there is not going to be a single unique answer for all use cases. Rather, we need to understand the above in context of various conditions: how does one pick the appropriate algorithm to match the needs of the use case?

So far, we focused on just one input and output map to assess the aggregation quality. Another interesting dimension may reveal itself when we consider a sequence of maps, aggregated in increasingly higher levels: how do we balance the quality of all the aggregations in the sequence? Moreover, it will be important to consider the stability of such a sequence, to avoid visual artifacts for interactive dot maps where the aggregation level adapts according to zoom level.

REFERENCES

[1] D. Cable. The racial dot map. Published online http://www.coopercenter.org/demographics/Racial-Dot-Map, 2013.
[2] Statistics Netherlands. CBS experimenting with dot maps. Online publication https://www.cbs.nl/en-gb/our-services/innovation/project/cbs-experimenting-with-dot-maps, 2017.
[3] M. Tennekes and E. de Jonge. Coloring Interactive Compositional Dot Maps. In *EuroVis 2016 - Posters*, 2016.