

Computing Stable Demers Cartograms^{*}

Soeren Nickel¹, Max Sondag², Wouter Meulemans², Markus Chimani³,
Stephen Kobourov⁴, Jaakko Peltonen⁵, and Martin Nöllenburg¹

¹ Algorithms and Complexity Group, TU Wien, Vienna, Austria
`soeren.nickel@tuwien.ac.at`, `noellenburg@ac.tuwien.ac.at`

² TU Eindhoven, Netherlands
`m.f.m.sondag@tue.nl`, `w.meulemans@tue.nl`

³ University of Osnabrück, Germany
`markus.chimani@uni-onsabrueck.de`

⁴ University of Arizona, Tucson, AZ, USA
`kobourov@cs.arizona.edu`

⁵ Tampere University, Finland
`jaakko.peltonen@tuni.fi`

Abstract. Cartograms are popular for visualizing numerical data for map regions. Maintaining correct adjacencies is a primary quality criterion for cartograms. When there are multiple data values per region (over time or different datasets) shown as animated or juxtaposed cartograms, preserving the viewer’s mental-map in terms of stability between cartograms is another important criterion. We present a method to compute stable Demers cartograms, where each region is shown as a square and similar data yield similar cartograms. We enforce orthogonal separation constraints with linear programming, and measure quality in terms of keeping adjacent regions close (cartogram quality) and using similar positions for a region between the different data values (stability). Our method guarantees ability to connect most lost adjacencies with minimal leaders. Experiments show our method yields good quality and stability.

Keywords: Time-varying data · Cartograms · Mental-map preservation

1 Introduction

Myriad datasets are georeferenced and relate to specific places or regions. A natural way to visualize such data in their spatial context is by cartographic maps. A choropleth map is a prominent tool, which colors each region in a map by its data value. Such maps have several drawbacks: data may not be correlated to region size and hence the visual salience of large vs small regions is not equal.

^{*} This research was initiated at NII Shonan Meeting 127 “Reimagining the Mental Map and Drawing Stability”. M. Sondag is supported by The Netherlands Organisation for Scientific Research (NWO) under project no. 639.023.20. S. Kobourov is supported by NSF grants CCF-1740858, CCF-1712119 and DMS-1839274. M. Nöllenburg is supported by FWF grant P 31119.

Moreover, colors are difficult to compare and not the most effective encoding for numeric data [23], requiring a legend to facilitate relative assessment.

Cartograms, also called value-by-area maps, overcome the drawbacks by reducing spatial precision in favor of clearer encoding of data values: the map is deformed such that each region’s visual size is proportional to its data value. Attention is then drawn to items with large data values and comparison of relative magnitudes becomes a task of estimating sizes – which relies on more accurate visual variables for numeric data [23]. This also frees up color as a visual variable. Cartogram quality is assessed by criteria [25] including **1. Spatial deformation:** regions should be placed close to their geographic position; **2. Shape deformation:** each region should resemble its geographic shape; **3. Preservation of relative directions:** spatial relations such as north-south and east-west should be maintained. **4. Topological accuracy:** geographically adjacent regions should be adjacent in the cartogram, and vice versa. **5. Cartographic error:** relative region sizes should be close to the data values. Criteria 1-4 describe geographical accuracy of the region arrangement. Maintaining relative directions also helps preserve a viewer’s spatial mental model [30] Criterion 5 (also called statistical error) captures how well data values are represented. Often techniques aim at zero cartographic error sacrificing other criteria.

Cartograms can also be effective for showing different datasets of the same regions, arising from time-varying data such as yearly censuses yielding temporally ordered values for each region, or from available measurements of different demographic variables that we want to explore, compare and relate, yielding a vector or set of values for each region. Visualizations for multiple cartograms include animations (especially for time series), small multiples showing a matrix of cartograms, or letting a user interactively switch the mapped value in one cartogram. See for example the interactive Demers cartogram accompanying an article from the New York Times⁶. In such methods, cartograms should be as similar as the data values allow: we thus want cartograms to be *stable* by using similar layouts. This helps retain the viewer’s mental map [22], supporting linking and tracking across cartograms. Thus, we obtain an important criterion with multivariate or time-varying data. **Stability:** for high stability, cartograms for the same regions using different data values should have similar layouts. The relative importance of the criteria depends on the tasks to be facilitated. Nusrat and Kobourov’s taxonomy of ten tasks [25] can also be considered with multiple cartograms. Many tasks focus on the data values. As such, a representation of a region of low complexity allows for easier estimation and size comparison.

Contribution. We focus on *Demers cartograms* (DC; [3]) which represent each region by a suitably sized square, similar to Dorling cartograms [9] which use circles. Their simplicity allows easy comparison of data values, since aspect ratio is no longer a factor, unlike, e.g., for rectangular cartograms [19]. However, as abstract squares incur shape deformation, in spatial recognition tasks the cartogram embedding as a whole must be informative, so the layout must op-

⁶ <https://archive.nytimes.com/www.nytimes.com/interactive/2008/09/04/business/20080907-metrics-graphic.html>, accessed June 2019.




Fig. 1. Cartograms displaying drug poisoning mortality, total GDP and population of the contiguous states of the US in 2016. The layout minimizes distance between adjacent regions. Lost adjacencies are indicated with red leaders. Color is used only to facilitate correspondences between the cartograms.

timize as much as possible the other geographic criteria: *spatial deformation*, *preservation of relative directions* and *topological accuracy*. We contribute an efficient linear programming algorithm to compute high-quality stable DCs. Our DCs have no cartographic error, satisfy given constraints on spatial relations, and allow trade-off between topological error and stability. Linear interpolation between different DCs yields no overlap during transformation. Lost adjacencies—satisfying a mild assumption—can be shown as minimal-length planar orthogonal lines. Fig. 1 shows examples. Experiments compare settings of our linear program to each other and to a force-directed layout we introduce (also novel for DCs); results show that our linear program efficiently computes stable DCs.

Related work. Cartogram-like representations date to the 1800s. In the 1900s most standard cartogram types were defined, including rectangular value-by-area cartograms [26] and more recent ones [13, 19]. The first automatically generated cartograms are continuous deformation ones [29] followed by others [12, 15]. Dorling cartograms [9] and DCs [3] exemplify the non-contiguous type representing regions by circles and squares respectively. Layouts representing regions by rectangles and rectilinear polygons have received much attention in algorithmic literature, see e.g. [1, 7, 11], and typically focus on aspect ratio, topological error and region complexity. Compared to DCs, rectilinear variants have higher visual complexity and added difficulty to assessing areas. No cartogram type can guarantee a both statistically and geographically accurate representation; see a recent survey [25]. Measures exist to evaluate quality of cartogram types and algorithms, see e.g. [2, 17].

There is little work on evaluating or computing stable cartograms for time-varying or multivariate data. Yet they are used in such manner, e.g., as a sequence of contiguous cartograms showing the evolution of the Internet [16].

DCs relate to contact representations, encoding adjacencies between neighboring regions as touching squares. The focus in graph theory and graph drawing literature lies on recognizing which graphs can be perfectly represented. Even the unit-disk case is NP-hard [5], though efficient algorithms exist for some restricted graph classes [8]. Klemz et al. [18] consider a vertex-weighted variant using disks, that is, with varying disk sizes. Various other techniques are similar

to DCs, using squares or rectangles for geospatial information. Examples include grid maps, see, e.g., [10] for algorithms and [21] for computational experiments. Recently, Meulemans [20] introduced a constraint program to compute optimal solutions under orthogonal order constraints for diamond-shaped symbols. We use similar techniques, but refer to Section 2 for a discussion of the differences.

2 Computing a single DC

First, we consider a DC for a single weight vector. We are given a set of weighted regions with their adjacencies, and a set of directional relations. We compute a layout realizing the weights with disjoint squares that may touch only if adjacent, so that directional relations are “roughly” maintained. We quantify the quality of the layout by considering the distances between any two squares representing adjacent regions. We show that the problem, under appropriate distance measures, can be solved via linear programming in polynomial time.

Formal setting. We are given an input graph $G = (R, T)$. For each *region* $r \in R$ we are given its centroid in \mathbb{R}^2 and its weight $w(r)$, the side length of the square that represents it in output. The graph has an edge in T if and only if the original regions are adjacent, thus their respective squares in the output *should* be adjacent as well. We are also given two sets H, V of ordered region pairs. A pair (r, r') is in H , if r should be horizontally separated from r' such that there exists a vertical line ℓ with the square of r being left of ℓ and r' to its right. Analogously, V encodes vertical separation requirements. If r and r' are adjacent, then (r, r') is either in H or in V (but not in $H \cap V$) and they *should* touch ℓ , otherwise we require a strict separation to avoid false adjacencies; we are given a minimum gap ε to ensure that this non-adjacency can be visually recognized.⁷ The sets H and V model the relative directions criterion for DCs and any two regions are paired in at least one of those sets. To ensure a DC exists satisfying the separation constraints, the directed graph $D = (R, H \cup V)$ must be a directed *acyclic* graph (DAG). We consider these relations transitive: if $(r, r') \in H$ and $(r', r'') \in H$, then this enforces that there exists a vertical line separating (r, r'') in any DC and thus (r, r'') is in H .

The output—a placement of a square for each region—can be stored as a point $P: R \rightarrow \mathbb{R}^2$ for each region, encoding the center of its square. A placement P is *valid*, if it satisfies the separation constraints of H and V . This implies all squares are pairwise interior disjoint (or fully disjoint for nonadjacent regions). We look for a valid placement where distances between non-touching squares of originally adjacent regions are minimized; this will be made more precise below.

Deriving separation constraints. The regions’ weights are given and their adjacencies and centroids easily derived, but separation constraints H and V are not. Various models can determine good directions or separation constraints [6]. We use the following model; it is symmetric and ensures constraints form a DAG.

⁷ In the implementation, ε is the minimum of the side length of the smallest region and 5% of the diagonal of the bounding box of the input regions R .

For two regions (r, r') represented by centroids, we check whether their horizontal or vertical distance is larger. In the former case, we add (r, r') to H if r is left of r' and (r', r) to H otherwise. In the latter case, we add the pair to V in the appropriate order. We call this the *weak setting*. We call constraints added in this setting *primary separation constraints*.

In the *strong setting*, we may add an extra constraint for nonadjacent region pairs whose bounding boxes admit both horizontal and vertical separating lines: if a pair has a primary separation constraint in H or V , we add a *secondary* separation constraint to V or H respectively.

Linear Program. We model optimal solutions to the problem via a polynomially-sized linear program (LP), which lets us solve the problem in polynomial time. For each $r \in R$, we introduce variables x_r and y_r for the center $P(r) = (x_r, y_r)$ of the square. For any originally adjacent regions $\{r, r'\} \in T$ we introduce variables $h_{r,r'}$ and $v_{r,r'}$ for the (non-negative) distance between two squares. For any two regions r, r' , we define shorthands: let $w_{r,r'} := \frac{(w(r)+w(r'))}{2}$ and let $gap_{r,r'} = \varepsilon$ if $\{r, r'\} \notin T$, and 0 otherwise.

$$\min \sum_{\{r,r'\} \in T} h_{r,r'} + v_{r,r'} \quad (1)$$

$$x_{r'} - x_r \geq w_{r,r'} + gap_{r,r'} \quad \forall (r, r') \in H \quad (2)$$

$$y_{r'} - y_r \geq w_{r,r'} + gap_{r,r'} \quad \forall (r, r') \in V \quad (3)$$

$$h_{r,r'} \geq \max\{(x_r - x_{r'}) - w_{r,r'}, (x_{r'} - x_r) - w_{r,r'}\} \quad \forall \{r, r'\} \in T \quad (4)$$

$$v_{r,r'} \geq \max\{(y_r - y_{r'}) - w_{r,r'}, (y_{r'} - y_r) - w_{r,r'}\} \quad \forall \{r, r'\} \in T \quad (5)$$

$$h_{r,r'}, v_{r,r'} \geq 0 \quad \forall \{r, r'\} \in T \quad (6)$$

The objective (1) minimizes a sum of the distances between regions with broken adjacencies in the L_1 metric. Constraints (2) and (3) ensure separation requirements by forcing square centers far enough apart. For nonadjacent regions, the *gap* function assures a recognizable gap of width ε between resulting squares. Constraints (4)–(6) bind distance variables h, v with positional variables x, y . Here (4) and (5) encode two linear constraints per line, one for each term in the ‘max’ function. As (1) minimizes the distances, it suffices to enforce lower bounds, hence the ‘ \geq ’ in the constraints. In an optimal solution, either one of the two versions, or the non-negativity constraint (6) will be satisfied with equality.

Improving the gaps. The above model has two minor flaws. *First*, two squares ‘touch’ even if they only do so at corners; we resolve this by adding ε to the right-hand side of (4) (or (5)) for vertically (or horizontally, respectively) separated region pairs in T . This allows $h_{r,r'} = 0$ ($v_{r,r'} = 0$), when squares share a segment at least ε long. *Second*, in the strong setting the LP asks for a minimum gap ε along both axes. This is not needed for visual separation, so we remove the gap requirement from the secondary separation constraint.

Fine-tuning the optimization criteria. The LP minimizes a sum of distances between adjacent regions. Cartogram literature emphasizes counting lost adjacencies between regions, not the distance between them. We prefer our measure since

1) there is a big difference if two neighboring countries are set apart by a small or large gap; 2) while the LP can be turned to an *integer* linear program to count lost adjacencies, it greatly increases computational complexity—optimizing for adjacencies is typically NP-hard, e.g., for disks [4, 5] or segments [14].

Our linear program typically admits several optimal solutions, due to translation invariance and since touching squares may slide freely along each other as long as they touch. We introduce a *secondary* term to the objective to nuance selection of better layouts, multiplied by a *small* constant to not interfere with the original (primary) objective. The secondary term optimizes preservation of relative directions between squares within the freedom of the optimal solution.

Consider regions r and r' . W.l.o.g., assume their original centroids are horizontally farther apart than vertically, and r is left of r' , so $(r, r') \in H$. We compute a directional deviation $d_{rr'} = |(y_r + \alpha(x_{r'} - x_r)) - y_{r'}|$, where α is the (finite) slope of the ray from r to r' in the input graph G . Similar to (4), the objective function will minimize $d_{rr'}$; we weigh this term more heavily for adjacent regions. We thus turn the above formula into two linear inequalities.

Alternatives exist for the secondary criterion: displacement from the original location helps find layouts maintaining many adjacencies for grid maps of equal-size squares [10, 21]. For each region we measure L_1 displacement from its *origin* (centroid of the original region in the geographic map) to the square center $P(r)$.

Comparison to overlap removal. A technique placing disjoint squares exists to remove overlap of diamond (45 degree rotated square) glyphs for spatial point data [20], asking to minimally displace varying-size diamonds to remove all overlap, constrained to keep orthogonal order of their centers. Rotating the scenario to yield squares does not yield axis-parallel order constraints but “diagonal” ones, different from our strong setting. A “weak order constraints” variant is mentioned, related to our LP in the weak setting, if we change our objective to one only optimizing displacement relative to original locations. Fig. 2 shows similarities and differences considering the feasibility area between two regions. Extensions in [20] can be applied in our scenario, e.g., reducing actively considered separation constraints by removing transitive relations (“dominance” in [20]). Time-varying data is briefly considered in [20], only conceptualizing a trade-off between origin-displacement and stability for artificial data; we discuss several optimization criteria, also focusing on adjacencies which are not considered in [20], use real-world data experiments, and compare to a baseline DC implementation to move beyond the limits of linear programming.

The lemma below matches an observation from [20] that carries over to our setting. It implies that cartograms for different weight functions but with the same constraints have a smooth and simple transition between any such DCs helping to retain the user’s mental map.

Lemma 1. *Let R be a set of regions with separation constraints H and V . Let A and B be two DCs for R , both satisfying H and V . Then, any linear interpolation between A to B also satisfies H and V and is thus overlap-free.*




Fig. 2. Feasibility area where r' may be placed w.r.t r , when r' is primarily to the right of r . (a) In terms of feasibility, our weak setting and weak order constraints in [20] coincide. (b) Feasibility using a rotated orthogonal order [20]. (c) Feasibility in the strong setting.

3 Computing stable DCs for multiple weights

The method can be extended for regions having multiple weights. We are given a set of weight functions $W = \{w_1, \dots, w_k\}$. We aim to compute a DC for each $w_i \in W$, i.e., positions $P_i(r)$ for each $r \in R$ and $w_i \in W$. If each weight function represents the same data semantic, say population size, at different times, we consider $W = \{w_1, \dots, w_k\}$ ordered by the k time steps; we call this setting *time series*. If each weight function represents measurements of different data semantics (possibly at the same time), say population and gross domestic product, we treat W as an unordered set; we call this setting *weight vectors*.

As we focus on cartogram stability over multiple datasets, we combine the weight functions into one LP that computes the set of DCs, with potentially different centers $P_i(r)$ for each region r and weight function w_i . This lets us add constraints and optimization objectives for stability. We change objective (1) and add constraints to minimize displacement between centers of the same region for different weight functions. We re-use notation in Section 2 with superscript i denoting respective variables for weight function $w_i \in W$.

$$\min \sum_{i=1}^k \sum_{\{r, r'\} \in T} (h_{r, r'}^i + v_{r, r'}^i) + \sum_{\{i, j\} \in I} \sum_{r \in R} (c_r^{i, j} + d_r^{i, j}) \quad (7)$$

$$c_r^{i, j} \geq \max\{(x_r^i - x_r^j), (x_r^j - x_r^i)\} \quad \forall r \in R, \{i, j\} \in I \quad (8)$$

$$d_r^{i, j} \geq \max\{(y_r^i - y_r^j), (y_r^j - y_r^i)\} \quad \forall r \in R, \{i, j\} \in I \quad (9)$$

Here set I contains index pairs of weight functions $\{w_i, w_j\}$ for which displacement should be minimized. For each $r \in R$, variables $c_r^{i, j}$ and $d_r^{i, j}$ measure the horizontal and vertical displacement between $P_i(r)$ and $P_j(r)$ due to (8) and (9).

For which weight functions to relate in I , we consider two options: 1) relate *all pairs* of functions so $I = \binom{W}{2}$, which is natural for weight vectors where an analyst may want to compare the DCs for any two weight functions; and 2) relating *consecutive* pairs in a predefined order of the functions so $I = \{(i, i+1) \mid 1 \leq i \leq k-1\}$, which is natural for time series. An alternative 3) for time series initially computes a DC for w_1 (e.g., minimizing displacement to region centroids

in the initial map) and then iteratively solves the LP for one DC and weight function w_i ($i \geq 2$), where we minimize the displacement only with respect to the previously solved DC for weight function w_{i-1} . Due to its restricted solution space 3) is expected to be faster to solve than 2) but with lower stability. In some scenarios another option 4) may be worthwhile: one weight function, say w_1 , may be considered central to the dataset and displacements are only considered relative to it, so I contains pairs $\{1, i\}$ for all $2 \leq i \leq k$.

Not all planar graphs can be represented using touching squares of any size. A real-world example is Luxembourg having three pairwise neighbors; the input graph G is a K_4 . Thus any DC may need to break some adjacencies. To show lost adjacencies we use *leaders* – orthogonal polylines connecting the two squares. We want leaders to have minimal length and low complexity which we can guarantee under mild assumptions: 1) leaders can coincide with square boundaries; 2) regions to be connected are *realisable*, i.e., a valid DC (with possibly different weights) exists for each pair of regions such that they are adjacent. Let $L_1^B(r_1, r_2)$ denote the minimal L_1 distance between squares of regions r_1 and r_2 in DC B . The following lemmas are proven in Appendix A in the full version [24] – the proof of the first is constructive and gives a simple $O(n^2)$ algorithm to compute all leaders.

Lemma 2. *Consider DCs with separation constraints H, V and two regions $\{r_1, r_2\} \in T$. Let (r_1, r_2) be a minimal pair in H or V . Then, in any DC B , there is a monotone leader ℓ between r_1 and r_2 with length $L_1^B(r_1, r_2)$.*

Lemma 3. *Let $\{r_1, r_2\} \in T$ and assume a DC A exists with r_1 and r_2 adjacent, from which H and V are derived in the strong setting. Then, for any DC B satisfying H and V , a leader ℓ exists between r_1 and r_2 with at most two bends.*

4 Experimental setup

We compare 18 variants of our linear programs with each other and to 4 variants of a baseline force-directed DC layout implementation, as described below.

Linear programs. We categorize our method according to three criteria: A) optimization term, B) method of deriving constraints, and C) how we deal with different time steps. For A) our linear program admits three primary optimization terms: *TOP* – distance between topologically adjacent regions; *CNT* – number of lost adjacencies; *ORG* – distance to the origin (region’s centroid in the geographic map). We use the indicated primary optimization term, complemented by the secondary constraint of maintaining relative directions. For B), separation constraints are deduced from the input map in one of two ways, S and W , matching the strong and weak case respectively. For C), we deal with different weight values (time series/weight vectors) in three ways called *stability implementations*: *CO* – we add an optimization term to minimize distance between layouts of all (complete) weight value pairs; (2) *SU* – we add an optimization term to minimize distance between layouts of successive weight values; (3) *IT* –

we iteratively solve a linear program including an optimization term to minimize distance to previously calculated layouts. We specify our methods by concatenating the three aspects in order, for example, TOP-S-SU indicates the linear program optimized for distances of topologically adjacent regions with strong separation constraints and with successive weight values linked.

Force-directed method. DCs are hard to track down in literature, especially regarding computation. To our knowledge, there is no common baseline for computing a DC; we introduce a simple one. As Dorling cartograms and DCs are similar [3] and Dorling cartograms use a force-directed method, we implement one here, too: FRC. For each pair of regions we define a disjointness force based on Chebyshev distance between their centers, which grows quadratically to push squares apart. We use the same desired distance as in Section 2 at which this force becomes zero. We also add a force for cartogram quality, either towards their original locations (FRC-O) or between adjacent regions (FRC-T). We initialize the process with map locations (U; unstable) or the result for previous weights (S; stable). See Appendix B in the full version [24] for more details.

Metrics: cartogram quality. Our algorithms inherently yield zero cartographic error, and shape deformation is constant over all possible DCs. To evaluate cartogram quality we use three metrics, each normalized between 0 and 1; smaller values are better. We measure topological accuracy as the number of lost adjacencies (**MADJ**) in each of the k computed layouts, normalized by the number of adjacencies $k|T|$. To measure preservation of relative directions (**MREL**) with respect to the input map, we use the Relative Position Change Metric [28] which captures the preservation of the spatial mental model (orthogonal order) in a fine-grained way. Each rectangle defines eight zones by extending its sides to infinite lines. Between a pair of input map regions (r, r') we consider fractions of the bounding box that fall into each zone; if bounding boxes overlap, we scale values so they sum to 1. We do the same between the corresponding squares in the cartogram layouts. The measure between two regions is half the sum over all absolute differences between fractions per zone; the value is in $[0, 1]$ but is not symmetric. Finally, we take the average over all pairs. For spatial deformation we measure distance to map origins (**MDIS**), average L_1 distance of each region r in the DC to its origin (centroid of r in the geographic map), normalized by dividing with the L_1 distance of the diagonal of the map.

Metrics: stability. We also want to assess stability, or layout similarity, between the DCs by two quality metrics, based on treemap stability metrics [28], interpreting DCs as special treemaps with added whitespace. The first is based on geometric distances between the layouts: the layout distance (**SDIS**) focuses on the change in position of the squares. The layout distance change function as presented by Shneiderman and Wattenberg [27] is the most common one. It measures Euclidean distance between rectangles r and r' . We take the average over all pairs, and normalize by dividing with the L_1 distance of the largest diagonal of the two DCs. The result is related to our optimization term for quality when dealing with multiple weights (see Section 3). The second metric, relative

directions between layouts (**SREL**), focuses on changes in relative directions; it is analogous to MREL, but compares two layouts instead.

Datasets. We run experiments on real-world datasets. For time-series data, we expect a gradual change and strong correlation between the different values. For weight-vectors data, we expect more erratic changes and less correlation. We use two maps with rather different geographic structures: the first (**World**) is a map of world countries, having mixed region (country) sizes in a rather unstructured manner; the second (**US**) is a map of the 48 contiguous US states, having relatively high structure in sizes of its states, with large states in the middle and along the west coast and many smaller states along the east coast. We collected five time series for the World and four for the US map of which the details are given in Appendix C in the full version [24]. We transformed these into a weight-vectors dataset by taking the values of 2016 for each of these time series, resulting in five weight vectors for the World map, and four for the US map.

The various datasets have different scales, and need be projected into a reasonable square size to compute a DC. We compute the diagonal Δ of the bounding box of the map. For a time-series dataset, we find the region r with maximal $w_i(r)$ for any i and scale values such that $w_i(r) = \Delta/4$. For a weight-vectors dataset, we do the same, but scale the values for each DC separately.

Running times. We ran the experiments using IBM ILOG CPLEX 12.8 to solve the (I)LP. We observe the following running times on a normal laptop: *-*-IT and FRC-O-* finished within seconds (USA) or a minute (World); *-*-{SU,CO} took around a minute (USA) or below 5 minutes (World); FRC-T-* was completed in minutes (USA) or hours (World). CNT-*-* is an *integer* linear program rather than a regular linear program (or force-directed method); its computational complexity is significantly higher, and intractable in many cases. Only CNT-*-IT variants were successfully solved, and only on the US map; for all other cases it ran out of memory (48 GB allocated).

5 Experimental results

We discuss results and four questions: 1) How much does the strong versus weak setting affect quality? 2) How much does stability implementation matter? 3) Which optimization criteria perform best? 4) What is the effect of separation constraints in our LP, compared to a force-directed method for DCs? Fig. 3 shows the result of two algorithms for the US. Appendix D and the supplementary video in the full version [24] show more DCs for different settings.

Strong versus weak setting. Fig. 4 shows the average metric values for the iterative variants, over all datasets and linear programs. We find that the strong case (additional separation constraints) reduces the error in relative direction for both cartogram quality and stability: the average score for MREL, including CNT variants where possible, reduces from 0.21 to 0.16; similarly, stability (SREL) decreases from 0.059 to 0.045 due to decreased movement freedom of the




Fig. 3. US election turnout data DC in 2016, by TOP-W-CO and CNT-W-IT.




Fig. 4. Bar chart of average metric scores, for IT settings of all linear programs and the FRC directed variants. We see similar effects when switching from the weak version to the strong version of the IT setting for all three optimization settings. We also see a strong effect when choosing different optimization settings for the IT setting. FRC is generally outperformed by the {TOP,ORG}-W-IT variants.

squares. This is at the expense of topological error (MADJ increases from 0.58 to 0.61) and origin displacement (MDIS increases from 0.16 to 0.17). The effect is present independent of optimization criterion and stability implementation though its strength varies. Effects remain noticeable but of varying strength when we control for type of dataset, except MDIS slightly decreases for US datasets (0.116 to 0.107) in the strong setting. We also see a clear difference between optimization terms (CNT, TOP, ORG), discussed later.

Stability implementation. In time-series datasets there is little difference in stability over the three settings: time series data change gradually over time so choosing which pairs to optimize does not have much influence. In weight-vectors datasets, even with only few weights per region (five for the World, four for the US), an effect becomes noticeable in the IT setting. CO and SU behave nearly identical, but this might be an artifact of only having a few weights per region. Compared to CO (and SU) setting, the iterative version scores better on MDIS (0.31 versus 0.26) but worse on the stability metric SDIS (0.084 versus 0.10). For weight-vectors datasets it is thus better to use the SU variant as this achieves better stability and is only slightly more expensive to compute compared to IT variants. The added complexity of CO does not seem to pay off.

Optimization criteria. We use three metrics for cartogram quality: MADJ and MDIS are optimized explicitly with the CNT and ORG objectives respectively, the third metric MREL corresponds to a secondary objective term. To compare the TOP/CNT/ORG objective terms, we consider the IT variant (see Fig. 4), as other stability implementations could not solve the CNT objective; still, we found similar patterns for the SU and CO cases.

For MADJ, CNT finds the optimal value (0.31) under the given constraints. TOP (0.57) does clearly better than ORG (0.70), somewhat in contrast to observations of [10, 21]: for grid maps, the MDIS metric that ORG optimizes is a good proxy for maintaining topology; our results suggest this is not so for DCs.

For MDIS and MREL metrics ORG performs best; for MDIS, CNT performs slightly better compared to TOP and vice versa for MREL. Thus, in terms of spatial quality, ORG seems a good objective, except for topological error – which is typically of primary concern for cartograms.

For stability metrics SDIS and SREL, ORG outperforms TOP which outperforms CNT. We explain it by inherent stability of the map which is the same for all DCs. CNT does poorly; it is fairly unconstrained for lost adjacencies whereas TOP aims to keep such pairs close.

ORG scores best on all metrics except MADJ; its MADJ score is high, losing 70% of adjacencies on average. In contrast, CNT optimizes the number of adjacencies, but is clearly worse on other metrics and is computationally expensive. There is thus a trade-off present between topological error and other quality aspects. TOP makes this trade-off, scoring reasonably on most metrics.

Comparison to FRC. Our linear programs enforce separation constraints which help maintain spatial relations and the spatial mental model; they are required for the linear program but not in general. To study their effect, we compare to FRC which does not enforce separation constraints; results are shown in Fig. 4. Comparing FRC-T and FRC-O variants, we see the same behavior as in the TOP versus ORG linear programs: FRC-O performs worse than FRC-T on ADJ, and better on the other metrics. Layout initialization trades off stability versus cartogram quality: FRC-*-S variants have better stability scores and worse quality scores compared to FRC-*-U.

As it has the fewest constraints, we compare ORG-W-IT to FRC methods: FRC-O-* are slightly worse or equal to ORG-W-IT on all metrics; FRC-T-* are worse than ORG-W-IT on all metrics except ADJ where it is a little better, but the number of adjacencies lost is still clearly higher compared to TOP-W-IT.

To conclude, in general we outperform FRC for the various metrics by an appropriate setting in our linear program. No single setting outperforms all FRC variants. The large difference with TOP-variants in terms of MADJ suggests TOP variants are a good choice for high-quality stable DCs.

6 Discussion and future work

We described a linear program to compute stable Demers cartograms, based on separation constraints and minimizing distance between adjacent regions.

It allows overlap-free transitions between weight functions and connecting lost adjacencies with short, low-complexity leaders. Experiments show it offers a good trade-off between topological error and other criteria. It outperforms basic force-directed layouts, though there is not a unique variant that does so, suggesting an interplay between separation constraints, optimization and quality metrics.

In future work we may consider stability in other cartogram styles, and perform human-centered comparisons in addition to computational ones, with methods implemented in interactive systems; such systems can, e.g., emphasize adjacent regions by drawing leaders (at all or more clearly) or link regions back to the geographic map. We focused on Demers cartograms, but there are many different styles of cartograms. Future work may also investigate stable variants of such other cartogram styles and quantitatively or qualitatively compare them.

References

1. Alam, M.J., Biedl, T., Felsner, S., Kaufmann, M., Kobourov, S.G., Ueckerdt, T.: Computing cartograms with optimal complexity. *Discrete & Computational Geometry* **50**(3), 784–810 (2013). <https://doi.org/10.1007/s00454-013-9521-1>
2. Alam, M.J., Kobourov, S.G., Veeramoni, S.: Quantitative measures for cartogram generation techniques. *Computer Graphics Forum* **34**(3), 351–360 (2015). <https://doi.org/10.1111/cgf.12647>
3. Bortins, I., Demers, S., Clarke, K.: Cartogram types. http://www.ncgia.ucsb.edu/projects/Cartogram_Central/types.html (2002)
4. Bowen, C., Durocher, S., Löffler, M., Rounds, A., Schulz, A., Tóth, C.D.: Realization of simply connected polygonal linkages and recognition of unit disk contact trees. In: Di Giacomo, E., Lubiw, A. (eds.) *Graph Drawing and Network Visualization (GD)*. LNCS, vol. 9411, pp. 447–459. Springer (2015). https://doi.org/10.1007/978-3-319-27261-0_37
5. Breu, H., Kirkpatrick, D.G.: Unit disk graph recognition is NP-hard. *Computational Geometry* **9**(1-2), 3–24 (1998). [https://doi.org/10.1016/S0925-7721\(97\)00014-X](https://doi.org/10.1016/S0925-7721(97)00014-X)
6. Buchin, K., Kusters, V., Speckmann, B., Staals, F., Vasilescu, B.: A splitting line model for directional relations. In: *Advances in Geographic Information Systems (SIGSPATIAL)*. pp. 142–151. ACM (2011). <https://doi.org/10.1145/2093973.2093994>
7. Buchin, K., Speckmann, B., Verdonschot, S.: Evolution strategies for optimizing rectangular cartograms. In: Xiao, N., Kwan, M.P., Goodchild, M.F., Shekhar, S. (eds.) *Geographic Information Science (GIScience)*. LNCS, vol. 7478, pp. 29–42. Springer (2012). https://doi.org/10.1007/978-3-642-33024-7_3
8. Di Giacomo, E., Didimo, W., Hong, S.H., Kaufmann, M., Kobourov, S.G., Liotta, G., Misue, K., Symvonis, A., Yen, H.C.: Low ply graph drawing. In: *Information, Intelligence, Systems and Applications (IISA)*. IEEE (2015). <https://doi.org/10.1109/IISA.2015.7388020>
9. Dorling, D.: Area cartograms: their use and creation. No. 59 in *Concepts and Techniques in Modern Geography*, University of East Anglia: Environmental Publications (1996), http://www.dannydorling.org/?page_id=1448
10. Eppstein, D., van Kreveld, M., Speckmann, B., Staals, F.: Improved grid map layout by point set matching. *International Journal of Computational Geometry & Applications* **25**(02), 101–122 (2015). <https://doi.org/10.1142/S0218195915500077>

11. Eppstein, D., Mumford, E., Speckmann, B., Verbeek, K.: Area-universal rectangular layouts. In: Symposium on Computational Geometry (SoCG). pp. 267–276. ACM (2009). <https://doi.org/10.1145/1542362.1542411>
12. Gastner, M., Newman, M.: Diffusion-based method for producing density-equalizing maps. Proceedings of the National Academy of Sciences of the United States of America **101**, 7499–7504 (2004). <https://doi.org/10.1073/pnas.0400280101>.
13. Heilmann, R., Keim, D., Panse, C., Sips, M.: Recmap: Rectangular map approximations. In: Information Visualization (InfoVis). pp. 33–40. IEEE (2004). <https://doi.org/10.1109/INFVIS.2004.57>
14. Hliněný, P.: Contact graphs of line segments are NP-complete. Discrete Mathematics **235**(1–3), 95–106 (2001). [https://doi.org/10.1016/S0012-365X\(00\)00263-6](https://doi.org/10.1016/S0012-365X(00)00263-6)
15. House, D.H., Kocmoud, C.J.: Continuous cartogram construction. In: IEEE Conference on Visualization. pp. 197–204 (1998). <https://doi.org/10.1109/VISUAL.1998.745303>
16. Johnson, T., Acedo, C., Kobourov, S., Nusrat, S.: Analyzing the evolution of the Internet. In: Eurographics Conference on Visualization (EuroVis). pp. 43–47. Eurographics Association (2015). <https://doi.org/10.2312/eurovisshort.20151123>
17. Keim, D.A., North, S.C., Panse, C.: Cartodraw: A fast algorithm for generating contiguous cartograms. IEEE Transactions on Visualization and Computer Graphics **10**(1), 95–110 (2004). <https://doi.org/10.1109/TVCG.2004.1260761>
18. Klemz, B., Nöllenburg, M., Prutkin, R.: Recognizing weighted disk contact graphs. In: Di Giacomo, E., Lubiw, A. (eds.) Graph Drawing and Network Visualization (GD). LNCS, vol. 9411, pp. 433–446. Springer (2015). https://doi.org/10.1007/978-3-319-27261-0_36
19. van Kreveld, M., Speckmann, B.: On rectangular cartograms. Computational Geometry **37**(3), 175–187 (2007). <https://doi.org/10.1016/j.comgeo.2006.06.002>
20. Meulemans, W.: Efficient optimal overlap removal: Algorithms and experiments. Computer Graphics Forum **38**(3), 713–723 (2019). <https://doi.org/10.1111/cgf.13722>
21. Meulemans, W., Dykes, J., Slingsby, A., Turkay, C., Wood, J.: Small multiples with gaps. IEEE Transactions on Visualization and Computer Graphics **23**(1), 381–390 (2017). <https://doi.org/10.1109/TVCG.2016.2598542>
22. Misue, K., Eades, P., Lai, W., Sugiyama, K.: Layout adjustment and the mental map. Journal of Visual Languages & Computing **6**(2), 183–210 (1995). <https://doi.org/10.1006/jvlc.1995.1010>
23. Munzner, T.: Visualization analysis and design. AK Peters/CRC Press (2014)
24. Nickel, S., Sondag, M., Meulemans, W., Chimani, M., Kobourov, S., Peltonen, J., Nöllenburg, M.: Computing stable Demers cartograms. CoRR **abs/1908.07291** (2019), <http://arxiv.org/abs/1908.07291>
25. Nusrat, S., Kobourov, S.: The state of the art in cartograms. Computer Graphics Forum **35**(3), 619–642 (2016). <https://doi.org/10.1111/cgf.12932>
26. Raisz, E.: The rectangular statistical cartogram. Geographical Review **24**(2), 292–296 (1934). <https://doi.org/10.2307/208794>
27. Shneiderman, B., Wattenberg, M.: Ordered treemap layouts. In: Information Visualization (InfoVis). pp. 73–78. IEEE (2001). <https://doi.org/10.1109/INFVIS.2001.963283>
28. Sondag, M., Speckmann, B., Verbeek, K.: Stable treemaps via local moves. IEEE Transactions on Visualization and Computer Graphics **24**(1), 729–738 (2018). <https://doi.org/10.1109/TVCG.2017.2745140>

29. Tobler, W.R.: A continuous transformation useful for districting. *Annals of New York Academy of Sciences* **219**, 215–220 (1973). <https://doi.org/10.1111/j.1749-6632.1973.tb41401.x>
30. Tversky, B.: Cognitive maps, cognitive collages, and spatial mental models. In: Frank, A.U., Campari, I. (eds.) *Spatial Information Theory (COSIT)*. LNCS, vol. 716, pp. 14–24. Springer (1993). https://doi.org/10.1007/3-540-57207-4_2