

About

This manual, the corresponding program and the corresponding data accompany the paper "[Small Multiples with Gaps](#)", published in the TVCG proceedings of IEEE VIS 2016.

These items are provided "AS IS" with no guarantee of functionality nor warranty of any kind. None of the authors or their host institutions are bound to provide support or updates of any kind, nor to accept liability in case of faulty or undesired results. By using them, you accept these terms. If you do not accept, do not use these items.

Starting the program

Step 1: run the pre-compiled jar file "DataExplorerSingle.jar" using Java Runtime Environment 8 or newer.

Step 2: go to the "/data/" subfolder (if it doesn't start there) and select one or more .zip archives to open these data. Select only files with the same prefix ("square", "flex0.8", "abstract" or "existing").

Step 3: go to the "/existing/" subfolder (if it doesn't start there) and select "existing.tsv" to open the data about existing layouts (useful mostly if you selected "existing" in step 2). You may also hit "cancel" to not load any existing layouts.

Inside the program

The program has five toplevel tabs: "Metrics Matrix" showing charts for all the measurements in the data set; "Trade-Off Matrix" to see how one optimization compares to others more directly; "Examples Matrix" is dedicated to exploring data about existing maps (only available if these data were loaded in step 3 above); "Map Statistics" showing statistics about the input maps as a PCP; "Layout Viewer" allowing you to inspect the maps, grids and assignments, with annotations of distortion.

General interaction

When focus is on a draw panel (clicking once may sometimes be needed):

- Spacebar: resets the view to neatly contain the current visualization
- Right mousebutton drag: pan the view
- Scrolling: zooming in/out
- Dragging with middle mouse button: zooming in to the dragged rectangle
- C: makes a screenshot in PNG format of the current view, saving it in the default folder.
- I: makes a screenshot in IPE format of the current view, saving it in the default folder.
- In each of the three "Matrix" tabs: Alt+Drag on columns/rows to order them! The note that clicking a cell in the matrix reorders both row and column simultaneously. You can release Alt after starting to drag. The "Optimizations" and "Measures" tab have a "Reset order" button to revert to the original order (same order as used in the tab itself).

Metrics Matrix

Initially, the screen may be mostly white and empty, because the viewport is misaligned. Hit spacebar to realign the view. There are a number of subtabs on the lefthand side, which allow us to control the visualization:

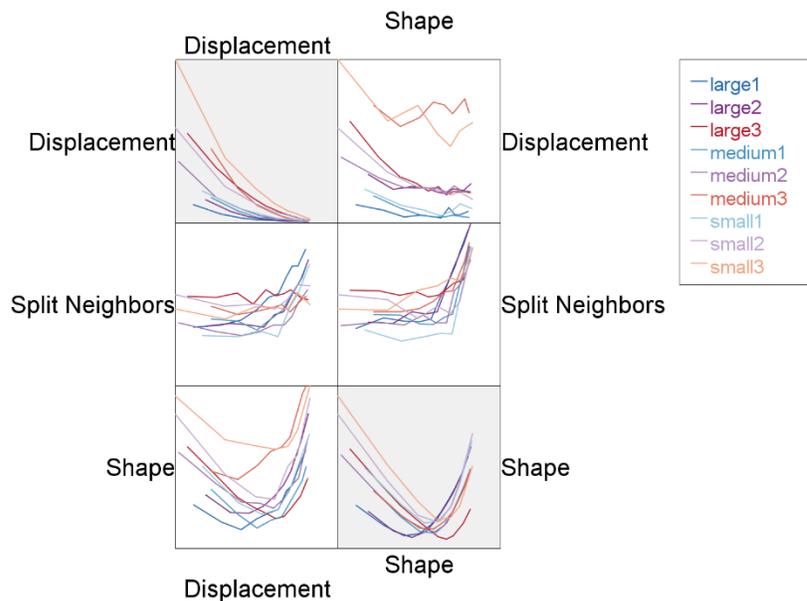
“Optimizations” allows you to select columns in the eventual matrix of charts, which will show you statistics about assignments when we choose to optimize for the selected measure.

Let’s select “Displacement” and “Shape” for now.

“Measures” allows you to specify the rows of the chart matrix, using mostly the same naming scheme as above, in addition to some additional timing measurements (which will typically not be of interest). Note that most measures have a “value” (V column of checkboxes) and a “normalized value” (N column of checkboxes). The former uses the pure values as measured, the latter normalizes these on a scale of 0% to 100%, where 100% is a (theoretic) worstcase value. Let’s select “Displacement”, “Shape” and “Split Neighbors”.

A dropdown allows you to select “AVERAGE” and “SUM”. This configures whether you look at the sum or average of all parts that make up a measure (e.g. the “SUM” or “AVERAGE” displacement of a region). It is probably best to leave this to “AVERAGE” to avoid introducing effects of the number of regions.

Now, click into the drawing pane on the right and hit “spacebar” again to fit the charts to the available area. You should see something like this:



Every square shows the maps in the data set (color legend and selecting/deselecting maps can be done under the “Maps” tab; note that you can press the colors in the “Maps” tab to change the color assignment) Columns are optimizations and rows and measures. Thus, the cell in column 1, row 2 shows us how “Split Neighbors” when we optimize “Displacement”.

Within one chart, we see the behavior as whitespace increases from left to right: the left boundary of the square would be 0% whitespace, the right boundary 100%. Vertically, it shows the quality of the measure: the bottom side is always 0: the best possible value; the topside is the maximum value in the entire row in case of “value” selections or 100% in case of “normalized value” selections. So, **lower is better!** This holds for all measures, even for those for which intuitively this should be the other way around, e.g. “Horizontal Alignment”: these have been inverted.

Note that if the row/column measures indicate the same measure, the border is very subtly thickened, to indicate that there are the same. (Roughly the diagonal of the matrix when selecting all measures).

The “Existing” tab allows you to control which existing maps from literature are shown in the statistics (visualized as crosses). Note that there is no “optimization” for these, so from left to right, (i.e. within the same row) these crosses are positioned equally. Note that this tab will not be available if no data was loaded about existing maps.

The “Interact” tab allows you to configure/toggle some basic interactions with the matrix of charts.

- “Highlight row/column”: As you hover over charts, the row/column you’re hovering will be given a light gray background, to make it easier to see which measures are shown in a given chart.
- As you move close to the points making up the individual trials, you will get to see a preview of this assignment in the leftbottom corner. The size of this preview can be configured by setting the value of “Preview size” to a different number of pixels.
 - The trial you hover over is by default also highlighted throughout the same column (different measures, but same optimization and thus the exact same assignment) but not highlighted in the same row (same measure, but different optimization and thus potentially different assignments). This can be toggled with “Highlight hover in column” and “Highlight hover in row” respectively. “Highlight hover in remainder” does this same effect throughout all other charts in the chart matrix.
 - You can move between trials in the grid also using the arrow keys. (You may have to click once in the draw area to set focus). Left/right decreases and increases the whitespace, up/down changes the map. If you hold Alt while pressing these keys, you can use these to jump to the adjacent row or column as you reach the end! If you hold down Shift, you’ll always directly jump a column or row.
 - If you want to further look at a particular assignment, Shift-click the trial or press Enter to go to the “Layout Viewer” directly, for your selected trial.
 - You can also see the entire lineup (for the currently hovered map and optimization criterion) by setting the configuration in the “Interact” tab.
- There is some (very) crude / elementary pattern detection implemented. You can set its sensitivity using the “sensitivity” slider, left being very sensitive, right being very UNsensitive to small changes in the pattern. Select the patterns you want to see using the checkboxes. Note that “NONE” is anything that couldn’t be matched to one of the other shapes. The slider at the bottom allows you to configure the transparency of the unselected shapes.

Examples Matrix

This view is similar in many ways to the “Metrics Matrix” view, but is focused on exploring the data on existing layouts and comparing them to what our optimization approaches yield. The Optimization/Measures/Existing subtabs on the lefthand side work identical to those in “Solution statistics”. The “Interact” tab, however, provides very different options.

First, there are three comboboxes, allowing you to control the type of visualization. The first controls how the measurements about the existing maps are displayed. “Hidden” does not display them, “Integrated” shows them integrated into each matrix element. Note that the values throughout a row (a measured statistic) are constant for each existing map. Hence, they can also be shown as a separate column. To do so, check “[Existing]” in the “Optimizations” subtab.

The second combo box decides on how to show the “comparator”, that is, the result of our optimization approach on the same grid (if found in the data set: load the existing.zip file in step 1 on startup to ensure this). Again, we can either choose to have these “Hidden” or visualized as a “Bar”. However, we can also select the two different views: “DIFF_BASE” shows the comparator as a rectangle indicating the difference between it and the existing map. “DIFF_ZERO” shows these same bars, but aligns them in the middle of a

matrix element, allowing easier comparisons between magnitudes and better highlighting the better/worse cases.

The third combo box decides on a color scheme for the comparator. If “Fixed” is selected, then comparator rectangles are all given the same color. “Equal” gives the comparator the same color as the existing map. “Dark Light” assigns a darker version of the existing-map color to comparators that perform worse than the existing map and a lighter version to those that perform better; “Light Dark” is its inverse. “Worse Better” assigns fixed colors to when the comparator is worse/better than the existing map.

The fixed colors mentioned above can be modified using the provided buttons. Moreover, a “fade” option is provided that fades the color to the color of the existing map, depending the magnitude of difference (zero difference yields exactly the same color).

Tradeoffs

This tab is again very similar to the “Metrics Matrix” and “Examples Matrix” tabs, in a way combining these. The basic principles of the graphic and interaction are nearly identical to “Examples Matrix”, but the difference lies in which solutions are being compared. Whereas the “Examples Matrix” allows you to compare an existing map to our computation solutions, this tab allows you to cross-compare our computed results.

Instead of an “Existing” subtab, there is now a “Grids” subtab that allows you to select which map-grid combinations you want to study. Note that there may be quite a large number of them! Scroll down to the bottom to select/unselect all. You can also hold down the shift button while checking/unchecking one of the grids to perform the operation on all grids of the same map.

Every column in the matrix is again an optimization criterion, as selected in the Optimizations subtab. Every row in the matrix is a measure as in the other toplevel tabs, as selected in the Measures subtab.

So what is it we see in one cell of the matrix? For each grid, we draw a comparison in the cell (similar settings available in the “Interact” subtab as for the “Examples Matrix” toplevel tab). The base for the comparison the solution according to the optimization (column). The comparator can be one of two things:

1. The solution according to the optimization that is being measured in the row (make sure “Compare to optimized” is checked in the “Interact” subtab). Note that the “diagonal” of the grid is going to not tell you anything in this particular view!
2. The solution according to one specific optimization (uncheck “Compare to optimized” and select which one you want using the combobox right below, in the “Interact” subtab).

Note that “better” and “worse” are defined in the sense of the comparator (determined using one of the two methods above) being better or worse than the base (determined by the column).

In particular, this allows us to investigate two things.

1. Using method 1 above, we can see when we optimize for a given criterion (within a column), how much we sacrifice in terms of all the other measures. That is, how much better could we have done if we really wanted to? E.g. how much of the “Shape” do we lose compared to optimum when we optimize for “Displacement”.

2. Using method 2 above, we can see how much we gain or lose in all measures between two solutions. E.g. if we compare optimizing “Shape” to “Displacement”, which tends to perform better in terms of “Split Neighbors”?

Note that “False Neighbors” is not one of the optimization criteria. So, in particular, with method (1), this row is going to be blank / not provide information.

Map Statistics

We collect some statistics about the input maps. These can be seen in a Parallel-Coordinates Plot [here](#). Some measures are aggregates of a per-region measurement (e.g. “Area”), and by default show only the coefficient of variation. However, you may select other summary statistics to be shown in the “Interact” subtab. Note that you can press the colored buttons in the “Maps” tab to change the color assignment.

Assignment viewer

Here you can view the various layouts computed in the experiments and the existing layouts. Under the “Input” tab, you find four dropdown boxes. The first allows you to select an existing assignment (if any were loaded). The next three work in concert to define one trial. The first of these selects the input map (e.g. “london”). The second allows you to specify the whitespace fraction (“London 80%” would be a grid where 80% of the cells remain unassigned). The third allows you to select the optimization criterion.

Note that if you click on a dropdown (for whitespace or optimization) and use the arrow keys up/down to toggle, the visualization is automatically updated.

Under the “Aux” tab you can select auxiliary elements to be displayed. Currently, you can display vectors from the regional centroid to the cell centroid to which it is assigned as well as some topological features. Note that these settings will carry over into previews of the other toplevel tabs.

If you Shift-click or hit Enter in this panel, you will be going back to the “Metrics Matrix” tab.