Weaving Emotional Response - Combining Art With Data Visualization

by Bradley Shanrock-Solberg, March 2015

Executive Summary

Background

Lia Cook is an artist who has done extensive work translating photographic images into woven tapestries. She became interested in the differences in emotional response to photographic prints, compared to woven translations. Working with neuroscientists she obtained help in designing a study to capture immediate emotional reactions of people who viewed wall-sized prints and weavings of the same image.

Having captured this information, Lia wished to translate it back into her art. Lia believed that her survey data could also be combined with the woven photo in a similar way, showing visually how people reacted to the actual image.

Issues

The people who helped set up the tests were unable to continue to work on the project and had only supplied some basic, overall statistical results and a mix of data files which had translated the paper surveys into electronic format. Lia did some early data visualizations on her own but was not entirely happy with the result. This was where I got involved.

My job was to finish the statistical analysis, advise on improvements to the data visualization charts and provide graphics Lia could incorporate into her own work as needed.

Results

I performed a factor analysis and found the significant statistical variations in the studies and provided a number of visualization options based on the data. This was packaged for reproducibility along with a companion brochure for future exhibits explaining the statistics and visualization choices. Chart generation was automated using R, allowing rapid turnaround of new chart requests or brochure changes as needed.

The Problem

Lia's initial data visualizations suffered from three distinct problems - incomplete statistical analysis, graphic choices that failed to communicate the nuances of the data and choices in chart design and colors that did not work well with the underlying prints.

Statistical Analysis

The experimental procedure and surveys were well designed, tracking carefully important elements such as whether a photo or weave was viewed first and capturing emotional reactions three different ways (a paragraph, a single word and 0-10 measures of emotional intensity and negative/positive reactions, in that order). Only after capturing the key emotional data and viewing both images did the individual add information about themselves. These factors included age, gender education, prior experience with art, weaving, scientific experiments or photography.

There were two studies, each with about 350 subjects. The subjects were evenly split between weave-first and photo-first viewing, but the studies were done in two locations (Houston and Pittsburgh) and each study used a different photo as the basis of the print and weave.

The data was transferred from paper to electronic format by several hands and several text formats, including an intern who was not involved in the study. Some of the surveys had formatting errors in the 0-10 scales, leading to a concern that some subjects might get the order reversed. One of the studies had a statistical analysis done of the entire data population with a couple basic graphs - a **matplot** looking much like a spider web and a basic **spot** chart.

The problem in this area revolved around pulling the data for both studies into a common format, performing a factor analysis on the data to see what if any qualities of a subject caused a different emotional response and producing clean data sets suitable for visualization charts.

Communicating the Data

Lia's initial visualizations were used as a foundation, as it was the charts in the limited study that had captured her imagination. I could see right away there were some problems in the initial approach. In the **matplots**, color was used in a way natural to the weaving process but not in a way useful for communicating the data, and thickness was not used at all. The **spots** graphs did a better job - using size and color to communicate data elements. The problem there was a lot of redundant chart legends and titles that drew attention away from both the image and the data.

The problem in this area was focused on bringing color and size into harmony, consistent across all of the chart choices. This would allow viewers to understand various statistical

comparisons in an exhibit by using a common theme and logic, reducing what is actually woven to the basic shapes and colors describing the response.

Colors, Shapes and Chart Graphics

In data visualization, there is always a tension between how much explanation is done in the margins and how much the data is allowed to speak for itself.

In art of this sort, there is also the concern of chart elements such as grid lines and legends drawing attention both from not just the data, but also the underlying image. In this project the choices are complicated by the fact that a basic grid has a certain resonance with weaving itself, and some approach with a grid that did not overpower the image was desired as an option.

Finally, the colors chosen must not clash with the image, or get lost in the image. Lia's color choices in the initial spots visualizations had the yellow washing out in lighter portions of the underlying image, and the red and blue elements washing out in darker portions unless she included a white border on the shapes.

The problem here is providing options that are consistent with letting the data tell the story, and which do not overpower the image, but also provide the artist with enough good choices to find the ones that have the greatest impact. Also, all choices must be possible with the technology on hand - electronic charts must be in a format usable by her graphics program, and the choices must be possible to weave, and look good when using thread instead of pixels.

Implementation and Analysis

Statistical Inference

The data cleansing exercise was typical for this kind of study - with everything in one or another form of text file, importing the data into R and sorting it out was not difficult. Some basic sanity checks on the data turned up transcription errors that required going back to the paper surveys.

The statistical analysis focused on T-tests of different factors, linear regression and analysis of variance with the most promising factors and finally some basic visualizations to see if the statistical variation was sufficiently different to be visible to a human when graphed.

The most important source of bias in the study was whether the subject viewed a **weave** or a **photo** first. The data had to be adjusted to correct for this data in order for the other relationships to appear clearly.

The Pittsburgh and Houston studies had a significantly different response, which is not terribly surprising since they had different images, although it is possible their different age mix also played a part.

There were 428 subjects in Houston, 396 subjects in Pittsburgh, people willing to both attend an art exhibit and participate in an experiment. When gender was indicated, both samples were a bit over 70% female. The Houston sample showed an older population than the Pittsburgh sample. Both had similar numbers for education, accounting for age.

Age	child	teen	20s	30s	40s	50s	60s	70s	80s	90+
Pittsburgh	21	34	79	42	44	80	51	13	1	0
Houston	13	20	52	50	38	80	89	28	4	1

What was most interesting is that both studies had an age cut-off where negative-positive response was significantly different among younger and older viewers. For Pittsburgh, this cutoff was at about age 50, for houston, the difference was seen at age 30. No significant difference in raw intensity was discovered in any subgroup of either test, nor was any other factor (such as gender, education, art experience etc) significant when age was included.

The core data for the visualizations focused on the entire populations of each study (for comparison between them) and populations within each study that fell into the major age group breakpoints.

Showing the Story in the Data

These are two of the early visualizations, **matplot** on the left, **spot** on the right. The matplot had the virtue of focusing entirely on the image and the lines formed by the data, but did not use color or line thickness to show the data.

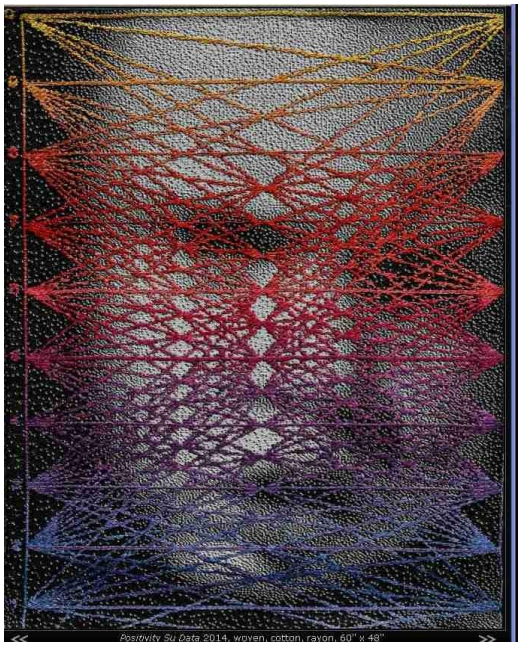

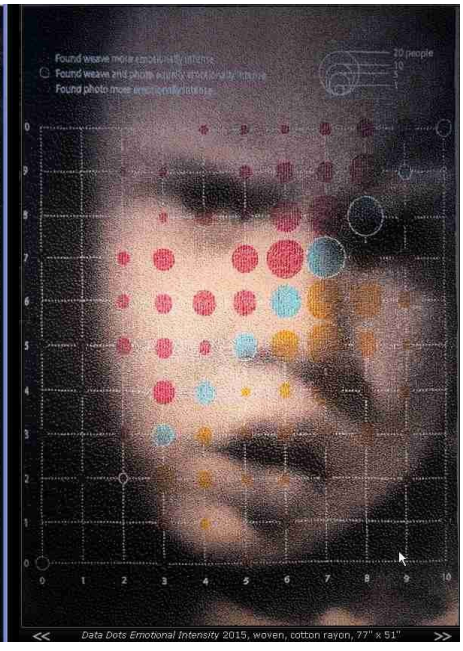

The spot graph applied meaning to the colors and size of the graphic elements, but then added a lot of distracting legend elements. The grid shown here blends

into the image, and Lia did another tapestry where the grid was dominant over the image.

My first approach was to simply apply the best of both consistently to all images.

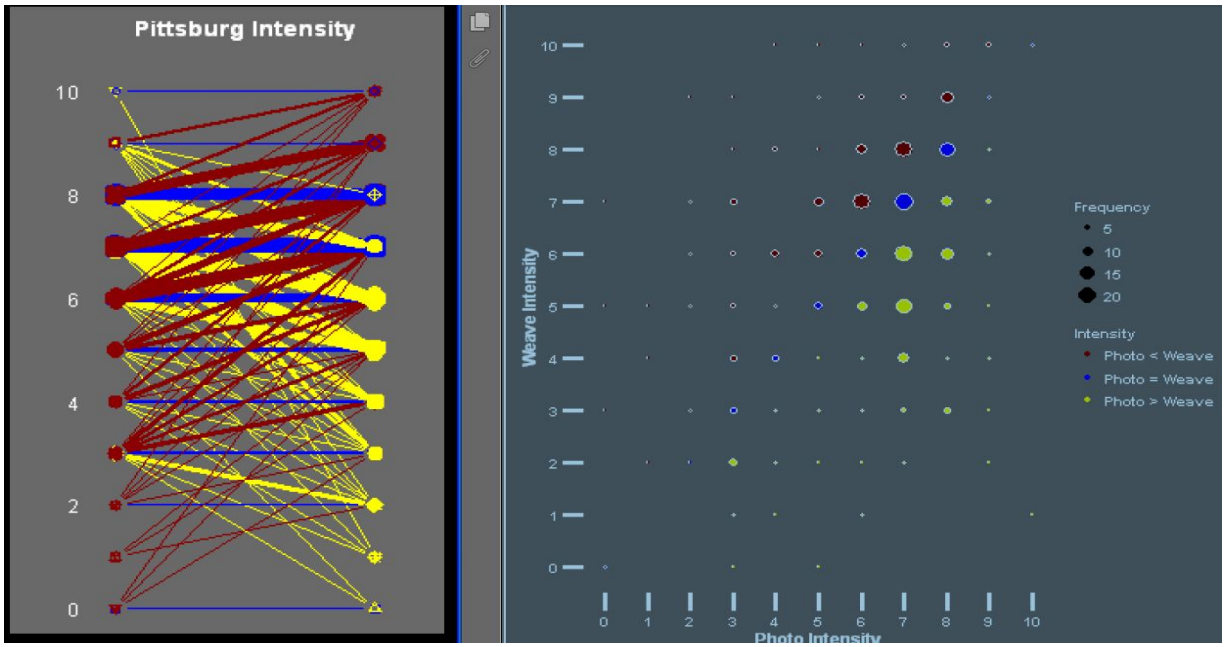

With these images color is used the same way in both graphs. The grid lines were removed from the **spot** graph and more of a legend was added to the **matplot**. I used several backgrounds to show that colors washed out, and that the three originally chosen colors probably were not ideal.

I retained the option to add grids of varying strength to the plot, as Lia had received positive feedback related to the grid based on her early visualizations. As a general rule I am not a fan of grids in data visualizations, as they often detract from the story the data can tell on its own, and the human eye will create a "grid" if the data is plotted on points, as in the image above. The lattice of a grid, however, has a resonance with the woven medium and that means the choice to include or exclude the grid is going to be a tension between the overall artistic effect, the focus on the data itself and the visibility of the original image. My favorite style if a grid is included is to weave it using the same thread used for the lighter portion of the photographic image. This puts the grid at the same level of importance as the photograph in the overall effect.

After reviewing these changes, Lia liked the consistent use of shape and color, and we felt that the legends, numbers and other explanations were also distracting. The meaning of the colors and line size/shape size could be explained in a companion document or as part of the exhibit. Experimenting with colored thread, Lia came up with a bright blue and bright red thread that would show clearly against both the dark and light weaves. Where the response was equal (blue in the above graphs) she would weave the red and blue together, giving a purple look from a distance or just use a bright purple thread. This means there is no need to outline the shapes in white (as in the above image) and opens the door to using shapes that are not solid.

The second round of visualizations incorporated these elements, and used the cleaned data populations derived from statistical inference, allowing comparisons and contrasts. I added an abstract "face" to show how the colors would overlay the image from a conceptual standpoint.

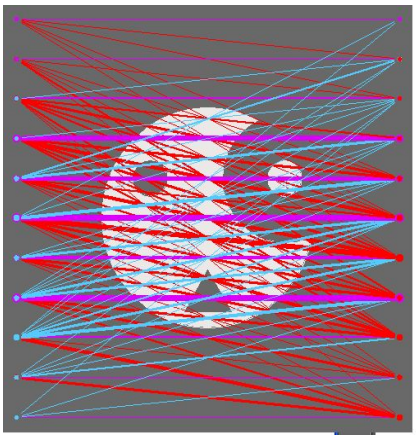

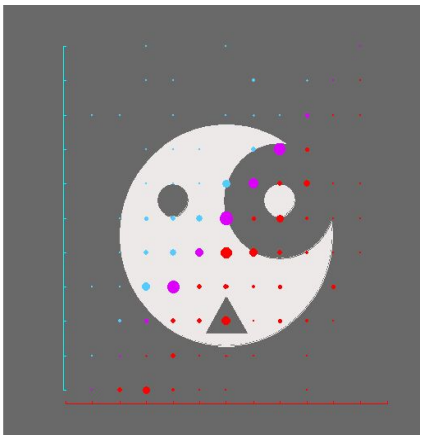

Finding colors to simulate her threads was fairly challenging - how brightness appears on a computer and a printer are not the same.

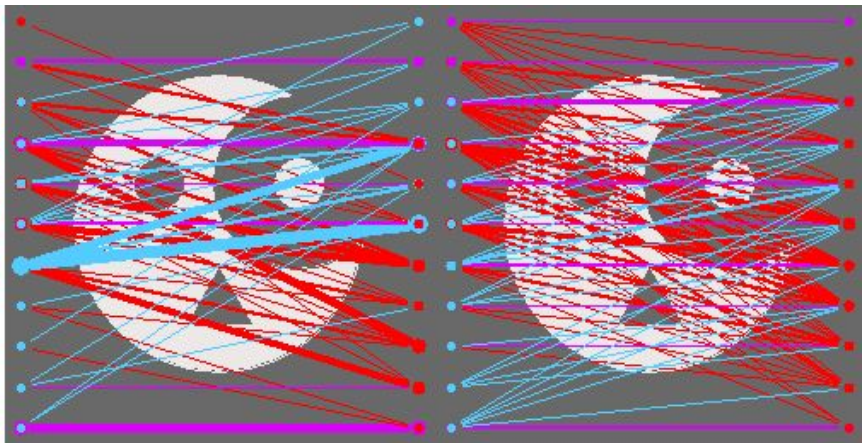

These graphs show the effects of the data, left is "under 30" and right is "over 30" in the Houston study. The less positive response (red) of the older contingent is very visible. For this kind of comparison, the line and shape sizes were pro-rated based on sample size.

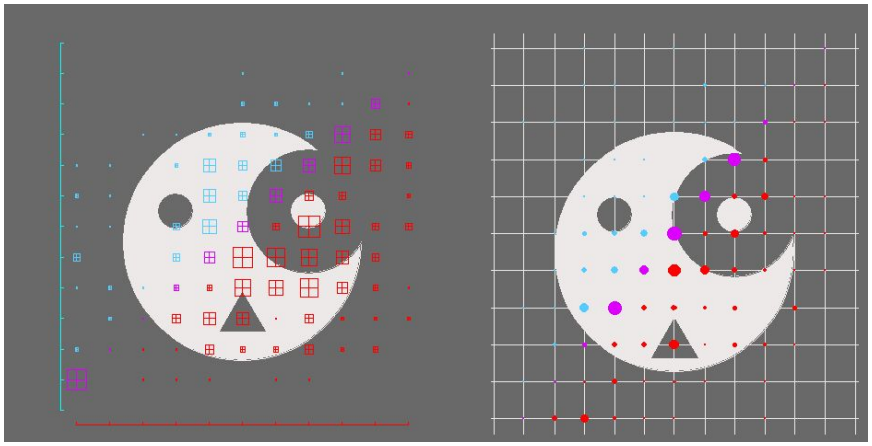

Graphs like these played with grid lines and shapes.

All of this round of graphs were generated using R functions crafted to vary background, shape, grid, color and data set.

While useful enough to show options, these graphs were not quite ready for Lia to blend them into her weaving process. The initial aspect ratio was based on the early data visualization tapestries, but they were kind of odd, and she had me switch to a more standard value (and I added a variable, so we can change it up as needed). She needed a clean .pdf with a

transparent background. I modified the program and sent her about 30 images to get her started.

Wordclouds

The study captured single word descriptions of the emotions felt by the subjects. As the main areas of statistical variation were in the negative-positive axis, I thought the words might prove interesting. Similar words were combined to make the wordcloud easier to read.

While superimposing wordclouds on a tapestry is technically challenging, Lia was interested in giving it a try.

Pittsburgh Weave More Positive

	Photo	Weave	Change
120	troubled	reflective	8
265	worried	interested	8
134	desire	fear	7
338	fear	peaceful	7
91	sad	evil	6
54	lost	wondering	5
183	butter	invigorating	5
282	sad	beautiful	5
280	wishing	angry	4
354	lonely	sad	4
387	uncertain	hope	4
396	reflective	angry	4

Results

This is to some extent an ongoing project, as Lia has a lot of other projects and is experimenting with these visualizations when she has time and will undoubtedly request new images going forward. At this point though, it is a turnkey process, and I have been given permission to show a few images which might become tapestries in time, to show how it all fits together. Compare these to the original visualizations earlier in the document.

This graphs on the right are Houston under 30 vs over 30,

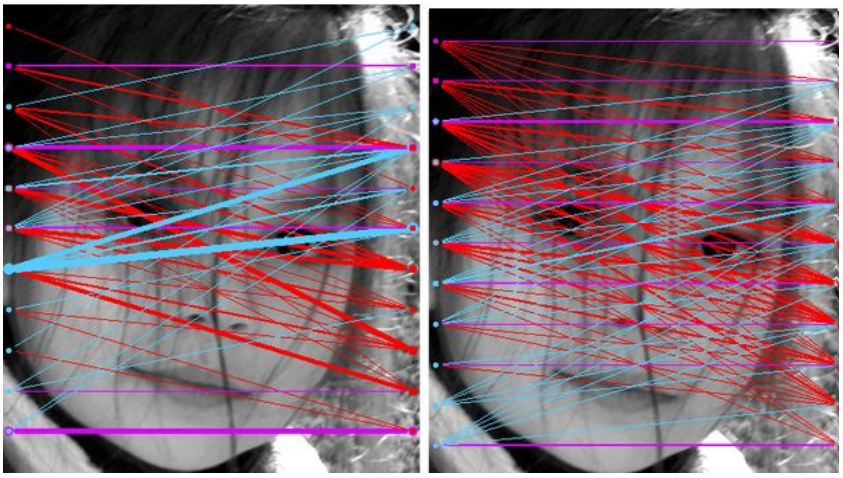

The three on the bottom are Pittsburgh full data set.

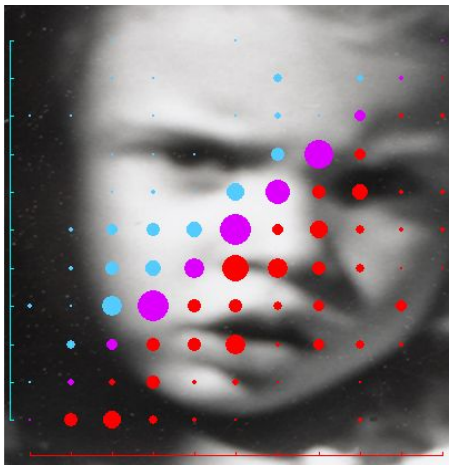

Concluding this case study is the first fruit of this effort, a small demonstration weave Lia created for a conference in April 2016. One of the quirks of the weaving process is that the image must be resized before it is used to program the loom. Lia provides an explanation:

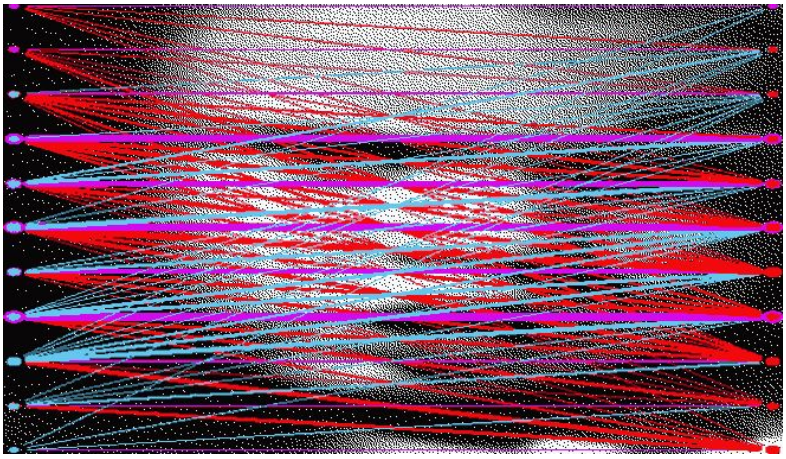

This is the prepared image before the weaves are assigned. In this particular case the final version will have 660 threads wide (warp on the loom) and 1140 threads high, (weft) woven at 66 weft threads per inch to end up with a total height of 17 inches. Every different kind of thread, thickness of thread used and type of weave structure used changes the number of threads per inch woven and therefore changes the aspect ratio. Most people do samples before they start weaving if they want the image to come out with the correct proportions.

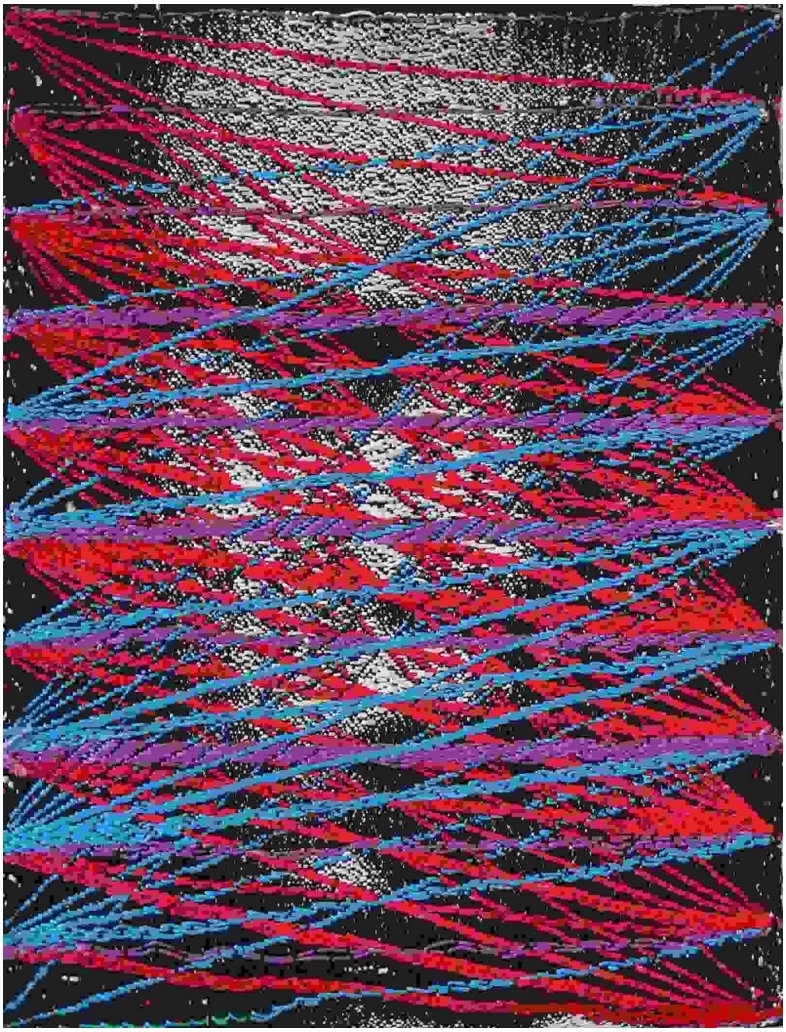

This weave is the Pittsburgh full data set **matplot**, showing positive (top=10) vs negative (bottom=0) response to **photo** (left side) vs **weave** (right side).

Blue lines indicate a preference for **weave**, which is why they all slope upward. Purple lines have equal preference, they are all horizontal. Red lines indicate a preference for **photo**, they all slope downward.

Line thickness is an indication of the percentage of people in the sample who had exactly the same response to photo vs weave.

Lia used purple thread for this weave because she preferred the blue line covering the purple to blending the red and blue threads.