

Four Types of Ensemble Encoding in Data Visualizations

Danielle Albers Szafir, Department of Computer Sciences, University of Wisconsin-Madison, Madison, WI

Steve Haroz, Department of Psychology, Northwestern University, Evanston, IL

Michael Gleicher, Department of Computer Sciences, University of Wisconsin-Madison, Madison, WI

Steven Franconeri, Department of Psychology, Northwestern University, Evanston, IL

Abstract

Ensemble encoding supports the rapid extraction of visual statistics about distributed visual information. Researchers typically study this ability with the goal of drawing conclusions about how such encoding extracts information from natural scenes. Here we argue that a second domain can serve as another strong inspiration for understanding ensemble encoding: graphs, maps, and other visual presentations of data. Data visualizations allow observers to leverage their ability to perform visual ensemble statistics on distributions of spatial or featural visual information to estimate actual statistics on data. We survey the types of visual statistical tasks that occur within data visualizations across everyday examples, such as scatterplots, and more specialized images, such as weather maps or depictions of patterns in text. We divide these tasks into four categories: *identification* of sets of values, *summarization* across those values, *segmentation* of collections, and estimation of *structure*. We point to unanswered questions for each category and give examples of such cross-pollination in the current literature. Increased collaboration between the data visualization and perceptual psychology research communities can inspire new solutions to challenges in visualization while simultaneously exposing unsolved problems in perception research.

Keywords: Ensemble Encoding, Data Visualization, Visual Search

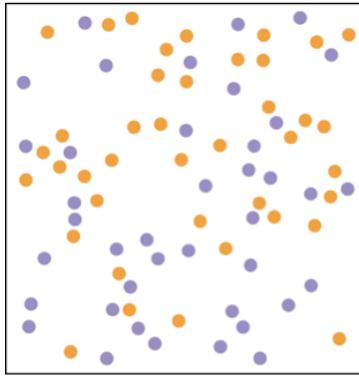
16 Introduction

17 Some types of visual information must be extracted from small numbers of objects at a time, such as complex object
18 identity (Wolfe, 1998) or spatial relationships (Franconeri, Scimeca, Roth, Helseth, & Kahn, 2012). Other types of
19 information can be extracted and combined in parallel from large numbers of objects at once, such as the average object
20 size (Ariely, 2001). A growing body of work seeks to understand such *ensemble encoding* of spatially distributed visual
21 information (see Whitney, Haberman, and Sweeny, 2014, and Alvarez, 2011, for surveys). Researchers typically study
22 this ability in order to draw conclusions about how ensemble encoding helps extract information from natural scenes.
23 For example, one might want to estimate the number of books on a shelf (Ross & Burr, 2010) or gauge the average
24 emotional expression within a crowd of people (Haberman & Whitney, 2007).

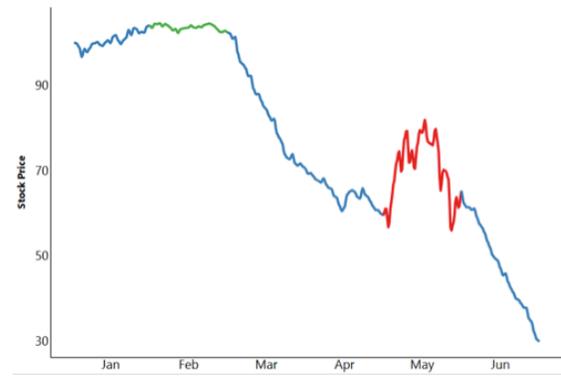
25 Here we argue for another domain that should serve as an equally exciting inspiration for understanding ensemble
26 encoding: visual presentations of data (e.g., maps, charts, & graphs). Data visualizations are ubiquitous to students,
27 scientists, and any broader audience that reads graphs, uses maps, or reads a newspaper. Visualizations communicate
28 patterns in data by mapping data dimensions to visual features (see Bertin, 1983, and Heer, Bostock, and Ogievetsky,
29 2010, for an overview). To illustrate, consider a scatterplot, which maps data values to spatial positions. For some
30 types of inspection, such as mapping symbols to a legend or knowing whether a particular data value is lower or higher
31 than another, we must serially inspect small numbers of data values at a time. But other types of information can be
32 extracted in parallel, such as the approximate mean position, or size, of an entire cloud of points (Figure 1a) or the
33 portion of a line graph with the highest variability (Figure 1b).

34 These judgments are ensemble judgments, and they merit more intense study both for their value as a case study
35 for understanding how ensemble processing works in the visual system and also for their practical importance within
36 information visualization. Information visualization research and practice has been previously inspired by research
37 in cognition and perception (see Ware, 2008 and 2013, and Healey and Enns, 2012, for surveys, and Rensink, 2014,
38 for a framework for reasoning about perceptions of visualization designs). However, existing work focuses on how
39 perception might inform visualization design. We instead aim to inspire a broader two-way conversation between
40 vision science and visualization—understanding how viewers estimate properties of visualized data offers potential
41 research directions for vision science, and this understanding can in turn inform more effective visualization designs.

42 Experiments in visualization have tried to quantify how effectively viewers perceive different properties encoded
43 using various visual features. However, much of this work focuses on single-value tasks, such as finding and estimating
44 values from individual datapoints (see Cleveland and McGill, 1984; Heer, Kong, and Agrawala, 2009; and Javed,
45 McDonnel, and Elmqvist, 2010, for examples). More recent work has begun to study ensemble processing in data
46 visualization, such as the construction of averages within a scatterplot (Gleicher, Correll, Nothelfer, & Franconeri,
47 2013), variance, range, and outliers in line graphs and heatmaps (Albers, Correll, & Gleicher, 2014), and estimation
48 and comparison of correlation (Harrison, Yang, Franconeri, & Chang, 2014; Rensink & Baldrige, 2010). Our goal
49 is to identify a broader set of such visualization tasks that benefit from ensemble encoding and to increase research
50 and discussion surrounding how these judgments work and how visual data displays can be better designed to support
51 them.



(a)



(b)

Figure 1: Even without explicitly showing statistics, a viewer can quickly and robustly observe that, on average, the orange dots have a higher Y value than the purple dots, or that there is more variance in May (highlighted in red) than February (highlighted in green).

52 In principle, the statistics that viewers perceive in a data visualization could be formally computed and shown
 53 directly to the viewer. However, visual extraction of statistics is often more attractive because formally computed
 54 statistics, which necessarily abstract over potentially critical patterns, are often insufficient to describe data. Try to
 55 imagine a scatterplot of a dataset that exhibits the following statistics: the X and Y variables both have a mean of 7.5
 56 and variance of 5, and the correlation coefficient of X and Y is 0.816. You are probably imagining that the underlying
 57 data look like the first plot in Figure 2. But any of the four datasets shown in Figure 2 would produce these statistics—
 58 all four datasets have identical means across both X and Y, variabilities across both X and Y, correlation coefficients,
 59 and linear regression formulas ($y = 3 + 0.5x$) (Anscombe, 1973). Yet each plot exhibits qualitatively different patterns.
 60 While there are increasingly complex statistics that could differentiate among these patterns, these statistics would not
 61 likely be run without the benefit of visual inspection to determine their necessity. Alternatively, attempting to provide
 62 statistical information explicitly, even through visual means, can quickly become cluttered and overwhelm the viewer,
 63 even for a small number of statistics (Figure 3). Visual estimation also provides a beneficial flexibility in terms of
 64 what data is being processed: viewers have direct control over the different subsets of the data they choose to compute
 65 statistics for.

66 If allowing observers to extract statistics and patterns about data with their visual system offers an alternative
 67 to explicitly providing raw statistics, then understanding the effectiveness of this process is critical. The benefits of
 68 efficiently estimating visual statistics provides both an application of, and challenges for, research on the perception of
 69 these features. How do the capabilities of the visual system match the needs of visual depictions of data? Conversely,
 70 the difficulties encountered in data visualization can challenge our understanding of perception. When a perceptual
 71 psychologist cannot answer the questions asked by a visualization designer, it shows the psychologist the gaps in their
 72 theory—what they did not realize that they did not know.

73 We organize our exploration of this synergy between perception research and visualization around two key ques-

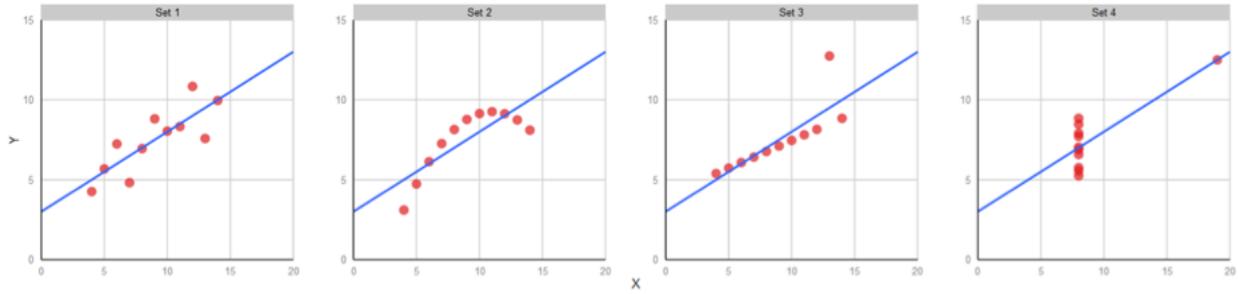


Figure 2: Common statistical abstractions may not capture potentially critical patterns in data. This example, from Anscombe (1973), shows four datasets that are identical across several common statistics (mean, variance, correlation, and linear regression), yet contain qualitatively different patterns. Visual inspection offers powerful and flexible processing of these differences, as well as rough approximations of statistics.

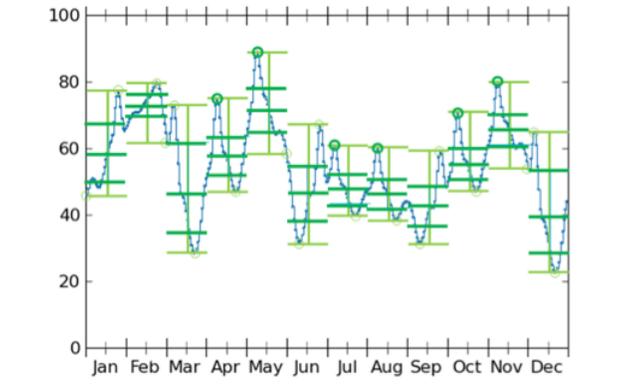


Figure 3: Providing explicit statistics (in this case, minimum, maximum, mean, variance, and outliers per month) can be overwhelming, even in a visual format.

74 tions: what visual 'statistics' can our perceptual system extract via ensemble encoding, and what potential needs in
75 visualization can these ensembles address? We can align visual statistics with visualization needs by understanding the
76 different kinds of tasks viewers might want to accomplish. In visualization, tasks are, informally, the visual operations
77 that people may want to perform with data, such as identifying points with high values or estimating the average of
78 a set of values. A flurry of recent efforts in the data visualization community propose taxonomies and typologies of
79 tasks (Roth, 2012; Schulz, Nocke, Heitzler, & Schumann, 2013; Shneiderman, 1996; Amar, Eagan, & Stasko, 2005),
80 creating abstractions that seek to help knowledge gained in one environment transfer to visualizations with differ-
81 ing contexts and details (see Brehmer and Munzner, 2013, for an extensive survey and comparison of prior efforts).
82 However, these taxonomies generally attempt to classify techniques used by designers rather than to understand how
83 properties of the data might be perceived in different designs.

84 In order to address the questions around ensemble encoding that bridge perception and visualization, we need a
85 categorization of visual tasks at the perceptual level: basic visual operations that serve as building blocks for more
86 complex analyses. In this paper, we introduce an organization of low-level tasks that require, or may require, ensemble
87 encoding into a framework of four categories: identification, summarization, segmentation, and structure estimation.
88 Figure 4 depicts these categories, as well as examples of each, for four common ways of visually depicting data
89 values (position, size, orientation, and color). Figure 5 demonstrates examples of these tasks applied to more complex
90 visualization systems. Understanding which combinations of visual feature and task are most effective is a critical
91 challenge. What statistics and patterns can we accurately extract, which are inaccurate, and which are systematically
92 biased? How does the choice of feature used to represent the data affect our ability to extract absolute values, statistics,
93 and patterns from datasets?

94 Both the perception and visualization research communities should explore and refine—or even completely reinvent—
95 the grid of tasks and features in Figure 4. For perception research, it holds a diverse set of unsolved problems, not only
96 for understanding ensemble encoding across different features and statistics, but also for revealing unsolved questions
97 surrounding visual search, multifocal attention, and visual comparison (see (Franconeri, 2013) for review of these
98 topics). For visualization research, it has the potential to produce concrete guidelines for optimizing the mapping of
99 visual features to data dimensions to support different tasks. In the following sections, we explore this grid, moving
100 serially among its columns, providing samples of relevant research on the perceptual issues related to each task, the
101 visualization applications that build on the task, and potential directions for future research. While by no means ex-
102 haustive, the sampling offers several potential research directions for perceptual psychologists that could also inspire
103 more effective visualization design.

104 We preface our argument with some caveats. We do not intend this categorization to be a final answer, but instead
105 the spark of a broader conversation. For example, we categorize outlier detection as an 'identification' task, but one
106 might also argue that it is a form of 'segmentation'. Our discussion of work related to each task, both from perceptual
107 psychology and data visualization, will not be exhaustive—instead, our goal is to provide a sampling of relevant
108 work from each community. Appendix A provides a table of additional visualization references for the discussed
109 tasks. Because of the need for brevity, some of the links that we draw may strike members of either community as

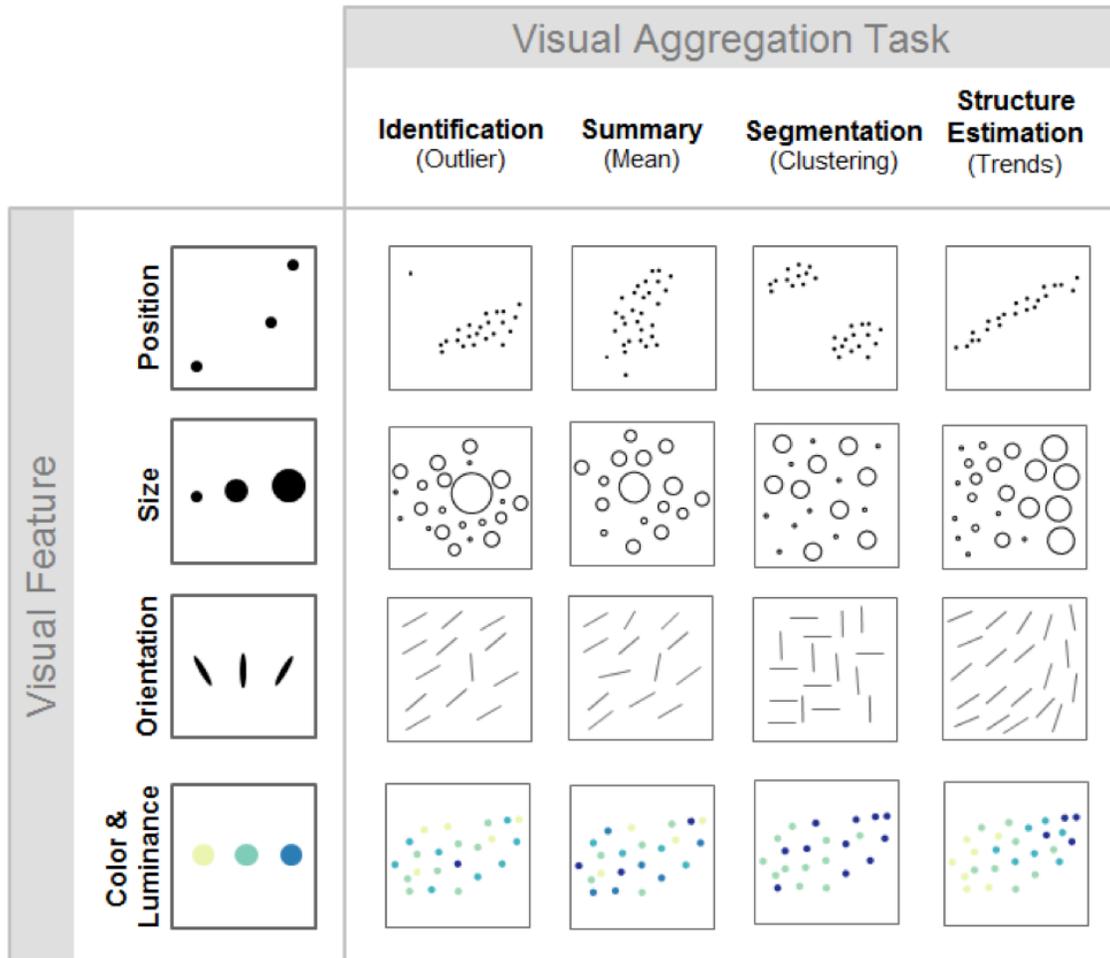
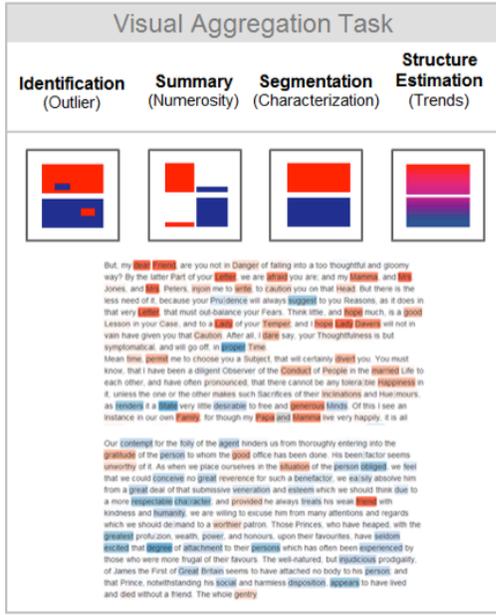
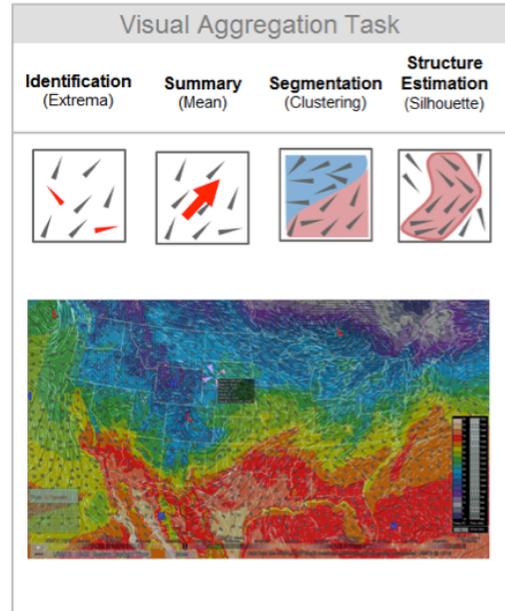


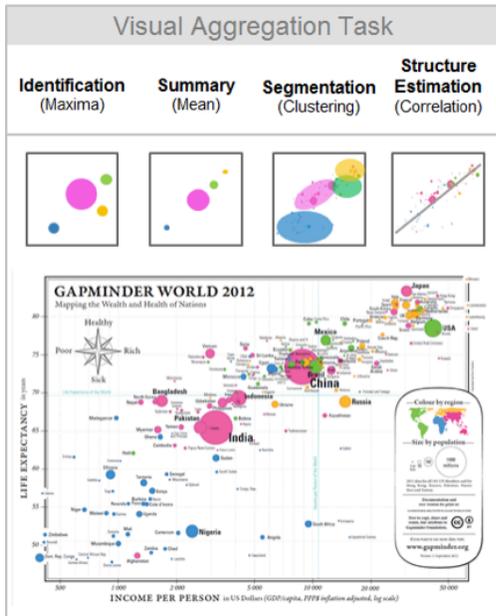
Figure 4: We identify four categories of visualization tasks (top) that require ensemble encoding of information spread throughout the visual field. The tasks can be performed on multiple visual features, but not necessarily with equal speed or efficiency. In visualizations, choosing which visual feature is mapped to each dimension of the dataset affects which tasks are most easily performed on which data dimensions (e.g., perhaps it is best to map size to the data dimension that will likely require a summary judgement, and position to the data dimension that will likely be segmented).



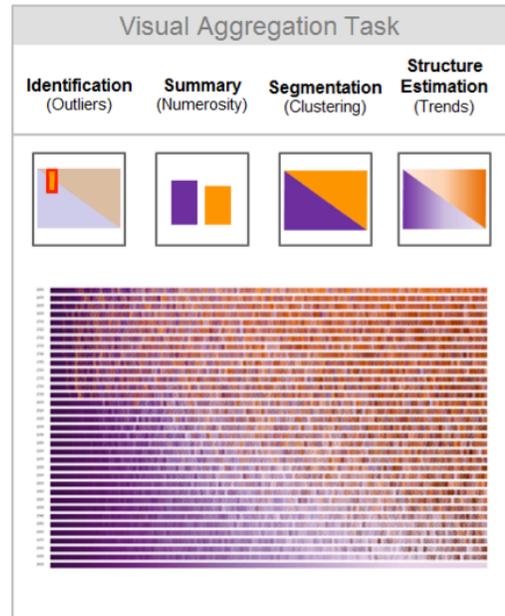
(a) Tagged text visualization uses color and position to allow users to identify linguistic patterns in a document (Alexander, Kohlmann, Valenza, Witmore, & Gleicher, 2014).



(b) Weather maps use color and orientation to visualize information about wind speeds, temperatures, and other meteorological data (Ware & Plumlee, 2013).

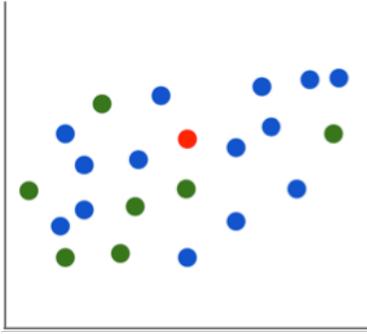


(c) GapMinder uses size, position, and color to reveal patterns in global demographics (Rosling, 2009).

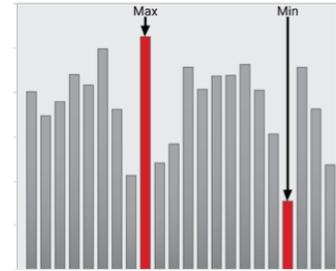


(d) Inspired by work on ensemble encoding, Sequence Surveyor depicts changes in the use of 170,000 words across 34 decades, locally permuting color to help the visual system construct ensemble summaries of noisy information (Albers, Dewey, & Gleicher, 2011).

Figure 5: In these example visualization applications, data is mapped to multiple visual features, such as (a, d) color and position, (b) color and orientation, and (c) position and size, to support a variety of analysis tasks. Understanding how efficiently these features communicate different kinds⁷of information can inspire effective visualization designs.



(a) Absolute value from color: Where are the red points?



(b) Extrema from length: What are the largest and smallest values in a bar chart?

Figure 6: Identification tasks require a viewer to locate a specific set of datapoints, such as (a) the class of points in a scatterplot that are labeled as red or blue, or (b) the minimum and maximum values that constitute the value range in a bar graph.

110 problematic. For example, we mention findings in data visualization that appear inconsistent with work in perceptual
 111 psychology (c.f. the discussion on mean position in scatterplots in 'Summary Tasks: Mean and Variance Estimation').
 112 A psychologist reading those sections may generate display constraints and confounding factors that could explain
 113 why the effect did not generalize, and may reflexively produce more precise guidelines that visualization designers
 114 could use to better predict when these effects will hold. We would be delighted by this response, as it highlights
 115 the importance of increased collaboration between visualization designers and perceptual psychologists. Finally, our
 116 review will focus on ensemble encoding, but given the blurriness of its definition and the need for it to interact with
 117 other types of visual processing, our review will include other related topics that are outside its strict definition, such
 118 as visual search, multifocal attention, feature-based attention, and shape recognition.

119 Identification Tasks

120 An identification task requires a viewer to isolate a specified subset of datapoints (Figure 6). Often, the viewer knows
 121 which values they seek, in which case the task requires identifying points that match given values. In other cases, the
 122 viewer may need to extract distributional information about a dataset to identify values that are important relative to
 123 the distribution, such as locating the minima, maxima, or median. In a third kind of identification task, the viewer must
 124 actively search through the dataset to detect outliers—values that are notably different from the rest of the dataset.

125 Absolute Value Identification

126 In absolute value identification tasks, viewers isolate points that match a specified data value, such as the red points
 127 (e.g., Figure 6a), a circle of a given size, or all points in a certain spatial region of the display.

128 The constraints on a viewer's ability to locate a *single* target with a known feature—red, circular or 2 units in

129 diameter—have been long-studied (Wolfe, 1998). As an example of one effect, finding a single target is more difficult
130 when that target is perceptually more similar to the distractors around it, and more difficult when there is more diversity
131 amongst the distractors (Duncan & Humphreys, 1989). This result is consistent with an established guideline for
132 designing visualizations: maximize the perceptual distance among the features that delineate data properties relevant
133 to the identification task, but minimize differences between features that map properties irrelevant to the task (Wickens
134 & Carswell, 1995).

135 There are also rich links between the visualization task of localizing the set of targets matching a known feature
136 —*all* of the red, circular, or 2-degree objects—and feature based attention, which allows a viewer to select multiple
137 objects with a visual ‘feature filter’ (e.g., Saenz, Buraças, & Boynton, 2003)). This feature filter is known to extend
138 broadly across a visual display, allowing the selection of large numbers of objects (or datapoints) at the same time
139 (Levinthal & Franconeri, 2011), and rules have been proposed governing how multiple features can, and cannot, be
140 logically combined (e.g., red AND square) (Nakayama, Silverman, et al., 1986; Huang & Pashler, 2007). Under-
141 standing how these filters operate across multiple features could inform visualizations that support identification tasks
142 acting across different dimensions of a dataset simultaneously. For example, in Figure 5b, understanding how well
143 viewers can combine color and orientation can help determine how effectively viewers can identify, for example, high
144 temperature regions where the wind blows towards the west.

145 When a set of points cannot be easily selected by visual features, it must instead be selected by their locations.
146 Perhaps a viewer needs to select points 4, 16, and 18 within a scatterplot because the text of their labels makes them
147 currently relevant. Here, research on attentional selection of multiple objects (Scimeca & Franconeri, 2014) explains
148 the ability to perform this task in visualization. As an example, there are limits to the number of locations that can be
149 selected (up to 7 or 8 in total), but this limit is closer to 3 or 4 in typical displays where objects become more tightly
150 packed (Scimeca & Franconeri, 2014). A recent collaboration between perception and visualization researchers has
151 shown that these selection constraints generalize to conditions similar to following points in a scatterplot through an
152 animated transition from one plot to another (Chevalier, Dragicevic, & Franconeri, 2014).

153 **Relative Value Identification**

154 While some absolute value identification tasks may not require ensemble encoding, relative value identification tasks
155 rely on distributional information about a dataset in order to identify datapoints with a pre-specified position within
156 that distribution. Because extraction of the distribution is needed to find relative values, even traditional visual search
157 tasks for relative values would seem to require an initial ensemble processing pass of a display before defining the
158 target to search for. Examples of relative value tasks include extracting the minimum or maximum value for the entire
159 set of data (e.g., the lowest data value) or within a subset of the data (e.g., the lowest red). In Figure 5b, for example,
160 an analyst might search for the range of wind directions in California. In Figure 5c, which two countries have the
161 largest populations, and which have the smallest? In a bar graph, the range of the data distribution might be revealed
162 by simultaneous visual selection of both the minimum and maximum values (Figure 6b). The strategy for locating
163 minima and maxima is unclear, though it may require ensemble encoding, as both minima and maxima are defined

164 with respect to all of the points in a collection.

165 There are other relative value identification tasks that beg for study by perceptual psychologists. In Figure 6b,
166 how well can you estimate the median value in the bar graph, and what perceptual process allows that judgment? One
167 heuristic could be to find the range, determine the imaginary horizontal line that hovers in the midpoint of that range,
168 and search for bars with tops near that area. That strategy works for non-skewed distributions, but fails when the data
169 are skewed. What perceptual strategies would be more robust, what downsides would they have, and how could they
170 be taught to graph readers? What graph designs would permit other strategies for finding the median—for example,
171 how could your abilities change for data plotted as color values instead of the positional and length values in Figure
172 6b? What happens when you ask all of these questions for the modal (most frequent) value, instead of the median?

173 **Outlier Identification**

174 Outlier detection tasks are defined by the need to identify targets that are different than others in the collection. They
175 are similar to relative value identification tasks, except that the position of the target datapoints within the distribution is
176 not specified *a priori*—viewers discover them while foraging through a dataset, allowing saliently different datapoints
177 to ‘pop out’ (Neisser, 1964; Prinzmetal & Taylor, 2006). This set of tasks reflects one of the strongest advantages
178 of using the visual system to compute statistics in visualization: cases where critical statistics are difficult to know
179 (and therefore cannot be mathematically computed) *a priori*. As a result, outlier identification provides a number of
180 opportunities for research in both perception and visualization.

181 When position is used to represent data, it is unclear what perceptual strategies allow viewers to determine when
182 a point in a scatterplot, or a bar in a bar graph, might be seen as an outlier. Studies of perceptual segmentation, as
183 discussed in the ‘Segmentation’ section of this paper, may offer insight into how a positional outlier may be identified
184 and the role ensembles might play in detecting these values. When data is instead plotted in a featural space, such as
185 when values are encoded with color in a heatmap, outliers that might be modeled by their perceptual salience (Itti,
186 2005; Itti & Koch, 2001). But what process might allow a viewer to detect outliers that are not prototypical extrema,
187 such as outliers in the middle of a widely-spaced bimodal color distribution? Identifying such outliers may rely more
188 strongly the power of ensemble encoding to extract information about the overall distribution of values.

189 Goals, context, tasks demands, and experience account for much of the variability in salience for natural scenes, but
190 whether this is still true in relatively simpler data displays remains to be tested. Work on attentional control (Serences
191 et al., 2005) and priming of features by recent experience (Chetverikov & Kristjansson, 2014) may contain important
192 insights for visualization designers, and the context of data visualization tasks could inspire perceptual psychologists
193 with new questions.

194 All three of the tasks described above can be performed either within an entire dataset or within a specific subset
195 of the data. This subset can be spatially defined: what is the minimum value in the left half of Figure 6b? What is
196 the pop out color in the first paragraph of the example shown in Figure 5a, or the upper-left corner of Figure 5d? The
197 subset can also be featurally defined: among the red circles in Figure 5c, are there any positional outliers at the bottom
198 of the display? In displays that simulate data visualizations, the spatially local level surprisingly does not appear to

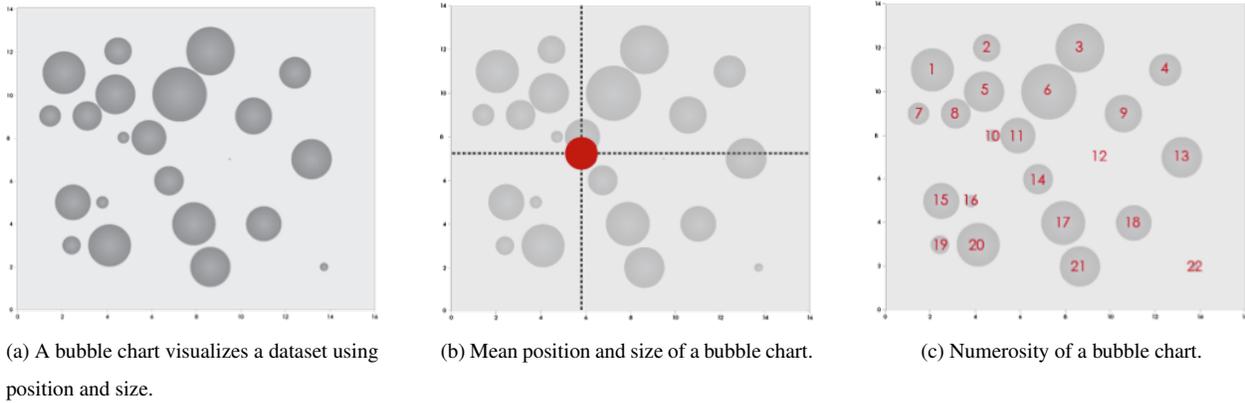


Figure 7: Summary tasks require viewers to estimate a value that summarizes a collection, such as its (b) mean and (c) numerosity.

199 play a stronger role in computing the salience of a potential outlier. Detecting the presence of an outlier depended on
 200 an item's global uniqueness rather than local uniqueness (Haroz & Whitney, 2012), implying the use of scene-wide
 201 variance rather than only local contrast.

202 Summary Tasks

203 Summary tasks require the viewer to extract properties that describe the collection in aggregate. In contrast to identifi-
 204 cation tasks, which extract subsets of objects, summary tasks create representative values, such as descriptive statistical
 205 measures. For example, a viewer may estimate the average height of a bar in a bar chart or the average position of
 206 points in a scatterplot (Figure 7b). While some summaries might overlap with value extraction, as when extracting a
 207 median, most summaries are not values from the set.

208 Mean and Variance Estimation

209 Estimating mean and variance are common summary tasks in visualization. Figure 8a depicts monthly stock prices
 210 as individual line graphs with data from one company colored red and data from the other colored blue. An analyst
 211 could estimate the mean orientation of the red and blue lines to compare how monthly stock prices change on average
 212 between the two companies, and orientation variance to compare the stability of stock values. Means can be efficiently
 213 computed for several visual features including size (Ariely, 2001; Chong & Treisman, 2003, 2005a, 2005b; Fouriezos,
 214 Rubinfeld, & Capstick, 2008), orientation (Choo, Levinthal, & Franconeri, 2012; Parkes, Lund, Angelucci, Solomon,
 215 & Morgan, 2001; Alvarez & Oliva, 2008; Bulakowski, Bressler, & Whitney, 2007), motion speed (Watamaniuk &
 216 Duchon, 1992) and direction (Watamaniuk, Sekuler, & Williams, 1989), brightness (Bauer, 2010), color (Webster,
 217 Kay, & Webster, 2014), and position (Hess & Holliday, 1992; Melcher & Kowler, 1999; Morgan & Glennerster,
 218 1991; Whitaker, McGraw, Pacey, & Barrett, 1996). Variance among values can be efficiently computed for orientation

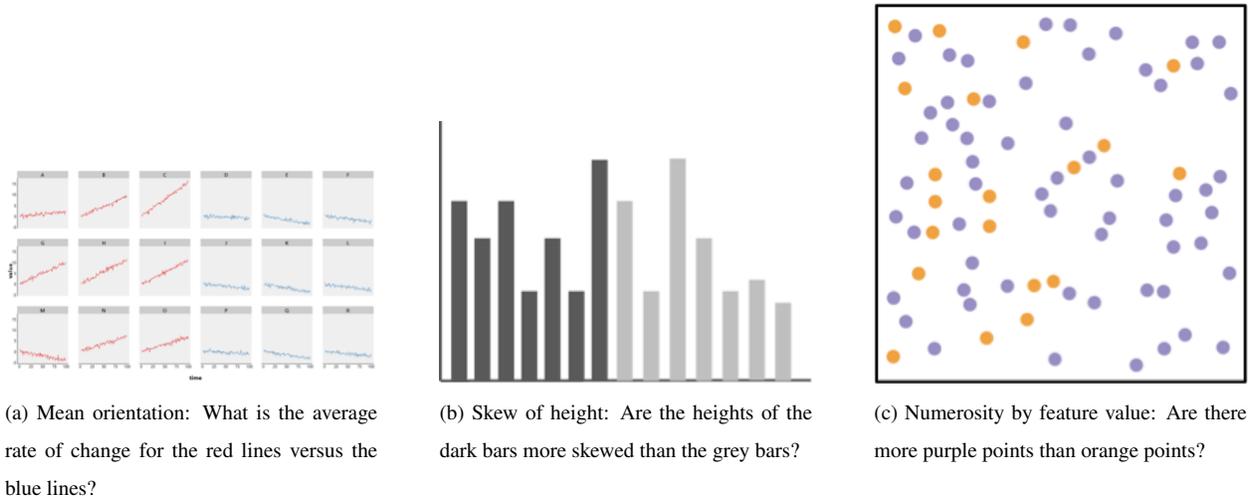


Figure 8: Estimating mean, skew, and numerosity are three types of summary tasks common to data visualization.

219 (Morgan, Chubb, & Solomon, 2008). Our ability to compute the mean of a collection is surprisingly robust in the
 220 face of other types of variability across collections, for irrelevant dimensions like spatial frequency (Oliva & Torralba,
 221 2006), density (S. C. Dakin, 2001; Chong & Treisman, 2005b), numerosity (S. C. Dakin, 2001; Chong & Treisman,
 222 2005b), temporal sequence (Chong & Treisman, 2005a), and distributional variance (S. C. Dakin, 2001).

223 Figure 7a provides a sample visualization where mean size and position can be rapidly extracted (see the red circle
 224 in Figure 7b). Figure 5c depicts a more complex visualization where mean size provides insight into demographic
 225 data. To compare average population size across different geographic regions, for example, you could identify circles
 226 of different colors and average the size of the resulting set. You could use a similar process to identify the average
 227 population size of low-income countries by spatially grouping the objects within the left third of the x-axis, and
 228 computing the average size of the resulting groups, allowing you to note that they tend to be small on average, with
 229 low size variability.

230 Recent work has tested the ability of viewers to estimate the mean value of collections within data visualizations.
 231 One study tested how well viewers could compare the mean position of two groups of points in a scatterplot (Gleicher
 232 et al., 2013), focusing on how the colors and shapes used to mark different datasets affected viewers' ability to compare
 233 their mean heights. Some results were consistent with intuitions from perceptual psychology. For example, using color
 234 to distinguish the two groups (making points orange vs. purple) led to higher accuracy in mean judgments compared to
 235 using shape (making points circles vs. triangles). But adding more diversity among the distractor classes (e.g., adding
 236 green objects to the orange and purple display) did not impair performance for comparing mean position between the
 237 two groups, as would be expected from previous work on visual search (Duncan & Humphreys, 1989). Adding more
 238 perceptual spacing between classes by combining cues (orange and circular vs. purple and triangular) surprisingly did
 239 not improve performance, contrary to findings from previous work on visual search (Duncan & Humphreys, 1989).
 240 Note that the underlying mechanism for such mean position judgments may or may not be an ensemble one, depending

241 on your definition of 'ensemble'. If the horizontal and vertical positions are truly averaged in the same manner as
242 other dimensions such as size or luminance, then the definition fits. But if the center is computed by shape recognition
243 heuristics that focus on a low-spatial frequency envelope (Harrison et al., 2014), then whether that counts as ensemble
244 processing depends on your definition.

245 Another set of studies tested how well viewers can estimate the mean and variance from visualizations of time
246 series data (Albers et al., 2014). These studies showed trade-offs between how accurately viewers can estimate mean
247 and variance (summary tasks) versus range and extrema (identification tasks) from data visualized using either color
248 or position. While each statistic could be extracted from both visual features, there was a salient difference between
249 the types of tasks best supported by each: mean and variance were more accurately extracted from data encoded
250 using color, whereas extrema and range were more accurately extracted from positional visualizations. These results
251 suggest different processing abilities for color and position—color may facilitate summation of values at low spatial
252 frequencies into a representation similar to a color histogram, while position may better represent shape boundary
253 properties. At the same time, the results show a trade-off for visualization design—color better supports summary
254 tasks while position better supports identification tasks.

255 People can also estimate the mean of a set of orientations (Parkes et al., 2001). Not all types of orientation are
256 averaged with the same precision: the average orientation of the boundary contours of a set of objects can be more
257 precisely extracted than the average orientation of their internal textures (Choo et al., 2012). In visualizations like
258 the weather map shown in Figure 5b, this predicts an improved ability to summarize wind directions in maps that use
259 oriented glyphs over maps that use oriented textures. Here, extracting a mean orientation across local regions has clear
260 utility for understanding how the general wind direction in cold regions (purple) differs from the wind direction in
261 warm regions (red). Ensemble processing of orientation is also useful in the stock market visualization in Figure 8a.
262 How accurately could a financial analyst determine the red company shows more variable performance than the other?
263 While the Weber fraction for variance—the point where differences in variance become indistinguishable—has been
264 studied for orientation (Morgan et al., 2008), there are far fewer studies of variance perception, compared to studies of
265 perceptions of average value.

266 **Distribution Statistics**

267 We can extract mean and variance from a collection of data points, but what about other aspects of a dataset's distri-
268 bution? For example, visualizations may reveal the skew and kurtosis of a dataset to a viewer. The bar chart depicted
269 in Figure 8b contains two datasets that differ in skew. In more complex visualizations of oceanographic data, skew
270 and kurtosis of sea surface height are important for making predictions about the movement and position of eddies
271 for applications in oil exploration, where eddies can damage off-shore drilling equipment (Hollt et al., 2014). Little
272 is understood about the accuracy or biases of our perception of these high-order statistics. One possibility is that the
273 visual system does not encode skew per se, but may approximate it after extracting a collection's centroid (S. Dakin
274 & Watt, 1997). Better insight into how these statistics can be inferred by the visual system can help in designing
275 visualizations to support a broader variety of statistical analyses on raw data.

276 Numerosity

277 Visualizations regularly require viewers to judge the numerosity of a set of data points (Figure 7c). This judgment
278 might be an estimate of an absolute number of data points—how many bubbles are in Figure 7a—or a comparison
279 between two or more values—are there more purple or orange points in the the scatterplot in Figure 8c? Numerosity
280 estimation is surprisingly robust in the face of variability in other dimensions, such as contrast, orientation, and density
281 (Burr & Ross, 2008). These findings align with results from visualization that show estimation of relative numerosity
282 is robust across color and orientation (Healey, Booth, & Enns, 1996).

283 However, changes in some featural dimensions, such as luminance, can bias numerosity estimates. Darker col-
284 lections can appear more numerous (Ross & Burr, 2010), so a visualization designer should be careful when using
285 luminance to differentiate collections of data, and the relative numerosity of those collections is relevant to the viewer.
286 Grouping objects using visual connection can cause viewers to underestimate the number of original parts (Franconeri,
287 Bemis, & Alvarez, 2009). A visualization designer working with network data (typically visualized as points con-
288 nected by lines), should be wary of the effect of these connections on number perception. The number of possible
289 simultaneous numerosity estimates may also be limited (Halberda, Sires, & Feigenson, 2006), which implies that the
290 number of visually distinct categories in a visualization should be limited if simultaneous numerosity estimation is
291 critical.

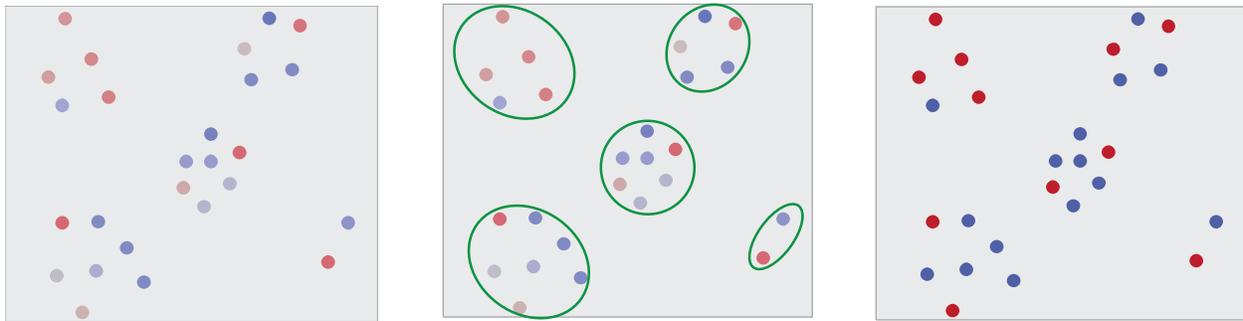
292 Understanding how the size of an item influences perceived numerosity could help reduce bias when size and
293 quantity visualize independent dimensions. Correll, Alexander, and Gleicher (2013) found that when comparing the
294 quantities of red and blue words in displays like Figure 5a, longer words biased viewers towards perceiving a higher
295 quantity. These results led to a new visualization approach that helps account for this bias: increasing the spacing
296 between letters in short words increases the overall length of the colored word, and improves numerosity estimation
297 in text displays.

298 Segmentation Tasks

299 Segmentation tasks require viewers to organize datapoints into subsets (Figure 9). Unlike identification tasks that iso-
300 late datapoints that adhere to specific constraints, these subsets are formed based on their similarity within some visual
301 dimension, typically either space (position) or a feature dimension (e.g. color or orientation). Ensemble processes
302 might guide these segmentation tasks by providing distributional information that allows detection of salient spatial or
303 featural clusters. For example, if luminance values formed a bimodal distribution, with one light mode and one dark
304 mode, it could signal two corresponding clusters of points.

305 Segmentation by Spatial Position

306 Segmentation is perhaps most intuitive when data are mapped to spatial position. Viewing Figure 9a, it is apparent that
307 there are five primary spatial groups (Figure 9b). Spatial segmentation helps viewers quickly form meaningful subset



(a) A multiclass scatterplot uses position, color, and luminance to encode data.

(b) Spatial clustering of a scatterplot

(c) Clustering a dataset by color

Figure 9: Segmentation tasks, such as (b) spatially or (c) featurally clustering data elements, require the viewer to visually segment the dataset into discrete clusters.

308 of related items within a dataset. For example, in Figure 5c, we see multiple spatial clusters that identify countries
 309 with similar demographic properties: a tight cluster of countries that share both high GDP per capita and high life
 310 expectancy in the upper right corner of the visualization, a looser cluster in the center with intermediate GDP per capita
 311 and life expectancy, and a scattered group of countries in the lower left with low GDP per capita and life expectancy.
 312 The importance of this segmentation process demands several explanations from perceptual psychologists, such as
 313 what counts as a ‘cluster’ in the visual field, how many clusters can be created, and what might bias this segmentation
 314 process?

315 Understanding visual grouping may be particularly useful in answering some of these questions. For example, spa-
 316 tial clustering should be largely based on the Gestalt grouping cue of ‘proximity’, and studies of proximity grouping
 317 suggest that it is a parallel and mandatory cue (Rock & Palmer, 1990). It also tends to dominate over other grouping
 318 cues, such as color (Oyama, 1961). While several clusters can be constructed simultaneously across a display, per-
 319 forming additional operations on these clusters, such as extracting the shape of each collection, can be a slow, or even
 320 strictly serial, operation (Trick & Enns, 1997).

321 An understanding of segmentation in data visualizations will also require an understanding of visual crowding.
 322 Items that are identifiable on their own can become indistinguishable or crowded when surrounding objects are too
 323 close (Whitney & Levi, 2011), and this problem worsens in the periphery (Pelli, Palomares, & Majaj, 2004). Crowding
 324 limits could contribute to the limit on the number of clusters that can be created in a visual display (Franconeri, Alvarez,
 325 & Cavanagh, 2013).

326 Segmentation by Features

327 A viewer can also cluster datapoints using featural similarity. Figure 9c depicts an alternative segmentation of the
 328 display in Figure 9a, using color value instead of spatial position. Figure 10 depicts an example where data can be
 329 clustered by orientation and color, ignoring their locations. Figure 5 depicts clustering by color and orientation in

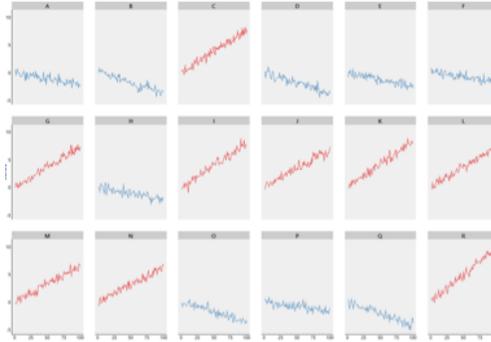


Figure 10: Segmentation tasks, such as clustering, can not only be accomplished with positional mappings, but also by featural mappings like color or orientation.

330 more complex displays: data in all four visualizations can be clustered according to color values, data in Figure 5b can
 331 be clustered by orientation, and data in Figure 5c can be clustered by size.

332 In some cases, using multiple visual features within a single visualization may make feature clusters harder to see in
 333 either feature dimension alone. For example, a visualization might use color hue to encode one property of a datapoint
 334 and luminance to encode a second. Luminance variation across these points might inhibit viewers' abilities to segment
 335 points that have similar hues. It may likewise be difficult to segment different points of different luminance levels if
 336 their hues are vastly different (Callaghan, 1984). However, some features may support more robust segmentation. For
 337 example, if both color hue and shape are used to encode different properties of a dataset, segmenting points based on
 338 shape is likely to be more challenging for shapes of different hues, whereas viewers can segment points of different
 339 colors regardless of shape (Callaghan, 1989). Understanding the role of ensemble encoding in segmentation may offer
 340 guidance for which features (or combinations thereof) can help viewers better identify divisions in visualized data.

341 **Structure Estimation Tasks**

342 Structure estimation tasks require viewers to extract patterns from sets of datapoints that are not always intuitively cap-
 343 tured by single statistics (Figure 11). Anscombe's Quartet (Figure 2) illustrates the importance of structure estimation
 344 tasks: the four datasets are identical across several statistics, yet have qualitatively different patterns. These patterns
 345 often require visualizations for a viewer to understand them.

346 **Trend Detection**

347 Detecting trends—the qualitative relationship between two variables—is perhaps the most ubiquitous form of structure
 348 estimation. The reader is likely most accustomed to trends between two variables mapped to position on a Cartesian
 349 grid, as in a scatterplot (Figure 11b). The trend that as X increases, Y increases, is immediately apparent and appears
 350 linear, as opposed to curved or U-shaped. The visual system is adept at comparing the relative strength of linear
 351 correlations (Rensink & Baldrige, 2010; Harrison et al., 2014) (Figure 11c). While the most common (and likely



(a) A scatterplot can map values to position and color.

(b) Position and color trends in a scatterplot.

(c) Correlation of position in a scatterplot.

Figure 11: Structure estimation tasks extract patterns from collection of values, such as (a) trends and (b) correlation. In the above scatterplots, these tasks can be computed across both position and color.

352 most powerful) visual mappings for representing the trend between two variables pair two spatial dimensions, feature-
 353 based depictions are also common when both spatial dimensions have already been mapped to other aspects of the
 354 data. For example, the bubble chart in Figure 12a may not contain an X-Y trend, but size clearly increases with the
 355 X value. Figure 12b depicts a trend that relies on neither spatial axis, but it is clear that as size increases luminance
 356 decreases.

357 More complex examples of such trends between a positional dimension and a featural dimension are depicted in
 358 Figure 5a and 5d—certain colors occupy certain spatial positions in each of these displays. In Figure 5a, this trend
 359 reflects that words common in novels (red) are more frequent at the beginning of the passage, whereas words asso-
 360 ciated with philosophical discussions (blue) are more frequent at the end (Alexander, Kohlmann, Valenza, Witmore,
 361 & Gleicher, 2014). Figure 5d visualizes the 2,000 most popular works per decade over the last 350 years, with each
 362 decade represented by a row and word popularity mapped to the X-axis. The figure shows that words that are popular
 363 in modern writing (purple) have slowly replaced those that were popular in earlier texts (orange) (Albers, Dewey, &
 364 Gleicher, 2011).

365 When trends are depicted across two spatial position axes, we have some idea of how they might be detected. For
 366 example, simple shape recognition networks might classify whether a cloud of points matches an oval, a curve, or a
 367 U-shaped object (e.g., Uttal & Tucker, 1977; Field, Hayes, & Hess, 1993), though some data suggest that such simple
 368 tricks may not be sufficient to explain performance, at least for certain kinds of trends (Rensink, 2014). The ability to
 369 identify these shapes within visualized data may help reveal complex relationships between variables. For example,
 370 mapping two variables to position will form a line if they are highly correlated or a parabola if there is a quadratic
 371 relationship between them. But when at least one feature dimension is involved, our understanding of how such
 372 patterns are recognized or encoded shrinks drastically. There is evidence that the visual system is capable of detecting
 373 featural correlation, sometimes as effectively as with positional correlation (Rensink, 2014), and that correlation may
 374 be estimated with similar accuracy across visualization designs using several different kinds of features (Harrison et al.,

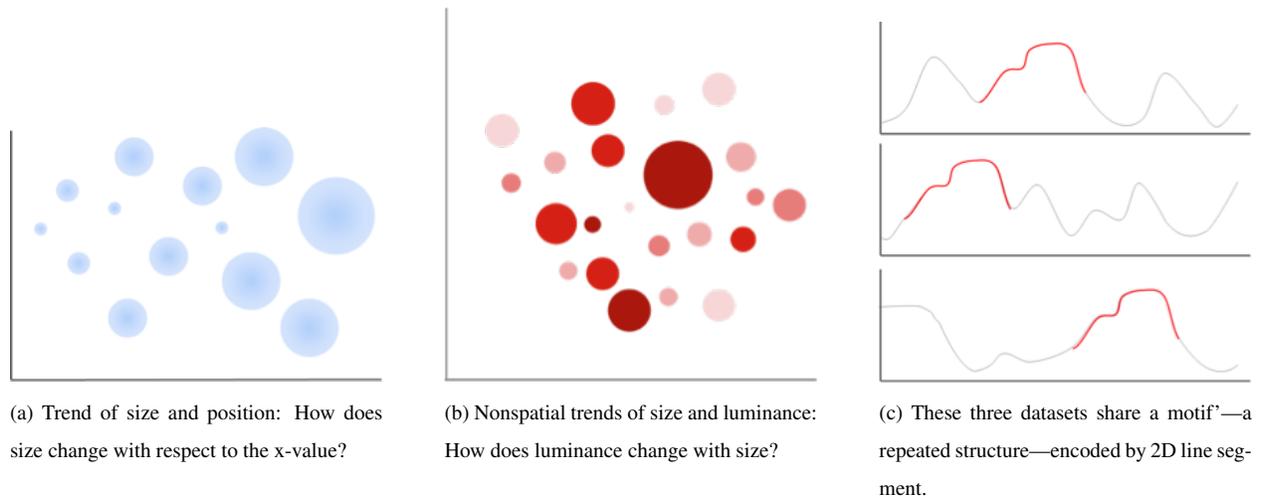


Figure 12: Trend and motif detection are two examples of structure recognition tasks in visualization.

375 2014). However, it is not clear whether this ability generalizes to other forms of trend detection, such as characterizing
 376 nonlinear relationships between features.

377 One possible strategy for characterizing featural trends is that a form of ensemble encoding extracts one or two
 378 feature distributions from the display, and they are compared. While there is some evidence that this cross-feature
 379 pattern detection may occur for orientation and size (Oliva & Torralba, 2006), the mechanism for this detection is
 380 unclear, as is whether it works for other feature combinations. We see this problem as a fertile one for perceptual
 381 psychologists to explore.

382 An innovative set of proposals suggest a more mechanistically precise alternative for detecting featural trends: that
 383 trends involving at least one feature dimension are processed by serial selection of certain feature values at a time
 384 (Huang & Pashler, 2002, 2007). For example, extracting a trend between luminance and size in a bubble chart might
 385 involve selecting dark and then light items, approximating the mean size at each of these luminance levels, and storing
 386 each mean size in memory for later comparison. Note that this method of structure estimation would be more similar
 387 to the 'segmentation' operations in the previous statement, rather than the other operations in this section. There is
 388 evidence for this type of serial processing in some kinds of visual comparison (Huang & Pashler, 2002) and visual
 389 grouping (Huang & Pashler, 2007; Levinthal & Franconeri, 2011), but there is a need for empirical work demonstrating
 390 that such a model could explain trend detection among, for example, the graphs shown in Figure 12.

391 **Similarity Detection**

392 Another aspect of structure estimation is determining similarity across different regions in a dataset, such as inferring
 393 how similar wind currents are across different geographic or temperature regions (Figure 5b). Similarity detection
 394 occurs at different scales: from the more holistic task of estimating the similarity of the shape of two line graphs, or the
 395 more local task of detecting small repeated patterns across two line graph (Figure 12c). Detecting repeated structures,

396 commonly called *motifs*, across a dataset is important in applications such as biology, where these patterns often
397 indicate blocks of genetic material with important biological functions that are conserved across different organisms
398 (Meyer, Munzner, & Pfister, 2009; Albers et al., 2011), or in time series data, where they often represent related events
399 (Lin, Keogh, Lonardi, Lankford, & Nystrom, 2004). Important motifs may appear among noise or other distortions
400 within a dataset or may be inverted in order. For example, in visualizing energy usage over time, an event (e.g. turning
401 on a device) may cause a drastic increase in energy usage. This motif indicates the event occurrence and the motif's
402 inversion may occur when the event ends (e.g. the device is turned off). How might the visual system detect similarity
403 between different collections in a visualization? How might it find small scale repeated patterns that form motifs, and
404 how do noise and inversion influence our ability to identify these patterns? How might the efficiency with which we
405 determine similarity change for different visual features?

406 Ensemble encoding may be an important part of computing similarity between visualized collections. For co-
407 located objects, the visual system might compute variance in a region (Morgan et al., 2008), and regions of low
408 variance indicate high local similarity. The visual system might identify motifs by detecting small regions with similar
409 ensemble statistics as a viewer scans a display. A potential strategy for estimating similarity across different clusters,
410 such as red and blue points in a scatterplot, might involve computing ensembles within spatial or featural clusters
411 (Corbett & Melcher, 2014) and then comparing those statistics between clusters (S. Dakin, 2014). This strategy
412 relies on comparing ensembles as opposed to detailed patterns to estimate similarity across different subsets of data
413 and correlates well with how viewers perceive similarity between pieces of artwork, another type of complex visual
414 scene—here, perceived similarity correlates with comparisons between mean luminances of corresponding spatial
415 regions in a painting (Graham, Friedenber, McCandless, & Rockmore, 2010).

416 **Distribution Shape**

417 Understanding the shape of a data distribution is important for a number of statistical inference tasks. In understanding
418 demographics data, a viewer may wish to characterize the distribution of wealth across a population. Alternatively,
419 they may wish to compare different attributes of a population, such as the distribution of age versus that of income.

420 Prior work provides evidence of an interaction between ensemble encoding processes and properties of featural
421 distributions, such as smoothness and range of variance (Utochkin & Tiurina, 2014). However, little is known about
422 how well the visual system might infer whether a distribution is uniform, Gaussian, multimodal, or some other classi-
423 fication. For positional representations, ensemble encoding seems to allow a viewer to readily perceive the mean and
424 variability of a collection, but other aspects of the distribution might be available as well (Figure 13). How well if
425 at all can the visual system perceive distribution shape? And how might encoding data with different visual features,
426 such as color or luminance, affect that ability?

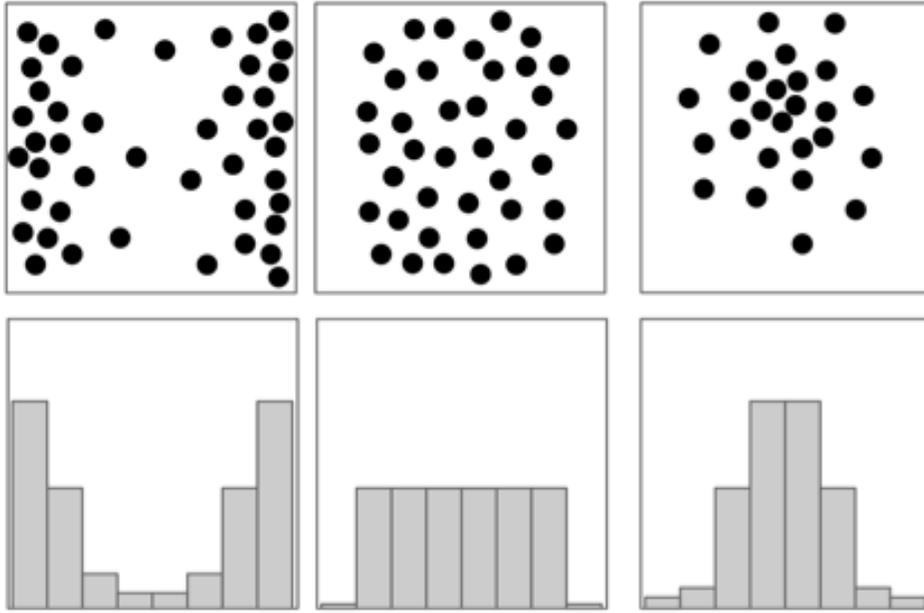


Figure 13: Visualizations may communicate important information about the distribution of values within a dataset, beyond simply mean and variance.

Conclusions

Data visualizations allow us to explore and analyze data using our visual system by mapping data values to spatial positions and visual features. Because viewers can use ensemble processing to efficiently extract statistical information from a dataset, a better understanding of these ensemble mechanisms could provide design guidelines for data displays that maximize a viewer’s ability to process data visually. The fact that many of these guidelines have yet to be firmly established—what types of visual features and displays facilitate what kinds of visual statistical decisions—reveals unsolved questions for the perceptual psychologist. These factors make the study of ensemble processing of data visualizations a fertile territory for collaboration between the perception and visualization communities.

To help organize the discussion at this interface, we have introduced a task categorization, and surveyed both past work and open problems for each category, across perception and visualization. The four categories of tasks are ubiquitous in data visualization: identification, summarization, segmentation, and structure estimation. A single example can clearly illustrate the importance of each of these tasks: tagged text visualization (Figure 5a). You can quickly *identify* outlier text tags in blue in the first paragraph. You can *summarize* that there are two to three dozen red tags in total. You can *segment* the major division between red tags in the top paragraph and blue in the bottom. You can *estimate structure* in the data to detect a red-to-blue trend that is systematically related to vertical position in the text.

In addition to ensemble encoding, research on the control of visual attention is relevant to all four task categories. In the examples that we have explored, we assume that viewers have perfect control over which subset of data they operate on—the datapoints on top, the red objects, the triangles, etc. In reality, visually selecting relevant datapoints can be difficult or noisy. For example, data displays often contain animations, motion, or transients that can distract

446 the viewer (Hollingworth, Simons, & Franconeri, 2010; Bartram, Ware, & Calvert, 2003) and may impair selective
447 ensemble processing of only the relevant visual features. Attentional control can be especially difficult when multiple
448 dimensions of data are depicted simultaneously. For example, a bar graph might map sales to color, profits to height,
449 and time to horizontal order. This visualization would present multiple dimensions of information via multiple visual
450 features simultaneously. Work on attentional control shows that when there is simultaneous variability in multiple
451 feature dimensions, the ‘wrong’ dimension can distract the viewer (Lustig & Beck, 2012). Some existing work ex-
452 plores attentional control in the context of data visualization (see Healey and Enns, 2012, for a survey), but many
453 questions of interest to both communities remain. Both perceptual psychology and data visualization would benefit
454 from a better understanding of whether our current conclusions about attentional control, which draw from a set of
455 laboratory tasks, apply to the more complex displays found in data visualization. Collaboration with data visualization
456 researchers brings this benefit to the perceptual psychologist more generally, as a way of testing whether knowledge
457 gained simplified displays and tasks is robust across new contexts.

458 Our categorization of tasks and links to relevant work in both communities are by no means intended to be exhaus-
459 tive – and it will not be the last word. Instead, our goal is to foster conversation between these communities around
460 ensemble phenomena. We find three themes particularly exciting.

461 First, we assume that the collection of visual processing abilities that we call ensemble encoding—processing
462 of information that can be extracted and combined in parallel from large numbers of objects at once—evolved and
463 developed to compute ‘statistics’ in the natural world. Those statistics are likely to be based on heuristics and other
464 ‘good-enough’ strategies that suffice for the natural world, but we know that some ensemble judgements introduce
465 biases in statistical inferences from data displays that would not be present in formal computed statistics (Sweeny,
466 Haroz, & Whitney, 2012). How common are those biases, which are potentially problematic among different visual
467 features, and how can data displays be designed to avoid them?

468 Second, there is research in data visualization that explores which dimensions allow the most precise extraction of
469 individual values (e.g., Cleveland & McGill, 1984). These studies have found that spatial position is precise, object
470 length is not quite as good, angular extent is bad, and color saturation is among the worst methods for precisely
471 representing individual data values. But these rankings change for ensemble encoding. For example, color depictions
472 can beat spatial position depictions when a viewer needs to analyze average values from a subset of raw data (Albers et
473 al., 2014). While inconsistent at first glance, we believe that this contrast may inspire a new requirement for dimensions
474 that lead to efficient ensemble processing - they may actually have to be *imprecisely* coded, so that their distributions
475 tend to overlap, leading to better representation of distributions as a whole. In contrast, dimensions that are precisely
476 coded may be tougher to combine, because representations of values tend to remain individuated (Franconeri et al.,
477 2013).

478 Third, we believe that collaborative work between these communities will help perceptual psychology researchers
479 define the set of operations that are possible via ensemble encoding. Currently, judgment of average value is the
480 dominant task given to participants in these studies. We hope that we have shown that other judgments, such as range,
481 median, skew, modality, or correlation (e.g., Rensink & Baldrige, 2010), would provide excellent testing grounds for

482 exploring whether these values are extracted via ensemble encoding or by combining ensemble encodings with other
483 visual strategies.

484 **Acknowledgements**

485 We thank Todd Horowitz for helpful advice. We also thank Liqiang Huang and Ron Rensink for their thoughtful
486 comments. This work was funded by NSF awards IIS-1162037 and IIS-1162067.

487 **References**

- 488 Albers, D., Correll, M., & Gleicher, M. (2014). Task-driven evaluation of aggregation in time series visualization. In
489 *Proceedings of the 32nd annual acm conference on human factors in computing systems* (pp. 551–560).
- 490 Albers, D., Dewey, C., & Gleicher, M. (2011). Sequence surveyor: Leveraging overview for scalable genomic
491 alignment visualization. *Visualization and Computer Graphics, IEEE Transactions on*, 17(12), 2392–2401.
- 492 Alexander, E., Kohlmann, J., Valenza, R., Witmore, M., & Gleicher, M. (2014). Serendip: Topic model-driven visual
493 exploration of text corpora. In *Visual analytics science and technology (vast), 2014 ieee conference on* (pp.
494 173–182).
- 495 Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in cognitive*
496 *sciences*, 15(3), 122–131.
- 497 Alvarez, G. A., & Oliva, A. (2008). The representation of simple ensemble visual features outside the focus of
498 attention. *Psychological science*, 19(4), 392–398.
- 499 Amar, R., Eagan, J., & Stasko, J. (2005). Low-level components of analytic activity in information visualization. In
500 *Information visualization, 2005. infovis 2005. ieee symposium on* (pp. 111–117).
- 501 Anscombe, F. J. (1973). Graphs in statistical analysis. *The American Statistician*, 27(1), 17–21.
- 502 Ariely, D. (2001). Seeing sets: Representation by statistical properties. *Psychological Science*, 12(2), 157–162.
- 503 Bartram, L., Ware, C., & Calvert, T. (2003). Moticons:: detection, distraction and task. *International Journal of*
504 *Human-Computer Studies*, 58(5), 515–545.
- 505 Bauer, B. (2010). Does stevens’ power law for brightness extend to perceptual brightness averaging? *The Psycholog-*
506 *ical Record*, 59(2), 2.
- 507 Bertin, J. (1983). Semiology of graphics: diagrams, networks, maps.
- 508 Brehmer, M., & Munzner, T. (2013). A multi-level typology of abstract visualization tasks. *Visualization and*
509 *Computer Graphics, IEEE Transactions on*, 19(12), 2376–2385.
- 510 Bulakowski, P. F., Bressler, D. W., & Whitney, D. (2007). Shared attentional resources for global and local motion
511 processing. *Journal of Vision*, 7(10), 10.
- 512 Burr, D., & Ross, J. (2008). A visual sense of number. *Current Biology*, 18(6), 425–428.

- 513 Callaghan, T. C. (1984). Dimensional interaction of hue and brightness in preattentive field segregation. *Perception*
514 *& psychophysics*, 36(1), 25–34.
- 515 Callaghan, T. C. (1989). Interference and dominance in texture segregation: Hue, geometric form, and line orientation.
516 *Perception & psychophysics*, 46(4), 299–311.
- 517 Chetverikov, A., & Kristjansson, Á. (2014). History effects in visual search for monsters: Search times, choice biases,
518 and liking. *Attention, Perception, & Psychophysics*, 1–11.
- 519 Chevalier, F., Dragicevic, P., & Franconeri, S. (2014). The not-so-staggering effect of staggered animated transitions
520 on visual tracking. *IEEE Transactions on Visualization and Computer Graphics*.
- 521 Chong, S. C., & Treisman, A. (2003). Representation of statistical properties. *Vision research*, 43(4), 393–404.
- 522 Chong, S. C., & Treisman, A. (2005a). Attentional spread in the statistical processing of visual displays. *Perception*
523 *& Psychophysics*, 67(1), 1–13.
- 524 Chong, S. C., & Treisman, A. (2005b). Statistical processing: Computing the average size in perceptual groups. *Vision*
525 *research*, 45(7), 891–900.
- 526 Choo, H., Levinthal, B. R., & Franconeri, S. L. (2012). Average orientation is more accessible through object
527 boundaries than surface features. *Journal of Experimental Psychology: Human Perception and Performance*,
528 38(3), 585.
- 529 Cleveland, W. S., & McGill, R. (1984). Graphical perception: Theory, experimentation, and application to the
530 development of graphical methods. *Journal of the American statistical association*, 79(387), 531–554.
- 531 Corbett, J. E., & Melcher, D. (2014). Characterizing ensemble statistics: mean size is represented across multiple
532 frames of reference. *Attention, Perception, & Psychophysics*, 76(3), 746–758.
- 533 Correll, M., Alexander, E. C., & Gleicher, M. (2013). Quantity estimation in visualizations of tagged text. In
534 *Proceedings of the sigchi conference on human factors in computing systems* (pp. 2697–2706).
- 535 Dakin, S. (2014). Seeing statistical regularities: Texture and pattern perception. In *Handbook of perceptual organiza-*
536 *tion*.
- 537 Dakin, S., & Watt, R. (1997). The computation of orientation statistics from visual texture. *Vision research*, 37(22),
538 3181–3192.
- 539 Dakin, S. C. (2001). Information limit on the spatial integration of local orientation signals. *JOSA A*, 18(5), 1016–
540 1026.
- 541 Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological review*, 96(3), 433.
- 542 Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local
543 association field. *Vision research*, 33(2), 173–193.
- 544 Fouriezos, G., Rubinfeld, S., & Capstick, G. (2008). Visual statistical decisions. *Perception & psychophysics*, 70(3),
545 456–464.
- 546 Franconeri, S. L. (2013). The nature and status of visual resources. *Oxford handbook of cognitive psychology*, 8481,
547 147–162.
- 548 Franconeri, S. L., Alvarez, G. A., & Cavanagh, P. (2013). Flexible cognitive resources: competitive content maps for

549 attention and memory. *Trends in cognitive sciences*, 17(3), 134–141.

550 Franconeri, S. L., Bemis, D., & Alvarez, G. (2009). Number estimation relies on a set of segmented objects. *Cognition*,
551 113(1), 1–13.

552 Franconeri, S. L., Scimeca, J. M., Roth, J. C., Helseth, S. A., & Kahn, L. E. (2012). Flexible visual processing of
553 spatial relationships. *Cognition*, 122(2), 210–227.

554 Gleicher, M., Correll, M., Nothelfer, C., & Franconeri, S. (2013). Perception of average value in multiclass scatterplots.
555 *Visualization and Computer Graphics, IEEE Transactions on*, 19(12), 2316–2325.

556 Graham, D. J., Friedenberg, J. D., McCandless, C. H., & Rockmore, D. N. (2010). Preference for art: similarity,
557 statistics, and selling price. In *Is&tspie electronic imaging* (pp. 75271A–75271A).

558 Haberman, J., & Whitney, D. (2007). Rapid extraction of mean emotion and gender from sets of faces. *Current*
559 *Biology*, 17(17), R751–R753.

560 Halberda, J., Sires, S. F., & Feigenson, L. (2006). Multiple spatially overlapping sets can be enumerated in parallel.
561 *Psychological science*, 17(7), 572–576.

562 Haroz, S., & Whitney, D. (2012). How capacity limits of attention influence information visualization effectiveness.
563 *Visualization and Computer Graphics, IEEE Transactions on*, 18(12), 2402–2410.

564 Harrison, L., Yang, F., Franconeri, S., & Chang, R. (2014). Ranking visualizations of correlation using weber’s law.
565 *IEEE Transactions on Visualization and Computer Graphics*.

566 Healey, C. G., Booth, K. S., & Enns, J. T. (1996). High-speed visual estimation using preattentive processing. *ACM*
567 *Transactions on Computer-Human Interaction (TOCHI)*, 3(2), 107–135.

568 Healey, C. G., & Enns, J. T. (2012). Attention and visual memory in visualization and computer graphics. *Visualization*
569 *and Computer Graphics, IEEE Transactions on*, 18(7), 1170–1188.

570 Heer, J., Bostock, M., & Ogievetsky, V. (2010). A tour through the visualization zoo. *Commun. ACM*, 53(6), 59–67.

571 Heer, J., Kong, N., & Agrawala, M. (2009). Sizing the horizon: the effects of chart size and layering on the graphical
572 perception of time series visualizations. In *Proceedings of the sigchi conference on human factors in computing*
573 *systems* (pp. 1303–1312).

574 Hess, R. F., & Holliday, I. E. (1992). The coding of spatial position by the human visual system: effects of spatial
575 scale and contrast. *Vision Research*, 32(6), 1085–1097.

576 Hollingworth, A., Simons, D. J., & Franconeri, S. L. (2010). New objects do not capture attention without a sensory
577 transient. *Attention, Perception, & Psychophysics*, 72(5), 1298–1310.

578 Holtt, T., Magdy, A., Zhan, P., Chen, G., Gopalakrishnan, G., Hoteit, I., . . . Hadwiger, M. (2014). Ovis: A framework
579 for visual analysis of ocean forecast ensembles.

580 Huang, L., & Pashler, H. (2002). Symmetry detection and visual attention: A binary-map hypothesis. *Vision research*,
581 42(11), 1421–1430.

582 Huang, L., & Pashler, H. (2007). A boolean map theory of visual attention. *Psychological review*, 114(3), 599.

583 Itti, L. (2005). Quantifying the contribution of low-level saliency to human eye movements in dynamic scenes. *Visual*
584 *Cognition*, 12(6), 1093–1123.

585 Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature reviews neuroscience*, 2(3), 194–203.

586 Javed, W., McDonnell, B., & Elmqvist, N. (2010). Graphical perception of multiple time series. *Visualization and*
587 *Computer Graphics, IEEE Transactions on*, 16(6), 927–934.

588 Levinthal, B. R., & Franconeri, S. L. (2011). Common-fate grouping as feature selection. *Psychological science*,
589 0956797611418346.

590 Lin, J., Keogh, E., Lonardi, S., Lankford, J. P., & Nystrom, D. M. (2004). Visually mining and monitoring massive
591 time series. In *Proceedings of the tenth acm sigkdd international conference on knowledge discovery and data*
592 *mining* (pp. 460–469).

593 Lustig, A. G., & Beck, D. M. (2012). Task-relevant and task-irrelevant dimensions are modulated independently at a
594 task-irrelevant location. *Journal of cognitive neuroscience*, 24(9), 1884–1895.

595 Melcher, D., & Kowler, E. (1999). Shapes, surfaces and saccades. *Vision research*, 39(17), 2929–2946.

596 Meyer, M., Munzner, T., & Pfister, H. (2009). Mizbee: a multiscale synteny browser. *Visualization and Computer*
597 *Graphics, IEEE Transactions on*, 15(6), 897–904.

598 Morgan, M., Chubb, C., & Solomon, J. A. (2008). A ‘dipper’ function for texture discrimination based on orientation
599 variance. *Journal of Vision*, 8(11), 9.

600 Morgan, M., & Glennerster, A. (1991). Efficiency of locating centres of dot-clusters by human observers. *Vision*
601 *research*, 31(12), 2075–2083.

602 Nakayama, K., Silverman, G. H., et al. (1986). Serial and parallel processing of visual feature conjunctions. *Nature*,
603 320(6059), 264–265.

604 Neisser, U. (1964). Visual search. *Scientific American*.

605 Oliva, A., & Torralba, A. (2006). Building the gist of a scene: The role of global image features in recognition.
606 *Progress in brain research*, 155, 23–36.

607 Oyama, T. (1961). Perceptual grouping as a function of proximity. *Perceptual and Motor Skills*, 13(3), 305–306.

608 Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded
609 orientation signals in human vision. *Nature neuroscience*, 4(7), 739–744.

610 Pelli, D. G., Palomares, M., & Majaj, N. J. (2004). Crowding is unlike ordinary masking: Distinguishing feature
611 integration from detection. *Journal of vision*, 4(12), 12.

612 Prinzmetal, W., & Taylor, N. (2006). Color singleton pop-out does not always pop out: an alternative to visual search.
613 *Psychonomic bulletin & review*, 13(4), 576–580.

614 Rensink, R. A. (2014). On the prospects for a science of visualization. In *Handbook of human centric visualization*
615 (pp. 147–175). Springer.

616 Rensink, R. A., & Baldridge, G. (2010). The perception of correlation in scatterplots. , 29(3), 1203–1210.

617 Rock, I., & Palmer, S. (1990). Gestalt psychology. *Sci Am*, 263, 84–90.

618 Ross, J., & Burr, D. C. (2010). Vision senses number directly. *Journal of Vision*, 10(2), 10.

619 Roth, R. E. (2012). Cartographic interaction primitives: Framework and synthesis. *The Cartographic Journal*, 49(4),
620 376–395.

- 621 Saenz, M., Buraças, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision*
622 *research*, 43(6), 629–637.
- 623 Schulz, H.-J., Nocke, T., Heitzler, M., & Schumann, H. (2013). A design space of visualization tasks. *Visualization*
624 *and Computer Graphics, IEEE Transactions on*, 19(12), 2366–2375.
- 625 Scimeca, J. M., & Franconeri, S. L. (2014). Selecting and tracking multiple objects. *Wiley Interdisciplinary Reviews:*
626 *Cognitive Science*.
- 627 Serences, J. T., Shomstein, S., Leber, A. B., Golay, X., Egeth, H. E., & Yantis, S. (2005). Coordination of voluntary
628 and stimulus-driven attentional control in human cortex. *Psychological Science*, 16(2), 114–122.
- 629 Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. In *Visual*
630 *languages, 1996. proceedings., ieee symposium on* (pp. 336–343).
- 631 Sweeny, T. D., Haroz, S., & Whitney, D. (2012). Reference repulsion in the categorical perception of biological
632 motion. *Vision research*, 64, 26–34.
- 633 Trick, L. M., & Enns, J. T. (1997). Clusters precede shapes in perceptual organization. *Psychological Science*, 8(2),
634 124–129.
- 635 Utochkin, I. S., & Tiurina, N. A. (2014). Parallel averaging of size is possible but range-limited: A reply to marchant,
636 simons, and de fockert. *Acta psychologica*, 146, 7–18.
- 637 Uttal, W. R., & Tucker, T. E. (1977). Complexity effects in form detection. *Vision Research*.
- 638 Ware, C. (2008). *Visual thinking for design*. Morgan Kaufman.
- 639 Ware, C. (2013). *Information visualization: Perception for design, 3e*. Morgan Kaufman.
- 640 Watamaniuk, S. N., & Duchon, A. (1992). The human visual system averages speed information. *Vision research*,
641 32(5), 931–941.
- 642 Watamaniuk, S. N., Sekuler, R., & Williams, D. W. (1989). Direction perception in complex dynamic displays: the
643 integration of direction information. *Vision research*, 29(1), 47–59.
- 644 Webster, J., Kay, P., & Webster, M. A. (2014). Perceiving the average hue of color arrays. *JOSA A*, 31(4), A283–A292.
- 645 Whitaker, D., McGraw, P. V., Pacey, I., & Barrett, B. T. (1996). Centroid analysis predicts visual localization of
646 first-and second-order stimuli. *Vision Research*, 36(18), 2957–2970.
- 647 Whitney, D., Haberman, J., & Sweeny, T. (2014). From textures to crowds: multiple levels of summary statistical
648 perception. *The New Visual Neurosciences., JS Werner and LM Chalupa, eds.(MIT, 2014)*, 685–709.
- 649 Whitney, D., & Levi, D. M. (2011). Visual crowding: a fundamental limit on conscious perception and object
650 recognition. *Trends in cognitive sciences*, 15(4), 160–168.
- 651 Wickens, C. D., & Carswell, C. M. (1995). The proximity compatibility principle: its psychological foundation and
652 relevance to display design. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 37(3),
653 473–494.
- 654 Wolfe, J. (1998). What can 1,000,000 trials tell us about visual search? *Psychological Science*, 9, 33–39.

Appendix A: Sample of Visualization Tasks

Task	Task Category	Visual Feature	Data Type	Sample References
Categorize datasets	Segmentation			Wehrend & Lewis, 1990
Categorize datasets	Segmentation	Line Height		Nourbakhsh & Ottenbacher, 1994
Characterize distribution	Structure Estimation			Amar, Eagan, & Stasko, 2005 Holt et al, 2014 Schulz et al, 2013 Sopan et al, 2014 Zhou & Feiner, 1998
Characterize distribution	Structure Estimation		Graph	Lee et al, 2006
Cluster	Segmentation			Aigner et al, 2008 Amar, Eagan, & Stasko, 2005 Buja, Cook, & Swayne, 1996 Hibino, 1999 Schulz et al, 2013 Tory & Moller, 2004 Wehrend & Lewis, 1990 Zhou & Feiner, 1998
Cluster	Segmentation		Graph	Lee et al, 2006
Cluster	Segmentation		Multiscale	Cui et al, 2006
Compute Derived Value	Summary			Amar, Eagan, & Stasko, 2005
Correlation	Structure Estimation	Position	2D	Doherty et al, 2007 Harrison et al, 2014 Rensink and Baldrige, 2010
Correlation	Structure Estimation			Amar, Egan, & Stasko, 2005 Harrison et al, 2014 Schulz et al, 201 Wehrend & Lewis, 1990 Zhou & Feiner, 1998
Correlation	Structure Estimation		Graph	Lee et al, 2006
Correlation	Structure Estimation	Angle	Categorical Data	Bendix, Kosara, & Hauser, 2005 Meyer, Munzner, & Pfister, 2009
Detect Ordering	Structure Estimation			Amar, Eagan, & Stasko, 2005
Extrema	Identification	Color	Time Series	Albers, Correll, & Gleicher, 2014 Fuchs et al, 2013
Extrema	Identification	Line Height	Time Series	Albers, Correll, & Gleicher, 2014 Fuchs et al, 2013
Extrema	Identification	Position	Time Series	Fuchs et al, 2013
Extrema	Identification	Height	Time Series	Fuchs et al, 2013
Extrema	Identification	Angle	Time Series	Fuchs et al, 2013
Extrema	Identification		Network	Lee et al, 2006

Grouping	Segmentation			Buja, Cook, & Swayne, 1996 Meyer, Munzner, & Pfister, 2009 Shneiderman, 1996 Zhou and Feiner, 1998
High-Order Statistics	Structure Estimation			Hollt et al, 2014
Identify Changes/Consistencies	Segmentation		Spatiotemporal	Andrienko, Andrienko, & Gatalisky, 2002 Heer & Robertson, 2007
Identify connected elements	Segmentation	Connectivity	Graph	Henry, Fekete, & McGuffin, 2007 Lee et al 2006 Tory & Moller, 2004
Identify connected elements	Segmentation	Connectivity	Sequences	Meyer, Munzner, & Pfister, 2009
Mean	Summary	Color	Time Series	Albers, Correll, & Gleicher, 2014 Correll et al, 2012
Mean	Summary	Line Height	Time Series	Albers, Correll, & Gleicher, 2014 Correll et al, 2012
Mean	Summary	Position	Time Series	Gleicher et al, 2013
Mean	Summary	Number	Number	Morris & Masnick, 2014
Motif Extraction	Structure Estimation		Time Series	Lin et al., 2004 Meyer, Munzner, & Pfister, 2009
Numerosity	Summary	Color	Text	Correll, Alexander, & Gleicher, 2013
Numerosity	Summary	Length	Text	Correll, Alexander, & Gleicher, 2013
Numerosity	Summary	Connected Points	Graph	Lee et al, 2006
Outliers	Identification	Color	Time Series	Albers, Correll, & Gleicher, 2014
Outliers	Identification	Line Height	Time Series	Albers, Correll, & Gleicher, 2014
Outliers	Identification			Amar, Egan, & Stasko, 2005 Buja, Cook, & Swayne, 1996 Elmqvist, Stasko, & Tsigas, 2008 Hibino, 1999 Schulz et al, 2013 Tory & Moller, 2004
Outliers	Identification		Graph	Lee et al, 2006
Outliers	Identification		Time Series	Lin et al, 2004
Pattern Identification	Structure Estimation		Time Series	Aigner et al, 2007 Aigner et al, 2008
Pattern Identification	Structure Estimation		Discrete	Borkin et al, 2011 Jerding & Stasko, 1995 Tory & Moller, 2004
Range	Identification	Color	Time Series	Albers, Correll, & Gleicher, 2014
Range	Identification	Line Height	Time Series	Albers, Correll, & Gleicher, 2014; Aigner et al, 2010

Range	Identification			Amar, Eagan, & Stasko, 2005
Relatedness/Similarity	Structure Estimation			Amar & Stasko, 2004 Hibino, 1999 Jerding & Stasko, 1995 Schulz et al, 2013 Shneiderman, 1996
Select Subsets	Segmentation			Shneiderman, 1996
Set Operations	Segmentation		Graph	Lee et al, 2006
Set Operations	Segmentation		High Dimensional	Elmqvist, Stasko, & Tsigas, 2008
Spatial Relationships	Structure Estimation		Spatial Data	Tory & Moller, 2004
Subsequence Similarity	Structure Estimation		Time Series	Lin et al. 2004 Meyer, Munzner, & Pfister, 2009
Trends	Structure Estimation	Line Height	Time Series	Aigner et al, 2010 Best, Smith & Stubbs, 2007 Fuchs et al, 2013 Meserth & Hollands, 1999
Trends	Structure Estimation	Color	Time Series	Aigner et al, 2010 Fuchs et al, 2013
Trends	Structure Estimation			Buja, Cook, & Swayne, 1996 Hibino, 1999 Schulz et al, 2013 Slingsby, Dykes, & Wood, 2009 Tory & Moller, 2004
Trends	Structure Estimation	Position	Time Series	Fuchs et al, 2013
Trends	Structure Estimation	Height	Time Series	Fuchs et al, 2013
Trends	Structure Estimation	Angle	Time Series	Fuchs et al, 2013
Trends	Structure Estimation	Angle	Categorical Data	Bendix, Kosara, & Hauser, 2005
Trends	Structure Estimation	Angle	Spatial Data	Livingston & Decker, 2011
Trends	Structure Estimation	Color	Spatial Data	Livingston & Decker, 2011
Variance	Summary	Color	Time Series	Albers, Correll, & Gleicher, 2014
Variance	Summary	Line Height	Time Series	Aigner et al., 2010 Albers, Correll, & Gleicher, 2014

Appendix A References:

Aigner, W., Bertone, A., Miksch, S., Tominski, C., & Schumann, H. (2007). Towards a conceptual framework for visual analytics of time and time-oriented data. *Simulation Conference, 2007 Winter*, 721-729.

- Aigner, W., Miksch, S., Muller, W., Schumann, H., & Tominski, C. (2008). Visual methods for analyzing time-oriented data. *IEEE Transactions on Visualization and Computer Graphics*, 14(1), 47-60.
- Aigner, W., Kainz, C., Ma, R., & Miksch, S. (2010). Bertin was Right: An Empirical Evaluation of Indexing to Compare Multivariate Time-Series Data Using Line Plots. *Computer Graphics Forum*, 30(1), 215-228.
- Albers, D., Correll, M., & Gleicher, M. (2014). Task-driven evaluation of aggregation in time series visualization. *Proceedings of the 32nd Annual SIGCHI Conference on Human Factors in Computing Systems*, 551-560.
- Amar, R., Eagan, J., & Stasko, J. (2005). Low-level components of analytic activity in information visualization. *IEEE Symposium on Information Visualization, 2005*, 111-117.
- Andrienko, N., Andrienko, G., & Gatalsky, P. (2002). Data and Task Characteristics in Design of Spatio-Temporal Data Visualization Tools. *Symposium on Geospatial Theory, Processing, and Applications*.
- Bendix, F., Kosara, R., & Hauser, H. (2005). Parallel sets: Visual analysis of categorical data. *IEEE Symposium on Information Visualization, 2005*, 133-140.
- Best, L. A., Smith, L. D., & Stubbs, D. A. (2007). Perception of linear and nonlinear trends: Using slope and curvature information to make trend discriminations 1, 2. *Perceptual and Motor Skills*, 104(3), 707-721.
- Borkin, M., Gajos, K., Peters, A., Mitsouras, D., Melchionna, S., Rybicki, F., Feldman, C., & Pfister, H. (2011). Evaluation of artery visualizations for heart disease diagnosis. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), 2479-2488.
- Buja, A., Cook, D., & Swayne, D. F. (1996). Interactive high-dimensional data visualization. *Journal of Computational and Graphical Statistics*, 5(1), 78-99.
- Correll, M., Albers, D., Franconeri, S., & Gleicher, M. (2012). Comparing averages in time series data. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1095-1104.
- Correll, M. A., Alexander, E. C., & Gleicher, M. (2013). Quantity estimation in visualizations of tagged text. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2697-2706.
- Cui, Q., Ward, M. O., Rundensteiner, E. A., & Yang, J. (2006). Measuring data abstraction quality in multiresolution visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 12(5), 709-716.
- Doherty, M. E., Anderson, R. B., Angott, A. M., & Klopfer, D. S. (2007). The perception of scatterplots. *Perception & Psychophysics*, 69(7), 1261-1272.
- Elmqvist, N., Stasko, J., & Tsigas, P. (2008). DataMeadow: a visual canvas for analysis of large-scale multivariate data. *Information Visualization*, 7(1), 18-33.
- Fuchs, J., Fischer, F., Mansmann, F., Bertini, E., & Isenberg, P. (2013). Evaluation of alternative glyph designs for time series data in a small multiple setting. *Proceedings of the SIGCHI Annual Conference on Human Factors in Computing Systems*, 3237-3246.
- Gleicher, M., Correll, M., Nothelfer, C., & Franconeri, S. (2013). Perception of average value in multiclass scatterplots. *IEEE Transactions on Visualization and Computer Graphics*, 19(12), 2316-2325.
- Harrison, L., Yang, F., Franconeri, S., & Chang, R. (2014). Ranking Visualizations of Correlation Using Weber's Law. *IEEE Transactions on Visualization and Computer Graphics*, 20(12), 1943-1952.
- Heer, J., & Robertson, G. G. (2007). Animated transitions in statistical data graphics. *IEEE Transactions on Visualization and Computer Graphics*, 13(6), 1240-1247.

- Henry, N., Fekete, J., & McGuffin, M. J. (2007). NodeTrix: a hybrid visualization of social networks. *IEEE Transactions on Visualization and Computer Graphics*, 13(6), 1302-1309.
- Hibino, S. L. (1999). Task analysis for information visualization. *Visual Information and Information Systems*, 139-146.
- Hollt, T., Magdy, A., Zhan, P., Chen, G., Gopalakrishnan, G., Hoteit, I., Hansen, C. D., & Hadwiger, M. (2014). Ovis: A Framework for Visual Analysis of Ocean Forecast Ensembles. *IEEE Transactions on Visualization and Computer Graphics*, 20(8), 1114-1126.
- Jerding, D. F., & Stasko, J. T. (1998). The information mural: A technique for displaying and navigating large information spaces. *IEEE Transactions on Visualization and Computer Graphics*, 4(3), 257-271.
- Lee, B., Plaisant, C., Parr, C. S., Fekete, J. D., & Henry, N. (2006, May). Task taxonomy for graph visualization. In *Proceedings of the 2006 AVI workshop on BEyond time and errors: Novel Evaluation Methods for Information Visualization*, 1-5.
- Lin, J., Keogh, E., Lonardi, S., Lankford, J. P., & Nystrom, D. M. (2004). Visually mining and monitoring massive time series. *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 460-469.
- Livingston, M. A., & Decker, J. W. (2011). Evaluation of trend localization with multi-variate visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), 2053-2062.
- Meyer, M., Munzner, T., & Pfister, H. (2009). MizBee: a multiscale synteny browser. *IEEE Transactions on Visualization and Computer Graphics*, 15(6), 897-904.
- Meserth, T. A., & Hollands, J. G. (1999). Comparing 2D and 3D displays for trend estimation: The effects of display augmentation. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 43(23), 1308-1312.
- Morris, B. J., & Masnick, A. M. (2014). Comparing data sets: Implicit summaries of the statistical properties of number sets. *Cognitive Science*, 39(1), 156-170.
- Nourbakhsh, M. R., & Ottenbacher, K. J. (1994). The statistical analysis of single-subject data: a comparative examination. *Physical Therapy*, 74(8), 768-776.
- Rensink, R. A., & Baldrige, G. (2010). The perception of correlation in scatterplots. *Computer Graphics Forum*, 29(3), 1203-1210.
- Schulz, H. J., Nocke, T., Heitzler, M., & Schumann, H. (2013). A design space of visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12), 2366-2375.
- Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. *IEEE Symposium on Visual Languages*, 336-343.
- Slingsby, A., Dykes, J., & Wood, J. (2009). Configuring hierarchical layouts to address research questions. *IEEE Transactions on Visualization and Computer Graphics*, 15(6), 977-984.
- Sopan, A., Freier, M., Taieb-Maimon, M., Plaisant, C., Golbeck, J., & Shneiderman, B. (2013). Exploring data distributions: Visual design and evaluation. *International Journal of Human-Computer Interaction*, 29(2), 77-95.
- Tory, M., & Moller, T. (2004). Rethinking visualization: A high-level taxonomy. *IEEE Symposium on Information Visualization, 2004*, 151-158.

Wehrend, S., & Lewis, C. (1990). A problem-oriented classification of visualization techniques. *Proceedings of the 1st Conference on Visualization 1990*, 139-143. IEEE Computer Society Press.

Zhou, M. X., & Feiner, S. K. (1998). Visual task characterization for automated visual discourse synthesis. *Proceedings of the SIGCHI Annual Conference on Human Factors in Computing Systems*, 392-399.