Vision, Modeling, and Visualization (2022)
J. Bender, M. Botsch, and D. Keim (Eds.)

Honeycomb Plots: Visual Enhancements for Hexagonal Maps

T. Trautner[1] M. Sbardellati[2] S. Stoppel[1,3] and S. Bruckner[1]

[1] University of Bergen, Norway
[2] TU Wien, Austria
[3] Wolftech Broadcast Solutions AS

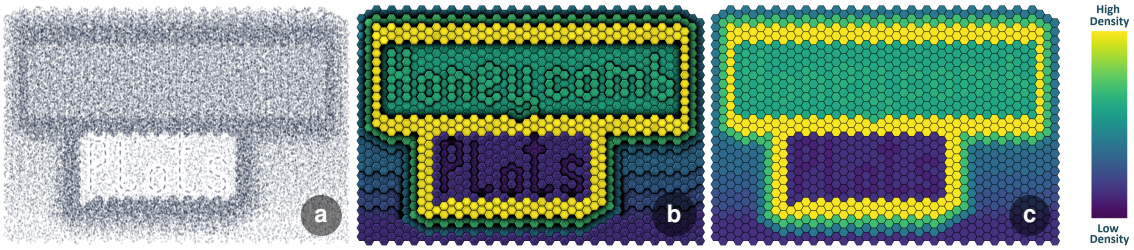

Figure 1: *Honeycomb plots (b) combine (a) scatter plots and (c) heat maps using per-tile densities. Through shading, ambient occlusion, and an implicit point-distribution encoding, the observer can explore data features that cannot be captured by either of the techniques alone.*

Abstract

Aggregation through binning is a commonly used technique for visualizing large, dense, and overplotted two-dimensional data sets. However, aggregation can hide nuanced data-distribution features and complicates the display of multiple data-dependent variables, since color mapping is the primary means of encoding. In this paper, we present novel techniques for enhancing hexplots with spatialization cues while avoiding common disadvantages of three-dimensional visualizations. In particular, we focus on techniques relying on preattentive features that exploit shading and shape cues to emphasize relative value differences. Furthermore, we introduce a novel visual encoding that conveys information about the data distributions or trends within individual tiles. Based on multiple usage examples from different domains and real-world scenarios, we generate expressive visualizations that increase the information content of classic hexplots and validate their effectiveness in a user study.

CCS Concepts

• ***Human-centered computing*** → ***Visualization techniques; Visualization theory, concepts and paradigms;***

1. Introduction

Hexplots represent a form of spatial aggregation usually applied to a large number of two-dimensional points, making them resistant to overplotting. They rely on a subdivision of the plane by a regular space-filling grid of hexagonal polygons. However, Cleveland and McGill [CM85] emphasize that the visual system may fail to detect quantitative information from geometric aspects of a visualization. For example, aggregation comes at the expense of the perceptibility of individual points, making outliers invisible [Dow14], and potentially obscures trends or clusters. Additionally, being able to distinguish between tile colors with at least a just-noticeable difference (JND) can be essential, especially when color describes quantitative data properties [Sto12; SAS14]. We, therefore, assessed potential disadvantages of current techniques by applying *algebraic visualization design* (AVD) [KS14]. Our strategy was inspired by McNutt [McN21], who highlights the advantages of this human operable and interpretable systematic framework. AVD analyzes how data changes affect the resulting visualization. Input for honeycomb plots are tabular 2D point coordinates. If the same data is displayed differently, this is referred to as a *hallucinator*. When changes in the data remain invisible, this is referred to as a *confuser*. Although both flaws cannot be avoided completely, they should be minimized. Therefore, we initially noted the following *confusers*:

C1: Point-data visualizations often assume equal density ranges. The underlying points are either sparse (scatter plots) or dense (density estimations and aggregations), making techniques reach their limits when data exhibit both simultaneously.

DOI: 10.2312/vmv.20221205

C2: Color-coded tiles containing a similar number of points may be perceived as equally dense, which would falsely correspond to the same visualization of actually different data.

C3: Heat maps only encode the quantity of aggregated points, so that uniform distributions within tiles cannot be distinguished from clusters or trends if their numbers of points match.

C4: Heat maps prevent the analysis of sparse features or outliers, allowing their position or arrangement to change although the aggregation, i.e., tile color, remains the same.

These confusers highlight the need for (C1) a novel hybrid visualization technique that preserves features of highly uneven distributions. Our goal was (C2) a visual encoding that overcomes problems with color coding of the first moment of statistics, i.e., the mean value describing average density as color, but also to (C3) capture statistical moments of higher order describing shape parameters, while (C4) strengthening the embedding of underlying points. Our contribution can be summarized as follows:

- An interpretation of hexagonal tiles as *relief mosaic* where ambient occlusion serves as a subtle aid to perceive nuanced color differences between neighboring discrete tiles.
- An extension of the visually encoded information content by incorporating the regression plane of the underlying densities per bin in the form of a per-tile *diamond cut*.
- A hybrid approach that blends point data with colored hexagonal tiles corresponding to an *amber inclusions* metaphor, enabling an exploration of trends and clusters also in sparse regions.
- A quantitative user study consisting of four typical hexplot tasks in which we compare our encodings to classic heat maps.

In addition to the listed contributions, we developed an efficient hexagonal aggregation algorithm for two-dimensional data points, based on arbitrary grid sizes, in real time on the GPU.

2. Related Work

Hexagonal binning as an aggregation technique was first introduced by Carr et al. [CLNL87] in 1987. They suggest the use of glyphs in the shape of hexagons whose size encodes the total number of data points within. Another technique called sunflower plots [CM84] uses flower-like glyphs in which petals represent aggregated points. A hybrid approach that visualizes individual data points as well as density estimations through binning are variable resolution bivariate plots, so-called varebi plots [HMS97]. Over the years, a variety of related techniques spread under similar names: *hexagon plots* or *hexplots*, *hexbins*, *hexagonal binning plots*, *hexagonal tiles*, *hexbin maps*, *hexagonal gridded maps*, and *hexagonal heatmaps*.

For statistical summaries in cartography, Carr et al. [COW92] emphasize the advantages of hexagon mosaic maps, i.e., hexagon grid-cell choropleth maps, extendable by information layers [PMAM16]. Battersby et al. [BSF16] investigate projection distortions of such grid structures resulting in either differently sized geographic areas as the basis for binning, or bin grid overlays of varying size. Example scenarios for spatio-temporal cartographic data contain the analysis of criminal activity [RRM15] or visitor flows in amusement parks when solving the VAST 2015 Mini-Challenge (MC1) [CGH*15]. Another recent approach [WK20] focuses on multivariate game metrics using hexbin maps, Wurman dots [PW66], and arrowheads as direction indicators. To explore data attributes at specific (geospatial) locations, so-called attribute blocks [Mil07], i.e., dynamically configurable regular arrays of "screen door" lenses, can be used. Possible alternatives are hexagonal cells as windows or magic lenses enabling the comparative visualization of multiple data sets [MHG10]. Other techniques [HKIH07] evaluate blending and weaving to encode multivariate information by color. Established guidelines derived from exploring the visual design space of multi-class point data [HCSG18] and different representations using similar approaches [JVDF19] evaluate blending, weaving, majority-based coloring, embedding of pie charts or bar charts as glyphs, etc. Furthermore, user studies analyzing user performance when comparing multiple heat maps [KAB*20] show that juxtaposed 2D heat maps work best for overview tasks, whereas stacked 3D heat maps (explored using stereo vision) are superior when reading and comparing single values. Although the two-dimensionality of honeycomb plots prevents occlusions and perceptual distortions, our spatialization cues could also be mapped to 3D geometry or digitally fabricated.

3. Honeycomb Plots

Hexplots, in contrast to choropleth maps [Dup26], enable a fairer spatial comparison of aggregated values since differently sized landmasses do not affect their interpretation. Although well suited for this, hexplots also limit the parameter space for additional visual encodings since *position*, *area*, *color*, *orientation*, and *shape*, as discussed by Munzner [Mun15], are already occupied. We, therefore, investigated potential spatialization cues, i.e., three-dimensional shape cues, applicable to 2D visualizations, preventing disadvantages of 3D visualizations such as viewpoint choice and view-dependent occlusion. Spatialization cues have a long tradition, for example, cushion treemaps [vWvdW99] and shaded Voronoi diagrams [TW01], enridged contour maps [vWT01], sunspot plots [TBSB20], and line weaver [TB21]. In all cases, luminance from shading expands the design space.

Perception studies analyze *lightness constancy*, i.e., if color and shadow can be separated by the human visual system. Szafir et al. [SSG16] analyze molecular visualizations and confirm that users can assign shadowed colors to corresponding unshadowed colors. Langer and Bülthoff [LB01] highlight that viewers assume light from top left and that a (convex) object is likely observed from above. Irani et al. [ISS04; ISS06] also showed that users prefer shaded visualizations which help them understand structures.

3.1. Relief Mosaic

Hexplots are well suited to spatially aggregate data points. However, in cases where minimal color differences are not merely attributable to falling within a margin of error, discrepancies between adjacent tiles are difficult to see. Inspired by locally adjusting heat maps [ZZW*21], we apply *ambient occlusion* (AO), i.e., encoding accessibility of a surface point [ZIK98], to spatial density which is represented by the tile height. We exploit the advantages of a regular hexagonal grid and calculate AO analytically which is, in such a well-defined scenario, most precise and interactive in real time.

As known from literature [Dut21; AMGA12], the AO integral of a surface point x occluded by a rectangle Q, defined by four corner points $q_{1,2,3,4}$, is calculated as follows:

$$AO(x,Q) = \frac{1}{2\pi}\sum_{i=1}^{3}\left(cos^{-1}\frac{v_i \cdot v_{i+1}}{\| v_i \| \cdot \| v_{i+1} \|}\right)\left(\frac{v_i \times v_{i+1}}{\| v_i \times v_{i+1} \|}\cdot N\right) \tag{1}$$

where $v_i = q_i - x$ corresponds to vectors pointing from x to the corner points of Q, and N is the surface normal of x. According to Quilez [Qui21], the double integral of the occluding surface can be reduced to a line integral of its perimeter, whose individual sides projected on the unit hemisphere correspond to the angle of the arc itself, which can be computed as the inverse cosine of two consecutive corner points. The contribution of the occlusion is then calculated as dot product between the surface normal at the point of occlusion and the normal of the triangle connecting the two corner points to the surface point. Unlike shadow mapping [Wil78], where a light source casts shadows, AO emphasizes the accessibility of a tile to its immediate neighbors. A point x can, therefore, only be occluded by a maximum of six rectangles corresponding to the adjacent side walls of potentially higher tiles. The total AO can, hence, be calculated as the sum of six individual occlusions:

$$AO(x,\mathbf{Q} = \{Q_{1,\cdots,6}\}) = \sum_{j=1}^{6} AO(x,Q_j) \tag{2}$$

The summarized AO factor can then be used as per-pixel darkening factor when coloring tiles. Figure 2 shows how AO can help understanding the spatial interpretation of scalars as a height field.

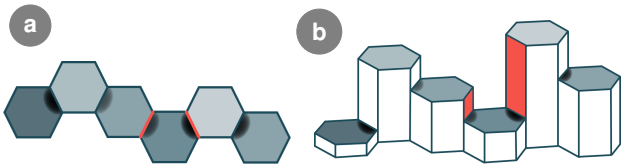

Figure 2: *Illustrations of how AO emphasizes structures using a viewpoint from (a) top-down and (b) "the side", including two red rectangular side walls that cause such an exemplary tile darkening.*

Figure 4(a-c) shows a synthetic example of a density increase from left to right, observable in (a) the scatter plot and (b) the heat map. The vertical trend in the center, however, can only be emphasized by (c) AO. Such an augmentation can be beneficial in two ways: first, to highlight color differences in continuous heat maps, and second, to visually distinguish more nuances than the number of colors that is available in a discrete heat map. For example, cartographers rarely use more than seven colors on choropleth maps [HB03]. Using a relief mosaic also shows that explaining the semantics of a heat map, i.e., which colors correspond to the maximum/minimum density, is now preattentive as it can be derived from the topological structure. Since darkening, however, might not always distort colors favorably, we conducted a user study (see Section 7) to evaluate its impact. We chose *viridis* as the heat map for all our visualizations, as it corresponds to the default in many standard tools and scientific environments, such as the python package matplotlib or R for statistical computing and graphics. It is perceptually uniform even when printed in black-and-white and was developed to improve readability for users with color deficiency or color blindness. Since none of our techniques focuses on absolute numbers, the color legend is only visible in the teaser Figure 1.

3.2. Diamond Cut

Classic hexplots rely on flat tiles, leading to the intuitive assumptions that densities within tiles are either uniform or neglectable. However, non-uniformly distributed points may form trends that are relevant for analysis. Inspired by hexagonally shaped glyphs [CLNL87] and Phoenixmaps [ZLG*21] as an abstraction of data points, we introduce a diamond cut metaphor. Similar to how a rough diamond is cut into a brilliant, we cut hexagonal pyramids so that their cut surface corresponds to the *regression plane* of the contained points, i.e., a fit plane that best approximates the underlying density distribution. This requires the computation of six intersections, one of which called p' can be calculated as follows:

$$p' = l_0 + l\frac{(p_0 - l_0)\cdot N}{l \cdot N} \quad \text{with} \quad l = \frac{l_0 - p_i}{\| l_0 - p_i \|} \tag{3}$$

wherein l_0 is the pyramid top, p_0 is the center point of a tile at the height of the average density, N is the normal, i.e., the regression plane orientation, and l is the normalized direction pointing from one of the six hexagonal grid corner points p_i to the pyramid top l_0. An illustration of this calculation is shown in Figure 3.

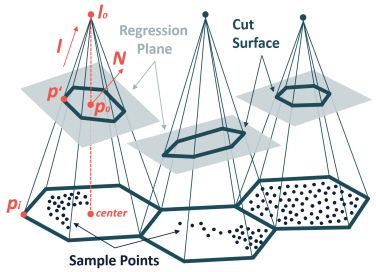

Figure 3: *Pyramid intersections with planes (tilted to the left/right and horizontally) that best approximate the underlying density.*

These slopes can be interpreted as diamonds pointing towards the steepest descent. The smaller and narrower a glyph gets and the more it is attracted to one side of the hexagon, i.e., the more it is repelled from another side, the steeper and more extreme the trend. An illustration of this is shown in Figure 5. Alternatively, inverting the pyramids, i.e., using concave cavities rather than convex spikes, without changing the orientation of the regression planes would instead produce glyphs reminiscent of *radar charts* where the attraction in one direction points towards the steepest ascent.

Figure 4(d-f) shows a data set containing three regions with different densities visible using (e) a heat map, but not in (d) the scatter plot. Using (f) a diamond cut, the regression plane within the tile can be inferred from the shape, shading, and outline of the truncated pyramid. To evaluate whether users understand this concept within 15-20 minutes, we conducted a user study (see Section 7).

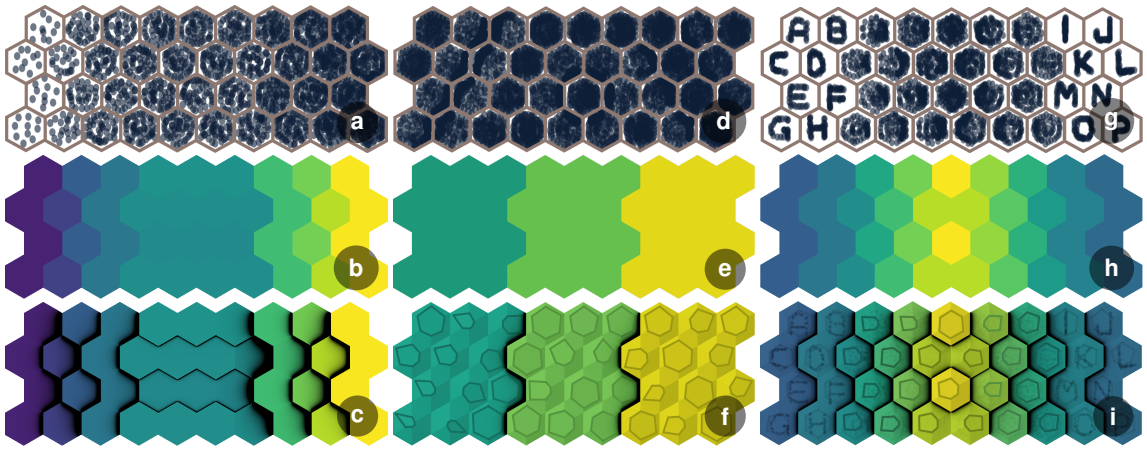

Figure 4: *Overview of three scenarios comparing scatter plots and heat maps to our techniques: scenario one (a-c) shows the applicability of the relief mosaic, scenario two (d-f) highlights advantages of the diamond cut, and scenario three (g-i) illustrates amber inclusions.*

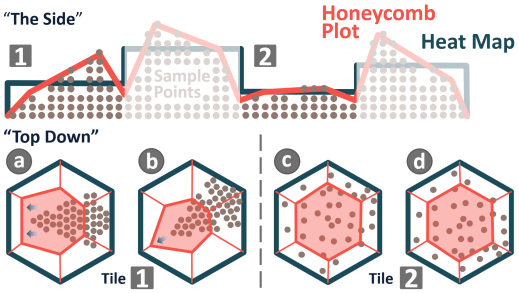

Figure 5: *Comparison of four neighboring tiles aggregating points viewed from "the side". The heat map (blue line) corresponds to the average density, whereas the diamond cut (red line) provides per-tile density distributions. When viewed "top down" as usual, tile (1) shows two possible decreases from right to left (a,b) and tile (2) exemplifies two feasible, almost uniform, distributions (c,d).*

3.3. Amber Inclusions

Regression planes are insightful but, unfortunately, cannot capture outliers. Similar to how Novotny and Hauser [NH06] preserved outliers in parallel coordinate plots (PCP), points within sparse regions could be analyzed as in scatter plots. Therefore, we follow approaches like sunspot plots [TBSB20] or Splatterplots [MG13] and combine hexplots with the underlying points based on the relative densities of tiles. We employ a per-tile opacity modulation mapping the transparency of a tile (encoding a relief mosaic and a diamond cut) to the interval $[0, \textit{density of the densest tile}]$. This results in a metaphor where flat and transparent amber shows enclosed insects rather than thick cloudy amber. We use the Porter-Duff [PD84] *over* operator which defines **A over B** as:

$$C = \frac{\alpha_A A + (1-\alpha_A)\,\alpha_B B}{\alpha_A + (1-\alpha_A)\,\alpha_B} \tag{4}$$

wherein C is the final pixel color, A the hexplot, B the scatter plot, α_A the normalized density/opacity, and α_B the scatter plot opacity. Figure 6 shows tiles with different densities using (a) amber inclusions, (b) their combination with diamond cut spatialization

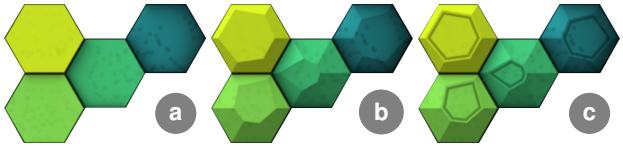

Figure 6: *(a) Blended tiles with (b) diamond cuts and (c) contours.*

cues, and (c) additional 2D contour polygons emphasizing the cut surface. Figure 4(g-i) shows a synthetic data set with particularly distinctive sparse features in the form of letters from A to P, which can be clearly seen in (g) the scatter plot. The (h) heat map captures the density increase in the center but obscures the distinctive letters. Using (i) amber inclusions, however, both the sparse features and the dense data characteristics are preserved. We, again, evaluated whether blending complicates the interpretation of the other encodings as part of our user study (see Section 7).

4. Usage Examples

In this section, we demonstrate the benefits of honeycomb plots using real-world usage examples. In each one, we show how classic hexplots can be enhanced to increase their information content.

4.1. US Tornadoes (1950 - 2019)

The NOAA's National Weather Service - Storm Prediction Center has been collecting tornado data [NOA21] since 1950. Figure 7 shows a honeycomb plot aggregating points corresponding to the starting position of a tornado between 1950 and 2019. In the center of the US, there is the so-called "Tornado Alley". The visualization, furthermore, reveals that most of Florida's tornadoes form along the west coast between Tampa Bay and Forth Myers. Nevertheless, there is a large number of tornadoes visible as continuous black dotted line along the east coast. This phenomenon describes so-called "waterspouts", which are extraordinary tornadoes forming over water. Additional indicators for this are the diamond cuts along the coast pointing towards the open sea, illustrating the highly decreasing number of tornadoes from this direction.

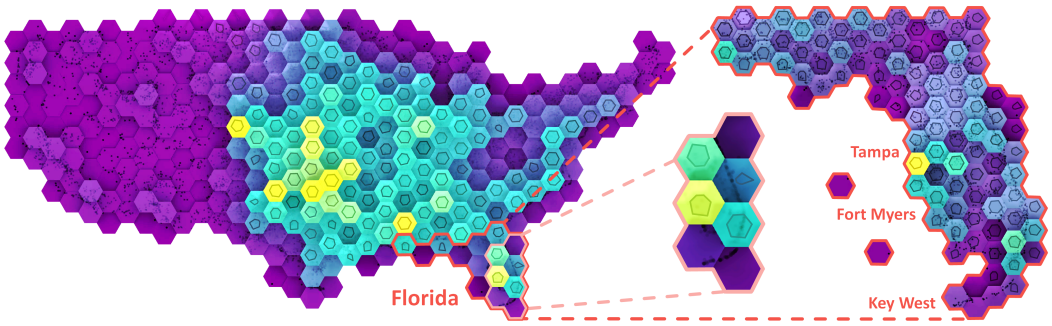

Figure 7: *Honeycomb plot of US tornadoes between 1950 and 2019.*

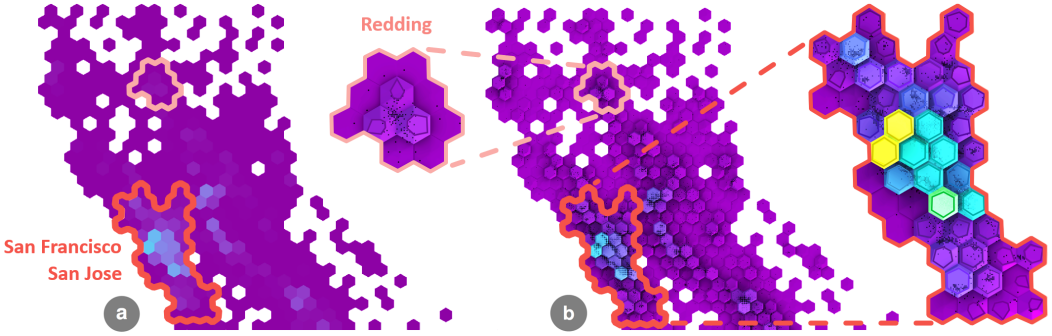

Figure 8: *Juxtaposition of (a) a heat map and (b) a honeycomb plot of the aggregated locations of houses originating from the 1990 US census. Both visualizations show agglomerated metropolitan areas along the west coast. Heat maps (a) alone, however, do not show underlying data points and the detection of minimal color differences depends on the user or hardware. Using (b) honeycomb plots, both are visualized as shown in the two enlarged very dense as well as sparse regions using an updated zoom-dependent heat map.*

4.2. California Housing Data (1990)

A well-known machine-learning data set is the California Housing data [Tor21] containing the longitude and latitude of houses, their age, the number of rooms, etc., collected during the 1990 US census. Figure 8 shows a hexagonal aggregation of their locations, where one point corresponds to the geographical position of a house. As highlighted by the heat map, the four largest metropolitan areas are located along the west coast: Los Angeles (4 M), San Diego (1.4 M), San Jose (1 M), and San Francisco (0.9 M). A heat map without blended points, however, fails to encode how houses are arranged, such as in San Francisco along the San Francisco Bay, and may lack emphasis on regions where density only varies slightly corresponding to minor changes in colors, such as in Redding (0.09 M), still being the largest city in California north of Sacramento. Using honeycomb plots, however, both data properties relevant for analysis and exploration are subtly highlighted.

4.3. Gender Equality Index EU-28 (2020)

This example illustrates a relief mosaic of data without underlying points. In many domains, e.g., cartography [dSL18; PMAM16; BSF16; COW92], it is common practice to rely on discretization. Figure 9 shows a simplified map of Europe with equally sized hexagonal countries, which supports a comparison regardless of landmasses. The color of each tile corresponds to the gender equality index, except purple countries which only provide geographical context. Additionally, each hexagon contains a picture of the prime minister (as of August 2021). The resulting visualization highlights a north-south disparity, as well as a west-east divide within Europe. Due to the coloring, countries like Sweden (SE), Denmark (DK), the Netherlands (NL), and France (FR) stand out. Unfortunately, when using the heat map alone, no difference can be seen between Estonia (EE), Latvia (LV), Lithuania (LT), Poland (PL), Czech Republic (CZ), Slovakia (SK), Hungary (HU), Romania (RO) and Croatia (HR). They appear to form a light-blue plateau of equal gender indices although they vary between 60.7 (EE) and 53.0 (HU), approximately 10 percentage points. Using a relief mosaic, three groups can be differentiated: Estonia and Latvia; Lithuania, Poland, Czech Republic, and Slovakia; and Hungary and Romania.

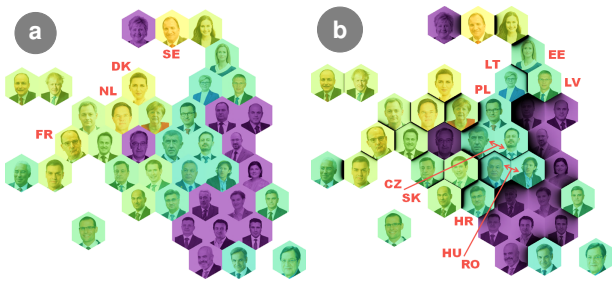

Figure 9: *Juxtaposition of (a) a heat map and (b) a relief mosaic visualizing a pseudo-spatially arranged hexbin map of Europe.*

We focused primarily on hexagonal grids, as these are among the most common types. A relief mosaic is, nevertheless, not limited to hexagons. Other shapes would equally be possible, e.g., deltoids, trapezoids, parallelograms, rectangles, triangles, etc.

5. Implementation

Our approach was implemented in C++ and OpenGL, and all calculations were performed in parallel on the GPU. Using ImGui [Cor21], all user-dependent parameters can be adjusted in real time. The five implementation steps necessary for honeycomb plots are listed below.

Scatter Plot: First, we create a scatter plot by rendering all points as semitransparent disks with additive blending. To prevent aliasing, the opacity of points decreases from their centers and their color as well as their radii can be modified using a GUI.

Kernel Density Estimation: Second, the density of the data set is calculated using kernel density estimation (KDE). As previously, we render points additively as disks, but now evaluate a two-dimensional Gaussian bell function centered at each sample point. This time, however, the point disks are enlarged to cover the Gaussian bell curves until their contributions are considered to be zero. Again, the GUI can be used to change the sigma of the Gaussian function depending on the granularity of the analysis.

Hexagonal Aggregation: Third, we aggregate all points by rendering them into a texture whose dimensions correspond to the number of rows and columns of the hexagon grid. After determining the tile in which a data point falls, a counter in the associated texture texel is increased using additive blending. The grid layout is adjustable using a GUI slider allowing for various grid resolutions.

Regression Plane: Fourth, using the central difference of per-pixel densities within each tile, originating from the KDE, we reconstruct per-pixel normals and sum them up in per-tile shader storage buffer objects (SSBO). Their normalized sum then corresponds to the regression-plane normal. Similarly, we average per-pixel densities as they define the height of a tile pyramid.

Compositing: Fifth, depending on the total number of points per tile, a heat map color is assigned. Next, optionally, the AO is calculated according to Equations 1 and 2 using the density, i.e., height, of neighboring tiles. If enabled, the diamond cut is calculated according to Equation 3 using the regression plane normal and the average density. Finally, the initial scatter plot is blended with the honeycomb plot using the *over* operator from Equation 4.

The source code of honeycomb plots is available on GitHub:
`https://github.com/TTrautner/HoneycombPlots.git`

6. Performance

The performance measurements listed in Table 1 were conducted on a desktop computer with an Intel Core i7-8700K CPU (3.7 GHz), 16 GB RAM, an NVIDIA GeForce RTX 2080 graphics card with 8 GB of texture memory, and a Windows 10 Home 64-bit operating system. Our analysis was based on the seven data sets used as paper figures ranging between 1,831 and 68,306 sample points aggregated into 36 to 1,242 hexagonal tiles. They correspond to four synthetically generated data sets and three real-world data sets. Table 1 additionally shows which of our techniques were enabled: relief mosaic (RM), diamond cut (DC), and amber inclusions (AI), indicated by check marks. Our focus was on the analysis of achieved frames per second (FPS) and has shown that all data sets can be rendered as well as interacted with in real time.

Figure	Points/ Tiles	RM	DC	AI	FPS 1280×720	FPS 1920×1080
1	68,306/ 1,242	✓	✓	✓	$80.11^{80.29}_{77.66}$	$26.29^{26.40}_{25.05}$
4c	2,024/ 38	✓			$629.24^{643.37}_{603.80}$	$433.79^{435.68}_{429.22}$
4f	3,244/ 36	✓	✓		$414.99^{417.94}_{412.62}$	$220.72^{222.26}_{218.78}$
4i	2,244/ 38	✓	✓	✓	$406.59^{407.78}_{401.27}$	$219.68^{220.35}_{216.23}$
7	66,388/ 329	✓	✓	✓	$36.79^{37.25}_{33.48}$	$18.73^{19.14}_{15.35}$
8	20,640/ 637	✓	✓	✓	$493.32^{497.22}_{483.75}$	$250.72^{251.28}_{247.38}$
9b	1,831/ 44	✓			$641.34^{649.28}_{625.18}$	$298.61^{303.63}_{296.84}$

Table 1: *Performance analysis of all seven visualized data sets.*

7. User Study

A major concern about honeycomb plots is their interpretability and the required learning effort. We, therefore, conducted a user study and compared standard heat maps (HM) to our three contributions, i.e., relief mosaic (RM), diamond cut (DC), and amber inclusions (AI), all exemplified in Figure 10. The website of the online survey as well as the collected results are included in the supplemental material. We investigated the following hypotheses:

H1: Relief Mosaic (RM) - ambient occlusion improves performance, i.e., correctness, in value estimation tasks.

H2: Diamond Cut (DC) - regression planes improve performance, i.e., correctness, in slope estimation tasks.

H3: Amber Inclusions (AI) - blending of points with tiles does not negatively affect other encodings, i.e., RM and DC.

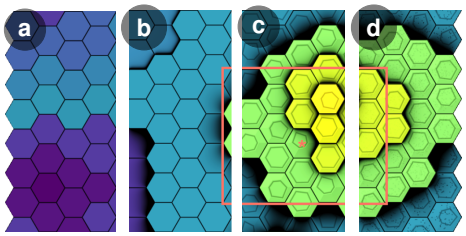

Figure 10: *Image collage comparing the four stimuli of a sample question used in our study. From left to right, the techniques expand on the previous stimuli: (a) HM, (b) RM, (c) DC, and (d) AI.*

7.1. Experiment Design and Questions

Our study, evaluating whether honeycomb plots can be understood within a short time, took 15 to 20 minutes. We chose a between-subject design and randomly assigned visualization techniques.

Stimuli:	HM vs. **RM**	HM vs. **DC**	HM vs. **AI**	RM vs. **DC**	RM vs. **AI**	DC vs. AI
Participants:	21 (11, 10)	20 (11, 9)	23 (11, 12)	19 (10, 9)	22 (10, 12)	21 (9,12)
Q1:	$\chi^2 = 8.10$ **p** $< .05^*$	$\chi^2 = 0.31$ $p = .57$	$\chi^2 = 5.09$ **p** $< .05^*$	$\chi^2 = 3.25$ $p = .07$	$\chi^2 = 0.20$ $p = .65$	$\chi^2 = 1.93$ $p = .16$
Q2:	$\chi^2 = 0.42$ $p = .51$	$\chi^2 = 15.55$ **p** $< .05^*$	$\chi^2 = 19.90$ **p** $< .05^*$	$\chi^2 = 14.48$ **p** $< .05^*$	$\chi^2 = 19.14$ **p** $< .05^*$	$\chi^2 = 1.33$ $p = .24$
Q3:	$\chi^2 = 3.03$ $p = .08$	$\chi^2 = 0.06$ $p = .80$	$\chi^2 = 0.11$ $p = .73$	$\chi^2 = 2.35$ $p = .12$	$\chi^2 = 1.74$ $p = .18$	$\chi^2 = 0.01$ $p = .91$
Q4:	$\chi^2 = 1.16$ $p = .28$	$\chi^2 = 17.27$ **p** $< .05^*$	$\chi^2 = 19.84$ **p** $< .05^*$	$\chi^2 = 15.78$ **p** $< .05^*$	$\chi^2 = 18.40$ **p** $< .05^*$	$\chi^2 = 0.75$ $p = .38$

Table 2: *Pairwise comparisons using a Kruskal-Wallis test with superior stimuli in bold and marked (*) significant p-values.*

	Q1:		**Q2:**		**Q3:**		**Q4:**	
	mean error	histogram	mean error	histogram	mean error	histogram	mean error	histogram
HM:	$\bar{m} = 39.3\%$		$\bar{m} = 72.7\%$		$\bar{m} = 9.0\%$		$\bar{m} = 93.9\%$	
RM:	$\bar{m} = 6.6\%$		$\bar{m} = 66.6\%$		$\bar{m} = 0.0\%$		$\bar{m} = 86.6\%$	
DC:	$\bar{m} = 33.3\%$		$\bar{m} = 3.7\%$		$\bar{m} = 7.4\%$		$\bar{m} = 0.0\%$	
AI:	$\bar{m} = 13.8\%$		$\bar{m} = 0.0\%$		$\bar{m} = 13.8\%$		$\bar{m} = 2.7\%$	

Table 3: *Overview of the mean errors and corresponding histograms with bins of 0, 1, 2, or 3 mistakes (from left to right).*

Inspired by Padilla et al. [PQMC17] and Kraus et al. [KAB*20], we chose Brehmer and Munzner's [BM13] visualization-tasks typology as basis for four questions:

Q1: *"Which tile has the highest/lowest density value within the highlighted box?"*

Q2: *"Which tile has the steepest/flattest change in density within the highlighted box?"*

Q3: *"Which tile has the highest/lowest density value, adjacent to the tile with the red asterisk (∗), within the highlighted box?"*

Q4: *"Which tile has the steepest/flattest change in density, adjacent to the tile with the red asterisk (∗), within the highlighted box?"*

We avoided asking questions about overall patterns in the data, as our techniques primarily improve the perception of individual tiles. Questions Q1 and Q2 focus on a global level, i.e., a larger highlighted region whose size varies. Questions Q3 and Q4 correspond to a local level, i.e., within the immediate vicinity of the tile with the red asterisk. Furthermore, two questions relate to density value estimations (Q1, Q3) and two to trends or density distributions (Q2, Q4). To answer the questions, participants may proceed differently, which then in turn corresponds to different tasks according to Brehmer and Munzner [BM13]. If, for question Q1, a participant first interprets the heat map and then searches for the corresponding value within the visualization, this represents a *locate* (target known, location unknown) and *identify* task. If the visualization is first examined and only afterwards compared to the heat map, this corresponds to an *explore* (target unknown, location unknown) and *identify* task. Since question Q2 requires a diamond glyph interpretation, it corresponds to an *explore* and *identify* task. Both Q3 and Q4 specify the location using a red asterisk. Again, Q3 may correspond to a *lookup* (target known, location known) or *browse* (target unknown, location known) and *compare* task, and Q4 again to a *browse* and *compare* task. Each question had to be answered using three different synthetic data sets, created by us, with varying difficulty, i.e., $4 \times 3 = 12$ questions. Finally, we evaluated binarily whether the correct tile was found or not.

7.2. Participants and Procedure

The online user study included 42 participants (11 HM, 10 RM, 9 DC, and 12 AI). Participants (32 male, 10 female) were aged from 16 to 69 ($\tilde{m}_{age} = 32$ years) and stated whether they were wearing glasses (19 yes, 23 no) or suffered from color vision deficiencies (1 with, 40 without, and 1 that did not know). They rated their familiarity with scatter plots ($\tilde{m}_{sp} = 4$), density estimations ($\tilde{m}_{de} = 2$), and heat maps ($\tilde{m}_{hm} = 4$) on a scale from $1 =$ *not familiar* to $5 =$ *very familiar*. At the beginning, the assigned visualization type was introduced, followed by a technical explanation of the survey. Subsequently, participants had to work through the 12 questions and finally answer general questions about themselves, were allowed to leave comments, and asked to rate from $1 =$ *strongly disagree* to $5 =$ *strongly agree* how aesthetically pleasing they found their assigned technique ($\bar{m}_{sp} = 3.23$, $\bar{m}_{hm} = 4.09$, $\bar{m}_{rm} = 4.7$, $\bar{m}_{dc} = 4.55$, $\bar{m}_{ai} = 4.5$).

7.3. Study Results

In the beginning, we evaluated whether the collected results were normally distributed, which was rejected by a Kolmogorov-Smirnov test at a 5% significance level. We, therefore, conducted pairwise comparisons of all techniques per question, shown in Table 2, using Kruskal-Wallis tests. Furthermore, Table 3 shows the mean error made by participants, including histograms (with bins from left to right for 0, 1, 2, or 3 mistakes) showing the number of errors made. Although we also measured times, these varied widely, which is why we did not analyze them further. For **H1**, we observed that the error rates for value estimation were generally low for RM, DC, and AI. On a global level (Q1), RM (6.6%) and AI (13.8%) were significantly lower in comparison to HM (39.3%). We also see that for question Q3, where a much smaller neighborhood had to be considered (comparing a tile to its six neighbors), there are no significant differences between the techniques, providing support for the assumption that color encoding alone is sufficient. In the case of DC (33.3%), the results were not significant. We suspect that this could be rooted in the fact that the introduc-

tion of a visual encoding that did not have any benefit for the task at hand may have confused some participants. For **H2**, we found strong evidence that our diamond cut metaphor is effective in facilitating slope encoding tasks, as both DC and AI performed significantly better than HM and RM on a global (Q2) as well as a local (Q4) level. In both cases, the participants were able to solve the corresponding tasks with very low error rates (DC: 3.7% for Q2 and 0.0% for Q4, AI: 0.0% for Q2 and 2.7% for Q4). With respect to **H3**, there were no significant differences between DC and AI in any of the questions (Q1-Q4), which supports our hypothesis that additional point encoding does not introduce negative side effects.

These findings demonstrate that while AO can provide substantial benefits for value estimation tasks, care must be taken when mixing it with other encodings that are not relevant to the task at hand. The results also suggest that participants were generally able to understand our diamond cuts based on the explanation in the study. We believe that this initial study provides support for the beneficial effects of incorporating spatialization cues into binning plots, while also demonstrating the clear need for particular care when combining multiple encodings based on the task at hand.

8. Discussion and Limitations

The focus of a relief mosaic is emphasizing color differences, making it unsuitable if these are unimportant or within a margin of error. Using a diamond cut for shape parameters also requires caution, especially if a grid layout is not implicitly defined since features may change if the grid is scaled, shifted, or rotated. Additionally, while we found it intuitive to adjust AO with a slider, future work could focus on automatic selections depending on, e.g., data variety. It would also be interesting to investigate to what extent AO alone can encode quantitative data. Future studies could examine if the sign of quantitative data, comparable to the neutral point of a classic heat map, can be encoded using AO or the pyramid geometry, i.e., tiles could have a convex or concave shape. We, furthermore, have not tested the effectiveness of amber inclusions but rather analyzed whether they negatively impact other encodings. Future research could quantitatively evaluate the analysis of outliers.

All our plots use the classic *over* operator, wherein the aggregated density influences the blending. Although this is well suited to highlighting sparse features, we have also experimented with other influence factors, e.g., point discrepancy. This would render tiles opaque when contained points are nearly evenly distributed, whereas tiles containing points close to each other or large empty areas are rendered transparently. In this work, we augmented hexplots with spatialization cues leading to $2\frac{1}{2}$D visualizations. We are, nevertheless, confident that honeycomb plots would also be well suited for 3D physicalization such as digital fabrication. It would be interesting to explore how tasks that are usually performed on hexagonal aggregation plots could be carried out by visually impaired or blind people using 3D prints of honeycomb plots.

9. Conclusion

In this work, we have presented novel visualization techniques for hexagonal binning plots. We proposed a relief mosaic that exploits ambient occlusion highlighting differences between similarly colored tiles, introduced a regression plane as a diamond cut that reveals point distributions, and emphasized trends or clusters in sparse regions using amber inclusions. We showed that our techniques can be implemented in real time on the GPU. Based on three usage examples, we generated expressive visualizations that increase the information content of hexplots, and showed, supported by a user study, that honeycomb plots can be understood and interpreted in short time.

References

[AMGA12] APOSTU, O., MORA, F., GHAZANFARPOUR, D., and AVENEAU, L. "Analytic ambient occlusion using exact from-polygon visibility". *Computers & Graphics* 36.6 (2012), 727–739. DOI: 10.1016/j.cag.2012.04.008 3.

[BM13] BREHMER, M. and MUNZNER, T. "A Multi-Level Typology of Abstract Visualization Tasks". *IEEE Transactions on Visualization and Computer Graphics* 19.12 (2013), 2376–2385. DOI: 10.1109/TVCG.2013.124 7.

[BSF16] BATTERSBY, S., STREBE, D., and FINN, M. "Shapes on a plane: Evaluating the impact of projection distortion on spatial binning". *Cartography and Geographic Information Science* 44.5 (2016), 1–12. DOI: 10.1080/15230406.2016.1180263 2, 5.

[CGH*15] CAKMAK, E., GÄRTNER, A., HEPP, T., et al. "Applying Visual Analytics to Explore and Analyze Movement Data". *Proc. IEEE Conference on Visual Analytics Science and Technology*. 2015, 127–128. DOI: 10.1109/VAST.2015.7347643 2.

[CLNL87] CARR, D. B., LITTLEFIELD, R. J., NICHOLSON, W. L., and LITTLEFIELD, J. S. "Scatterplot Matrix Techniques for Large N". *Journal of the American Statistical Association* 82.398 (1987), 424–436. DOI: 10.2307/2289444 2, 3.

[CM84] CLEVELAND, W. S. and MCGILL, R. "The Many Faces of a Scatterplot". *Journal of the American Statistical Association* 79.388 (1984), 807–822. DOI: 10.2307/2288711 2.

[CM85] CLEVELAND, W. S. and MCGILL, R. "Graphical Perception and Graphical Methods for Analyzing Scientific Data". *Science* 229.4716 (1985), 828–833. DOI: 10.1126/science.229.4716.828 1.

[Cor21] CORNUT, O. *ImGui*. https://github.com/ocornut/imgui. Accessed: August. 2021 6.

[COW92] CARR, D. B., OLSEN, A. R., and WHITE, D. "Hexagon Mosaic Maps for Display of Univariate and Bivariate Geographical Data". *Cartography and Geographic Information Systems* 19.4 (1992), 228–236. DOI: 10.1559/152304092783721231 2, 5.

[Dow14] DOWNEY, A. *Think Stats: Exploratory Data Analysis*. O'Reilly Media, 2014, 91–95 1.

[dSL18] DE SOUSA, L. M. and LEITÃO, J. P. "HexASCII: A file format for cartographical hexagonal rasters". *Transactions in GIS*. 22.1 (2018), 217–232. DOI: 10.1111/tgis.12304 5.

[Dup26] DUPIN, C. *Tableau des arts et métiers et des beaux-arts: présenté pour servir à propager l'institution des cours de géométrie et de méchanique appliquées aux arts, dans les villes de la France*. Bachelier, Paris, 1826 2.

[Dut21] DUTRÉ, P. *Global Illumination Compendium - The Concise Guide to Global Illumination Algorithms*. https://people.cs.kuleuven.be/~philip.dutre/GI/. 2021 3.

[HB03] HARROWER, M. and BREWER, C. A. "ColorBrewer.org: An Online Tool for Selecting Colour Schemes for Maps". *The Cartographic Journal* 40.1 (2003), 27–37. DOI: 10.1179/000870403235002042 3.

[HCSG18] HEIMERL, F., CHANG, C., SARIKAYA, A., and GLEICHER, M. "Visual Designs for Binned Aggregation of Multi-Class Scatterplots". *CoRR* abs/1810.02445 (2018) 2.

[HKIH07] HAGH-SHENAS, H., KIM, S., INTERRANTE, V., and HEALEY, C. "Weaving Versus Blending: a quantitative assessment of the information carrying capacities of two alternative methods for conveying multivariate data with color". *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), 1270–1277. DOI: 10.1109/TVCG.2007.70623 2.

[HMS97] HUANG, C., MCDONALD, J., and STUETZLE, W. "Variable-Resolution Bivariate Plots". *Journal of Computational and Graphical Statistics* 6.4 (1997), 383–396. DOI: 10.1080/10618600.1997.10474749 2.

[ISS04] IRANI, P., SLONOWSKY, D., and SHAJAHAN, P. "The Effect of Shading in Extracting Structure from Space-Filling Visualizations". *Proc. International Conference on Information Visualisation*. 2004, 209–216. DOI: 10.1109/IV.2004.1320146 2.

[ISS06] IRANI, P., SLONOWSKY, D., and SHAJAHAN, P. "Human Perception of Structure in Shaded Space-Filling Visualizations". *Information Visualization* 5.1 (2006), 47–61. DOI: 10.1057/palgrave.ivs.9500113 2.

[JVDF19] JO, J., VERNIER, F., DRAGICEVIC, P., and FEKETE, J.-D. "A Declarative Rendering Model for Multiclass Density Maps". *IEEE Transactions on Visualization and Computer Graphics* 25.1 (2019), 470–480. DOI: 10.1109/TVCG.2018.2865141 2.

[KAB*20] KRAUS, M., ANGERBAUER, K., BUCHMÜLLER, J., et al. "Assessing 2D and 3D Heatmaps for Comparative Analysis: An Empirical Study". *Proc. Conference on Human Factors in Computing Systems*. 2020, 1–14. DOI: 10.1145/3313831.3376675 2, 7.

[KS14] KINDLMANN, G. and SCHEIDEGGER, C. "An Algebraic Process for Visualization Design". *IEEE Transactions on Visualization and Computer Graphics* 20.12 (2014), 2181–2190. DOI: 10.1109/TVCG.2014.2346325 1.

[LB01] LANGER, M. S. and BÜLTHOFF, H. H. "A Prior for Global Convexity in Local Shape-from-Shading". *Perception* 30.4 (2001), 403–410. DOI: 10.1068/p3178 2.

[McN21] MCNUTT, A. "What are Table Cartograms Good for Anyway? An Algebraic Analysis". *Computer Graphics Forum* 40.3 (2021), 61–73. DOI: 10.1111/cgf.14289 1.

[MG13] MAYORGA, A. and GLEICHER, M. "Splatterplots: Overcoming Overdraw in Scatter Plots". *IEEE Transactions on Visualization and Computer Graphics* 19.9 (2013), 1526–1538. DOI: 10.1109/TVCG.2013.65 4.

[MHG10] MALIK, M. M., HEINZL, C., and GRÖLLER, E. "Comparative Visualization for Parameter Studies of Dataset Series". *IEEE Transactions on Visualization and Computer Graphics* 16.5 (2010), 829–840. DOI: 10.1109/TVCG.2010.20 2.

[Mil07] MILLER, J. R. "Attribute Blocks: Visualizing Multiple Continuously Defined Attributes". *IEEE Computer Graphics and Applications* 27.3 (2007), 57–69. DOI: 10.1109/MCG.2007.54 2.

[Mun15] MUNZNER, T. "Visualization Analysis and Design". CRC Press, 2015, 95–116. DOI: 10.1201/b17511 2.

[NH06] NOVOTNY, M. and HAUSER, H. "Outlier-Preserving Focus+Context Visualization in Parallel Coordinates". *IEEE Transactions on Visualization and Computer Graphics* 12.5 (2006), 893–900. DOI: 10.1109/TVCG.2006.170 4.

[NOA21] NOAA'S NATIONAL WEATHER SERVICE - STORM PREDICTION CENTER. *Severe Weather Database Files (1950-2019)*. https://www.spc.noaa.gov/wcm/. Accessed: August. 2021 4.

[PD84] PORTER, T. and DUFF, T. "Compositing Digital Images". *ACM SIGGRAPH Computer Graphic* 18.3 (1984), 253–259. DOI: 10.1145/964965.808606 4.

[PMAM16] POLISCIUC, E., MACAS, C., ASSUNCAO, F., and MACHADO, P. "Hexagonal Gridded Maps and Information Layers: A Novel Approach for the Exploration and Analysis of Retail Data". *Proc. SIGGRAPH ASIA Symposium on Visualization*. 2016, 1–8. DOI: 10.1145/3002151.3002160 2, 5.

[PQMC17] PADILLA, L., QUINAN, P., MEYER, M., and CREEM-REGEHR, S. H. "Evaluating the Impact of Binning 2D Scalar Fields". *IEEE Transactions on Visualization and Computer Graphics* 23.1 (2017), 431–440. DOI: 10.1109/TVCG.2016.2599106 7.

[PW66] PASSONNEAU, J. R. and WURMAN, R. S. "Urban Atlas : 20 American Cities". (Boston) M.I.T. Press, 1966 2.

[Qui21] QUILEZ, I. *Useful maths - Box ambient occlusions*. https://www.iquilezles.org/www/articles/boxocclusion/boxocclusion.htm. Accessed: November. 2021 3.

[RRM15] ROTH, R., ROSS, K., and MACEACHREN, A. "User-Centered Design for Interactive Maps: A Case Study in Crime Analysis". *International Journal of Geo-Information* 4 (2015), 262–301. DOI: 10.3390/ijgi4010262 2.

[SAS14] STONE, M., ALBERS SZAFIR, D., and SETLUR, V. "An Engineering Model for Color Difference as a Function of Size". *Proc. Color and Imaging Conference*. Vol. 6. 2014, 253–258 1.

[SSG16] SZAFIR, D. A., SARIKAYA, A., and GLEICHER, M. "Lightness Constancy in Surface Visualization". *IEEE Transactions on Visualization and Computer Graphics* 22.9 (2016), 2107–2121. DOI: 10.1109/TVCG.2015.2500240 2.

[Sto12] STONE, M. "In Color Perception, Size Matters". *IEEE Computer Graphics and Applications* 32.2 (2012), 8–13. DOI: 10.1109/MCG.2012.37 1.

[TB21] TRAUTNER, T. and BRUCKNER, S. "Line Weaver: Importance-Driven Order Enhanced Rendering of Dense Line Charts". *Computer Graphics Forum* 40.3 (2021), 399–410. DOI: 10.1111/cgf.14316 2.

[TBSB20] TRAUTNER, T., BOLTE, F., STOPPEL, S., and BRUCKNER, S. "Sunspot Plots: Model-based Structure Enhancement for Dense Scatter Plots". *Computer Graphics Forum* 39.3 (2020), 551–563. DOI: 10.1111/cgf.14001 2, 4.

[Tor21] TORGO L. - UNIVERSITY OF PORTO. *California Housing Data (1990)*. https://www.dcc.fc.up.pt/~ltorgo/Regression/cal_housing.html. Accessed: August. 2021 5.

[TW01] TELEA, A. and WIJK, VAN, J. "Visualization of Generalized Voronoi Diagrams". *Proc. IEEE TCVG Symposium on Visualization*. 2001, 165–174. DOI: 10.1007/978-3-7091-6215-6_18 2.

[vWT01] Van WIJK, J. and TELEA, A. "Enridged Contour Maps". *Proc. IEEE Visualization*. 2001, 69–543. DOI: 10.1109/VISUAL.2001.964495 2.

[vWvdW99] Van WIJK, J. and van de WETERING, H. "Cushion Treemaps: Visualization of Hierarchical Information". *Proc. IEEE Symposium on Information Visualization*. 1999, 73–78. DOI: 10.1109/INFVIS.1999.801860 2.

[Wil78] WILLIAMS, L. "Casting Curved Shadows on Curved Surfaces". *ACM SIGGRAPH Computer Graphic* 12.3 (1978), 270–274. DOI: 10.1145/965139.807402 3.

[WK20] WALLNER, G. and KRIGLSTEIN, S. "Multivariate Visualization of Game Metrics: An Evaluation of Hexbin Maps". *Proc. Annual Symposium on Computer-Human Interaction in Play*. 2020, 572–584. DOI: 10.1145/3410404.3414233 2.

[ZIK98] ZHUKOV, S., IONES, A., and KRONIN, G. "An Ambient Light Illumination Model". *Proc. Rendering Techniques*. 1998, 45–55. DOI: 10.1007/978-3-7091-6453-2_5 2.

[ZLG*21] ZHAO, J., LIU, X., GUO, C., et al. "Phoenixmap: An Abstract Approach to Visualize 2D Spatial Distributions". *IEEE Transactions on Visualization and Computer Graphics* 27.3 (2021), 2000–2014. DOI: 10.1109/TVCG.2019.2945960 3.

[ZZW*21] ZENG, Q., ZHAO, Y., WANG, Y., et al. "Data-driven Colormap Adjustment for Exploring Spatial Variations in Scalar Fields". *IEEE Transactions on Visualization and Computer Graphics* Pre-Print (2021). DOI: 10.1109/TVCG.2021.3109014 2.