Vision, Modeling, and Visualization (2018)
F. Beck, C. Dachsbacher, and F. Sadlo (Eds.)

Correlated Point Sampling for Geospatial Scalar Field Visualization

Riccardo Roveri[1], Dirk J. Lehmann[2], Markus Gross[1] and Tobias Günther[1]

[1]ETH Zurich, Switzerland [2]University of Magdeburg, Germany

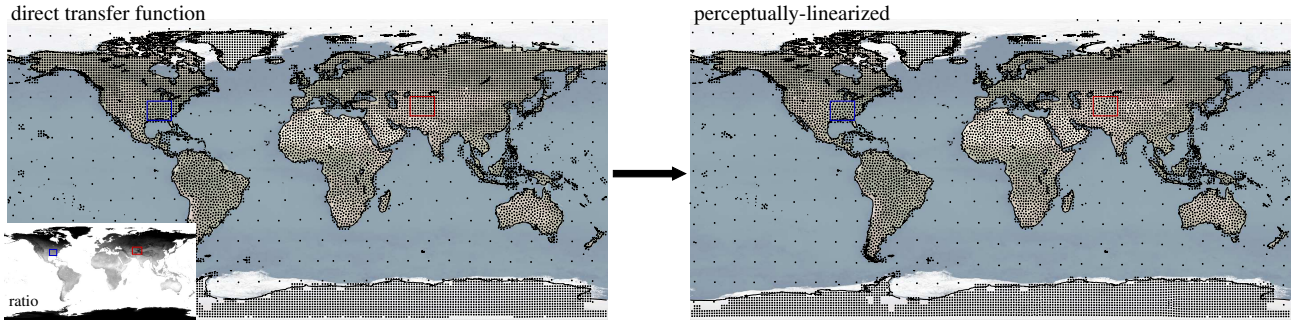

Figure 1: *Visualizations of land surface temperature using a point distribution interpolation from a regular grid (low temperature) to blue noise (high temperature). In this work, we perform a perceptual linearization of the transfer function. Left: an interpolation without perceptual linearization. Right: an interpolation with perceptual linearization. Note the differences in the indicated blue and red boxes.*

Abstract

Multi-variate visualizations of geospatial data often use combinations of different visual cues, such as color and texture. For textures, different point distributions (blue noise, regular grids, etc.) can encode nominal data. In this paper, we study the suitability of point distribution interpolation to encode quantitative information. For the interpolation, we use a texture synthesis algorithm, which paves the path towards an encoding of quantitative data using points. First, we conduct a user study to perceptually linearize the transitions between uniform point distributions, including blue noise, regular grids and hexagonal grids. Based on the linearization models, we implement a point sampling-based visualization for geospatial scalar fields and we assess the accuracy of the user perception abilities by comparing the perceived transition with the transition expected from our linearized models. We illustrate our technique on several real geospatial data sets, in which users identify regions with a certain distribution. Point distributions work well in combination with color data, as they require little space and allow the user to see through to the underlying color maps. We found that interpolations between blue noise and regular grids worked perceptively best among the tested candidates.

This is the authors preprint. The definitive version is available at http://diglib.eg.org/.

1. Introduction

The visualization of scalar fields is a fundamental aspect of visualization. In geospatial data analysis, the scalar fields are given on maps, i.e., each data value has a spatial embedding. Examples are found in meteorology [Wan14], oceanology [WPS*16], environmental studies [LMH*15] and urban planning [FPV*13]. In the context of multi-variate data analysis, there is a keen interest in displaying multiple scalar fields at the same time in order to assess their correlation. Thereby, the encoding of the mixture of two scalar fields is a challenging problem, since we are not only interested in the mixing ratio, but also in the scalar value itself. Typically, these properties are visualized with visual attributes that can encode quantitative data, e.g., color. In addition, texture-based approaches may be used that model varying densities of hatching patterns. However, textures are typically used to encode nominal data, for instance through the display of different point distributions, such as points on a regular grid or a particular noise pattern such as blue noise or green noise.

The recent texture synthesis approach by Roveri et al. [ROG17] is able to interpolate between different point distributions, such as regular patterns and noise patterns, as illustrated in Fig. 1. Their

method paves the path towards the encoding of quantitative data using mixtures of point distributions. Thereby, the types of point distributions form a nominal dimension, e.g., rain and snow, and their interpolation encodes the mixing ratio. The density of the points constitutes a quantitative dimension, such as the amount of precipitation. The perceptual interpretation of point distribution interpolations has not yet been studied. Thus, in this paper, we perceptually linearize the interpolation between three pairs of uniform point distribution interpolations, namely from regular grids to blue noise, hexagonal grids to blue noise, and regular grids to hexagonal grids. Which of the three will perform best? To answer these questions, we perform three user studies. In the first study, we estimate the perceptual transfer function from the actual mixing ratio to the perceived mixing ratio. The second study is used to validate the obtained transfer functions from the first study using an independent set of participants and to determine the remaining errors and individual variations. In the third study, we apply the best perceptually linearized transfer function to real geospatial data in order to experimentally assess the accuracy of the interpretation with and without legends.

2. Related Work

Perception. Since visualization is concerned with the precise and effective communication of information, research on cognitive processing received a significant amount of attention. First steps in identifying and sorting through common visual tasks, with links to preattentive vision and perception performance were laid out by Cleveland and McGill [CM84]. The research in this area spans several aspects, targeting the understanding of the cognitive system, the evaluation of existing visualization techniques, as well as the application of the concepts to craft novel and more effective encodings. For instance, with their studies on the impact of the context on the perception of contrast, Chubb et al. [CSS89] contributed to the understanding of the cognitive processes. Similarly, Bartram et al. [BCS11] assessed the effect of color and transparency on overlaid grids. In terms of analysis of existing techniques and as a framework for the development of new methods, Albuquerque et al. [AEM11] laid out perception-based visual quality measures. Perception-based evaluations of specific visualization techniques was applied in the field of high-dimensional data visualization, including the evaluation of projection methods by Etemadpour et al. [EMdSP*15], the study of parallel coordinates by Johansson et al. [JFLC08], as well as the assessment of scatterplots in terms of multi-class perception errors [GCNF13] and correlation indices [SBLC17]. Perception is also considered in computer graphics, e.g., for the evaluation of surface reflectance models [WAKB09].

Geospatial Visualization. The visualization of geospatial data is often at the juncture between scientific visualization and information visualization. On the one hand, there is the spatial embedding with its direct application to scientific problems, such as meteorology [Wan14], oceanology [WPS*16], or environmental studies [LMH*15]. On the other hand, the data is often multi-variate and can therefore profit from a direct embedding of information visualization techniques to effectively convey more information and therefore insight. Wood et al. [WDSC07] concentrated on interactivity in geospatial visualizations. Nocke et al. [NSBW08] gave an overview of common geospatial visualization techniques for climate data visualization. A summary of information visualization techniques in climate research was given by Tominski et al. [TDN11]. A practical guide through visualization and analysis software for meteorological data analysis was given by Wang [Wan14]. Dübel et al. [DRST14] reviewed 2D and 3D techniques for spatial data visualization. For a thorough entry to the visualization of cartograms, we refer to Nusrat and Kobourov [NK16], and for an introduction to spatially-embedded traffic data we refer to Chen et al. [CGW15]. An approach to multi-variate data visualization is the encoding with glyphs, which is summarized by Borgo et al. [BKC*13].

Point Distributions. In computer graphics, many approaches have been proposed to synthesize point distributions with target characteristics. In particular, point distributions have mostly been used for element placing [ROG17], anti-aliasing [Uli88, HSD13] and artistic applications, such as stippling [Fat11]. Above all, the blue noise pattern has been studied, where the points are randomly distributed but have a minimum distance between each pair [Fat11, Uli88, HSD13]. These distributions are usually generated by tiling patches of points [Ost07] or by optimizing their positions (possibly by even adding and removing new points) [Llo82, BSD09]. Adaptive density (especially needed for stippling, for example) has been achieved with many methods [LWSF10, dGBOD12]. On the other hand, most work focusing on the correlation of the distributions (i.e, the arrangement of the points that define a pattern) handled the case for a single given constant pair-wise correlation model [ZHWW12, OG12]. Based on their previous patch-based texture synthesis approach [ROM*15], the work of Roveri et al. [ROG17] is the first technique to synthesize general point distributions with adaptive correlation and density. Point distributions have been used also for visualization. Tang et al. [TQWZ06] proposed a texture synthesis using natural images for data visualization. Points have been used in other aspects of information visualization, for instance as building block of space-filling tree maps [SHS11]. In scientific visualization, point distributions were used for spot noise [VW91] and glyph placement [GWG*11].

Point Distribution Interpolation. Point distributions are characterized by two main properties: the density and the arrangement of the sampling points. The latter is defined as the pair-wise correlation, and is responsible for the different patterns that point distributions can assume. Blue noise and regular distributions are common examples used in practice. While many approaches exist to analyze and synthesize point distributions with adaptive densities [LWSF10, dGBOD12], the synthesis methods presented in [ROG17] are the first to handle distributions with adaptive correlations as well, i.e. to generate spatially varying patterns. In our work, we decided to adopt the discrete texture synthesis method from [ROG17], which allows for varying densities and can handle the special case of non-ergodic sampling patterns, thus allowing us to utilize regular patterns.

3. Evaluation of Point Distribution Interpolation

We study the suitability of point distribution interpolations for the encoding of geospatial scalar data, which has two use cases:

1. Encode the mixture of two scalar fields, showing the maximum, minimum, sum or average scalar value of the two fields (via density) and the mixing ratio by the point pattern.

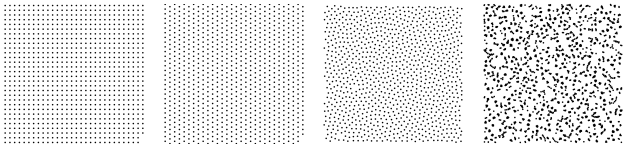

(a) *Regular* **(b)** *Hexagonal* **(c)** *Blue noise* **(d)** *Green noise*

Figure 2: *Different point distributions (later used as reference).*

2. Encode a single scalar field by the point pattern and keeping the density constant, which shows the minima and maxima with different point distributions.

First of all, this requires the perceptual linearization of the point distribution interpolation, the validation and error estimation of the linearization, and the assessment of the user performance in real-world geospatial data. This raises a number of questions: Which point distributions are suitable to encode a data class? How accurately can humans perceive the mixing ratio between two distributions and how much does this depend on the chosen point distributions? How well can users perform identification tasks in geospatial data? To answer these questions, we performed three user studies. In the following, we elaborate on the protocols and report our findings.

3.1. Selection of Point Distributions

First, we need to select appropriate point distributions. Four of the most commonly found point distributions are displayed in Fig. 2, namely regular grids, hexagonal grids, blue noise and green noise. Not every point distribution is suitable to encode information.

Uniform vs. Non-Uniform. An essential requirement is that the point distribution has a uniform density. If this was not the case, the synthesized point set would show varying density. Since we intend to encode scalar information by the density, a varying density of the point distribution itself would be misinterpreted as a variation of the underlying scalar field. The density of green noise is spatially-varying, which thus disqualifies it as a suitable candidate. In the following, we analyze regular grids, hexagonal grids and blue noise, as all three are uniform and thus fulfill the requirement.

Interpolation Gradients. A direct application of the texture synthesis method of Roveri et al. [ROG17] to the three candidate point distributions (regular grid, hexagonal grid and blue noise) results in the linear gradients shown in Fig. 3. It is yet unclear, how these gradients are visually perceived. For instance, will the result that is linearly interpolated at $t = 0.5$ be perceived as a half-way blend between the two distributions? We answer this by a user study.

3.2. Perceptual Linearization

The first study is dedicated to the perceptual linearization of the transition between two given point distributions.

Study Design. We studied the three interpolations shown in Fig. 3. For this, we conducted a user study with 12 graduate students who have a computer graphics and computer vision background. The study comprised 33 tasks (11 for each interpolation), printed on paper. For each task, we showed two point distributions and an

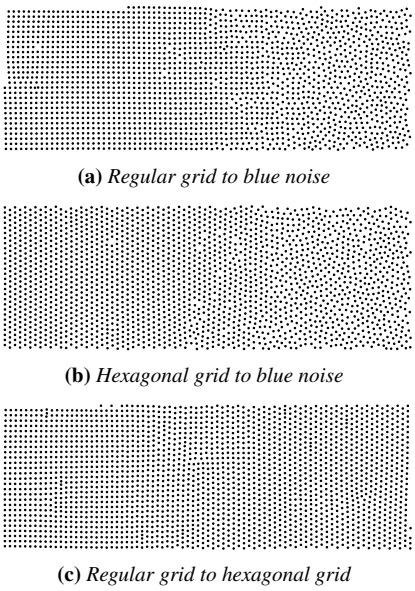

(a) *Regular grid to blue noise*

(b) *Hexagonal grid to blue noise*

(c) *Regular grid to hexagonal grid*

Figure 3: *Direct gradients for the interpolations between the three distributions. These are the direct results, when applying the texture synthesis of Roveri et al. [ROG17].*

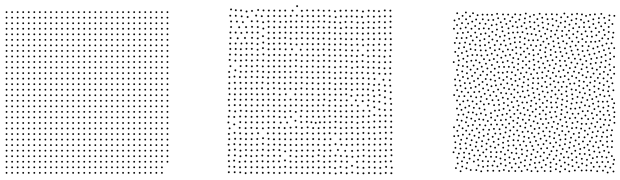

(a) *regular grid, 0%* **(b)** *to estimate by user* **(c)** *blue noise, 100%*

Figure 4: *Example for a task in the user study on the perceptual linearization. The user is shown two reference distributions (a) and (c), and is asked to estimate their mixing ratio in (b). The solution to this particular example is shown at the end of the paper.*

arbitrary interpolation between them. The users were asked to give an arbitrary estimate of the mixing ratio, e.g., 40% of the reference distribution on the left and 60% of the reference distribution on the right. An example for such a task is shown in Fig. 4. The tested mixing ratios between the two distributions were uniformly sampled and shown in a random permutation. Users were not given any additional aid, such as a legend or an example, since we were interested in their unbiased opinion. Users were not allowed to change their previous answers, since this would have resulted in an ordering and a calibration. To reduce the memory effect, we alternated between the three different interpolations. All tasks were completed within 10–30 minutes.

Study Results. A plot of the computed mixing ratio and the perceived mixing ratio is shown in Fig. 5 for all three interpolations. The transfer functions are non-linear and vary greatly among the three interpolations. As shown in Table 1, the least deviation occurred for interpolations with the blue noise pattern, which gives hope that those can be used for an encoding of information. The root-mean-square errors are in the order of 19%. It should be noted that in a real scenario, a legend would be given, which can be used

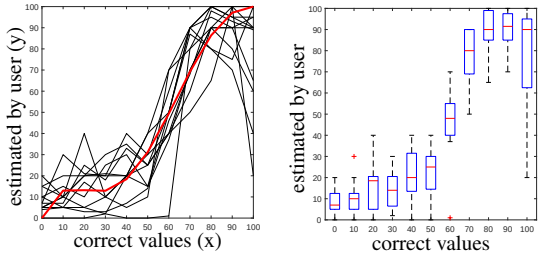

(a) *Regular grid to blue noise, transfer function* $y(x) = 14.7059x^6 - 36.1918x^5 + 22.0716x^4 + 5.4334x^3 - 6.6478x^2 + 1.6301x - 0.0007$

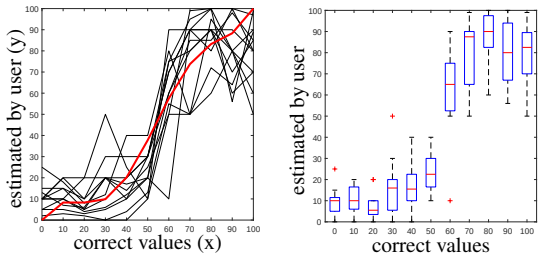

(b) *Hexagonal to blue noise, transfer function* $y(x) = 17.7483x^6 - 31.8233x^5 + 4.4327x^4 + 17.3818x^3 - 8.0492x^2 + 1.3114x - 0.0009$

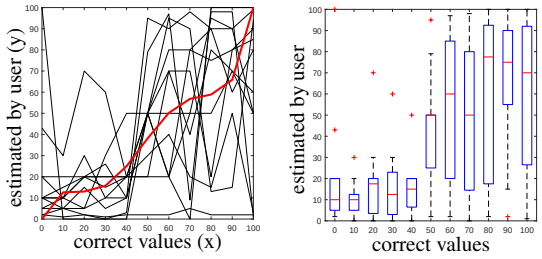

(c) *Regular to hexagonal grid, transfer function* $y(x) = 7.8431x^6 + 6.5832x^5 - 43.9637x^4 + 43.0286x^3 - 14.6691x^2 + 2.1788x - 0.0006$

Figure 5: *Estimated transfer functions after the first user study (in %). The transfer functions are non-linear and vary among the three tested interpolations. The least-squares fit of the perceptually linearizing transfer functions is shown in red for all interpolations.*

Regular to Blue Noise	Hex. to Blue Noise	Regular to Hex.
19.39	19.28	33.13

Table 1: *Root-mean square errors (in %) of the point distribution interpolations without a perceptual linearization (Study 1).*

by the user to calibrate and compare. Our goal is to find a transfer function for this legend so that the changes along the legend are perceptually linear, which maximizes the perceived distance everywhere. The interpolation from a regular grid to a hexagonal grid showed with a root-mean-square error of 33% the largest deviation among the participants. It is also apparent that in this case the answers were not monotone, which is an indicator for a randomness in the decision process. Without any reference, e.g., an example of a half-way blending, and without any prior experience, we expected a certain deviation among the participants.

Given the data of the user study, we estimated the transfer function for each of the three interpolations by fitting a monotone polynomial

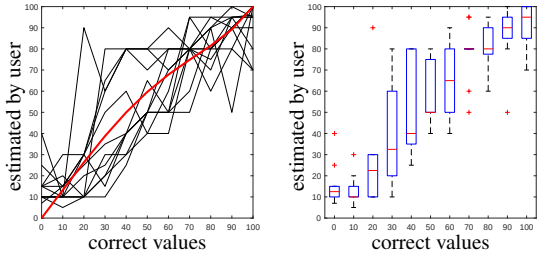

(a) *Regular grid to blue noise*

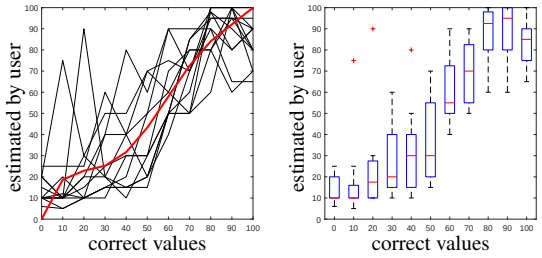

(b) *Hexagonal grid to blue noise*

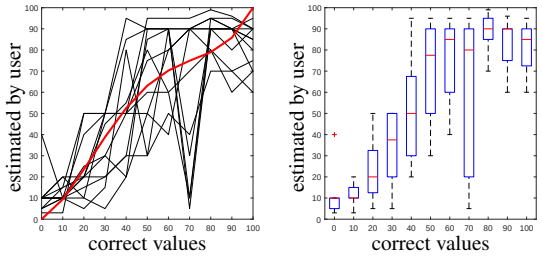

(c) *Regular grid to hexagonal grid*

Figure 6: *Validation of the linearized transfer function with a second independent study. The least-squares fit of the perceptually linearizing transfer functions is shown in red. The obtained transfer functions are almost linear.*

with end-point interpolation in a least-squares approximation sense. The transfer functions are listed and visualized in red in Fig. 5.

3.3. Determination of Accuracy

In the previous section, we determined transfer functions that perceptually linearize the interpolations among point distributions. Next, we determine the accuracy of our linearizations by performing a validation study with an independent group of participants.

Study Design. In the validation study, we repeated the experiments from Section 3.2, i.e., the tasks in Fig. 4, with a small modification. This time, we uniformly sampled the linearized values, which results in mixing ratios that exhibit a maximal perceived distance. Again, the participants reported their individually perceived mixing ratios for the given point distribution interpolations. The participants of the second study were 12 graduate students with a background in computer graphics and computer vision.

Study Results. In Table 2, we report the accuracy of the perceptually linearized transfer functions. Compared to the direct transfer

Regular to Blue Noise	Hex. to Blue Noise	Regular to Hex.
17.18	17.77	21.28

Table 2: *Root-mean square errors (in %) of the point distribution interpolations with a perceptual linearization (Study 2).*

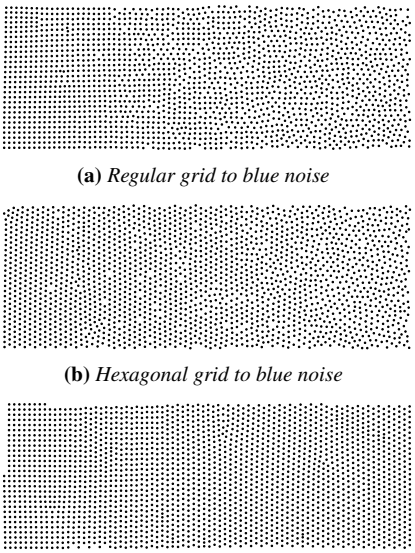

(a) *Regular grid to blue noise*

(b) *Hexagonal grid to blue noise*

(c) *Regular grid to hexagonal grid*

Figure 7: *Perceptually linearized gradients for the interpolations between the three tested distributions.*

functions, the errors reduced consistently. Fig. 6 shows a plot of the predicted mixing ratio and the perceived mixing ratio. For the interpolations with blue noise in (a) and (b), the average response curve (in red) is almost linear, which is evidence for a good prediction. As indicated by the individual responses (in black) and also visible in the box plots, the error is still high. Unfortunately, some users accidentally confused the mixing ratios and entered a mixing ratio of 20%, instead of 80%. Those errors are strongly noticeable outliers and they are reflected in the root-mean-square error. To test whether the remaining deviation is attributed to the missing legend, we perform the last user study both with and without legend.

Based on these results, we can make the following recommendation: An interpolation between regular grids and blue noise is the most effective and most accurate interpolation between two point distributions. The interpolation between hexagonal grids and blue noise follows closely, and the interpolation between regular grids and hexagonal grids is by far the most difficult and inaccurate one. The perceptually linearized gradients for the three tested point distribution interpolations are reported in Fig. 7.

3.4. Application to Geospatial Data

In the following, we apply the best linearized point interpolation model to the visualization of real geospatial scalar fields, i.e., the interpolation from blue noise to a regular grid. We applied the method to temperature and vegetation maps. To assess the effectiveness of the point interpolations, we designed a third user study.

Geospatial Data. For our tests, we used the three geospatial scalar fields that are shown in Fig. 8. Figs. 8a and 8b show the vegetation index in March and September 2017. This index estimates the amount of green plant life on all of Earth's land surface. We encode the ratio between green plant life in March and in September to visualize the change in vegetation over the year. Due to the changing seasons, we expect to find variations in the temperate and subtropic geographical zones. The second data set in Fig. 8c shows the average land surface temperature at day time. We observed the monthly average for February 2017. For this data set, we encode minima and maxima with different distributions.

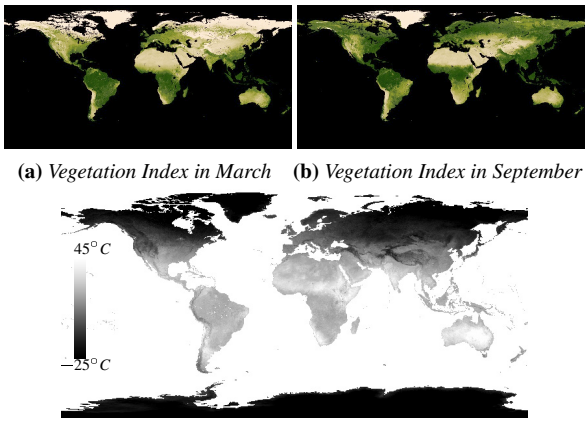

(a) *Vegetation Index in March* **(b)** *Vegetation Index in September*

(c) *Land surface temperature*

Figure 8: *The three geospatial scalar fields used in the third user study. The data sets are courtesy of NASA's Earth Observatory.*

	w/o Gradient	w/ Gradient
Direct transfer function	32.34	25.42
Perceptually-linearized	18.65	17.12

Table 3: *Root-mean square errors (in %) for the experiments with real geospatial data (Study 3).*

Study Design. Chubb et al. [CSS89] observed that texture patterns are perceived differently, based on the patterns in the neighborhood. In the third study, we want to see how strongly this effect influences the perception of spatially-varying point distributions in practice. We applied the point distribution interpolation to the aforementioned scalar fields and manually highlighted regions with nearly constant mixing ratio with a red ellipse. The users were asked to identify the constant mixing ratio of the selected region. In total 12 maps were generated, half of them with the direct transfer functions and the other half with the perceptually-linearized transfer functions. We split 23 graduate students into two separate groups. One group was assigned the maps without a gradient as demonstrated in Fig. 9a. The other group was given maps with a discrete gradient as shown in Fig. 9b. We hypothesized that showing a gradient should increase the accuracy by a certain degree. In the tasks, regions with a ratio of 20%, 40%, 50%, 60%, 70% or 90% were shown.

Study Results. Figs. 10a and 10b compare the user responses for the direct and perceptually-linearized cases in the absence of a gradient for help. The box plots show the errors to the ground truth. An optimal solution is therefore a deviation of zero. The sign implies an overestimation or underestimation. The direct case consistently underestimates the ground truth by up to 50%, whereas

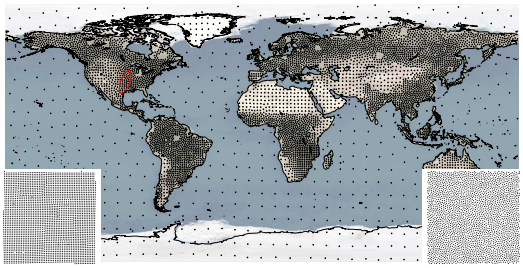

(a) *Identification of a marked region* ***without gradient*** *in a point distribution interpolation that uses the* ***direct transfer functions****.*

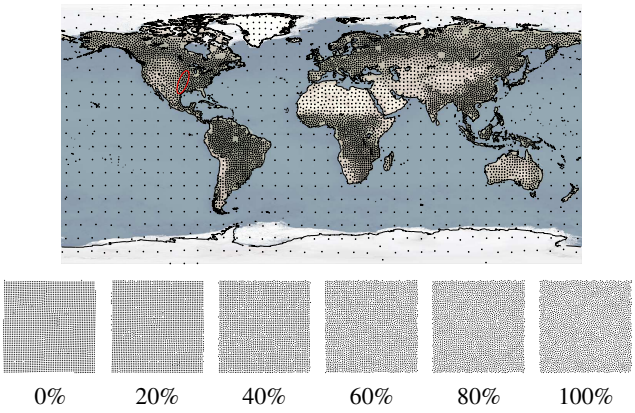

(b) *Identification of a marked region* ***with gradient*** *in a point distribution interpolation that uses the* ***perceptually-linearized transfer functions****.*

Figure 9: *Examples of the task in the third user study. The user is asked to identify the mixing ratio of the area that is highlighted by a red circle. In 9a no gradient is given and in 9b a gradient is shown. The point density displays the maximum density of the scalars. The solution to these tasks are shown at the end of the paper.*

the perceptually-linearized case deviates around zero. Still, the error margins are noticeable. How much of these errors is attributed to the lack of a gradient can be seen by comparing with Figs. 10c and 10d, which report the results with the gradients. With gradients, the mean error decreases both in the case of direct transfer functions, as well as in the case of perceptually-linearized transfer functions.

The root-mean-square errors are listed in Table 3. Adding a gradient helped both the direct and the perceptually-linearized transfer functions. The perceptually-linearized transfer functions clearly outperformed the direct transfer functions. One possible reason for the remaining residuals is the influence of the neighborhood region [CSS89], which might alter the texture perception. Another possibility is that the highlighted regions were too small to identify the pattern properly by the limited number of points inside the region. The observed deviations with the perceptually-linearized method have almost zero mean, which indicates that the perceptual linearization shifted the response curve into the right direction, but the interpersonal variation was too high to deliver accurate results.

We draw the conclusion that interpolations of point distributions should be used carefully. Users always exhibit a certain individual bias. Point distribution interpolations can serve well as a supporting information. Fortunately, the point sets cover only a small fraction of the image, leaving much space for combinations with other visual cues. For instance, point distribution interpolation can be well combined with color maps that can be placed in the background.

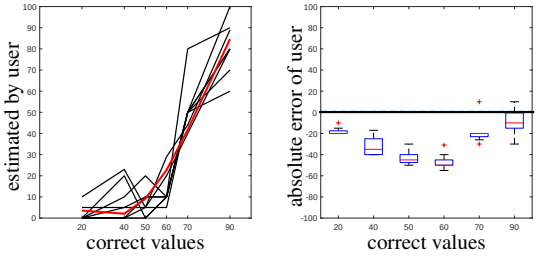

(a) ***Direct*** *transfer function* ***without gradients****.*

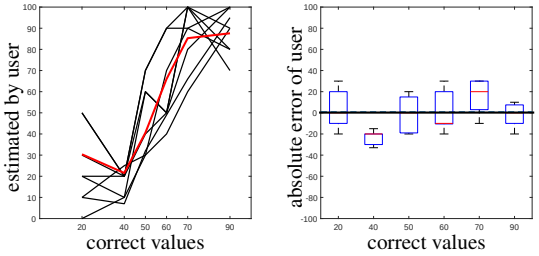

(b) ***Perceptually-linearized*** *transfer function* ***without gradients****.*

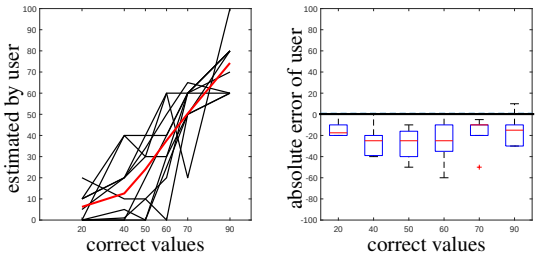

(c) ***Direct*** *transfer function* ***with gradients****.*

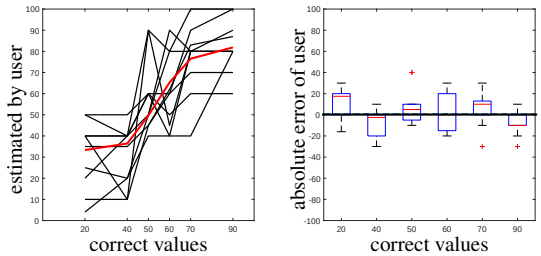

(d) ***Perceptually-linearized*** *transfer function* ***with gradients****.*

Figure 10: *Results of the final user study, in which the perceptually-linearized transfer function from regular to blue noise was applied to geospatial data. Left: user responses, right: box plot of the errors.*

4. Discussion

Point Density. The perceptual linearization was conducted for a constant point density, assuming that the density has no influence on the mixing ratio. It remains to be explored whether the linearization is different for other constant densities and how a continuous change in the density affects the perceived mixing ratio. We believe that there is a difference, but that this difference might be small for slowly varying (low-frequency) scalar fields.

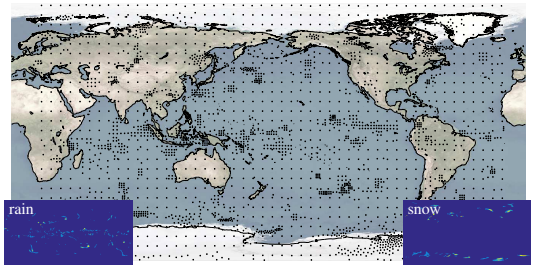

Figure 11: *Point distribution interpolation from regular (rain) to blue noise (snow) applied to a mixture of two scalar fields with high-frequency details. The structures are too small to be captured adequately with distinguishable patterns.*

Sparse Representation. Visualization techniques can be categorized into sparse and dense approaches, based on their coverage of the observation space. Visualizations using point distributions are inherently *sparse*. All sparse techniques have in common that they encode only discrete parts of a continuous field. Thus, small details between two discrete elements are lost. In case of point distributions, this distance is slightly bigger, since multiple points are perceived together as a particular pattern. This should be kept in mind, when selecting point distributions for visualization. Fig. 11 gives an example for a scalar field with high frequency details that cannot be captured with point distributions at this scale. In this case, users could zoom in, giving the observation space more area on the screen and therefore a bigger canvas to place points.

Outliers in the Study. Unfortunately, about two participants of each study produced very inconsistent results, as they accidentally stated the inverse mixing ratio. We detected and removed the obviously wrong answers in the cases where the ground truth was the minimum or maximum value, since for those the answer is usually impossible to guess wrong. An improved user study design would ask redundant questions to rule out inconsistent answers.

User Feedback. We asked the users for feedback throughout all studies and were particularly interested in the strategies that they applied to identify the mixing ratio. Multiple participants reported that they looked for the uniformity of points along horizontal and vertical lines for regular and hexagonal grids. Since both, regular and hexagonal grids have horizontal point alignments in common, the participants had less visual cues to identify the differences between the two. In fact, the only difference is the difference between a 90 degree alignment (for regular) or a 60 degree alignment (for hexagonal). This difference was apparently difficult to assess, as visible in the study results in Fig. 5c.

Performance. For the sample generation, we followed the texture synthesis approach of Roveri et al. [ROG17]. The input example distributions contained about 1000 samples, and the synthesized output ones up to 11k samples, for the geospatial data. In those cases, one iteration of the method took up to 1 minute with our unoptimized code, and 50 iterations were performed (convergence was always reached with less iterations). We used a PC with an Intel i7-3770K CPU. Due to its local nature, the method can be easily parallelized, thus substantially better performance can be achieved.

5. Conclusions

In this paper, we assessed the suitability of point distribution interpolation using a texture synthesis approach [ROG17] for the encoding of mixing ratios of two scalar fields in the context of geospatial data. First, we perceptually linearized the perceived mixing ratio with a user study in which the participants were asked to estimate the mixing ratios of given point distribution interpolations, including regular grid to blue noise, hexagonal grid to blue noise and regular grid to hexagonal grid. Second, we confirmed the linearizing transfer functions in a second study with an independent group of participants and we estimated the accuracy of the point distribution encodings in the absence of a legend or examples. This allowed us to design interpolation gradients along which the perceived distance is constant, which in turn can be used as legends. We found that the interpolation between regular grids and blue noise is the most accurate. Finally, we applied the best linearized transfer function to real geospatial data and determined the usefulness in an identification task. We found that the perceptual linearization reduced the mean errors, but still showed some deviation due to individual personal biases. Due to its limited accuracy, we conclude that point distribution interpolation should be used carefully. We see its role in augmentation and support of other visual encodings such as color maps to which it offers an additional visual cue.

We concentrated on the linearization of point distribution encodings, which opens many directions for future work. In the future, we would like to compare point distributions with other (possibly redundant) encoding techniques, such as color. Note that our approach is orthogonal to that and can be used to encode additional attributes. Though, the influence of color on the point perception would be interesting to investigate. It also remains to be explored whether other uniform point distributions exist that are easy to distinguish by humans. If this is the case, the interpolations could be extended to barycentric interpolation between three or even more classes. Another open question is the influence of point density and point size on the perceived mixing ratio. While we concentrated on quantitative data, it is also yet unknown how many ordinal classes can be reliably identified and how large the areas need to be, cf. [War10]. It would also be interesting to explore, whether point distributions could be used to design projections of high-dimensional data [LT16], as users are able to find outliers in the regularity of patterns.

Acknowledgements

We thank NASA's Earth Observatory (NEO) for providing the data shown in Fig. 8. The average land surface temperature (at day time) was provided by Jesse Allen, using data courtesy of the MODIS Land Group. The vegetation index is courtesy of Jesse Allen and Reto Stockli, using data provided by the MODIS Land Science Team. In addition, we thank the European Centre for Medium Range Weather Forecasts (ECMWF) for providing the ERA-Interim reanalysis simulations [DUS*11], which we used in Fig. 11.

Solution Fig. 4. The used mixing ratio was 50%. However, users were likely to predict a lower value.

Solution Fig. 9. The ground truth mixing ratio is 70% in both images. With the direct transfer function more regularity is visible, inclining users to select a lower value.

References

[AEM11] ALBUQUERQUE G., EISEMANN M., MAGNOR M.: Perception-based visual quality measures. In *IEEE Conference on Visual Analytics Science and Technology (VAST)* (2011), IEEE, pp. 13–20. 2

[BCS11] BARTRAM L., CHEUNG B., STONE M. C.: The effect of colour and transparency on the perception of overlaid grids. *IEEE Trans. Vis. Comput. Graph. 17*, 12 (2011), 1942–1948. 2

[BKC*13] BORGO R., KEHRER J., CHUNG D. H., MAGUIRE E., LARAMEE R. S., HAUSER H., WARD M., CHEN M.: Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics (State-of-the-Art Reports)* (2013), pp. 39–63. 2

[BSD09] BALZER M., SCHLÖMER T., DEUSSEN O.: Capacity-constrained point distributions: A variant of lloyd's method. *ACM Trans. Graph. 28*, 3 (July 2009), 86:1–86:8. 2

[CGW15] CHEN W., GUO F., WANG F. Y.: A survey of traffic data visualization. *IEEE Transactions on Intelligent Transportation Systems 16*, 6 (Dec 2015), 2970–2984. 2

[CM84] CLEVELAND W. S., MCGILL R.: Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of American Statistical Association 79*, 387 (1984), 531–554. 2

[CSS89] CHUBB C., SPERLING G., SOLOMON J.: Texture interactions determine perceived contrast. *Proc Natl Acad Sci, 86* (1989), 9631–9635. 2, 5, 6

[dGBOD12] DE GOES F., BREEDEN K., OSTROMOUKHOV V., DESBRUN M.: Blue noise through optimal transport. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 171:1–171:11. 2

[DRST14] DÜBEL S., RÖHLIG M., SCHUMANN H., TRAPP M.: 2D and 3D presentation of spatial data: A systematic review. In *2014 IEEE VIS International Workshop on 3DVis (3DVis)* (Nov 2014), pp. 11–18. 2

[DUS*11] DEE D. P., UPPALA S. M., SIMMONS A. J., BERRISFORD P., POLI P., KOBAYASHI S., ANDRAE U., BALMASEDA M. A., BALSAMO G., BAUER P., BECHTOLD P., BELJAARS A. C. M., VAN DE BERG L., BIDLOT J., BORMANN N., ET AL.: The ERA-Interim reanalysis: configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society 137*, 656 (2011), 553–597. 7

[EMdSP*15] ETEMADPOUR R., MOTTA R., DE SOUZA PAIVA J. G., MINGHIM R., DE OLIVEIRA M. C. F., LINSEN L.: Perception-based evaluation of projection methods for multidimensional data visualization. *IEEE Transactions on Visualization & Computer Graphics (Proc. IEEE Information Visualization)* (2015). 2

[Fat11] FATTAL R.: Blue-noise point sampling using kernel density model. *ACM Trans. Graph. 30*, 4 (July 2011), 48:1–48:12. 2

[FPV*13] FERREIRA N., POCO J., VO H. T., FREIRE J., SILVA C. T.: Visual exploration of big spatio-temporal urban data: A study of new york city taxi trips. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (Dec 2013), 2149–2158. 1

[GCNF13] GLEICHER M., CORRELL M., NOTHELFER C., FRANCONERI S.: Perception of average value in multiclass scatterplots. *IEEE Transactions on Visualization and Computer Graphics 19*, 12 (2013), 2316–2325. 2

[GWG*11] GOLDAU M., WIEBEL A., GORBACH N. S., MELZER C., HLAWITSCHKA M., SCHEUERMANN G., TITTGEMEYER M.: Fiber stippling: An illustrative rendering for probabilistic diffusion tractography. In *IEEE Symposium on Biological Data Visualization (BioVis)* (2011), IEEE, pp. 23–30. 2

[HSD13] HECK D., SCHLÖMER T., DEUSSEN O.: Blue noise sampling with controlled aliasing. *ACM Trans. Graph. 32*, 3 (2013), 25:1–25:12. 2

[JFLC08] JOHANSSON J., FORSELL C., LIND M., COOPER M. D.: Perceiving patterns in parallel coordinates: determining thresholds for identification of relationships. *Information Visualization 7*, 2 (2008), 152–162. 2

[Llo82] LLOYD S.: Least squares quantization in pcm. *Information Theory, IEEE Transactions on 28*, 2 (Mar 1982), 129–137. 2

[LMH*15] LEE T. M., MARKOWITZ E. M., HOWE P. D., KO C.-Y., LEISEROWITZ A. A.: Predictors of public climate change awareness and risk perception around the world. *Nature climate change 5*, 11 (2015), 1014–1020. 1, 2

[LT16] LEHMANN D. J., THEISEL H.: General projective maps for multidimensional data projection. *Computer Graphics Forum (Proc. Eurographics) 35*, 2 (2016). 7

[LWSF10] LI H., WEI L.-Y., SANDER P. V., FU C.-W.: Anisotropic blue noise sampling. *ACM Trans. Graph. 29*, 6 (Dec. 2010), 167:1–167:12. 2

[NK16] NUSRAT S., KOBOUROV S.: The state of the art in cartograms. In *Proceedings of the Eurographics / IEEE VGTC Conference on Visualization: State of the Art Reports* (Goslar Germany, Germany, 2016), EuroVis '16, Eurographics Association, pp. 619–642. 2

[NSBW08] NOCKE T., STERZEL T., BÖTTINGER M., WROBEL M.: Visualization of climate and climate change data: An overview. *Digital earth summit on geoinformatics* (2008), 226–232. 2

[OG12] ÖZTIRELI A. C., GROSS M.: Analysis and synthesis of point distributions based on pair correlation. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 170:1–170:10. 2

[Ost07] OSTROMOUKHOV V.: Sampling with polyominoes. *ACM Trans. Graph. 26*, 3 (July 2007). 2

[ROG17] ROVERI R., ÖZTIRELI A. C., GROSS M.: General point sampling with adaptive density and correlations. *Computer Graphics Forum (Proc. Eurographics) 36*, 2 (2017), 107–117. 1, 2, 3, 7

[ROM*15] ROVERI R., ÖZTIRELI A. C., MARTIN S., SOLENTHALER B., GROSS M.: Example based repetitive structure synthesis. *Comput. Graph. Forum 34*, 5 (Aug. 2015), 39–52. 2

[SBLC17] SHER V., BEMIS K. G., LICCARDI I., CHEN M.: An empirical study on the reliability of perceiving correlation indices using scatterplots. *Computer Graphics Forum (Proc. EuroVis)* (2017). 2

[SHS11] SCHULZ H.-J., HADLAK S., SCHUMANN H.: Point-based visualization for large hierarchies. *IEEE Transactions on Visualization and Computer Graphics 17*, 5 (2011), 598–611. 2

[TDN11] TOMINSKI C., DONGES J. F., NOCKE T.: Information visualization in climate research. In *15th International Conference on Information Visualisation* (July 2011), pp. 298–305. 2

[TQWZ06] TANG Y., QU H., WU Y., ZHOU H.: Natural textures for weather data visualization. In *Information Visualization, 2006. IV 2006. Tenth International Conference on* (2006), IEEE, pp. 741–750. 2

[Uli88] ULICHNEY R.: Dithering with blue noise. *Proceedings of the IEEE 76*, 1 (Jan 1988), 56–79. 2

[VW91] VAN WIJK J. J.: Spot noise texture synthesis for data visualization. *ACM Siggraph Computer Graphics 25*, 4 (1991), 309–318. 2

[WAKB09] WILLS J., AGARWAL S., KRIEGMAN D., BELONGIE S.: Toward a perceptual space for gloss. *ACM Trans. Graph. 28* (September 2009), 103:1–103:15. 2

[Wan14] WANG Y. Q.: Meteoinfo: Gis software for meteorological data visualization and analysis. *Meteorological Applications 21*, 2 (2014), 360–368. 1, 2

[War10] WARE C.: *Visual thinking: For design*. Morgan Kaufmann, 2010. 7

[WDSC07] WOOD J., DYKES J., SLINGSBY A., CLARKE K.: Interactive visual exploration of a large spatio-temporal dataset: Reflections on a geovisualization mashup. *IEEE Transactions on Visualization and Computer Graphics 13*, 6 (2007), 1176–1183. 2

[WPS*16] WOODRING J., PETERSEN M., SCHMEISSER A., PATCHETT J., AHRENS J., HAGEN H.: In situ eddy analysis in a high-resolution ocean climate model. *IEEE Transactions on Visualization and Computer Graphics 22*, 1 (Jan 2016), 857–866. 1, 2

[ZHWW12] ZHOU Y., HUANG H., WEI L.-Y., WANG R.: Point sampling with general noise spectrum. *ACM Trans. Graph. 31*, 4 (July 2012), 76:1–76:11. 2