

Pixnostics: Towards Measuring the Value of Visualization

Jörn Schneidewind*
University of Konstanz
Germany

Mike Sips†
Stanford University
USA

Daniel A. Keim‡
University of Konstanz
Germany

Abstract

During the last two decades a wide variety of advanced methods for the Visual Exploration of large data sets have been proposed. For most of these techniques user interaction has become a crucial element, since there are many situations in which an user or an analyst has to select the right parameter settings from among many or select a subset of the available attribute space for the visualization process, in order to construct valuable visualizations that provide insight into the data and reveal interesting patterns. The right choice of input parameters is often essential, since suboptimal parameter settings or the investigation of irrelevant data dimensions make the exploration process more time consuming and may result in wrong conclusions. In this paper we propose a novel method for automatically determining meaningful parameter- and attribute settings based on the Information content of the resulting visualizations.

Our technique called Pixnostics, in analogy to Scagnostics[1] automatically analyses pixel images resulting from diverse parameter mappings and ranks them according to the potential value for the user. This allows a more effective and more efficient visual data analysis process, since the attribute/parameter space is reduced to meaningful selections and thus the analyst obtains faster insight into the data. Real world applications are provided to show the benefit of the proposed approach.

Keywords: Visual Data Exploration, Visualization techniques, Visual Analytics

Index Terms: I.6.9 [Visualization]: Information Visualization—Visualization Techniques and Methodologies

1 Introduction

A wide variety of advanced Visual Exploration and Visualization methods have been proposed in the past. These techniques have proven to be of high value in supporting researchers and analysts to obtain insight in large data sets and to turn raw data into useful and valuable knowledge by integrating the human in the exploration process. However, with the increasing volume and complexity of today's data sets, new challenges for Visualization techniques arise. To keep step with the growing flood of information, Visualization techniques are getting more sophisticated, e.g. by integrating automated analysis methods or providing new visualization metaphors as proposed in the context of Visual Analytics [2].

But this also means, that visualization techniques are getting more complex, forcing the user to set up many different parameters to adjust the mapping of attributes to visual

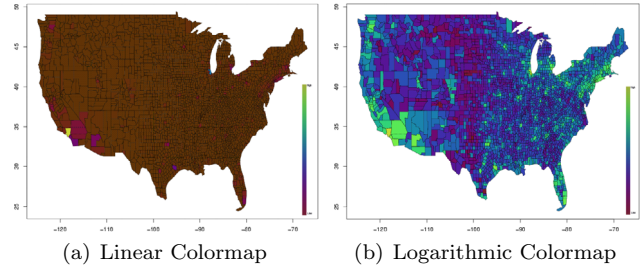


Figure 1: A typical application scenario – The visual analysis of a census data set involves different normalizations to a color scale; Although both visualizations are based on exactly the same input data, the right figure provides more insight since a logarithmic color scale is more suitable for the underlying data distribution

variables on the display space. In classical data exploration, playing with parameters to find a promising parameter setting is an important part of the exploration process, but with the increasing number and diversity of the parameters it becomes more and more difficult to determine a good parameter setup, which is vital for insightful visualizations.

For example, if we have 50 attributes (or attribute dimensions) and 4 parameters for the visual mapping, as e.g. in Pixel Bar Charts [3] employed in Section 4.2, then we have over 5 million possible mappings and it is very unlikely to find useful ones interactively.

Suboptimal parameter settings or the investigation of irrelevant data dimensions make the exploration process tedious and an interactive search impossible. In general, finding a good parameter setup is a challenging task for the analyst, since it is often not clear what is the best parameter setting for a given task, due to the huge parameter- and attribute space [4].

A simple application scenario is shown in Figure 1. The figure shows two choropleth maps visualizing USA population density data at county level. The two maps are based on the same input data, but created with two different parameter settings. More precisely, in the left figure a linear color mapping was chosen, in the right figure a logarithmic color mapping was chosen. It is easy to see that the linear mapping provides much less insight to the data than the logarithmic mapping, because the data is highly non-uniformly distributed. For instance, very high populated areas around Los Angeles, Chicago or Manhattan cause uniform dark colors for the remaining USA and it is almost impossible to see fine structures or differences in population density among them. In practice, the analyst does not know a priori which normalization function is best suited for a given dataset and he may test some preferred ones. Of course, there are typically much more parameters that have to be selected.

But the growing data complexity does not allow such playing with data by hand anymore. Therefore the paper aims at supporting the user in finding promising parameter se-

*schneide@inf.uni-konstanz.de

†ms@pixelmap.org

‡keim@inf.uni-konstanz.de

tups from the available parameter space to speed up the exploration process. We present a framework that employs automated analysis methods to detect potentially useful parameter settings for a given pixel-based visualization technique and an associated input data set, with respect to a given user task like Clustering or Outlier detection.

2 Problem Definition

In many application scenarios analysts have to deal with large parameter spaces when using visualization techniques to explore large data sets. These parameters control the visual encoding of the data, including the selection of attributes from the input data, the selection of the color scale, algorithm parameters, the selection of visual variables and so on. The problem is that the optimal parameter setting for a given task is often not clear in advance, which means that the analyst has to try multiple parameter settings in order to generate valuable visualizations. Since such selections can hardly be done manually, the integration of automated methods to support the analyst has been recognized as an important research problem in the context of Visual Analytics [2].

2.1 Parameter space of the Visualization Process

The challenging task in generating expressive visualizations is to find an adequate visual encoding of the input data set in the data display space. The visual encoding depends on the input data, the employed visualization technique and the visual variables given by a particular parameter setting. In classical data exploration, the mapping to visual variables described by Bertin [5], such as position (x, y), size, or color is manually controlled by the user. More precisely, the central process of visualization V can be described as

$$I(t) = V(D, S(P), t)$$

where data D is transformed using a specification S into a time varying image $I(t)$, according to the work proposed by Wijk [6]. To simplify matters, we aim at determining initial parameter settings for non-animated visualizations, therefore the time t can be excluded from our considerations. (Note that the complexity of the problem would be boosted exponentially, if we take time into consideration.)

Based on the formula above, the data set D is given as input data within a database environment with $D = (d_1, \dots, d_n)$. In the presentation of our approach, we focus mainly on demographic data provided by the US Census Bureau [7]. The data set contains, for example, information about US education levels, crime rates, housing or household incomes on different levels of detail (country, states, counties, block level). Typical exploration tasks focus on the extraction of information about housing neighbourhoods for particular areas within the US including the identification of correlations between statistical parameters like household income, house prices, education levels and crime rates.

The employed visualization technique V is defined by the application scenario and is given by the user. In this paper we use novel pixel-based techniques such as Pixel Bar Charts [3] and Jigsaw maps [8] as examples to explain our new idea.

Pixel Bar Charts (described in Section 4.2) have 4 parameters and more to adjust the visual encoding. The technique uses 1 parameter to separate the data into bars, 2 parameters for ordering of pixels in x and y direction and 1 parameter for color coding. Obviously, the manual selection of useful parameter settings does not scale with practical data sets, since the number of parameter settings which the user may

try is limited.

In our setting the input data D and the visualization technique V is given by the user and $P = \{P_i = (p_i^1, \dots, p_i^m)\}$ as instance of a parameter setting generating image $I(S(P_i))$ is determined by the system.

2.2 Limits and Problem Complexity

Most visualization techniques can handle less attributes than provided by the dimensionality of the input data set, so in the visualization step potentially useful attribute combinations must be selected from the data. As an example, we consider a data analyst who wants to use the Pixel Bar Chart technique to analyze real world customer purchase data. Such data sets typically include at least 20 dimensions (attributes), including name of item, price of item, name/id of the customer, status and so on. The number of possible parameter settings $P_i = (p_i^1, \dots, p_i^m)$ that control the visual encoding and therefore the number of possible images I_{P_i} is defined by the number of different attribute combinations of size m from the available number of dimensions d , given as:

$$|\{I(S(P_i)) = V(D, S(P_i), t)\}| = \frac{d!}{(d-m)!}$$

For the Pixel Bar Chart example, we may specify $m = 4$ attributes at once from the input data with 20 dimensions, which would result in 116,280 mappings. This number may be increased by additional parameters like different colormaps or different scalings. The equation above also shows that increasing dimensionality boosts the parameter space exponentially. If we have 50 attributes the number of possible mappings is 5.5 million. Then the analyst faces the problem, how to determine interesting subsets from the available data dimensions for visual analysis, that reveal interesting relationships.

2.3 Analysis Objectives

The aim of this paper is to provide techniques to prune the available parameter space to relevant settings to reduce the costs of creating insightful visualizations by filtering potentially relevant visualizations from irrelevant ones.

In [9] a semi-automated technique to search through visualization parameter space with applications in surface texturing is presented that focuses especially on perceptual and aesthetic concerns. The basic idea of this approach is to employ a genetic algorithm to guide a human-in-the-loop search through the parameter space.

In [1, 10, 11] graph theoretic approaches to solve these problems for scatterplots were proposed. This work called Scagnostics highly influenced our work. The goal of the mentioned approaches was to find interesting attribute relationships by creating scatterplot matrices from the data and then analyze each scatterplot, which reveals a relationship between 2 attributes, for certain properties using graph theoretic methods. The basic idea is to construct geometric graphs based on the data points of each scatterplot and then to compute relevance measurements from these graphs. For example, properties of the convex hull and the minimal spanning trees of the scattered points are used for outlier or cluster analysis.

With our approach we extend this idea to a broader set of visualization techniques. We provide analysis functions to analyse both patterns in the data using data analysis techniques as well as the patterns contained in the images by using image analysis techniques. Although data mining methods are commonly used for data analysis, only little research

has been done on analyzing visualizations with respect to their information content using image analysis methods.

In this paper we focus on pixel-based visualization techniques in which every data point corresponds to a pixel. Since we deal with pixel images instead of scatterplots, we call our approach Pixnostics instead of Scagnostics.

Our idea is to employ image analysis and retrieval methods to compute measurements based on the properties of the image pixels like color and pixel neighbourhoods. A very promising way to extract information about the content of an image is Shannon’s entropy measure [12]. It is frequently used in image processing and analysis. In [13] an entropy based approach for image cluster analysis is proposed, a more general approach for image retrieval using entropy is introduced in [14]. We adapt these techniques to measure the relevance of images I_i resulting from visualizations $V(D, S(P_i))$ with respect to a given task T .

In classical visual exploration the user visually analyzes a collection of data items to find answers to various questions (analysis tasks), whereas in our framework different analysis tasks are supported by evaluating properties of the resulting visualizations. Since the applied methods highly depend on the selected task, it is one of the major challenges to identify the most relevant tasks, identify their relationship to visual properties in the resulting image I and finally to find adequate analysis functions for each of the tasks, i.e. to find good predictors for these properties in images. Unique properties in images are homogenous areas, color outliers, edges and segments etc. Previous studies on visualization design proposed a range of different analysis goals and tasks [15] [16] [17] [18] [19]. They propose individual taxonomies of information visualizations using different backgrounds and models, so that users and analysts can quickly identify various techniques that can be applied to their domain of interest. Based on the proposed approaches, generic tasks like *identify*, *locate*, *cluster*, *associate*, *compare*, *correlate*, *match* and *sort* can be identified. In the following we describe the individual steps of Pixnostics where we mainly focused on analysis functions for clustering tasks.

3 The Pixnostics approach

Our Pixnostics approach follows a three step process based on the current task-at-hand:

- **Analytical Filtering and Pruning** of the set of possible images $\{I(S(P_i))\}$ by analyzing the parameter space P_i . The aim is to extract useful attribute selections and useful parameter settings automatically (candidate set CS),
- **Image Analysis** of the remaining candidate set CS – generating visualizations using the determined candidate attributes
- **Ranking and Output** of the candidate set CS – providing a ranking of candidate images I_{CS}

In the following we describe the individual steps in more details.

3.1 Step 1: Analytical Filtering and Pruning

In practice, the number of attributes is greater than the capabilities of most visualization techniques. The first step of our Pixnostics approach is therefore to determine relevant relationships among the different attributes analytically. We use data mining techniques, more precisely, correlation analysis, partial matching techniques and cluster analysis to accomplish this step.

3.1.1 Correlation Analysis

Attributes that are correlated may be interesting for detailed analysis, because they may reveal relevant impact relationships. Therefore we employ correlation analysis to find groups of correlated attributes. First, we determine the pair-wise global correlations among all measurements as given by Pearson’s correlation matrix. Pearson’s correlation coefficient r between bivariate data A_{1i} and A_{2i} with $(i = 1, \dots, n)$ is defined as

$$r = \frac{\sum_{i=1}^n (A_{1i} - \bar{A}_1)(A_{2i} - \bar{A}_2)}{\sqrt{\sum_{i=1}^n (A_{1i} - \bar{A}_1)^2 \sum_{i=1}^n (A_{2i} - \bar{A}_2)^2}} \quad (1)$$

where \bar{A}_1 and \bar{A}_2 are the means of the A_{1i} and A_{2i} values, respectively.

If two dimensions are perfectly correlated, the correlation coefficient is 1, in case of an inverse correlation -1. In case of a perfect correlation, we can omit one of the attributes since it contains redundant information. In Figure 2 an example

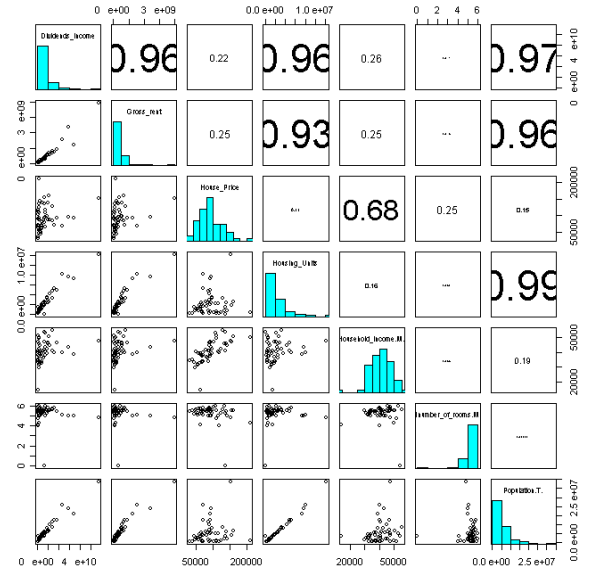


Figure 2: Identifying correlations in census housing data on US state level: Besides trivial correlations (e.g. Population and Number of House Units), some interesting correlations reveal e.g. between population and gross rent (because of demand and supply effects)

from the census housing data on US state level is shown, correlation coefficients for pairs of attributes are shown in the upper right half of the matrix, histograms in the diagonal show the data distribution. The figure clearly shows that states with high total population have high gross rents (0.96) or that Median Household incomes are correlated with Median House prices per state (0.68). The analyst may now investigate such relations in more detail. In most cases, however the correlations are not perfect and we are interested in high correlation coefficients and select sets of highly correlated attributes to be visualized.

An available alternative for adjacently depicting similar dimensions is to use the normalized Euclidean distance as a

measure for global similarity Sim_{Global} defined as:

$$Sim_{Global}(A_i, A_j) = \sqrt{\sum_{i=0}^{N-1} (b_i^1 - b_i^2)^2} \quad (2)$$

where $b_i^j = (a_i^j - \min(A_j)) / (\max(A_j) - \min(A_j))$

The global similarity measure compares two whole dimension such that any change in one of the dimensions has an influence on the resulting similarity. The defined similarity measure allows it to determine groups of similar attributes for the following visualization. Since in general, computing similarity measures is a non-trivial task, because similarity can be defined in various ways and for specific domains, special measures may be included for specific tasks.

3.1.2 Cluster analysis

In order to perform a visual analysis, it is important to have the possibility to partition the value ranges appropriately. Cluster analysis can help to do this based on the characteristics of the data instances. The cluster analysis may, for example, find out that the data instance of a data set may be partitioned into different groups, which may be then independently analysed using visualization techniques. Since attribute parameter values may be continuous (sales amount) or categorical values (item name) the clustering approach has to take these properties into account. There are a large number of clustering methods which have been proposed in the literature ([20] presents a nice overview). In the Pixnostics prototype we employed the *k-means* approach [21], one of the most popular approaches, since it is easy to implement and provides good results.

3.1.3 Classification Analysis

In some applications, for example in visual root-cause analysis, the goal of the data exploration is to understand the relationship between data attributes and some specific target attribute, e.g. which attributes have an influence on the target attribute. The task is to find the attributes which are best predicting the outcome of the target attribute. A well-known heuristic for this task is the GINI index [22], which is commonly used in decision tree construction.

Given a target attribute (e.g. a business metric) A_T which is partitioned into a disjoint set of k classes (e.g. accept, reject) or value ranges (e.g. large, medium, small) denoted by C_1, \dots, C_k , ($B = \bigcup_{i=1}^k C_i$), then the GINI index of an attribute A which induces a partitioning of A into A_1, \dots, A_m is defined as

$$InfoGain_{GINI}(A_T, A) = \sum_{i=1}^m \frac{|A_i|}{|A_T|} GINI(A_i) \quad (3)$$

where

$$GINI(A_i) = 1 - \sum_{j=1}^k \left[\frac{|C_j|}{|A_i|} \right]^2$$

The *InfoGain* is determined for all attributes and attribute combinations and the attributes with the highest *InfoGain* with respect to the target attribute A_T are chosen for visualization. These attributes are best predicting the outcome of the target attribute and therefore they may be relevant for detailed analysis.

3.2 Step 2: Image Analysis

Once we have selected candidate parameter settings $P_i, i = 1, \dots, max$ based on promising attribute selections, where max is the number of parameter settings, we generate visualizations by computing all possible mappings of the candidate parameter set to visual variables, and then determine the Information content of the resulting images using image analysis.

In the first step we generate and store I_{P_i} as a matrix U of scalars representing gray scale values, the pixel-matrix representation with $U = (u_{i,j}), i \in [0, \dots, I_{width}], j \in [0, \dots, I_{height}]$ (color images are converted to gray values). The next step is to process the image, to generate some measurements of its information content and thus return the relevance of this image with respect to a given task T . Numerous image processing operations $\sigma()$ exist in the literature, e.g. for Image Segmentation, Edge Detection, Image Denoising or Image Inpainting [23].

The base for our analysis is the distribution of gray values within the image. Thus we are interested to know the pixel

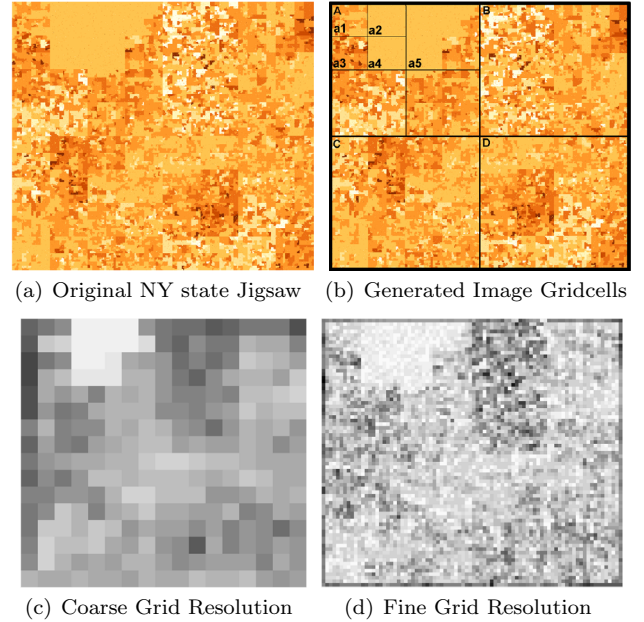


Figure 3: Basic idea of grid based information content: Based on the entropy values for certain grid resolutions, measurements for the relevance of the image are generated. Darker gray levels correspond to higher entropy values.

distribution H in certain areas of the Image as a function of gray levels g . Image Histograms are an efficient way to reach this goal. The histogram of the 2D image $g(U)$ can be seen as a 1D function $H[g]$ where the independent variable is the gray value g and the dependent variable is the number of pixels H with that level. We can then use the histogram properties to make assumptions about the information contained in the image. For example, if most pixels in an image are contained in a small range of gray levels, the image can be seen as redundant since it provides little new information and thus the underlying parameter setting would not lead to insightful visualizations. If there are too many different gray levels, the image represents noise and it is not likely that it contains relevant information. An image with a bimodal histogram (i.e. a histogram with two peaks) may contain clusters and may be relevant for visual exploration.

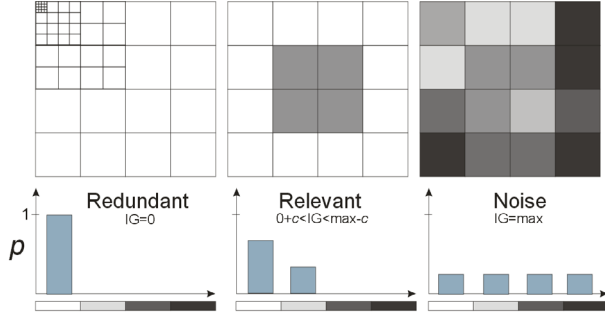


Figure 4: Information content (IG) of different gray level images. From an analyst's point of view, interesting images should have an Information content in a certain range c between 0 and IG_{max}

Since all pixels in the image must have some gray value in the allowed range, the sum of populations of the histogram bins must be equal the total number of image pixels N :

$$N = \sum_{g=0}^{g_{max}} H(g)$$

where g_{max} is the maximum gray value ($g_{max} = 255$ for an 8-bit quantizer). The histogram function is equal to the scaled probability distribution function $p(g)$ of gray levels in that image:

$$p(g) = \frac{1}{N} H(g) \text{ with } \sum_{g=0}^{g_{max}} p(g) = 1$$

Based on the probability distribution we can now compute a measure for the information content of the image. In general any function $\sigma()$ can be used, but a common way of doing this is Shannons Entropy [12], which is equal to the minimum number of bits which are required to store the image. If the probability of gray level g in the image $u(i, j)$ is represented as $p(g)$, the definition of the quantity of information in the image is:

$$E(g) = - \sum_{g=0}^{g_{max}} p(g) \log_2(p(g)) \quad (4)$$

From this definition, it is easy to show that the maximum information content E is obtained if each gray level has the same probability; in other words, a flat histogram corresponds to maximum information content. The minimum information content $E = 0$ is obtained if the image contains only one gray level. Since minimal information content means redundancy and maximum information content means information overload or noise, the interesting images should have an information content in between, e.g. in a task dependent range c shown in Figure 4

Alternatively we use the standard deviation $stdev$ as a measure of spread of gray levels g in a given image I with N as the number of different gray levels in the image:

$$stdev(I, g) = \sqrt{\frac{1}{N} \sum_{i=1}^N (g_i - \bar{g})^2} \quad (5)$$

Since we want to analyze the images not only in whole, but also find interesting local patterns in the image, we use a regular grid to separate the image in regular grid cells and then apply the methods mentioned above to compute

values for each grid cell. The computation of the information content of each cell is identical to the methods described above.

The only difference is that from the individual grid values we then compute a single relevance value for each image, described in the next section. To adapt our method to given application scenarios, we do not only use a fixed grid-resolution, but a hierarchy of grid cells, efficiently implemented by a quadtree data structure [24].

Input: Candidate Set

$$C = \{\{A_D^1, \dots, A_D^h\}, \{P_1, \dots, P_l\}\}$$

with Candidate Data Attributes A_D^1, \dots, A_D^h ($h < n$),

Candidate Parameter Settings P_1, \dots, P_l : ($l < k$),

Visualization V with Specification S

Performed Task T

Output: Ranking Scores $R(\{I(S(P_i)) = V(\{A_D^1, \dots, A_D^h\}, S(P_1, \dots, P_l))\})$

Procedure Visualization Analysis

Generate $\{I_i = V(\{A_D^1, \dots, A_D^h\}, S(P_1, \dots, P_l))\}$

for $i \leftarrow 1$ **to** $|C|$ **do**

 ComputeRegularGridonImage(I_i)

for each GridCell $GC(I_i)$ **do**

 ComputeEntropyValues $f(GC(I_i))$

end

for all Entropy values $f(GC(I_i))$ **do**

$R = \text{ComputeRankingScore}(f(GC(I_i)), T)$

end

end

return Ranking Scores $R(I_i)$

Algorithm 1: Entropy-based Image Analysis

3.3 Step 3: Ranking and Output to the User

Having a function f , such as $E, stdev$, which measures the information content of an image, we can now compute rankings from the candidate parameter sets with respect to a given user task T . For effective data analysis we can identify two major tasks that are common in most analyses processes, namely outlier analysis, including the search for local outliers or values of interest (find all counties or cities that have similar household income or unexpected household income), and cluster analysis (find areas with similar statistical parameters). In our prototype framework we provide ranking functions for both tasks and show how we applied it to real world data sets. Of course, the user may also use other ranking functions for specific tasks, which can be easily integrated into the Pixnostics framework.

3.3.1 Computing the Global Ranking Score

The first step in processing the image is to superpose a regular grid on the image. Each grid cell may then again be separated into grid cells at finer resolution until a sufficient resolution is reached as shown in Figure 3 (b).

We initialize each regular grid cell $GC(I_i)$ with its information content score $f(GC(I_i))$. Then we start to merge regular grid cells $GC_k(I_i)$ and $GC_l(I_i)$ that have similar content scores, to larger cells. The new content score $f'(GC(I_i))$ is determined using local term weighting $f(GC_{common}(I_i)) = l(f(GC_k(I_i)) + f(GC_l(I_i)))$. The term weight function l is defined over the set of all regular grid cells $\{GC(I_i)\}$. We

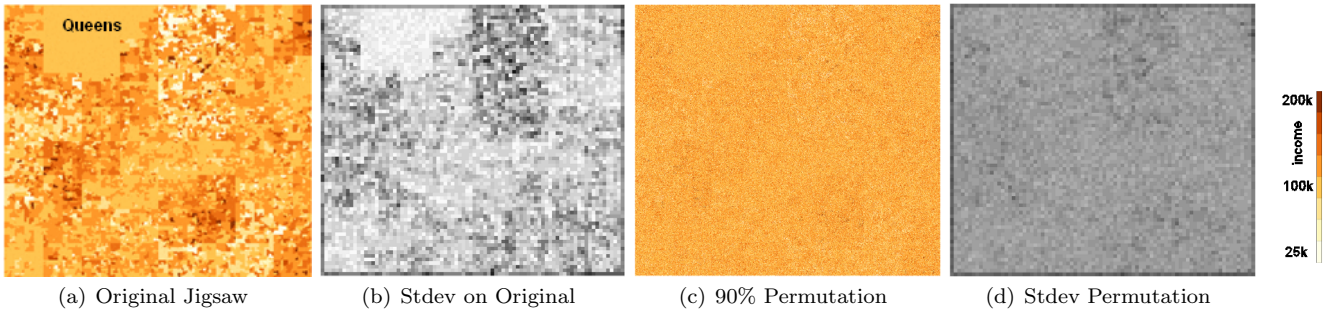


Figure 5: Visualization of Information content: Jigsaw maps are generated from NY median household income, darker colors correspond to higher income. Gray levels show the information content of image sections, darker gray levels correspond to higher information content. The permuted image has significant higher information content, which indicates bad clustering properties

use this measure to investigate clustering properties, where we are looking for images with higher information content f on coarser grid resolutions and with lower information content f at finer grid resolutions. The technique is similar to Single Linkage Clustering. We order on each resolution level the grid cells according to their value f , starting with the finest resolution. Then we expand grid cells with similar low content scores and sum up their weight while enlarging the grid cells. Since the spread of gray levels is typically higher if the size of the grid cell increases, we use cell size as an weighting factor. Note that the local term weight l directly depends on the given task T . That means, the user just needs to provide useful weight functions to get a relevance measure of the image for the given task. Well-known and widely used weight functions are binary operators, logarithmic or augmented normalized term frequency.

In outlier analysis, the goal is to compute a ranking of a collection of images in such a manner that images showing outliers should have higher global ranking scores. The inverse normalized term frequency is common choice for outlier analysis. The normalized term frequency l is defined as the logarithm of the sum of two information content scores $f(GC_k(I_i))$ and $f(GC_l(I_i))$ normalized by the total number of regular grid cells $|\{GC(I_i)\}|$.

$$l(f(GC_k(I_i)) + f(GC_l(I_i))) = \log\left(\frac{|\{GC(I_i)\}|}{f(GC_k(I_i)) + f(GC_l(I_i))}\right)$$

4 Evaluation and Application

To show the usefulness of our approach we applied the Pixnostics technique to generate Jigsaw maps and Pixel Bar Charts. The proposed experiments show how Pixnostics can steer the visual exploration process in an unsupervised manner, to increase the efficiency of the exploration process and to actively support the analyst to reduce the effort of getting insight from the data.

4.1 Census data Jigsaw Maps

Our first application example analyzes US census data, in particular median household income for the state of New York on block level. We generated visualizations using Jigsaw maps [8], a pixel-based space filling technique. The basic idea is to map the census data into the 2D plane in such a way that properties like locality and clusters in the data are preserved by using a space filling curve.

To verify our proposed techniques, we generated a Jigsaw map from the New York state census median household income data on block level which should preserve the clusters in the data (clusters of areas with high/low income) and






Figure 6: Census data Jigsaw unsorted (top) und sorted (line by line starting at top left corner with most relevant) by ranking function based on Entropy and clustering task (bottom)



(a) Ranking of images generated from promising attribute selections. Each Box shows a thumbnail of a Pixel Bar Chart based on a certain attribute selection, ordered by information content (desc)



(b) Ranking after image analysis: The top 25 images are shown, with a fixed target attribute as splitting attribute for the bars

Figure 7: Pixel Bar Chart showing top 25 results after image analysis using entropy measure. It is easy to see that the bar in the middle ("reject" parts) show significant differences in comparison to the 2 other bars.

their spatial location shown in Figure 5(a). As the figure shows, in the Queens area households with a income around \$100,000 are clustered. Then we permute the pixels in the data at different permutation rates. This should of course destroy, or at least reduce the clustering/locality properties. Now we apply our automated analysis function based on the clustering task, that ranks the underlying figures according to their clustering properties. The original jigsaw should of course be the image with the highest rank, since it provides the best clustering properties. The more permutation in the image, the lower should the relevance of the image, i.e. its rank, be.

Figure 6 shows the experimental results. The upper figure shows the unordered input data set, a set of Jigsaw images. It is easy to visually identify images with good clustering properties, i.e. images having a cluster with low income in the upper left corner surrounded by high income areas. In the lower figure, the result after the analysis step is shown. It is easy to see that figures with good clustering properties are ranked first, while images containing more noise have lower relevance. To determine the ranking, either the Entropy or the Standard deviation of the pixel gray levels in combination with the regular grid cell hierarchies are employed. Figure 5 shows the rationale for the ranking. An image that provides a good clustering has areas with very low Entropy or low Stdev of gray levels while the complete figure does not necessarily have low Entropy. Therefore we start with a fine grid and determine the information content of each cell like shown in Figure 5(b). Then we hierarchically compare neighbouring grid cells similar to Single Linkage Clustering and try to extend clusters. Finally we aggregate the Information content of the clusters and order the images according to their Information content values.

4.2 Pixel Bar Charts

Pixel bar charts[3] are derived from regular bar charts. The basic idea of a pixel bar chart is to present the data values directly instead of aggregating them into a few data values by representing each data item by a single pixel in the bar chart.

The detailed information of one attribute of each data

item is encoded into the pixel color and can be accessed and displayed as needed. To arrange the pixels within the bars one attribute is used to separate the data into bars and then two additional attributes are used to impose an ordering within the bars along the x and y axes. A Pixel Bar Chart can be seen as a combination of traditional bar charts and x-y diagrams.

Although Pixel Bar Charts have been successfully applied to explore large data sets, the analyst has to choose selections of attributes for separation, ordering and color coding of data points from the underlying data manually, according to his analysis tasks. On one hand, this is time consuming since he has to try multiple parameter settings even those that do not reveal interesting patterns, on the other hand he may overlook interesting patterns since only a few attribute combinations can be analyzed manually. To face this problem we combined Pixnostics and Pixel Bar Charts, to guide the analyst through the exploration process and indicate potentially interesting parameter settings.

We applied our approach to a production data example. The data set contains data from an assembly line, in particular measurements from different stages of the assembly line like cast temperatures, part measurements and the quality of the output. All in all the data set contains 22 attributes. The output parts are classified into 3 groups: accept, reject, rework. Parts that are grouped "accept" pass the quality check, "rework" parts need to be reworked to pass the quality check and "reject" parts must be rejected because of defects. The analysis of such data is an important task in order to reduce rejected parts and thus to reduce production cost. Using Pixel Bar Charts, the analyst faces to problem of how to find groups of attributes that may influence the quality of the output. There are 175560 combinations possible to choose 4 attributes as visual variable from 22, even if the target attribute "Quality" is fixed for separation of the bars, there are still over 9000 combinations for selection of 3 attributes out of 22, which can't be checked manually. Therefore we first apply our automated analysis tools to determine attributes that most influence the "Quality" variable, using correlation and classification analysis. Of course we can additionally prune all parameter settings




Figure 8: Pixel Bar Chart constructed from the output result of Pixnostics (from the Chart in the upper right corner in Figure 7 (b)).

where "Quality" is not involved, since these will not bring us any new insight.

From the remaining combinations we either generate images and order them by information content directly, like shown in Figure 7 (a) or we can filter Pixel Bar Charts where the target attribute is fixed as splitting attribute and select the most valuable ones from them, as shown in Figure 7 (b). The figure shows the 25 most relevant Pixel Bar Charts having "Quality" as splitting attribute. Note that the left bar shows parts that are "rework", the middle bar shows "reject" parts and the right bar show "accept" parts. It is easy to see that the "reject" bars look significantly different than the rest. The analyst may now select a single image from the provided images, and a Pixel Bar Charts is created from this selection as shown in Figure 8. The analyst can now easily discover relevant patterns by visual based root cause analysis. In the image the color shows the temperature of a particular casting mold and the ordering in y direction shows the duration of the part at this stage. It is easy to see that the casting mold had a significantly higher temperature for "reject" parts, which is a potentially reason for a damaged part. In this manner the analyst may investigate further high ranked images, which provides a more efficient way of visual analysis than manual feature selection.

5 Conclusion and Future work

Integrating automated analysis methods into the Visual Exploration process is an important challenge in the age of massive data sets and has been recognised as a major research area in the context of Visual Analytics. Therefore the aim of this paper is to show how unsupervised analysis functions can help to speedup the Visual Exploration process by supporting the user with task driven relevance functions for a more effective data analysis. The basic idea of the proposed method is to measure the relevance of the resulting visualization with respect to input parameters and user tasks and to provide a ranking of potentially useful initial visualizations. This helps the analyst to focus on relevant parts of the data and relevant parameter settings and could lead to an improved exploration process. We provided a formal definition of our work and showed how the technique can be used with Jigsaw's and Pixel Bar Charts. Future work will focus on the improvement of the proposed technique and its application to a variety of visualization techniques, including geometric and iconic techniques. Another important issue is the evaluation of the proposed analysis functions, where bridging the gap between human expressions of an analysis task and the corresponding analysis functions in Pixnostics is one of the major topics for further research.

Acknowledgement

This work was supported by the Max Planck Center for Visual Computing and Communication (MPC-VCC).

References

- [1] Leland Wilkinson, Anushka Anand, and Robert Grossman. Graph-theoretic scagnostics. In *INFOVIS '05: Proceedings of the 2005 IEEE Symposium on Information Visualization*, page 21, Washington, DC, USA, 2005. IEEE.
- [2] J.J. Thomas and K.A. Cook. Illuminating the path: Research and Development agenda for Visual Analytics. In *IEEE*, pages 79–86, October 2005.
- [3] D. A. Keim, M. C. Hao, U. Dayal, and M. Hsu. Pixel Bar Charts: A visualization technique for very large multi-attribute data sets. *Information Visualization*, 1(1):20–34, 2002.
- [4] Robert Spence. *Information Visualization*. ACM Press Books, Pearson Education Ltd., UK, 2001.
- [5] J. Bertin. *Semiology of graphics*. University of Wisconsin Press, Madison, Wisconsin, 1983.
- [6] Jarke J. van Wijk. The value of visualization. In *Proceedings of IEEE Visualization*, pages 79–86, Minneapolis, MN, October 2005.
- [7] US Census Bureau, <http://www.census.gov>, 2003.
- [8] Martin Wattenberg. A note on space-filling visualizations and space-filling curves. In *INFOVIS*, page 24, 2005.
- [9] Donald H. House, Alethea Bair, and Colin Ware. On the optimization of visualizations of complex phenomena. In *IEEE Visualization*, page 12, Minneapolis, MN, 2005.
- [10] J.W. Tukey. *Exploratory Data Analysis*. Addison Wesley Publishing, Reading, MA, 1977.
- [11] J.W. Tukey and P.A. Tukey. Computing graphics and exploratory data analysis: An introduction. In *Proceedings of the Sixth Annual Conference and Exposition: Computer Graphics 85*. Nat. Computer Graphics Assoc., 1985.
- [12] M.D. Esteban and D. Morales. A summary of entropy statistics. *Kybernetika*, 31(4):337–346, 1995.
- [13] M. Koskela, J. Laaksonen, and E. Oja. Entropy-based measures for clustering and som topology preservation applied to content-based image indexing and retrieval. *17th Int. Conference on Pattern Recognition*, 2:1005–1009, 2004.
- [14] John Zachary, S. S. Iyengar, and Jacob Barhen. Content based image retrieval and information theory: a general approach. *J. Am. Soc. Inf. Sci. Technol.*, 52(10):840–852, 2001.
- [15] B. Shneiderman. The eye have it: A task by data type taxonomy for information visualizations. In *Proc. IEEE Conference on Visual Languages*, pages 336–343, 1996.
- [16] Peter R. Keller and Mary M. Keller. *Visual Cues - Practical Data Visualization*. IEEE Press, 1st edition, 1993.
- [17] Stephen M. Casner. Task-analytic approach to the automated design of graphic presentations. *ACM Trans. Graph.*, 10(2):111–151, 1991.
- [18] Ed H. Chi. A taxonomy of visualization techniques using the data state reference model. In *INFOVIS*, pages 69–76, 2000.
- [19] D. Keim. Designing pixel-oriented visualization techniques: Theory and applications. *Transactions on Visualization and Computer Graphics*, 6(1):59–78, Jan–Mar 2000.
- [20] A. Hinneburg and D. A. Keim. Clustering techniques for large data sets From the past to the future. In *KDD '99: Tutorial notes of the fifth ACM SIGKDD*, pages 141–181, New York, NY, USA, 1999. ACM Press.
- [21] T. Kanungo, D. Mount, N. Netanyahu, C. Piatko, R. Silverman, and A. Wu. An efficient k-means clustering algorithm: analysis and implementation, 2002.
- [22] J. M. Zytkow. Types and forms of knowledge (patterns): decision trees. In *Handbook of data mining and knowledge discovery*, Oxford University Press, Inc., 2002.
- [23] Steve Mann. *Intelligent Image Processing*. John Wiley and Sons, November 2 2001.
- [24] R.A. Finkel and J.L. Bentley. Quad trees: A data structure for retrieval on composite key. *Acta Informatica*, 4(1):1–9, 1974.