

Information Visualization

<http://ivi.sagepub.com/>

Visual exploration of classification models for various data types in risk assessment

Malgorzata Migut and Marcel Worring

Information Visualization 2012 11: 237 originally published online 16 February 2012

DOI: 10.1177/1473871611433715

The online version of this article can be found at:

<http://ivi.sagepub.com/content/11/3/237>

Published by:



<http://www.sagepublications.com>

Additional services and information for *Information Visualization* can be found at:

Email Alerts: <http://ivi.sagepub.com/cgi/alerts>

Subscriptions: <http://ivi.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://ivi.sagepub.com/content/11/3/237.refs.html>

>> [Version of Record](#) - Jun 27, 2012

[OnlineFirst Version of Record](#) - Feb 16, 2012

[What is This?](#)

Visual exploration of classification models for various data types in risk assessment

Information Visualization
11(3) 237–251
© The Author(s) 2012
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1473871611433715
ivi.sagepub.com



Malgorzata Migut and Marcel Worring

Abstract

In risk assessment applications well-informed decisions need to be made based on large amounts of multi-dimensional data. In many domains, not only the risk of a wrong decision, but also of the trade-off between the costs of possible decisions are of utmost importance. In this paper we describe a framework to support the decision-making process, which tightly integrates interactive visual exploration with machine learning. The proposed approach uses a series of interactive 2D visualizations of numerical and ordinal data combined with visualization of classification models. These series of visual elements are linked to the classifier's performance, which is visualized using an interactive performance curve. This interaction allows the decision-maker to steer the classification model and instantly identify the critical, cost-changing data elements in the various linked visualizations. The critical data elements are represented as images in order to trigger associations related to the knowledge of the expert. In this way the data visualization and classification results are not only linked together, but are also linked back to the classification model. Such a visual analytics framework allows the user to interactively explore the costs of his decisions for different settings of the model and, accordingly, use the most suitable classification model. More informed and reliable decisions result. A case study in the forensic psychiatry domain reveals the usefulness of the suggested approach.

Keywords

Visual analytics, interactive visual exploration, decision boundary visualization, multi-dimensional space, classification

Introduction

Risk assessment is important in applications such as disaster management, security and human health^{1,2}. Decisions have to be made based on the analysis of multi-dimensional data. The risk assessment expert has a difficult task to understand such data, make decisions based on it and, moreover, to foresee the consequences of those decisions. The crucial part of the decision-making process is determining the optimal balance between the different costs incurred by the decision.

An example domain is the field of forensic psychiatry, where the risk of criminal behavior of patients with psychiatric disorders has to be predicted. It is of high social importance that a suitable trade-off is achieved between incorrectly assessed re-offending criminal patients and incorrectly assessed non-re-offending ones. Clearly, these

cases have a completely different impact on the society and the patients themselves. Another example is from the medical domain, where a certain disorder, for instance a genetic liver disorder, is to be diagnosed. The consequences of both types of mistake, diagnosing the healthy patients and not diagnosing the ill patients, may be fatal.

Intelligent Systems Lab Amsterdam, University of Amsterdam,
Science Park 904, 1098XH Amsterdam, The Netherlands

Corresponding author:

Malgorzata Migut, Intelligent Systems Lab Amsterdam, University of Amsterdam, Science Park 904, 1098XH Amsterdam, The Netherlands.
Email: mmigut@gmail.com

Neither a fully automatic nor a pure expert knowledge-driven assessment is the answer to the challenging requirements of risk assessment. Combining automated data analysis, in particular machine learning, and interactive visualization techniques with user expertise in the specific domain has proven itself useful in various applications¹⁻³. These techniques could help the domain expert gain insight to the problem and make well-informed decisions. This visual analytics approach, as formally described by Keim et al.¹, integrates the strength of the expert and the strength of the machine. Visual analytics is also promising for risk assessment.

In visual analytics, many different methods are used to visualize multi-dimensional datasets such as scatter-plots, heatmaps, parallel coordinates and parallel sets⁴. The visualizations should relate to the domain-specific understanding framework of the expert. Therefore, they should represent the data in the attributes themselves and not some (non-linear) combinations. Furthermore, the different types of data, including numerical and ordinal, should be visualized in a consistent manner. To compensate for the potentially limited expressiveness of such simple visualizations, all visual elements should be tightly linked to reveal their inter-relationships⁵.

To utilize the knowledge of an expert, the data elements should be assigned a meaning. However, it is often difficult to interpret the visualized numerical or ordinal values for each data element. Images related to the data elements could provide an expert with insightful associations. In forensic psychiatry, the photographs of the patients could be used to help an expert recall the information of the patient, expanding the scope of the dataset. In the medical example, the radiology images of the liver allow the expert to directly relate measures derived from an image to the actual content of the image.

Next to the visualization of the data, the core element in the decision-making process is an accurate and transparent predictive model. Although machine learning techniques may exhibit excellent performance, it is often difficult to obtain intuitive understanding of the induced model. The transparency of the model is, however, essential for the user to understand his decisions. Moreover, intuitive understanding of the classifier obtained is important not only for the validation of the model, but also to deepen the insight into the domain and its underlying processes. In order to help a user understand the intrinsics of a domain, it would be favourable to obtain a visual comprehension of what the classifier looks like in the multi-dimensional space. Many visualization techniques, such as receiver operating characteristics (ROC) curves or precision and recall graphs, are used to present the results of the classifier⁶. However, visualization of the results alone does not give insight to the model.

One of the most informative characteristics of the classification model is the decision boundary. It can be easily visualized for both 2D and 3D datasets⁶. This

allows a user to identify the areas belonging to different classes for the displayed dimensions as imposed by the classifier. Thus, it contributes to better comprehensibility of the model.

To allow better understanding of the model in relation to the dataset, we could also take advantage of the expert's knowledge about data elements. Using the visual representation (images) of the data elements, as mentioned earlier, can trigger immediate associations⁷. In particular, visualizing the critical data elements responsible for the cost change of the classifier's output would contribute to awareness in cost selection.

Based on the above considerations, we propose an interactive risk assessment framework. We interactively couple visualizations of different data types, the visualizations of the decision boundary and classifier performance, and associative images representing the data. This provides an expert with the possibility of visually exploring the classifier and the costs of classification for different trade-offs.

This paper is organized as follows. In the next section we present an overview of existing systems combining interactive visualizations with machine learning and their context of use. In the subsequent section, we propose visualizations for different types of data and the classifiers, in relation to the individual data elements. From there, we show how to make the visualization of the performance curve more informative by linking it with the visualizations of the data. Following the section describing classifiers' performance visualization, we propose an interactive risk assessment system and show its application in forensic psychiatry.

Related work in visual analytics

In recent years the idea of combining classification with interactive visualization has attracted a great deal of interest. A framework for visual analytics was formally described by Keim et al.¹. It proposes a tightly coupled system, with controlled interaction between the automated data analysis and visualization elements. Moreover, it allows users to steer the visualization process and to actively participate in the classification process. Several of such tightly integrated approaches have been proposed in different usage contexts. Building upon Keim's framework, Yu et al.⁸ propose a smooth interface between data mining and visualization for multimedia data in social and behavioural studies. They visualize all intermediate and final results of data mining, in terms of found patterns, allowing the user to obtain new insights and develop more hypotheses about the data. Ankerst et al.⁹ propose the DataJewel architecture coupling a visual, an algorithmic and a database approach for temporal data mining. The system focuses on the improvement of the discovery of useful patterns and not on how these patterns are obtained. Interactive construction of decision tree classifiers has been proposed in references

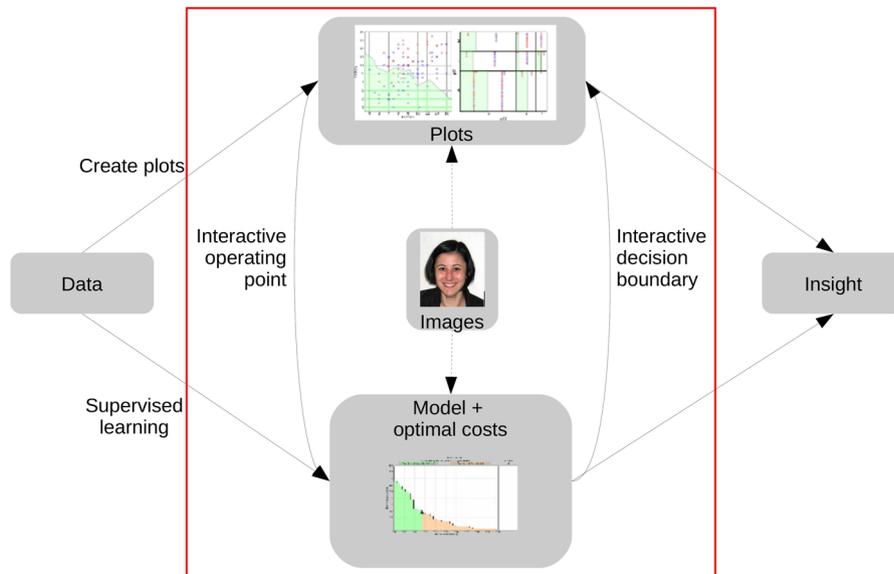


Figure 1. The risk assessment framework proposed in this paper, following the general visual analytics framework proposed by Keim et al.¹. The elements within the rectangle indicate the focus of this paper. For privacy reasons we illustrate the usage of images in the framework with the photographs of the authors of this paper, rather than the real data.

10–12. In these methods the user can interactively select the splitting attribute from the dataset visualization. Then, the current decision tree is visualized and the user can proceed with expanding the decision tree. Fuchs et al.³ describe a way of improving and analysing a classifier. Starting from an initial hypothesis, created with linking and brushing, the user steers a heuristic search algorithm to look for alternative hypotheses.

Our work shows how the classifiers' results, in terms of the correctly classified and misclassified data elements, can be visualized and explored, with particular emphasis on the classification costs and the data elements with the highest risk factor. We show how existing techniques can be integrated intelligently to support risk assessment, by combining and structuring them in a highly interactive way, as illustrated in Figure 1. This work is an extended version of our work previously published in reference 13.

Visualizations for exploration of classification models

The main visualization components constituting our system are the visualizations of the data, the classifier and the performance curve. The choice of the techniques we use to facilitate these visualizations is dictated by several requirements that emerge from the specifics of the risk assessment problem. Therefore, we start this section with the description of the problem and we explain our choice of techniques.

Problem description

Let us first formalize the problem. Assume a training dataset with n objects represented by feature vectors in

a p -dimensional space. Risk assessment is a two-class problem, with a positive and a negative class. We consider the case where the complex features are known individually to the expert, so the dimensionality of the space is limited. Feature vectors may contain variables of different data types, in particular numerical, ordinal and nominal. A classifier trained on the dataset results in a decision boundary in the p -dimensional space. An object, which, according to this decision boundary, is classified as positive, is called a true positive if the actual value is also positive. It is termed a false positive if the actual value is negative. The object classified as negative is called a true negative if the actual value is negative and a false negative if the actual value is positive.

One of our tasks is to visualize the model. The risk assessment context imposes several requirements. In particular, our requirements emerged from the collaboration with risk assessment researchers in forensic psychiatry. However, we feel they are quite generic. First of all, the expert has deep knowledge of the individual features and can only understand the model if it is related to those. It implies that for each type of data comprising the feature vector, we need to find a suitable data visualization technique, which would also allow visualizing an interpretable decision boundary. Moreover, the visualization should support the user's interest in each individual data element. The expert should be able to immediately determine the result of the trained model for each individual data element and its relation to the decision boundary.

The visualization of the data itself should be simple and preferably should employ well-known techniques familiar to the expert. We could minimize the number of visualized dimensions, making the exploration potentially

easier for the expert, using one of many dimension-reduction techniques. This is, however, not desirable for several reasons. Those methods have been designed to create 2D views by finding the most interesting 2D planes. However, their notion of ‘interest’ has a specific and restricted interpretation. Thus, the method might fail in finding all of the planes of interest. Furthermore, and more importantly for our purpose, methods that compute 2D projections of multi-dimensional spaces are often beyond the understanding capabilities of the expert. Moreover, the expert should be stimulated to make use of his knowledge about the data that is not contained in the data themselves. Therefore, a meaning has to be assigned to the data elements in such a way that it has an associative role for the expert.

Numerical data

Data visualization. Keim¹⁴ provided a taxonomy of multidimensional visualizations. The categories listed include standard 2D/3D displays, geometrically transformed displays, iconic displays, dense pixel displays, and stacked displays.

Our requirement of simplicity, familiarity among users and visualizing the attributes themselves accommodating high visual clarity brings us to the scatterplot¹⁵. This implies trading in some very good features of more advanced visualization techniques, but it does fulfil the risk assessment requirements. In a 2D scatterplot, data elements are drawn as points in the Cartesian space defined by two graphical axes defined by the actual attribute values. Therefore, they accommodate the understanding framework of the user. Scatterplots are also frequently used and are basic building blocks in statistical graphics and data visualization¹⁶. Multidimensional visualization tools that feature scatterplots, such as Spotfire¹⁷, XmdvTool¹⁸, Tableau/Polaris¹⁹ and GGobi²⁰, typically allow mapping of data dimensions to graphical properties such as point colour, shape and size.

However, the number of dimensions that a single scatterplot can reliably visualize is considerably less than is found in many realistic datasets. Therefore, multiple scatterplots should be visualized in sequence for all combinations of the dimensions, which can then be explored by the user. This approach yields little structure to the visual exploration and provides no relations between the data dimensions. To that end, multiple plots can be arranged in a scatterplot matrix¹⁶ and we can link the multiple visualizations of combinations of two dimensions in order to reveal relationships between them⁵ through interactions, as described in Matrix overview section.

Model visualization. In the scatterplot we use colour and shape to visually express the results of classification. Colour expresses the original class membership of the data. Shape indicates whether or not the data element is

misclassified by the classifier. As mentioned in the introduction, the decision boundary is one of the most informative characteristics of the classifier. The visualization of multi-dimensional decision boundaries is, however, a difficult problem. Several attempts have been made^{10,21,22}. Those methods do not relate the visualization of the decision boundary to the meaningful axes for the user. In ref.²¹ the authors visualize the support vector machine (SVM) classifier using a projection-based tour method. The authors show visualizations of histograms of the data-predicted class, visualization of the data, and the support vectors in 2D projections. They weigh up the plane coordinates to choose the most important features for the classification. Poulet¹⁰ displays histograms of the data distribution according to the distances to the boundary and a set of linked scatterplot matrices or parallel coordinates for SVM. Hamel²² uses self-organizing maps to visualize results of the SVM. These examples are all applied to SVM, with the exception of Poulet’s method¹⁰, which is also applicable to other classifiers such as decision trees or regression lines. The discussed approaches apply to multi-dimensional data, but are mostly specific to one type of classification model. What is needed is a uniform approach applicable to any classification method.

As we use scatterplots to visualize the data, we propose to visualize the multi-dimensional decision boundary in the series of scatterplots. However, the projection to 2D of the multi-dimensional decision boundary onto these plots would result in a loss of separating information. Therefore, we propose to use an approximation of the decision boundary, using Voronoi tessellation.

In short, a Voronoi diagram can be described as follows. Given a set of points (referred to as nodes), a Voronoi diagram is a partition of the space into regions, within which all points are closer to some particular node than to any other node (see Figure 2). Two Voronoi regions that share a boundary are called Voronoi neighbours. We apply the Voronoi diagram to a combination of two dimensions used by the classifier. All the data objects are used as nodes to make the Voronoi diagram. The boundaries of the Voronoi regions corresponding to neighbours belonging to different classes, according to the labels assigned by a classifier, form the decision boundary. Such a representation of class separation for two given features is a piece-wise linear approximation of the actual decision boundary as imposed by the multi-dimensional classifier (see Figure 2). The distances between the actual object and the decision boundary are not preserved, but class membership is. To indicate the actual distances to the decision boundary in the multi-dimensional space, the posterior probabilities for each data element could be calculated and visualized. To simplify the cognitive task of recognizing on which side of the decision boundary a data element is located, we not only show the decision boundary as a line, but also fill the regions corresponding to one of the classes with

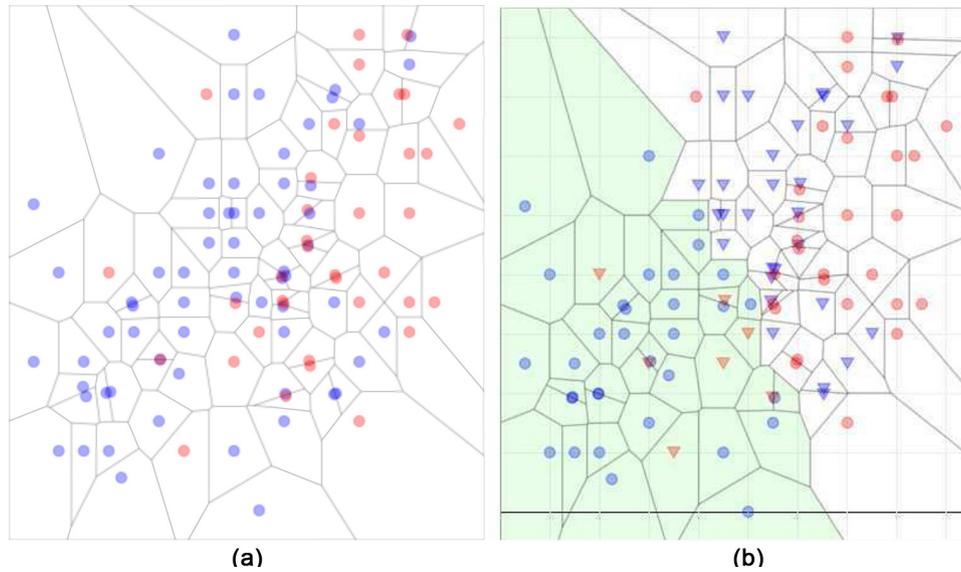


Figure 2. Visualization of a 2D dataset with two classes. (a) The scatterplot represents two numerical features and the Voronoi diagram of this space. (b) The scatterplot represents two numerical features, the Voronoi diagram of this space and the results of the nearest mean classifier for this dataset. For visual clarity, the classes are assigned different colours. Further, the misclassified data elements are assigned triangles and the correctly classified elements are assigned circles. The approximated decision boundary follows the Voronoi cells’ boundaries, which separate elements belonging to different classes as imposed by the multi-dimensional classifier. The regions belonging to one class are coloured green, whereas the regions belonging to the other class are not coloured.

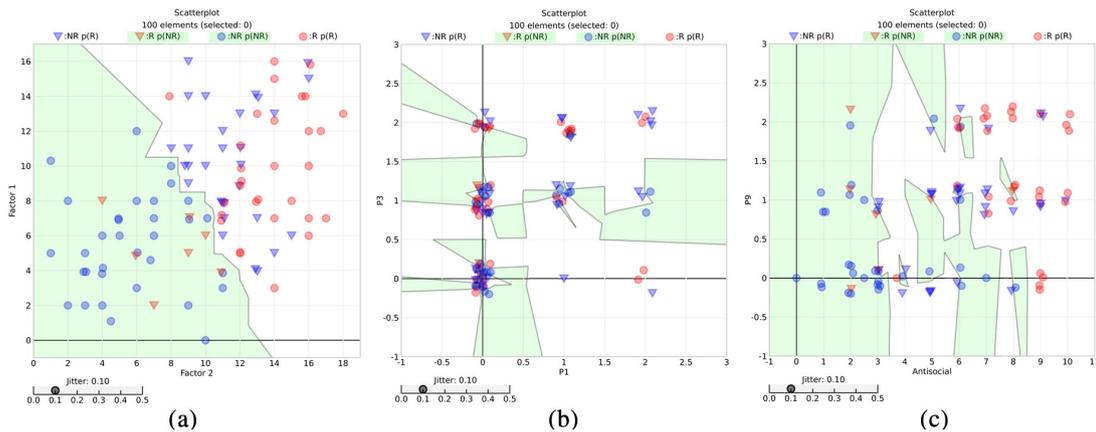


Figure 3. Visualization of the three data types using scatterplots and the Voronoi-based decision boundary: (a) numerical data on the x- and y-axes, (b) ordinal data on both axes, (c) x-axis with a numerical variable, y-axis with an ordinal variable. As indicated by our experts, (b) and (c) are not good representations for understanding the relation between the data and the classifier. We will consider alternatives in Figure 4.

a colour. In this way the expert immediately knows to which class an element has been assigned by the classifier. Figure 3(a) shows our implementation of the two variables’ scatterplot with the Voronoi-based approximation of the multi-dimensional decision boundary.

Categorical data

Data visualization. Categorical data can be either ordinal, in which the attributes are ranked categories (e.g. age group), or nominal, in which the categories do not have

an inherent ordering (e.g. sex). For the ordinal data an ordering can always be imposed, either randomly or according to user preferences. Without loss of generality we focus only on the visualization of the ordinal data.

There are several approaches to visualizing ordinal data. As we have chosen to visualize numerical data using series of scatterplots, we aim to visualize the ordinal data in a consistent way. During an initial evaluation of the system, we concluded that it is undesirable to visualize ordinal data and the model using scatterplots. Because of the occlusion of the individual data elements, the

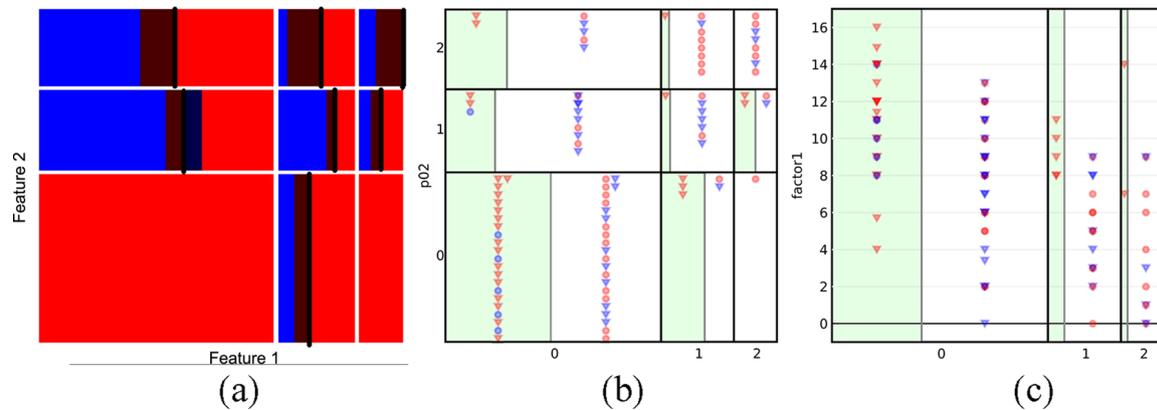


Figure 4. The visualization of the data and the decision boundary for different data types. (a) An example of the mosaic plot for ordinal data, as used at first. Each combination of the ordinal value pairs gives rise to one large rectangle. The sub-rectangles from left to right indicate the proportion of TNs, FPs, TPs and FNs belonging to that block. (b) After initial evaluation the plot has been modified to support the intrinsic characteristics of the risk assessment application, as indicated by the experts. A boundary is drawn to divide the sub-rectangles according to the number of individual elements being located at the different sides of the decision boundary. The data elements belonging to each of the sub-rectangles are visualized and arranged in a structure, but their order in this structure is random; (c) shows the modified mosaic-scatterplot for combined attributes. The ordinal variable is plotted on the x-axis and the numerical variable on the y-axis. The width of the sub-rectangles for the categories is proportional to the number of individual elements belonging to that category. The decision boundary is plotted in each of the sub-rectangles and placed according to the number of elements classified to each class. The individual elements are plotted within the sub-rectangle according to their numerical value. The ordering of the individual elements on the y-axis is therefore of importance.

Voronoi-based decision boundary does not provide insight into how the classifier separates the data for those ordinal features.

Therefore, we use mosaic plots instead, which are graphical displays devoted to examining the relationship among two or more categorical variables. First proposed by Hartigan et al.²³, they are now used in visualization toolkits such as Mondriaan²⁴ and Manet²⁵. Based on an analogy to the scatterplots we visualize two attributes at a time, where ranked categories are assigned to the vertical and horizontal axes. This results in a series of plots for all combinations of categories.

In a mosaic plot, each possible combination of categories is a cell that is represented by a rectangle proportional in area to the number of data elements in it. A mosaic plot shows quickly whether particular combinations of cases dominate the dataset. The appearance of a mosaic plot depends on the order in which the categorical variables are drawn. The categories of the first variable divide the horizontal axis into columns whose widths are proportional to their numbers. The categories of the second variable divide up the vertical axis similarly. To simplify, we make an assumption that the two variables are independent. Therefore, the separating gaps between all levels line up²⁶; otherwise, this would yield a visualization that would not be consistent with scatterplots for numerical data.

Model visualization. For each cell of the mosaic plot, we want to present information about the classification results. To make the visualization of the classification

results consistent with the scatterplot, we first divide each cell into parts according to the classifier. Therefore, in each cell a division occurs, equivalent to a decision boundary in the scatterplot, corresponding to the number of elements present on both sides of the decision boundary. Further, we show, as in a scatterplot, the blocks containing misclassified elements on both sides of the decision boundary. At first we used the following methodology. In each block we visualized the proportion of true positives (TPs), true negatives (TNs), false positives (FPs) and false negatives (FNs) according to the labels assigned by the classifier. We first assigned a colour to each class, according to the original labels. Parts of the block containing correctly classified elements were visualized according to the original class colour. Parts containing misclassified elements were visualized in a darker tint of the original class colour. In this way we made an analogy to different shapes of the visualized points in the scatterplots. Figure 4(a) shows such a mosaic plot for two ordinal features.

From the feedback we collected during the initial evaluation (described in Evaluation section) it became apparent that the most important element of decision boundary visualization is the relation of the individual data elements to the decision boundary, as described in Problem description section. The mosaic plot we originally proposed does not correspond to this. Therefore, we propose a modified mosaic plot that plots the original elements in the rectangles representing all the categories' combinations of two ordinal features. As there is

no natural order of the data within a sub-rectangle, the individual elements within each sub-rectangle are organized in a fixed structure, with random ordering. Data points are assigned a colour according to the ground truth class membership, and a shape according to the class membership as assigned by the classifier. As we are primarily interested in the correctly and misclassified samples in risk assessment, we assign distinctive shapes to correctly predicted data points, and a different shape to wrongly predicted data points. Originally, we divided each of the mosaic plot sub-rectangles in four regions. As the individual elements can now indicate the class membership and the result of the classifier, we no longer require that. Therefore, we show only the proportion of the elements on different sides of the decision boundary. We also fill the regions corresponding to one of the classes, as predicted by the classifier, with a colour. This is consistent with the visualization of the decision boundary for the scatterplot.

Combined data types

We use separate techniques to visualize 2D combinations of numerical attributes and ordinal attributes. However, it would be desirable to visualize the relation between a numerical and an ordinal attribute. As indicated by the experts, the visualizations of combined attributes were not clear enough using scatterplots (see Figure 3). In order to use a visualization of the data and the model consistent with the scatterplot and the mosaic plot, we combine the mosaic plot visualization of the ordinal attribute on one axis with the numerical attribute visualized on the orthogonal axis. In fact, we use the same modified mosaic plot that we used for two ordinal variables, together with the individual data elements and the model visualized in each of the sub-rectangles. In the mosaic plots for combined types of data, however, the order of the individual elements visualized on the numerical axis does have a meaning. There are no more sub-rectangles per category of each combination of two variables, but only sub-rectangles for each category for one ordinal feature. On the other axis, which represents the numerical variable, the individual elements are plotted as in the scatterplot, according to the scale (see Figure 4(c)). The individual elements are assigned a colour and a shape in the same manner as in the scatterplot and the mosaic plot. The decision boundary is plotted in the sub-rectangles on the ordinal axis, in a similar manner to the mosaic plot. The regions corresponding to the negative class, as assigned by the classifier, are filled with a colour.

Matrix overview

Obviously the visualization of multi-dimensional data by means of 2D plots will result in many scatterplots that are neither easily nor clearly arrangeable. To deal with

the large number of visualized plots we use the widely known static scatterplot matrix and order the dimensions based on their importance, as indicated by the experts. Alternatively, the dimensions could be ordered according to some interesting criterion. Ankerst et al.²⁷ propose placing the similar dimensions next to each other, after deriving the similarity metrics. Elmqvist et al.¹⁵ follow this method, defining the similarity measure as an absolute correlation of pair-wise dimensions to order the rows and the dissimilarity measure as an inverse of the absolute correlation to order the columns. Tatu et al.²⁸ developed methods to find the best scatterplots showing the separating classes.

Those techniques would be indispensable when the number of dimensions increases and it is impossible to show all the combinations of dimensions in the matrix. Then only the most 'interesting' dimensions could be displayed at first, where 'interesting' is defined by one of the above-mentioned methods. Alternatively, one of many feature selection algorithms could be applied to identify the features most relevant in the model training process. We could also develop a criterion based on the model visualization for the numerical data, where good views should show good class separation. The minimal overlap of classes would then denote the importance of dimensions. The order of rows and columns could also be changed manually by the user in an interactive way, as used in *Polaris*¹⁹, *XmdvTool*¹⁸ and *ScatterDice*¹⁵.

To use the space of the scatterplot matrix more efficiently we follow the example of Heer et al.²⁹. We remove the meaningless plots on the diagonal, which are the plots of the variable with itself, and replace them with the names of variables plotted in the matrix. The matrix overview is shown in Figure 5.

Performance visualization

General performance curves, such as precision and recall graphs or ROC curves, capture the ranking performance of the binary classifier as its discrimination threshold is varied. The ROC curve, often used in medicine, visualizes the trade-offs between hit rate and false alarm rate³⁰. The hit rate is the rate of positive instances classified correctly and the false alarm rate is the negative instances wrongly classified as positives. In information retrieval the precision and recall curve is often used. It depicts the trade-off between the proportion of relevant instances in the result and the proportion of relevant instances with respect to the whole dataset. Viewed in a different way, the precision is the rate of correctly classified positive instances and the recall is the rate of correctly classified positive instances of all the instances classified as positive. In risk assessment applications the balance between the misclassification in each class is of interest. Therefore, the performance curve should depict the trade-off between the classifier's errors for both classes.

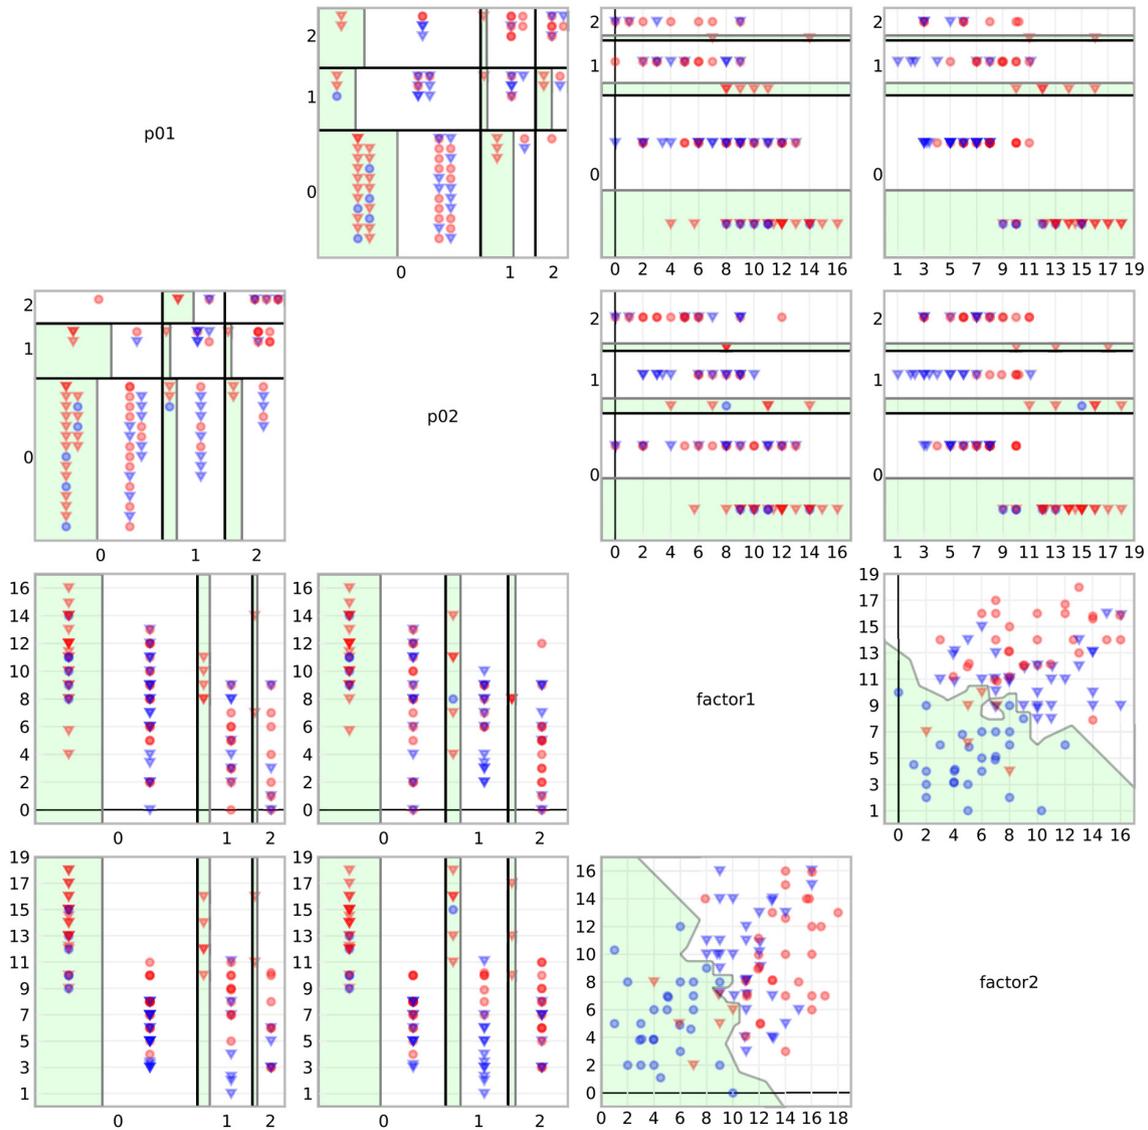


Figure 5. Example of a matrix overview showing four dimensions of a dataset. The overview matrix contains scatterplots for numerical attributes and adjusted mosaic plots for two ordinal attributes and mixed attribute types.

The curve, as we use it, represents the trade-off between the false positives and false negatives. If FP is the number of incorrectly classified negative data items and NP the total number of positive examples, then the false positive rate, FPr, of a classifier is:

$$FPr = \frac{FP}{NP} \tag{1}$$

The false negative rate, FNr, of a classifier is:

$$FNr = \frac{FN}{NN} \tag{2}$$

where FN is the incorrectly classified positive items and NN is the total number of negative items. On the performance graph FNr is plotted on the y-axis and FPr is plotted on the x-axis. These statistics vary with a threshold on the classifier’s continuous outputs.

We note several points on the curve. The upper left point (0,1) represents the classifier that misclassifies all positive data elements; the lower right point (1,0) represents the classifier that misclassifies all negative data elements. The point (0,0) represents perfect classification and the line $y = x$ represents the random guess. Such a performance curve allows visual comparison of error trade-offs of a set of classifiers. The trade-off of the current classifier is visualized by means of an operating point.

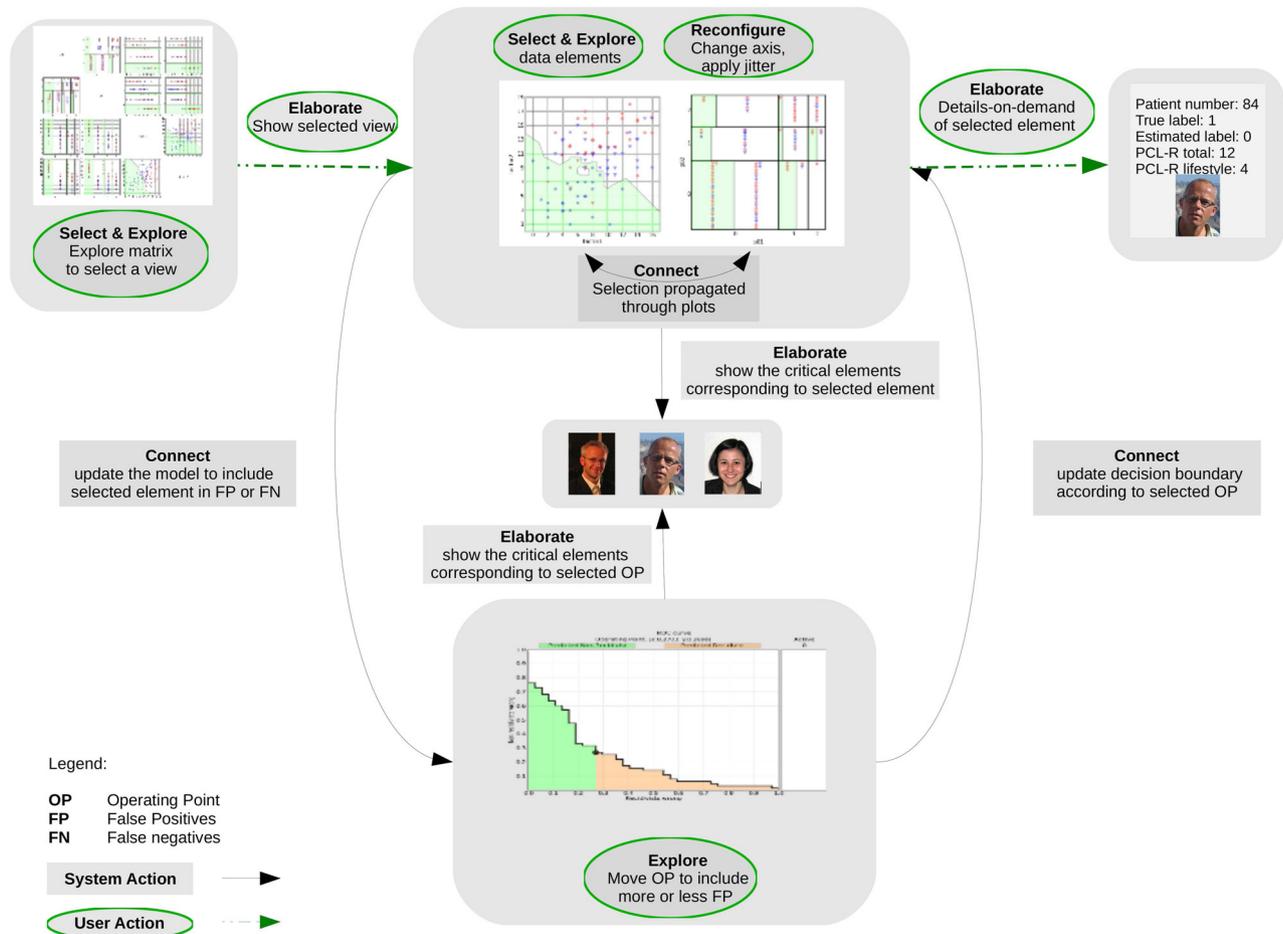


Figure 6. Our framework for a visual risk assessment system, elaborating upon Figure 1, with the proposed visualization solutions and interaction techniques. Both system and user actions are indicated.

Images for inspection of critical data elements

An expert and his knowledge are an integral part of the visual analytics-based risk assessment process. To take advantage of his knowledge in an effective way, we suggest triggering the associations he or she has with the data elements. Purely from the visualization of the data through scatterplots and mosaic plots, or any other graphical displays, an expert is not always able to recall the meaning of a particular data item. To assign a meaning to the data elements, we propose to use images that refer to the data in some way. For the datasets consisting of features of patients, the photos of the patients could be used. For the database consisting of the symptoms of a certain liver disorder, the radiological images could be used. Assigning images to the data elements might have a ‘trigger’ effect on an expert’s mind, stirring conscious and unconscious memories about certain data elements. By using images we refer to both episodic and semantic memory of the experts³¹. In particular, for characteristic data elements the associations can be strong⁷.

We distinguish two cases in which the images play a role in the exploration of data and classification results. First, when the images are not used as data in the classification, they are used to support the visual exploration and have merely an associative role. When viewing the images in this case one should focus not on the physical properties of the object on the images, but on the associations that these objects invoke. Second, measured characteristics of the images might also be a part of a feature vector used to train the classifier.

In the context of risk assessment we would like to focus the attention of an expert on the most important data elements in the cost selection process. We consider the most important elements as those which are assigned a different label by the classifier when changing the classifier’s trade-off. We term these ‘critical elements’ and propose to visualize them instantly whenever the classification model changes. Figure 6 shows how we visualize the images for the critical data elements, occurring in between adjacent operating points. In forensic psychiatry, for example, the expert examines his patients and therefore possesses extensive knowledge about them. Adding

the images of the patients to the data visualization could be helpful in the exploration of the dataset and the classifier. As the data elements are visualized in relation to the classifier, and the expert can associate them with actual patients, he knows exactly how particular patients relate to the classification model. He can also view the group of patients that is in the direct neighbourhood of the recognized patient, and therefore associate other patients with him. Therefore, the image functions as an orientation point for the general group of patients, which might be unknown to the expert. One might argue that by showing the images we introduce a bias, as we allow free associations of the expert that cannot be justified and are definitely not objective. However, the expert makes unconscious associations all the time. Therefore, we can take advantage of it and structure steering of the exploration process in this way. Moreover, we put constraints on expert actions concerning the change of the classification model based on image associations. This is described in Interaction techniques section.

Interactive risk assessment

To support risk assessment, the techniques we described in the previous sections need to be combined into a highly interactive visual analytics framework. In this section we describe how we couple the elements in order to design an interactive tool for data and classifier exploration, including in particular the trade-off selection in the context of risk assessment.

Interaction techniques

Through the interactive data exploration the expert can gain insight to not only the data but also the classification process. Therefore, several interaction techniques are integrated into our method, forming a minimum set of interaction techniques for interactive risk assessment. We choose the techniques from the seven categories of interaction techniques proposed in the study by Yi et al.³². They propose a well-structured way to handle interactivity based on the user's functional needs. We distinguish between the interaction techniques that are strictly user based and interactions that are triggered by the user but result in predefined system actions. Figure 6 shows the possible interaction techniques in the system.

Select. The select interaction technique aims to make the elements of interest visually distinctive. An expert can highlight these elements in the scatterplot or in a rectangle in the mosaic plot, resulting in a colour change of the selected element. This allows the user to keep track of elements of interest. The expert can also deselect the element if he or she is no longer interested in it.

Explore. Explore interaction techniques enable users to examine a different subset of data cases. First, all the plots in the matrix overview can be explored in the main window. Second, the direct-walk technique allows users to smoothly move the viewing focus from one position in the information structure to another by 'a series of mouse movements'. This occurs when moving the operating point on the performance curve or by including certain data elements in a specific class.

Connect. Connect is used to highlight relationships between data items and show hidden relations between relevant items. In our system we connect all the visual displays. First of all, the visualizations of data elements are connected. All the elements selected on one scatterplot are instantly selected on all the other displayed plots. Thus, it allows seeing the relations between the decision boundary and selected elements for all the dimensions visualized. Moreover, we connect the plots with the performance curve, illustrated in Figure 7. This connection is twofold.

On the one hand, we use an interactive operating point on the performance curve connected with the visualizations of the model on the plots. An operating point on the performance curve represents the current classifier. Without knowledge of the domain this is often chosen by optimizing the equal error rate (i.e. FP = FN). In the context of risk assessment, the expert is not always interested in the lowest combined classification error. Therefore, we allow the user to change the position of the operating point on the performance curve. In this way the user can choose the suitable trade-off for his application. The resulting FP rate and FN rate initiate an update of the classifier. As the visual displays are connected, the expert instantly observes what effect the change in trade-off has on the classifier for all data plots. Moving the operating point also triggers the display of the images of the critical data elements. Those elements are, moreover, highlighted on the scatterplots and the mosaic plots.

On the other hand, we allow the user to update the operating point, and therefore the classification model, through selection of data elements in the plots. The expert is allowed to choose a data element and requires the classifier to assign it to a different class. The model is then instantly updated, together with the corresponding trade-off on the performance curve. The images representing critical elements are automatically displayed and critical data elements are highlighted on all the plots. It is important to note that an expert can only change the decision boundary as a whole by selecting a different operating point. So forcing, for example, a false negative to become a positive would also introduce new false positives. In such a way the expert can explore whether a change of the model for the particular data element results in a

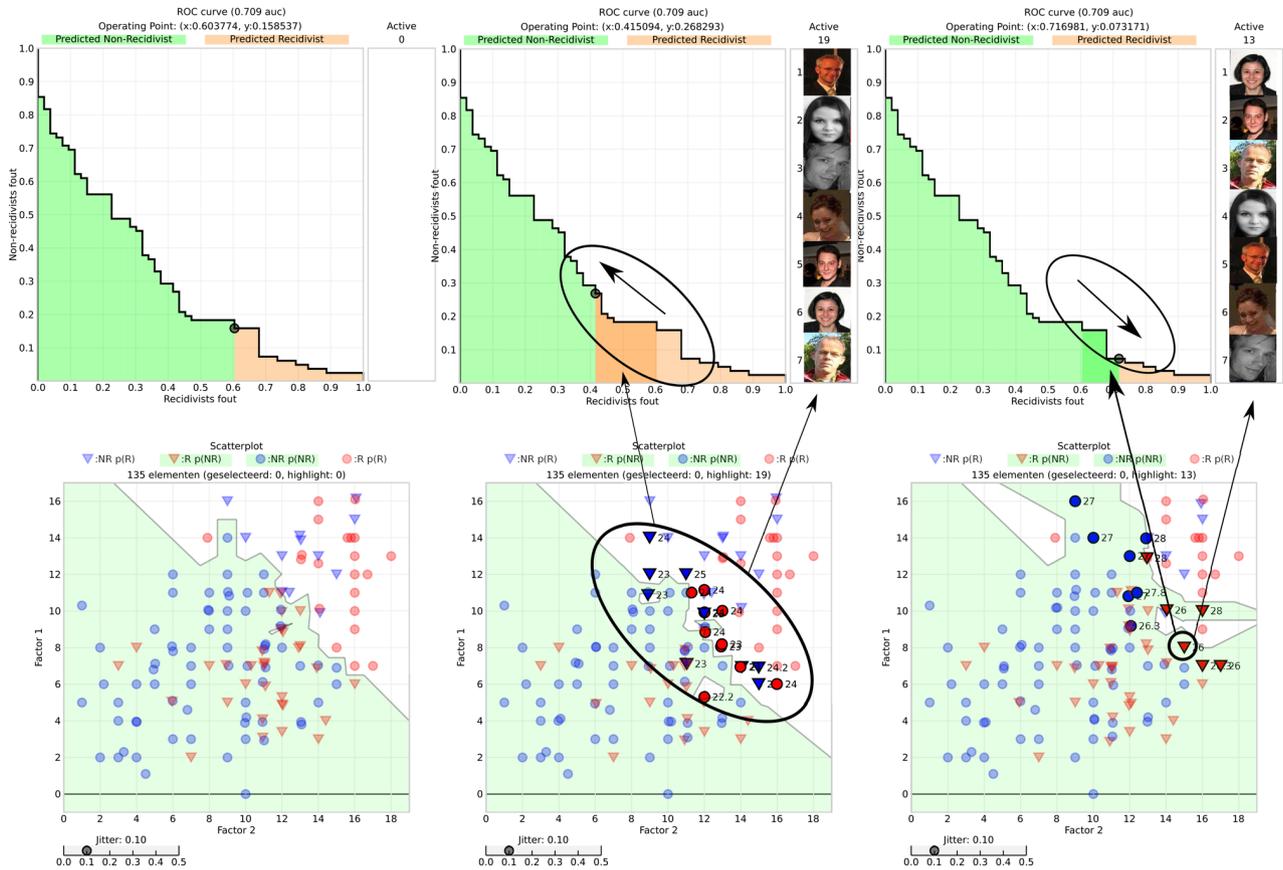


Figure 7. Example of the connect interaction. The figures at the left show the initial state of the system, with the initial operating point on the performance curve and corresponding visualization of the model for the scatterplot. The figures in the middle show the state of the system when an expert manipulates the operating point to include fewer true negatives. The figures on the right represent the state of the system when an expert includes a data element into false positives. The updated model is visualized in the scatterplot (using a Voronoi-based approximation of the multi-dimensional decision boundary), together with the critical elements highlighted in the scatterplot, and represented as photos. On the ROC curve the updated operating point is visible, as well as the rate of change on the x-axis.

serious cost change and which other data elements would be affected, without compromising the classifier.

Elaborate. Elaborate allows the user to see more details, and hence adjusting the level of abstraction of the data representation. We use two techniques in this category.

Details-on-demand allows users to examine all the information of a particular data point. The values of all attributes of that data point are displayed together with the true and estimated label. Moreover, the critical elements are displayed when the user aims at placing such an element on another side of the decision boundary. This is a system action that is always triggered by the explore interaction technique. Another technique that triggers details-on-demand is exploration of the operating point on the performance curve, which results in the visualization of images representing the critical data elements.

Zooming in/out allows users to simply change the scale of the representation so that the user can see a detailed

view of a part of the dataset. First, the plots in the matrix overview can be viewed in detail on a resized display. Second, within each plot one can zoom in to view a selected part of a plot in detail.

Reconfigure. The reconfigure interaction technique stands for changing the spatial arrangement of data representations. This gives experts different views on the data. We use instantiations of two interaction techniques belonging to this category, also implemented in, for example, Spotfire¹⁷. First, we allow the user to change the attributes presented on the x- and y-axes in the plots. Changing the attributes assigned to the x- and y-axes changes the set of attributes or variables to be examined among the entire dataset, so it eventually changes relationships between data items and provides different perspectives. Second, we provide the user with a jitter technique. Many data elements might have the same values for the visualized numerical features, resulting in

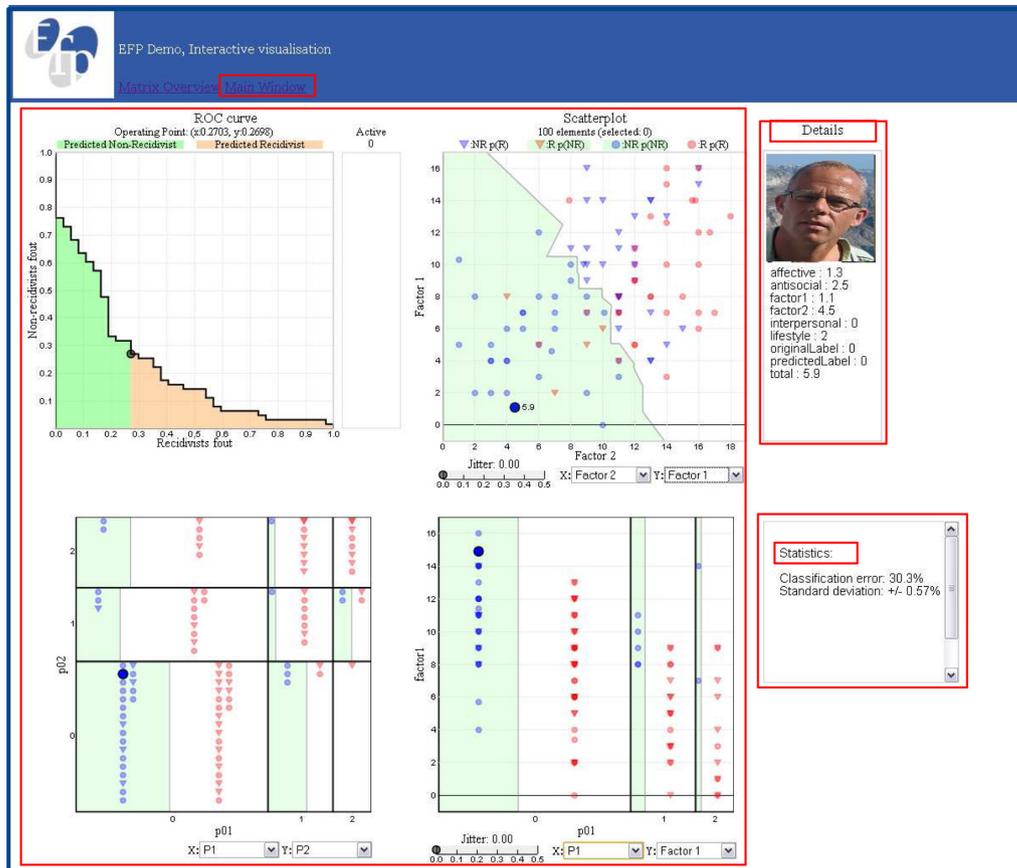


Figure 8. Screenshot of the system.

overlapping items on the scatterplots. To avoid occlusion, jitter can be applied, where the position of each item is randomly shifted by a small spatial increment. In this way the occluded data items are uncovered and the user can see the individual data elements and their actual densities.

Visual design

The system is implemented using ProtoVis³³. ProtoVis is a free and open-source, JavaScript and scalable vector graphic-based toolkit for web-native visualizations. The classifiers are precomputed in Matlab. The system we propose consists of several visual components. The implementation of the system is shown in Figures 5 and 8. The main components of the system are:

- *Overview window*: serves as an overview tool. Miniature versions of the individual plots are shown as the cells of the matrix. The plots are visualized together with the decision boundary and are connected.
- *Main window*: displays three plots currently chosen by the expert using axis selection. The performance curve and critical elements are also part of the main window.

- *Statistics window*: displays the statistical information on the classification model for the current interactive operating point, such as the relative errors, or amount of TP and TN. This is part of the main window.
- *Details window*: includes the information on a particular interactively selected data element, including the image associated with this element. This is part of the main window.

A case study of interactive risk assessment in forensic psychiatry

As indicated earlier, one of the practical fields where interactive risk assessment could be of great help is forensic psychiatry. In the Dutch legal system forensic experts often have to advise the court whether a sentenced, mentally disordered criminal should be released back to society or should be treated in a closed mental institution. Most importantly, the forensic expert has to assess the risk of re-offending of such a patient. Obviously, from the perspective of the safety of the society, releasing a potential recidivist has greater consequences than keeping a potential non-recidivist imprisoned. However, the release of the patient to society is an integral part of the treatment. Therefore, informed choices have to be made

Table 1. Performance of the selected classifiers for PCL-R dataset obtained with 10-fold cross-validation, with three repeats

Classifier	Cross-validation error (%)	Standard deviation
Nearest mean	30.3	± 0.5
Logistic regression	31.3	± 1.5
Fisher	31.0	± 1.0
Linear support vector machine	31.3	± 1.0
Decision tree	33.6	± 4.04

to minimize the risk of making severe mistakes, but the costs of mistakes in disadvantage to the patient should also be avoided.

Case study

As a case study we use a dataset of forensic psychiatric patients in the Netherlands provided by the Expertise Center for Forensic Psychiatry (EFP). The dataset consists of 100 male offenders who, at the time of the alleged crime, suffered from a mental disorder and received what is called a 'disposal to be involuntarily admitted to a forensic psychiatry hospital on behalf of the state' (TBS order). The TBS orders have been terminated based exclusively on the professional expertise of the clinicians. Each patient is assigned a class label indicating whether he has been convicted for a new crime after his TBS order has been terminated. Of the 100 defendants, 37 were convicted again, whereas 63 are non-recidivists. The 20 ordinal features are the scores of the PCL-R (Psychopathic Checklist-Revised) test. The seven numerical features are summed combinations of the PCL-R outcomes. Patients were retrospectively scored with these risk assessment measures and recidivism data was retrieved from the documentation of the Dutch Ministry of Justice. To visualize the patients we assigned a red colour to recidivists and blue to non-recidivists. As we cannot publish the real photos for this dataset for reasons of privacy, we use images of the authors and their colleagues as substitutes in this paper. To illustrate the usage of the system we choose to focus on the results of one classifier.

Table 1 presents the 10-fold cross-validation error, over three repeats for selected classifiers trained on 27 features. The difference between them seems statistically insignificant. We choose to visualize the nearest mean classifier as it performs well and has a low standard deviation.

Usage scenario

We now demonstrate how our system can be used for visual exploration of the nearest mean classifier trained on the criminal patients' dataset. We describe a basic

scenario. An expert, Analyst X, gets a newly developed classification model. Analyst X is asked to explore the model applied to patient data with which she is familiar.

First, the dataset is loaded and the model applied to it. All possible combinations of dimensions are visualized together with the decision boundary for the current classifier output in the matrix in the overview window. The classification error and the performance curve with the current operating point are displayed. Figure 8 shows a screenshot of the system that is now available for the expert.

Analyst X first looks at the overview window to get an overall feeling for the visualization of the dataset. She moves to the elaborate view (main window), where three currently selected plots from the matrix (overview window) are displayed.

First, she explores the connected plots. She selects two interesting-looking data elements on the scatterplot. Both elements are patients who did not re-offend, but the model classifies them as recidivists. On all the visible plots in the main window the patients are highlighted. They are located in the areas representing high values of the attributes. She seeks an explanation as to why those patients had such high attribute values but did not re-offend. Analyst X elaborates on these two patients and asks for their details. When she looks at the photos she realizes that one of the patients had very good re-socialization conditions and a strong network of people taking care of him, which could prevent him from re-offending. For the other patient, however, she does not seem to have an explanation.

Then she starts exploring the performance curve. She moves along the curve, varying the operating point. The operating point is connected to all the plots, so she observes the instant changes in the model visualization in the main window. She is also observing the pop-up photos of critical patients. While she moves the operating point to include fewer false negatives, as she is not satisfied with so many recidivists being misclassified, she notices among the images a patient who is familiar to her. She asks for the details. The scores indicating the violent behaviour of the patient are, indeed, not very high, indicating an average risk of re-offence. However, she recognizes Patient X and realizes that he is very unreliable at taking his medication. Therefore, the risk of re-offending when he leaves the clinic, with no one to control his medication use, would increase. She requires the model to include that patient in the true positives. She explores the result on the connected plots. Both the model and the performance curve are updated when the operating point is updated. Based on this limited exploration she finds the updated model suitable for use in her clinic to help in making more informed decisions.

A presentation of the system in use, clarifying the exploration process, can be found in the accompanying video material (<http://staff.science.uva.nl/~migut/publications.html>).

Evaluation

Evaluation of the visual analytics tools is a challenging task. It is different from evaluating either visualization-only systems or human–computer interaction-based system³⁴. The difficulty lies in the primary functionality of the system, which is exploratory analysis³⁵. The decision-making process is complex and the system has a comprehensive set of options. Only the experts can thoroughly verify the usability and usefulness of the system. However, this would lead to domain-specific improvements. For the purpose of this paper we limited ourselves merely to initial feedback from the experts and how it influenced the choice of methods and the design. We limited the structure of the evaluation to several crucial aspects:

- overall impression of the usefulness of the system;
- familiarity with the visualization techniques and comprehension of those; and
- comprehension of the interaction structure.

An evaluation took place in a forensic psychiatry clinic in the Netherlands. The informal evaluation meeting with nine employees took the form of a cognitive walk-through. It consisted of several parts. First, all the separate elements of the system were explained, focusing on the exploration of the dataset, exploration of the classification model, and exploration of the costs of classification. Second, an interactive demo was given. The participating experts could explore the system as a group by offering suggestions on what they wanted to explore. Third, a discussion took place and all the participants completed a questionnaire. The questionnaire had several questions, mainly referring to separate components of the system, their interactive combination and their usefulness in the forensic psychiatry setting. Through the questionnaire the experts assessed each of the visualization components separately, which triggered many questions about the data quality and the interpretation of the visual representation of the data.

First impressions of the experts were positive. The users could directly acknowledge the usefulness of the interactive visualizations and analysis for research purposes in a forensic psychiatry setting. The experts were not used to any visual representation of the data, and they immediately saw the potential of discovering new patterns, which was not possible with raw data. Special attention was given to the use of images of the patients. The expert, who actually had seen the system before without images and tried to make sense of simple visualizations, indicated clearly the additional value of using the images. That is also in agreement with their daily practice where each dossier of a patient includes a photograph. Moreover, the system allows comparing different patients by associating the photos. This was acknowledged to be a great improvement and of high value. Further, from the initial feedback

it became clear that the use of 2D plots of the ordinal data (see Figure 4(a)), in their original design, did not contribute to the understanding of the data. On the contrary, it was considered by all users as confusing. Even when the concept of a mosaic plot was explained and understood, the supportive value of them could not be confirmed. Therefore, we discussed with the users how this could be improved, and, as a result, we added the individual elements to the mosaic plot and combined plot visualization.

Conclusion

This paper proposes an interactive approach to risk assessment. The main idea is to provide an expert with a framework tightly integrating interactive visual exploration with machine learning and taking into account the demanding requirements of risk assessment applications. In order to support an expert in the visual exploration of the dataset and the classification model we propose to couple visualizations of numerical and ordinal data, classifier visualization, performance visualization and associative images representing meaningfully the data elements. In particular, we focus on the cost selection of the classification model. We have shown how the existing techniques can be tightly coupled in a structured and highly interactive way. We have applied the proposed framework to the forensic psychiatry domain, and verified chosen techniques based only on generic feedback. In this way, the framework remains general and applicable to other domains. The system could, for example, successfully be applied to support decision-making in domains such as diagnosing liver disorder in human health or detecting dangerous behaviour in security management.

Acknowledgments

This research is supported by the Expertise Center for Forensic Psychiatry, The Netherlands. The authors would like to thank Jan van Gemert for the re-implementation of the original system in Protovis.

Funding

This research is financially supported by the Expertise Center for Forensic Psychiatry, The Netherlands.

Declaration of conflicting interests

The authors declare that they do not have any conflicts of interest.

References

1. Keim D, Mansmann F, Schneidewind J, et al. Visual analytics: scope and challenges. *Lect Notes Comput Sci* 2008; 4404: 76–90.
2. Thomas JJ and Cook KA. *Illuminating the path: the research and development agenda for visual analytics*. Los Alamitos, CA: IEEE CS Press, 2005.

3. Fuchs R, Waser J and Gröller ME. Visual human + machine learning. *IEEE Trans Visual Comput Graphics* 2009; 15(6): 1327–1334.
4. Bendix F, Kosara R and Hauser H. *Parallel sets: visual analysis of categorical data*. In: *INFOVIS '05: proceedings of the IEEE symposium on information visualization*, Minneapolis, MN, 23–25 October 2005.
5. Collins C and Carpendale S. VisLink: revealing relationships amongst visualizations. *IEEE Trans Visual Comput Graphics* 2007; 13(6): 1192–1199.
6. Duda RO, Hart PE and Stork DG. *Pattern classification*. New York: Wiley-Interscience Publication, 2000.
7. Ware C. *Visual thinking for design*. Burlington, MA: Morgan Kaufmann Publishers, 2008.
8. Yu C, Zhong Y, Smith T, et al. Visual data mining of multimedia data for social and behavioral studies. *Inform Visual* 2009; 8(1): 5–70.
9. Ankerst M, Kao A, Tjoelker R, et al. DataJewel: integrating visualization with temporal data mining. *Lect Notes Comput Sci* 4404: 312–330.
10. Poulet F. Towards effective visual data mining with cooperative tools. In *Visual Data Mining - Theory, Techniques and Tools for Visual Analytics*, Simeon Simoff, Mihael Boehlen, Arturas Mazeikas (eds.), Vol 4404, pp. 389–406, Lecture Notes in Computer Science, Springer-Verlag, 2008.
11. Ankerst M, Ester M and Kriegel HP. Towards an effective cooperation of the user and the computer for classification. In *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '00)*. ACM, New York, NY, USA, pp. 179–188.
12. Van Den Elzen S and Van Wijk JJ. Baobab view: interactive construction and analysis of decision trees. In: *Proceedings of the IEEE symposium on visual analytics science and technology (IEEE VAST)*, Providence, RI, 23–28 October 2011.
13. Migut MA and Worrying M. Visual exploration of classification models for risk assessment. In: *Proceedings of the IEEE symposium on visual analytics science and technology (IEEE VAST)*, Salt Lake City, UT, 24–28 October 2010, pp. 11–18.
14. Keim DA. Information visualization and visual data mining. *IEEE Trans Visual Comput Graphics* 2002; 8(1): 1–8.
15. Elmqvist N, Dragicevic P and Fekete JD. Rolling the dice: multidimensional visual exploration using scatterplot matrix navigation. *IEEE Trans Visual Comput Graphics* 2008; 14: 1141–1148.
16. Cleveland WA and McGill ME. *Dynamic graphics for statistics*. New York: CRC Press Inc., 1988.
17. Spotfire Inc. <http://www.spotfire.com> (2007). Accessed on 23rd of January 2012.
18. Ward MO. XmdvTool: integrating multiple methods for visualizing multivariate data. In: *VIS '94: proceedings of the conference on Visualization*, Washington, DC, October 1994, pp. 326–333. Los Alamitos, CA: IEEE Computer Society Press.
19. Stolte C, Tang D and Hanrahan P. Polaris: a system for query, analysis, and visualization of multidimensional relational databases. *IEEE Trans Visual Comput Graphics* 2002; 8(1): 52–65.
20. Swayne DF, Lang DT, Buja A, et al. GGobi: evolving from XGobi into an extensible framework for interactive data visualization. *Comput Stat Data Anal* 2003; 43(4): 423–444.
21. Caragea, D., Cook, D., and Honavar, V. (2001). Gaining Insights into Support Vector Machine Classifiers Using Projection-Based Tour Methods. In: *Proceedings of the Conference on Knowledge Discovery and Data Mining (KDD 2001)*, August 26–29, San Francisco, CA, USA. pp. 251–256. ACM Press.
22. Hamel L. Visualization of support vector machines with unsupervised learning. In: *Proceedings of 2006 IEEE symposium on computational intelligence in bioinformatics and computational biology*, Toronto, Canada, 28–29 September 2006.
23. Hartigan JA and Kleiner B. Mosaic for contingency tables. In: *Computer science and statistics: proceedings of the 13th symposium on the interface* (ed. WF Eddy), 1981, pp. 268–273. New York: Springer.
24. Theus M and Urbanek S. *Interactive graphics for data analysis: principles and examples (computer science and data analysis)*. Boca Raton, FL: Chapman & Hall/CRC, 2008.
25. Unwin A, Hawkins G, Hofmann H, et al. Interactive graphics for data sets with missing values: MANET. *J Comput Graph Stat* 1996; 5(2): 113–122.
26. Theus M and Lauer SRW. Visualizing of loglinear models. *J Comput Graph Stat* 1999; 8(3): 396–412.
27. Ankerst M, Berchtold S and Keim DA. Similarity clustering of dimensions for an enhanced visualization of multidimensional data. In: *INFOVIS '98: proceedings of the 1998 IEEE symposium on information visualization*, 1998, pp. 52–60. Washington, DC: IEEE Computer Society. Atlantic City, NJ, USA, October 2009, pp. 59–66.
28. Heer J, Bostock M and Ogievetsky V. A tour through the visualization zoo. *Commun ACM* 2010; 53(6): 59–67.
29. Fawcett T. An introduction to ROC analysis. *Pattern Recogn Lett* 2006; 27(8): 861–874.
30. Tulving E. Episodic memory: from mind to brain. *Annu Rev Psychol* 2002; 53: 1–25.
31. Yi JS, Kang J, Stasko J, et al. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Trans Visual Comput Graphics* 2007; 13(6): 1224–1231.
32. Bostock M and Heer J. Protovis: a graphical toolkit for visualization. *IEEE Trans Visual Comput Graphics* 2009: 1121–1128.
33. Plaisant C. 2004. The challenge of information visualization evaluation. In *Proceedings of the working conference on Advanced visual interfaces (AVI '04)*. ACM, New York, NY, USA, 109–116.
34. Perer A and Shneiderman B. Integrating statistics and visualization for exploratory power: from long-term case studies to design guidelines. *IEEE Comput Graphics Appl* 2009; 29: 39–51.