

Visual Exploration of Classification Models for Risk Assessment

Malgorzata Migut*

Intelligent Systems Lab Amsterdam
University of Amsterdam

Marcel Worrington†

Intelligent Systems Lab Amsterdam
University of Amsterdam

ABSTRACT

In risk assessment applications well informed decisions are made based on huge amounts of multi-dimensional data. In many domains not only the risk of a wrong decision, but in particular the trade-off between the costs of possible decisions are of utmost importance. In this paper we describe a framework tightly integrating interactive visual exploration with machine learning to support the decision making process. The proposed approach uses a series of interactive 2D visualizations of numeric and ordinal data combined with visualization of classification models. These series of visual elements are further linked to the classifier's performance visualized using an interactive performance curve. An interactive decision point on the performance curve 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 context 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 and make more informed and reliable decisions. A case study on data from the Forensic Psychiatry domain reveals the usefulness of the suggested approach.

Keywords: Visual Analytics, Interactive Visual Exploration, Decision Boundary Visualization, Multi-dimensional Space, Classification

1 INTRODUCTION

Risk assessment is important in applications such as disaster management, security, medicine [1, 2] and forensics. 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 foresee the consequences of those decisions. The crucial part of the decision making process is determining the optimal balance between the cost of the different decisions.

An example domain is the field of Forensic Psychiatry where the risk of criminal behavior of patients with psychiatric disorders has to be predicted. In this domain 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, wrong prediction of re-offending has a completely different impact than wrong prediction of non re-offending, both for the society and for 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 sort of mistakes, diagnosing the healthy patients

and not diagnosing the ill patients, may be fatal.

Neither a fully automatic, nor a pure expert knowledge driven assessment is an 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 been indicated to be useful in various applications [1, 2, 3]. These techniques could help the domain expert gain insight in the problem and make well-informed decisions. This Visual Analytics approach, as formally described by Keim [1], integrating the strength of the expert and the strength of the machine is also promising for risk assessment.

In Visual Analytics, many different methods are used to visualize multi-dimensional datasets such as scatterplots, heatmaps, parallel coordinates and parallel sets [4]. The visualizations should relate to the domain specific understanding framework of the expert. Therefore, they should represent the data in the real attribute values and not some (non-linear) projections. Furthermore, the different types of data, including numeric 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 linked to each other revealing the relationships between them [5]. Moreover, to utilize the knowledge of an expert the data elements should be assigned a meaning. 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 field the exploration of radiology images of the liver when diagnosing a certain disorder allows 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 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 domain intrinsics, it is favorable to obtain a visual comprehension of what the classifier looks like in the multi-dimensional space. To enhance the understanding of the model, many visualization techniques are used to present the results of the classifier, such as ROC curves or Precision and Recall graphs [6]. However, visualization of the results alone does not give insight in the model.

One of the most informative characteristics of the classification model is the decision boundary. It can be easily visualized for two and three dimensional datasets [6]. 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 expert knowledge about data elements. Using the images representing the data elements, as mentioned earlier, would trigger immediate associations between the data and the model. Especially if visualizing the critical data elements responsible for the cost change of the classifiers output would contribute to awareness in cost selection.

*e-mail: M.A.Migut@uva.nl

†e-mail: M.Worrington@uva.nl

In this work we propose an interactive risk assessment framework. We interactively couple visualizations of different data types, decision boundary and classifier performance visualization and associative images representing the data to provide an expert with a possibility to visually explore 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 data elements. From there we show how to make the visualization of the performance curve more informative when linked with the visualizations of data. Following the section describing classifiers performance visualization, we propose an interactive risk assessment system and show its application in Forensic Psychiatry.

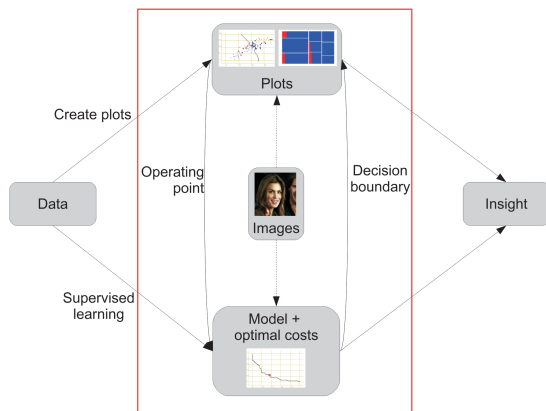


Figure 1: The risk assessment framework proposed in this paper following the general Visual Analytics framework proposed by Keim[1]. The elements within the rectangle indicate the focus of this paper. For privacy reasons we illustrate the usage of images in the framework with publicly available photographs of celebrities as an example, rather than the real data.

2 RELATED WORK

In recent years the idea of combining classification with interactive visualization has gained a lot of interest in the literature. A framework for Visual Data Mining was formally described by Keim [1]. It proposes a tightly coupled system, with system flows between the automated data analysis and visualization elements. It moreover employs user interaction to steer the visualization process and to actively participate in the classification process. Several of such tightly integrated approaches have been proposed aiming at different usage contexts. Building upon Keim’s framework, Yu et al. [7] propose a smooth interface between data mining and visualization for multimedia data in social and behavioral studies. All intermediate and final results of data mining, in terms of found patterns, are visualized allowing the user to obtain new insights and develop more hypotheses about the data. Ankerst proposes the DataJewel architecture [8] 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 [9, 10]. 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. A way of improving and analyzing a classifier is described in [3]. Starting from an initial hypothesis, created with linking and brushing, the user steers a heuristic search algorithm to look for alternative hypothesis generation.

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 intelligently be integrated to support risk assessment, by combining and structuring them in a highly interactive way, as illustrated in figure 1.

3 VISUALIZATION OF CLASSIFICATION MODELS

The main visualization components constituting our system are the visualizations of the data, the classifier and the performance. 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 motivate our choice of techniques.

3.1 Problem description

Let us formalize the problem. Assume a training dataset with n objects represented by feature vectors with numerical data in a p -dimensional space. Here we consider the two class problem (with a positive and negative class) and low dimensional spaces which are typical for detection problems in risk assessment. Feature vectors contain the variables of different data types. A classifier is trained on the dataset, resulting in a p -dimensional decision boundary. An object classified as positive is called a true positive if the actual value is also positive and is called 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 false negative if the actual value is positive.

One of our tasks is to visualize the model. Several requirements are imposed that are specific in a risk assessment context. First of all the expert understands the data in the real feature values and can therefore easily understand the model, only if related to the real feature values. The visualization of the data itself should be simple and should preferably employ well known techniques, which the expert is familiar with. To minimize the amount of visualized dimensions and therefore make the exploration potentially easier for the expert we could use one of many dimension reduction techniques. This is however not desirable for several reasons. Those methods have been designed to find the most interesting 2D planes and to create 2D views that best summarize a multi-dimensional dataset. However their notion of ‘interest’ has specific and restricted interpretation. They might not find all the planes being interesting for the expert in the context of risk assessment. Furthermore, and more importantly for our purpose, methods that compute 2D projections of multi-dimensional spaces are beyond the understanding 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 itself. Therefore, a meaning has to be assigned to the data elements in such a way that it has an associative role for the expert.

3.2 Numeric data

3.2.1 Data visualization

A taxonomy of multidimensional visualizations is provided by [11]. 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 real values of dimensions accommodating high visual

clarity, brings us to the scatterplot [12]. That implies trading in some very good features of more complex visualization techniques, but it does fulfill the requirements. As in a 2D scatterplot data elements are drawn as points in the Cartesian space defined by two graphical axes defined by the real attributes values, therefore accommodating the understanding framework of the user. Scatterplots are also frequently used. They are basic building blocks in statistical graphics and data visualization [13]. Multidimensional visualization tools that feature scatterplots, such as Spotfire [14], XmdvTool [15], Tableau/Polaris [16], GGobi [17], typically allow mapping of data dimensions also to graphical properties such as point color, shape, and size.

However, the number of dimensions that a single scatterplot can reliably visualize is considerably less than many realistic datasets. Therefore, a series of scatterplots should be visualized for all combinations of the dimensions, where all the dimensions can be explored by the user. However this approach yields little structure to the visual exploration and provides no relation between the data dimensions. Multiple plots can be arranged in a scatterplot matrix [13] and we can link the multiple visualizations of combinations of two dimensions in order to reveal relationships between them [5] through interactions.

3.2.2 Model visualization

We use scatterplots to visualize the results of a classification model. We use the color and the size to visually express the results of classification. Color expresses the original class membership of the data. Size indicates whether the data element is misclassified by the classifier, if so it is assigned twice the regular size.

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 to visualize decision boundaries for multi-dimensional data [9, 18, 19]. However, those methods do not relate the visualization of the decision boundary to the data projections of the axis that are meaningful for the user. In [18] authors visualize the support vector machine classifier (SVM) using 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 and weighting the plane coordinates to choose the most important features for the classification. Poulet [9] 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. Further, Hamel [19] uses self-organizing maps to visualize results of the SVM. These examples of methods to visualize the results of the classifiers are all applied to SVM, with the exception of [9] who proposes a method that can be applied to other classifiers like 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. Since we use scatterplots to visualize the data we propose to visualize the multi-dimensional decision boundary in series of scatterplots. However, the projection to 2D of the multi-dimensional decision boundary onto this data projections would result in loss of separating information. Therefore, we propose to use an approximation of the decision boundary, using Voronoi tessellation.

3.2.3 Voronoi-based approximation of decision boundary

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 neighbors. 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 neighbors 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 piecewise 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 elements as returned by the classifier could be visualized.




Figure 2: Visualizations of the results of a nearest mean classifier for a two-class dataset. (a) Scatterplot for the numerical data with the Voronoi diagram together with an approximated decision boundary following the Voronoi cells boundaries (thick solid line). For visual clarity the classes are assigned different colors and the misclassified data elements are assigned twice the size of the correctly classified elements; (b) Mosaic plot for ordinal data. Each combination of the ordinal value pairs gives rise to one large rectangle. The subrectangles from left to right indicate the proportion of TN, FP, TP and FN belonging to that block.

3.3 Ordinal data

3.3.1 Data visualization

There are several approaches to visualize ordinal data, where attributes are ranked categories. Since we have chosen to visualize numeric data using series of scatterplots, we visualize the ordinal data in a consistent way. Therefore, we use mosaic plots which are graphical displays to examine the relationship among two or more categorical variables. First proposed by [20] they are now used in visualization toolkits such as Mondriaan [21] and Manet [22]. Based on an analogy to the scatterplots we visualize two attributes at the time, where ranked categories are assigned to the vertical and horizontal axis. This results in a series of plots for all combinations of categories.

In a mosaic plot each possible combination of categories is a cell and 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 data set. 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 are lined-up [23]. Another option would be to allow one of the attributes to be dependent on the other, but this would yield a visualization which would not be consistent with scatterplots for numeric data.




Figure 3: Example of a matrix overview showing 6 dimensions of a dataset. The overview matrix contains scatterplots for numerical attributes and mosaic plots for ordinal attributes.

3.3.2 Model visualization

For each block of the mosaic plot we want to present information about the classification results. We use the following methodology. In each block we visualize the proportion of TP, TN, FP and FN according to the labels assigned by the classifier. First we assign a color to both classes, according to the original labels. Parts of the block containing correctly classified elements are visualized according to the original class color. Parts containing misclassified elements are visualized in the darker tint of the original class color. In this way we make an analogy to different sizes of the visualized points in the scatterplots. To make the visualization of the classification results even more consistent with the scatter plot, we first divide each block into parts as divided by the classifier. Therefore, in each block a division occurs which could be considered a decision boundary in the scatterplot. Further we show, as in a scatterplot, the blocks containing misclassified elements on both sides of the decision boundary. Figure 2 shows a mosaic plot for two categorical ranked features.

3.4 Matrix overview

Obviously the visualization of multi-dimensional data by means of 2D plots will result in many scatterplots which are neither easy nor clearly arrangeable. To deal with the large amount of visualized plots we use the widely known scatterplot matrix and order the dimensions according to a suitable criteria. The dimensions could be ordered according to some interestingness criterion. Ankerst [24] proposes to place the similar dimensions next to each other, after deriving the similarity metrics. In [12] Elmqvist et al. follow this method, defining the similarity measure as absolute correlation of pairwise dimensions to order the rows and a dissimilarity measure inverse of the absolute correlation to order the columns. Tatu [25] developed methods to find the best scatterplots showing the separating classes.

We could also develop a criterion based on the model visualization for the numeric data, where good views should show good class separation. The minimal overlap of classes would then be the notion of importance of dimensions. In the risk assessment context the dimensions could be ordered according to the order which is natural for the expert. The order of rows and columns could be changed manually by the user in an interactive way, as used in Polaris [16], XmdvTool [15], ScatterDice[12]. A notion of interestingness is then not included, however the understanding framework of the expert is supported. To use the space of the scatterplot matrix more efficiently we follow the example of [26]. We remove the 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 3.

3.5 Performance visualization

In 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 Receiver Op-

erating Characteristics (ROC) curve, often used in medicine, visualizes the trade-offs between hit rate and false alarm rate [27]. The Precision and Recall curve often used in information retrieval depicts the trade off between the fraction of retrieved documents relevant to the search and the fraction of the documents relevant to the query successfully retrieved. 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. The curve, as we use it, represents the trade-off between the False Positives and False Negatives. Let FP be the number of incorrectly classified negative data items and NP the total number of positive examples.

The false positive rate, FPr of a classifier is:

$$FPr = \frac{FP}{NP} \quad (1)$$

The false negative rate, FNr of a classifier is:

$$FNr = \frac{FN}{NN} \quad (2)$$

where FN are the incorrectly classified positive items and NN 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 mis-classifies all positive data elements, the lower right point (1, 0) represents the classifier that mis-classifies 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.

3.6 Images 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 to trigger the associations he/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 of an organ 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. In particular, for characteristic data elements the associations can be strong. We distinguish two cases where 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 not focus on the physical properties of the object on the images, but on the associations that these objects invoke. Secondly, 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. As most important we consider those elements which are assigned a different label by the classifier when changing the classifier's trade-off. We call those critical elements and propose to visualize them instantly whenever the classification model changes. Figure 4 shows how we visualize the images for the critical data elements, occurring in between adjacent operating points. In Forensic Op-

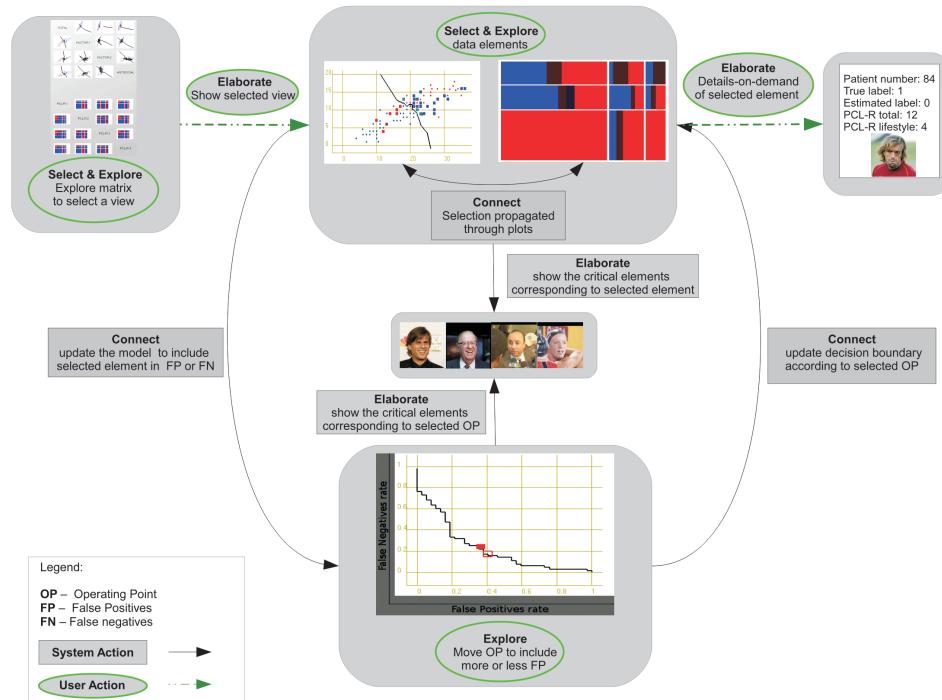


Figure 4: Our framework for a visual risk assessment system, elaborating upon figure 1, with the proposed visualization solutions and interaction techniques. Both system and user actions are indicated.

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. Since the data elements are visualized in relation to the classifier and the expert can associate them with actual patients he exactly knows how particular patients relate to the classification model. He also can view the group of patients that is in the direct neighborhood of the recognized patient and therefore associate other patients with him. Therefore, the image is a sort of 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, since we allow free associations of the expert that can not 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 section 4.1.

4 INTERACTIVE RISK ASSESSMENT

The purpose of this work is to support risk assessment. The techniques we described in the previous sections are well established methods. Combining them into a highly interactive Visual Analytics framework, makes them suitable for the risk assessment task. In this section we describe how we couple the elements in order to design an interactive tool for data and classifier exploration, particularly including the trade-off selection in the context of risk assessment.

4.1 Interaction Techniques

Through the interactive data exploration the expert can not only gain insight into the data but also into the classification process.

Therefore, several interaction techniques are integrated into our method which form a minimum necessary set of interaction techniques for the interactive risk assessment. We chose the techniques from the 7-categories of interaction techniques proposed in the study of [28]. They propose a well structured way to handle interactivity based on the users' functional needs. We make a distinction between the interaction techniques that are strictly user based and interactions that are triggered by the user, but result in pre-defined system actions. Figure 4 shows the possible interaction techniques in the system.

4.1.1 Select

The selection interaction technique aims at making the elements of interest visually distinctive. An expert can highlight the element of interest in the scatterplot or a rectangle in the mosaic plot, resulting in a color change of the selected element. This allows to keep track of elements of interest. The expert can also de-select the element if he is no longer interested in it. In the overview matrix the currently viewed plots are selected with the navigation lens.

4.1.2 Explore

Explore interaction techniques enable users to examine a different subset of data cases. In the matrix overview an expert can explore all the plots using the navigation lens. Direct-Walk allows users to smoothly move the viewing focus from one position in the information structure to another by "a series of mouse movements". That occurs when moving the operating point on the performance curve and by including certain data elements into a specific class. This is described in detail in the subsequent section.

4.1.3 Connect

Connect is used to highlight relationships between data items and show hidden relations between relevant items. In our system we

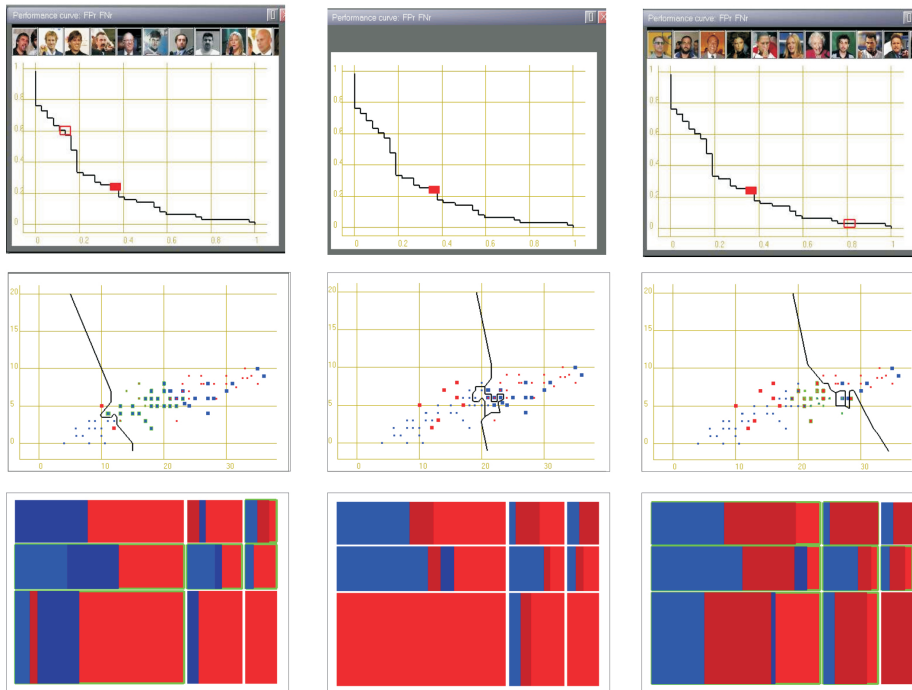


Figure 5: Example of the connect interaction. Figure in the middle shows the initial state of the system, with initial operating point on the performance curve and corresponding visualization of the model for scatter- and mosaic plot. Figure on the left shows the state of the system when an expert manipulates the operating point to include more False Negatives and on the right to include more False Positives. The updated model is visualized in the scatterplot and mosaic plot, together with the critical element represented as photos and highlighted in all visual components.

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 to see 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 5. That connection is twofold.

On the one hand we propose to 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 interactive operating point is then directly connected to the classification model. This means that the classifier is updated according to the new FP rate and FN rate. Since the visual displays are connected we instantly observe what effect the change of trade-off has on the classifier. The updated classifier for the adjusted operating point is directly visible 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 system highlights the combination of categories they belong to on the mosaic plot.

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 require 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 positive would also introduce new False Positives. In such a way he can explore whether a change of the model for the particular data element results in a serious cost change and which other data elements are affected.

4.1.4 Elaborate

Elaborate allows the user to see more details, hence to adjust the level of abstraction of the data representation. We use two techniques belonging to 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 can be displayed, when the user aims at placing that element on another side of the decision boundary ("pre-defined" 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. In the matrix overview an expert can move the lens which will directly result in displaying the selected choice of plots.

4.2 Visual design

The system we propose consists of several visual components. The implementation of the system is shown in figure 6. The main components of the system are:

- *Overview Window*: serves both as an overview and a navigation tool. Miniature versions of the individual plots are shown as the cells of the matrix, and the lens that is used for navigation is placed above the plots that are currently displayed in the main window. The plots are visualized together with the decision boundary.
- *Main Window*: displays three plots currently chosen by the expert using lens navigation in the matrix overview. The performance curve is 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 error, amount of TP and TN.
- *Details Window*: includes the information on a particular interactively selected data element, including the photograph associated with this element.

5 APPLICATION OF INTERACTIVE RISK ASSESSMENT TO 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 the society or should be treated in a closed mental institution. In fact, 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 the society is an integral part of the treatment of mentally ill patients. Therefore informed choices have to be made to minimize the risk of making severe mistakes but the costs of mistakes in disadvantage of the patient should also be avoided.

5.1 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) [29]. The dataset consists of 100 male offenders who, at the time of the alleged crime, suffered from 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 7 numeric features are the summed combinations of the PCL-R outcomes. Patients were retrospectively scored with this risk assessment measures and recidivism data was retrieved from the documentation of the Dutch Ministry of Justice. To visualize the patients we assign a red color to recidivists and blue to non-recidivists. As we can not use the real photos for this dataset due to privacy reasons. As an example we use a set of images of celebrities downloaded from the internet. These images are used to illustrate merely the design of the system and how the images are interactively coupled to the other visual elements. To illustrate the usage of the system we choose to focus only on the results of one classifier. Table 1 presents the 10-fold cross-validation error, over 3

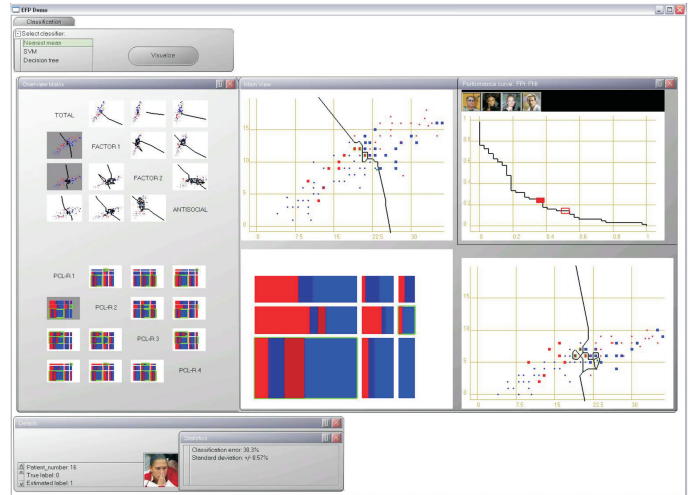


Figure 6: Screenshot of the system.

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.

Table 1: Performance of the selected classifiers for PCL-R dataset obtained with 10-fold cross-validation, with 3 repeats.

Classifier	Cross-val error %	std
Nearest Mean	30.3	± 0.5
Logistic Regression	31.3	± 1.5
Fisher	31.0	± 1.0
Support Vector Machine	31.3	± 1.0
Decision Tree	33.6	± 4.04

5.2 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, has been provided with a newly developed classification model. Analyst X is asked to explore the model applied to the patient data which she is familiar with.

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 current operating point are displayed. Figure 6 shows the screenshot of the system that is now available for the expert.

Analyst X first looks at the Overview Window to get an overall feeling of 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 they are classified 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 questions herself whether it can be explained why those patients have such high attributes values and they still 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 a very good re-socialization conditions and strong network of people taking care of him, that 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 less False Negatives, as she is not satisfied with so many recidivists being misclassified, she notices among the images a patient that is familiar to her. She asks for the details. The scores indicating the violence behavior of the patient are indeed not very high indicating the average risk of re-offence. However she recognizes Patient X and realizes he is very bad in taking his medication. Therefore the risk of re-offending when he leaves the clinic and has no one to control his medication use, would increase. She requires the model to include that patient into True Positives. She explores the result on the connected plots, where the model is updated and the performance curve, where the operating point is updated. Based on that limited exploration she finds the updated model suitable for use in her clinic.

6 DISCUSSION AND FUTURE WORK

This paper proposes an interactive approach to risk assessment. The main idea is to provide an expert a framework tightly integrating interactive visual exploration with machine learning taking into account the demanding requirements of risk assessment applications. We propose to couple visualizations of numeric and ordinal data, classifiers visualization, performance visualization and associative images representing meaningfully the data elements in order to support an expert in the visual exploration of the dataset and the classification model. 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. The most important next step is to conduct systematic empirical evaluation of the system. We plan to use multiple rounds of evaluation with the end users to test the usefulness of the approach and identify the areas of improvement. However user evaluation is difficult for such broad tasks as visual exploration, so we anticipate performing time consuming qualitative studies.

7 ACKNOWLEDGMENTS

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

REFERENCES

- [1] D. A. Keim, F. Mansmann, J. Schneidewind, J. Thomas, and H. Ziegler. Visual analytics: Scope and challenges. pages 76–90, 2008.
- [2] J.J. Thomas and K.A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE CS Press, 2005.
- [3] R. Fuchs, J. Waser, and M.E. Gröller. Visual human+machine learning. *IEEE TVCG*, 15(6):1327–1334, October 2009.
- [4] F. Bendix, R. Kosara, and H. Hauser. Parallel sets: Visual analysis of categorical data. In *INFOVIS '05: Proceedings of the Proceedings of the 2005 IEEE Symposium on Information Visualization*, 2005.
- [5] C. Collins and S. Carpendale. VisLink: Revealing relationships amongst visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6), 2007.
- [6] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley-Interscience Publication, 2000.
- [7] C. Yu, Y. Zhong, T. Smith, I.n Park, and W. Huang. Visual data mining of multimedia data for social and behavioral studies. *Information Visualization*, 8(1):56–70, 2009.
- [8] M. Ankerst, D. H. Jones, A. Kao, and C. Wang. Datajewel: Tightly integrating visualization with temporal data mining. *ICDM Workshop on Visual Data Mining*, 2003.
- [9] F. Poulet. *Towards Effective Visual Mining with Cooperative Approaches*. Springer-Verlag, Berlin, Heidelberg, 2008.
- [10] M. Ankerst, M. Ester, and H. P. Kriegel. Towards an effective cooperation of the user and the computer for classification. In *Knowledge Discovery and Data Mining*, pages 179–188, 2000.
- [11] D. A. Keim. Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):1–8, 2002.
- [12] N. Elmqvist, P. Dragicevic, and J.D. Fekete. Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *IEEE Transactions on Visualization and Computer Graphics (Proc. InfoVis 2008)*, 14:1141–1148, 2008.
- [13] W.A. Cleveland and M. E. McGill. Dynamic graphics for statistics. *Statistics/Probability Series*, 1988.
- [14] Spotfire Inc. Spotfire. <http://www.spotfire.com>, 2007.
- [15] M. O. Ward. Xmdvtool: integrating multiple methods for visualizing multivariate data. In *VIS '94: Proceedings of the conference on Visualization '94*, pages 326–333. IEEE Computer Society Press, 1994.
- [16] C. Stolte, D. Tang, and P. Hanrahan. Polaris: A system for query, analysis, and visualization of multidimensional relational databases. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):52–65, 2002.
- [17] D. F. Swayne, D. T. Lang, A. Buja, and D. Cook. Ggobi: evolving from xgobi into an extensible framework for interactive data visualization. *Computational Statistics and Data Analysis*, 43(4):423–444, 2003.
- [18] D. Caragea, D. Cook, and V. G. Honavar. Gaining insights into support vector machine pattern classifiers using projection-based tour methods. In *KDD '01: Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 251–256, 2001.
- [19] L. Hamel. Visualization of support vector machines with unsupervised learning. In *Proceedings of 2006 IEEE Symposium on Computational Intelligence in Bioinformatics and Computational Biology*, 2006.
- [20] A. Unwin, G. Hawkins, H. Hofmann, and B. Siegl. Mosaic for contingency tables. *Computer Science and Statistics: Proceedings of the 13th Symposium on the interface*.
- [21] M. Theus and S. Urbanek. *Interactive Graphics for Data Analysis: Principles and Examples (Computer Science and Data Analysis)*. Chapman & Hall/CRC, 2008.
- [22] A. Unwin, G. Hawkins, H. Hofmann, and B. Siegl. Interactive graphics for data sets with missing values: Manet. *Journal of Computational and Graphical Statistics*, 5(2):113–122, 1996.
- [23] M. Theus and S.R.W. Lauer. Visualizing of loglinear models. *Journal of Computational and Graphical Statistics*, 8(3):396–412, 1999.
- [24] M. Ankerst, S. Berchtold, and D.A. Keim. Similarity clustering of dimensions for an enhanced visualization of multidimensional data. In *INFOVIS '98: Proceedings of the 1998 IEEE Symposium on Information Visualization*, page 52, Washington, DC, USA, 1998. IEEE Computer Society.
- [25] A. Tatu, G. Albuquerque, M. Eisemann, J. Schneidewind, H. Theisel, M. Magnor, and D. Keim. Combining automated analysis and visualization techniques for effective exploration of high-dimensional data. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology (IEEE VAST)*, pages 59–66, Atlantic City, New Jersey, USA, 10 2009.
- [26] J. Heer, M. Bostock, and V. Ogievetsky. A tour through the visualization zoo. *Commun. ACM*, 53(6):59–67, 2010.
- [27] T. Fawcett. An introduction to ROC analysis. *Pattern Recogn. Lett.*, 27(8):861–874, 2006.
- [28] J.S. Yi, J. Kang, J. Stasko, and J. Jacko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231, 2007.
- [29] H. Hildebrand, B.L. Hesper, M. Spreen, and H.L.I. Nijman. *The value of structured risk assessment and the diagnosis of Psychopathy (in Dutch)*. Expertise Center for Forensic Psychiatry, 2005.