VRPM: An R Package to Visualize Risk Prediction Models

Van Belle V.[*1,2], Van Calster B.[3], and Van Huffel S.[1,2]

[1]Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, KU Leuven, Leuven, Belgium
[2]iMinds Medical IT, Leuven, Belgium
[3]Department of Development and Regeneration, KU Leuven, Leuven, Belgium

February 5, 2016

1 Abstract

Nomograms are graphical tools to illustrate the working of regression models. For non-linear and interaction effects, line-based nomograms are difficult to interpret. An alternative nomogram based on colors that is able to deal with non-linear and interaction effects in an appropriate way was recently proposed. This article describes how the *VRPM* package can be used to produce these plots for "glm","coxph","mfp", "multinom" and "ksvm" objects. Additional plots summarizing the risk prediction process for specific observations are provided as well.

2 Introduction

Risk prediction models become more available for non-scientific and non-statistical users. People are confronted with them in daily life due to the increase in patient consent and informed decision making in medicine. In the future, the broad introduction of smart watches to track lifestyle, health and fitness levels will only increase their use. A major concern in this field will be the explanation of associated risk predictions, a problem that has proven to be difficult in the past: clinicians tend to use overly simple models, by preference score systems, since the interpretation of these methods is straightforward.

However, in a lot of cases, more complex models are necessary to obtain adequate risk predictions. Explanation of how the risk prediction is obtained, based on measured features becomes therefore increasingly important. The *rms* package (Harrell, Jr, 2015) enables to illustrate the working of binary and ordinal logistic regression models (Harrell, 2001) by means of nomograms. An illustration of a nomogram and a thorough discussion on how they are created can be found in Iasonos et al. (2008). Nomograms were originally proposed as tools to calculate risk prediction without the need for a computer. Since personal computers and smartphones have become very powerful, use of paper based representation to calculate this risk is no longer necessary. Figure 1 gives an example of a nomogram for a logistic regression model, where the admission status into a graduate school is modeled by means of `gre` (Graduate Record Exam scores), `gpa` (grade point average) and `rank` which encodes the prestige of the undergraduate school that was attented (`rank`=1 is the highest prestige, `rank`=4 the lowest). The model is taken from the Institute for digital research and education. The length of the lines corresponding to each predictor indicate the range of the contributions to the linear predictor in the training data, and as such indicate the importance of the predictor. Originally, the nomogram was intended to draw vertical lines

*vanya.vanbelle@esat.kuleuven.be

from the predictor values to the points scale (top axis). For a student with a gre=410 this would yield 30 points. The resulting points for all predictors then need to be added to obtain the total score, which is then converted to a risk prediction by drawing a vertical line from the total points axis to the chance axis (bottom). More important than being an offline risk calculator, the nomogram offers an interpretation of the impact of the different predictors in the represented risk prediction process. A closer look at the plot learns that a higher gre and a higher gpa result in higher points (top axis) and as such in higher estimated chances (bottom axis). Having attended an undergraduate school with less prestige yields a very low point and reduces the chance of being admitted.

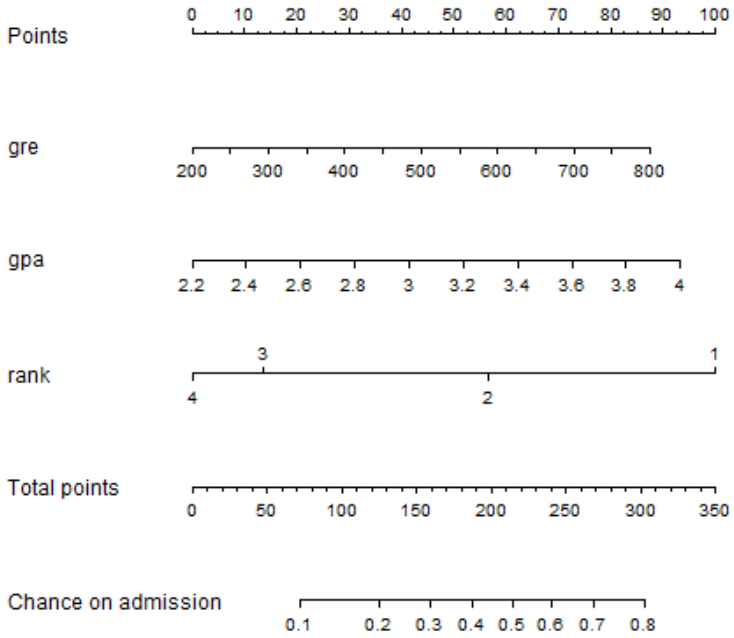

Figure 1: Nomogram for a logistic regression model.

An extension of this type of nomogram for multinomial logistic regression models with only three output levels is proposed in Ardoino et al. (2014). The latter work provides R sample code, but the proposed plot is not extendable to more than three output levels.

A drawback of the nomograms discussed so far is found when non-linear transformations of the predictors or interactions between continuous predictors are necessary to achieve a well performing risk prediction model. In the first case, the lines corresponding to the predictors with non-monotonic transformations need to be replicated. For medical applications this is often not a problem, but in other domains with more complex transformations, interpretation of the resulting plot becomes less straightforward. To allow the representation of interactions between two continuous predictors, one of the predictors needs to be categorized. This can potentially lead to loss of information and when one predictor is present in multiple interaction effects, interpreting the global model becomes very cumbersome.

Representation of Support Vector Machines (SVMs) classifiers (Vapnik, 1998) is even more difficult. These models project the predictors into a higher dimensional space to enable a linear decision bound in the latter space. To accomplish this, a kernel is used such that the transformation of the predictors can not be known in general. Particularly when using the popular Radial Basis Function (RBF) and polynomial kernels, it is unknown how the predictors influence the estimated outcome of the SVM model. Attempts have been made to make these methods more transparent by using decomposable kernels and extending the use of nomograms to SVMs (Jakulin et al., 2005). Drawbacks of this approach are that it is restricted to non-conventional kernels, non-linearities are represented by 2-dimensional plots and interactions between continuous predictors

need categorization.

In Van Belle and Van Calster (2015) an alternative representation of risk prediction models was proposed: a color based nomogram. In this new type of nomogram, the color encodes the points, similar to the length of the lines in the traditional nomogram. As a result, non-linearities and interactions can be represented appropriately. In Van Belle et al. (b) this approach was extended towards multinomial logistic regression and in Van Belle et al. (a) a new approach was proposed to represent SVMs with the RBF and polynomial kernel by means of these color based nomograms.

The contribution of the *VRPM* package is to gather all functionality to represent logistic regression models (binary and multinomial using the *nnet* package (Venables and Ripley, 2002b)), survival models and SVMs (*kernlab* (Karatzoglou et al., 2004)) using a linear, RBF and polynomial kernel by means of a color based nomogram. Additional functions are provided to summarize the risk prediction process for specific observations, to summarize the different risk estimates for multinomial logistic regression models and to evaluate the correctness of the approximation of the SVM models used to create the nomogram.

3 Color based nomograms for logistic regression

In order to build a nomogram from a risk prediction model, the risk prediction process needs to be written in the following form

$$\hat{p} = h\left(\sum_{p=1}^{d} f^{(p)}\left(x^{(p)}\right)\right), \tag{1}$$

or

$$\hat{p} = h\left(\sum_{p=1}^{d} f^{(p)}\left(x^{(p)}\right) + \sum_{p=1}^{d}\sum_{q=p+1}^{d} f^{(p,q)}\left(x^{(p,q)}\right)\right), \tag{2}$$

when interactions are allowed. Here the estimated risk $\hat{p}$ is formed as a monotonic function $h(\cdot)$ of a linear predictor or latent variable, that is found as the sum of contributions depending on one ($x^{(p)}$) or two predictors ($x^{(p,q)}$), with $x^{(p)}$ the pth predictor of input x. For a binary logistic regression model with only linear main effects, the function $h(\cdot)$ is a sigmoid function and the contributions are $f^{(p)}(x^{(p)}) = \beta^{(p)}x^{(p)}$, with $\beta^{(p)}$ the estimated coefficient corresponding to predictor $x^{(p)}$. The nomogram in Figure 1 relates points, visualized by means of the length of the line to $\beta^{(p)}x^{(p)}$, for each predictor $x^{(p)}, p = 1, \ldots, d$. The sum of all these points is converted to a risk by means of the bottom two scales. The relation between both scales represents the function $h(\cdot)$.

The color based nomogram is explained using the example of admission to a school. The fitted model coefficients of the binary logistic regression model are reported in Table 1. The corresponding color based nomogram is illustrated in Figure 2. In general, the color corresponding to a value a of the pth predictor $x^{(p)}$ encodes $\beta^{(p)}a$ points, analogous to the length in the nomogram in Figure 1. The package allows more flexibility by means of the option zerolevel. A zerolevel equal to "zero" results in $\beta^{(p)}a$ points. To enable a plot that is most similar to traditional nomograms (where the contributions are shifted such that contributions always range from zero to a positive value), zerolevel can be put to "min". Again, points corresponding to all predictors (and predictor pairs in the presence of interactions) need to be added. Conversion from the resulting score to the estimated risk is performed using the bottom most color bar, representing the function $h(\cdot)$. The following code results in Figure 2.

```
> library("VRPM")
> mydata <- read.csv("http://www.ats.ucla.edu/stat/data/binary.csv")
> mydata$rank <- factor(mydata$rank)
> fit.lr <- glm(admit ~ gre + gpa + rank, data = mydata, family = "binomial")
> obs1=data.frame(gre=408,gpa=3.87,rank=factor(2))
```

Table 1: Coefficients (β) and odds ratios (OR) for the binary logistic regression model to estimate admission.

	β	OR	2.5 %	97.5 %
(Intercept)	-3.990	0.019	0.002	0.167
gre	0.002	1.002	1.000	1.004
gpa	0.804	2.235	1.174	4.324
rank2	-0.675	0.509	0.272	0.945
rank3	-1.340	0.262	0.132	0.512
rank4	-1.551	0.212	0.091	0.471

```
> colplot(fit.lr,coloroptions=3,zerolevel="min",obs=obs1,filename="lrmin",
+                  risklabel="Chance on admission")
```

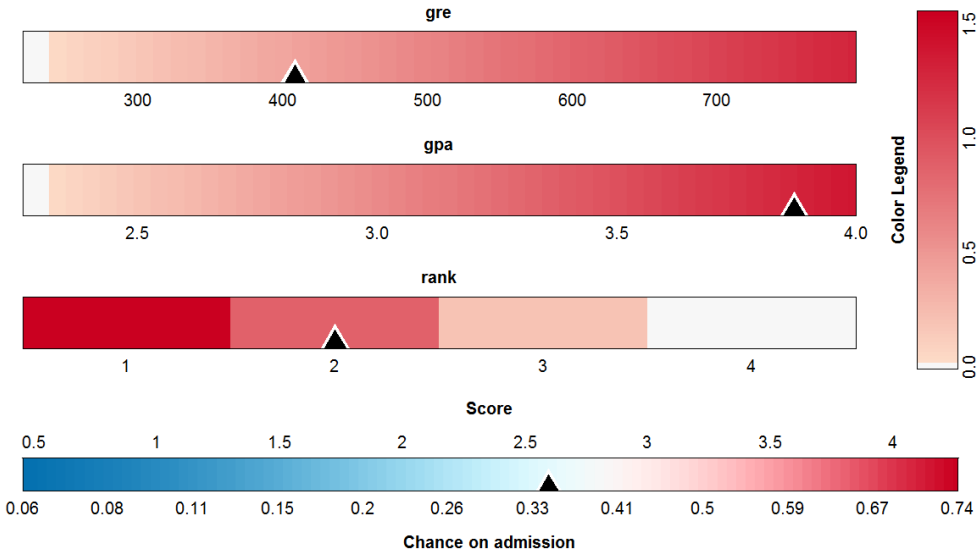

Figure 2: Colored nomogram for a logistic regression model. The triangles indicate the predictor values of one specific observation. For each bar corresponding to a predictor, the color is used to obtain a point using the color legend at the right. Addition of all these points results in the score. Using the conversion bar at the bottom, this score is converted to a risk estimate.

For this example it was chosen to use a zerolevel=”min” and to add an observation to the graph. The predictor values and risk estimation for this observation are indicated by means of triangles.

The representations shown in Figures 1 and 2 are alternatives as they provide the same information. The colplot function has an option coloroptions enabling the user to choose between 5 different color maps. Further options include the choice of the shift of the contributions (zerolevel=c(”zero”,”min”,”mean”,”median”)), addition of an observation (obs) and the filename, among others. See the manual for more options and details Van Belle (2015).

4 Color based nomograms for Cox proportional hazard models

For Cox proportional hazard models (Cox, 1972), similar nomograms can be created. When no time point at which the estimated survival should be reported is provided, the median survival time within the training data is used as time point. The following code trains a Cox model including multivariable fractional polynomials (see *mfp* (original by Gareth Ambler and modified by Axel Benner, 2015; Royston and Sauerbrei, 2008)) and plots a nomogram to predict survival at 500 days (Figure 3). In this case, the contributions are shifted such that the median of each contribution equals zero (zerolevel="median"). In combination with a diverging color map (selected by setting coloroptions=3), the median contributions will have a white color, higher and lower contributions will be presented in red and blue, respectively. This representation instantly reveals the skewness of the predictors posnodal, prm, tumsize and esm.

```
> library(mfp)
> data(GBSG)
> fit.coxph <- mfp(Surv(rfst, cens) ~ age+fp(tumsize)+fp(posnodal)+fp(prm)+fp(esm)+
+                                  menostat+tumgrad, family = cox, data = GBSG)
> colplot(fit.coxph,coloroptions=3,zerolevel="median",time=500,filename="coxmfp")
```

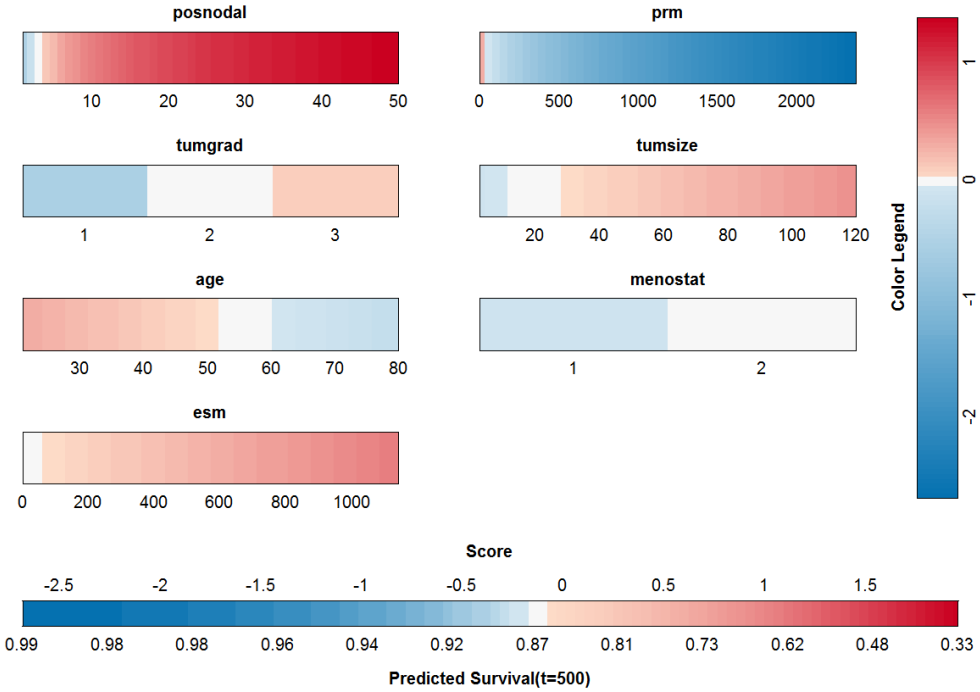

Figure 3: Colored nomogram for a Cox proportional hazard model, including non-linearities trained by means of the *mfp* package.

5 Color based nomograms for multinomial logistic regression

For multinomial logistic regression models, one colored nomogram is created for each output level, except for the reference level. These plots indicate how the linear predictor for each level is obtained, starting from the predictors. The linear predictor for the reference level equals zero by definition, such that a figure is not needed in this case. Since the linear predictor for one output level is not sufficient to obtain a risk prediction, these nomograms do not contain a convertion from the score to a risk estimate. Instead k additional plots, with k the number of output levels, are provided, each of which illustrate the estimation process of the risk of one output level. For the reference level, the predicted risk $\hat{p}_k$ is calculated as

$$\hat{p}_k = \frac{1}{1 + \sum_{l=1}^{k-1} \exp\left(\mathrm{lp}_l\right)}, \tag{3}$$

with lp_l the linear predictor of the lth outcome level. It can be seen that this equation is in the same format as equation (1), when taking $h(\cdot) = \frac{1}{1+\mathrm{score}}$, and score $= \sum_{\mathrm{l}=1}^{\mathrm{k}-1} \exp\left(\mathrm{lp}_\mathrm{l}\right)$.

For non-reference output levels, the estimated risk $\hat{p}_l$ is given by

$$\hat{p}_l = \hat{p}_k \exp\left(\mathrm{lp}_l\right), \tag{4}$$

or in the format of equation (1)

$$\hat{p}_l = \exp\left(\ln\left(\hat{p}_k\right) + \mathrm{lp}_l\right). \tag{5}$$

The code below illustrates how these plots can be generated for the prediction of marital status based on the gender, age, weight and height of a person.

```
> library(nnet)
> library(VGAMdata)
> data(xs.nz)
> marital.nz <- xs.nz[,c("marital","sex","age","height","weight")]
> mydata <- marital.nz[complete.cases(marital.nz),]
> fit.mlr <- multinom(marital ~ sex + age + height + weight, data = mydata,model=TRUE)
> # for multinimial logistic regression, a vector of risk labels needs to be made
> # and provided to the colplot function
> outnames=colnames(fitted(fit.mlr))
> labels=c(paste("Linear predictor for",outnames[-1]),paste("Predicted chance of being",
+                 outnames))
> # visualize the model: more than one plot is generated in the current directory
> obs2=mydata[45,]
> colplot(fit.mlr,filename="multinom",coloroptions=3,zerolevel="zero",risklabel=labels,
+                 obs=obs2)
```

Since more than one colored nomogram is created, showing all these on one page becomes difficult. The package therefore provides a function `HTMLsummary` that creates an HTML page containing all these colored nomograms. Note that the filename provided to the `colplot` function should match the filename provided to the `HTMLsummary` function. The HTML page is obtained by the following code:

```
> HTMLsummary(fit=fit.mlr, summaryfile="mysummary", filename="multinom",
+                  title="Global summary of a multinomial logistic regression model
+                  estimating the risk on different marital statuses.")
```

The resulting HTML page is available as supplementary information (mysummary.html).

5.1 Alternative plot to represent the risk of non-reference output level

An alternative plot to illustrate the risk calculation process for the non-reference outcome levels is given in Figure 4. It is based on equation (4), where the risk on a non-reference outcome level is expressed in function of the estimated risk on the reference level and the linear predictor of the level the risk of which needs to be predicted. The plot is generated when using the function colplot and saved with extension `wing`. The red lines indicate how the risk for a specific observation is obtained. In this case, the estimated risk on the reference level and the linear predictor for the outcome level of interest is 0.29 and 2.4, respectively. Drawing a vertical line until the crossing with the curve representing an estimated risk on the reference outcome of 0.29, followed by drawing a horizontal line from this point to the y-axis, results in an estimated risk on the outcome of interest, i.e. 0.04. This approach can be repeated for all non-reference outcomes. However, separate plots are necessary for each outcome level.

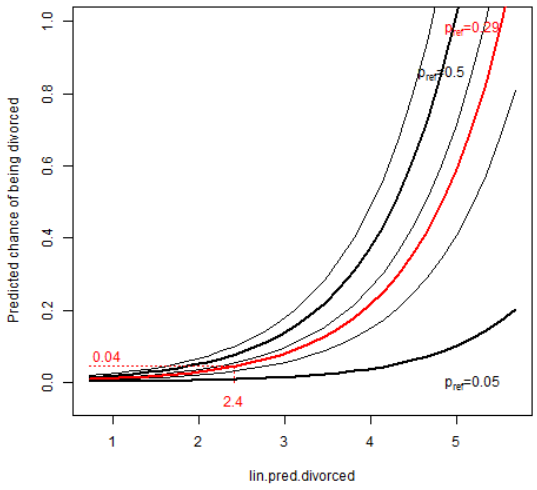

Figure 4: Alternative graph to represent the risk estimation process for a non-reference outcome level in multinomial logistic regression.

The HTMLsummary function will always provide a summary of these alternative plots, in addition to the colored nomograms.

5.2 Summary plots for specific observations within multinomial logistic regression models

In cases for which a specific observation is provided to the colplot function for a multinomial regression model, additional observation summaries as illustrated in Figure 5 are generated. The first summary indicates the estimated risks for all outcome levels and relates them to the prevalence of the training data (baseline risk). It also reports the relative risks to guide the user in revealing which risks have increased most compared to the baseline. The second summary relates the

estimated risks for this observation to the range of risk estimates observed in the complete training set.

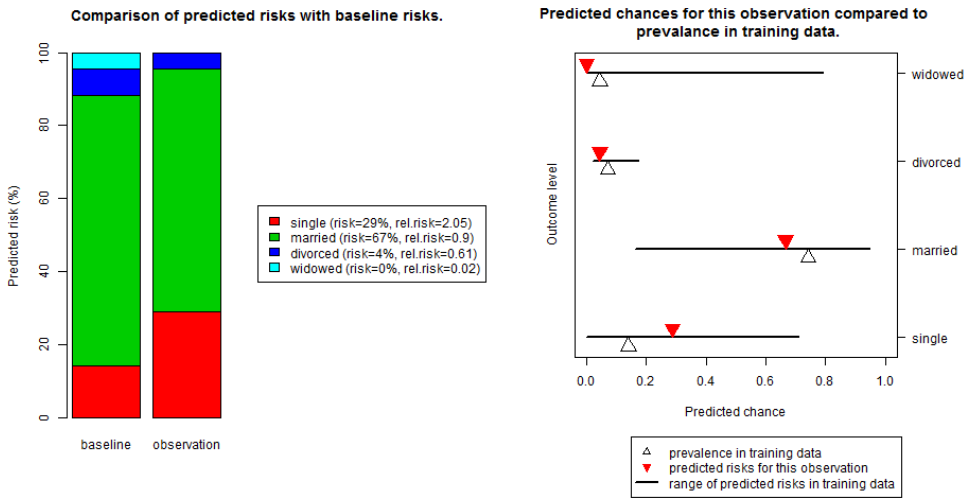

(a) (b)

Figure 5: Summary plots for one specific observation. The plots are created whenever an observation is provided to the colplot function to represent a multinomial logistic regression model. (a) Stacked bar plot comparing the baseline risk in the training set with the estimated risks for this specific observation. The relative risks (defined as the risk of this observation divided by the average risk in the training data) are presented as well. (b) Indication of baseline risks (white triangles) and estimated risks for this observation (red triangles). The black lines indicate the range of the estimated risks within the training set.

When an observation is provided to the colplot function, subsequently calling the HTMLsummary function will also provide the plots in Figure 5.

6 Color based nomograms for support vector machines

The *VRPM* package provides the possibility to create color based nomograms for SVMs when using three types of kernels: (i) linear kernels, (ii) polynomial kernels, (iii) RBF kernels. For a linear kernel, the color based nomogram is an exact representation of the SVM model. For polynomial and RBF kernels the representation of the SVM model by means of color based nomogram is based on the decomposition of these kernels in terms depending on one or two predictors. As such, only an approximation of the SVM model can be represented since interaction effects involving more than two predictors have contributions to the latent variable that are not considered in the representation. See Van Belle et al. (a) for more details regarding this decomposition. The method of Platt (1999) is used to obtain probabilities from the SVM model. The code below illustrates how an SVM model with RBF kernel can be represented. Note that the colplot function is not provided with the ksvm model itself. This would result in a representation that is solely based on the support vectors, and might not cover the complete input space. It is therefore necessary to first use the preplotperf function, with the ksvm model, the training data and the column number containing the output as arguments. The plot resulting from this code is presented in Figure 6.

```
> library(kernlab)
> data(iris)
```

```
> levels(iris$Species)[levels(iris$Species)=="setosa"] <- "other"
> levels(iris$Species)[levels(iris$Species)=="virginica"] <- "other"
> names(iris)=c("SL","SW","PL","PW","Species")
> # RBF kernel
> model <-ksvm(Species ~ ., data = iris,prob.model=TRUE,kpar=list(0.03),C=10)
> newmodel=preplotperf(model,iris,indy=5,zerolevel="median")
> colplot(newmodel,filename="svm",zerolevel="median",coloroptions=3)
```

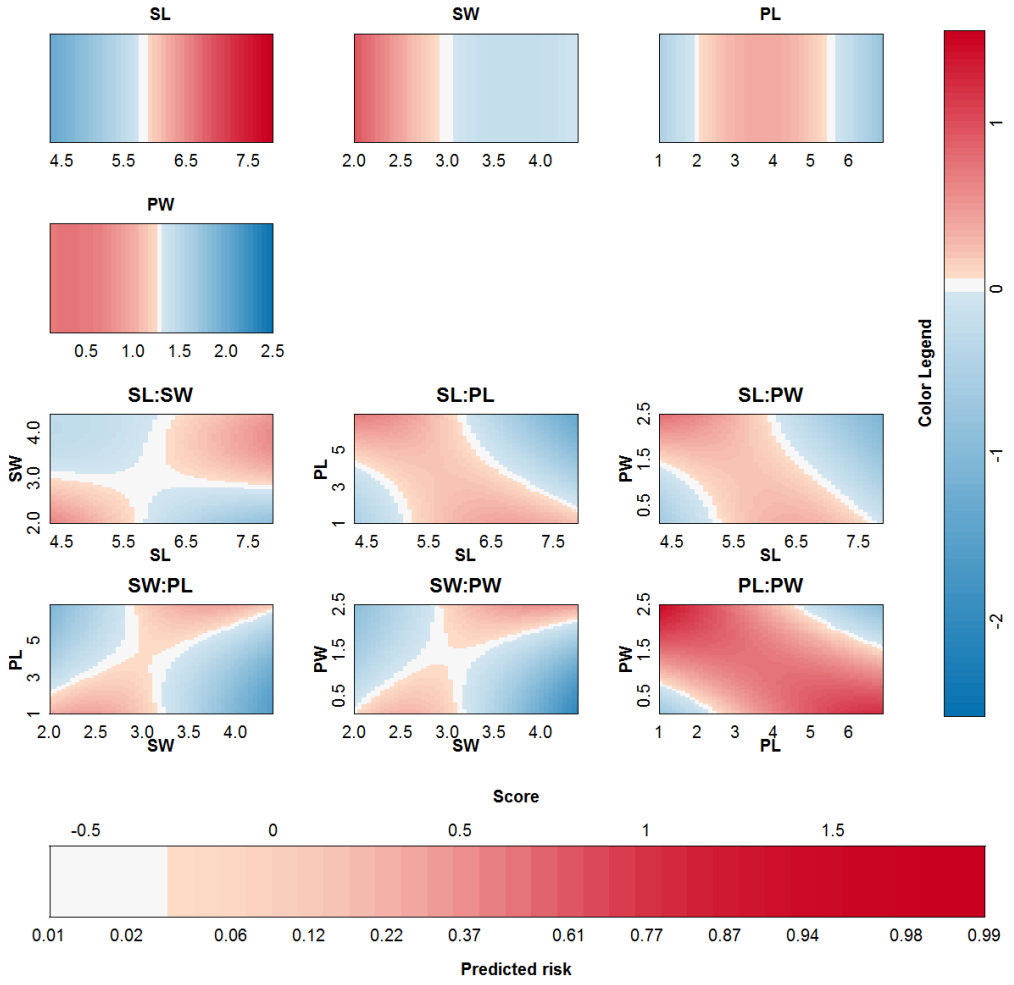

Figure 6: Colored nomogram for a support vector machine classifier with an RBF kernel. Due to the use of the RBF kernel, contributions are found in all main and all two-way interaction effects.

The choice of a diverging color map (coloroptions=3) and zerolevel="median", maps the median of the estimated risks in the training data to a white color in the bottom most bar of the plot. As such, the plot instantly reveals that the estimated chances in the training data have a skewed distribution.

6.1 Diagnostic plots

Using the *VRPM* it is possible to generate a colored nomogram for support vector machines using a linear, polynomial or Radial Basis Function kernel. However, in most cases, the created plot

will visualize an approximation of the SVM model. For a linear kernel the plot is exact. For a polynomial kernel, the plot is only exact when the degree of the kernel does not exceed 2 or when the input data has only 1 or 2 dimensions. For the RBF kernel, the plot is only exact for data with less than 3 dimensions. It is therefore extremely important to investigate the correctness of the approximation before interpreting any nomogram created from this approximation. The *VRPM* provides some diagnostic plots to aid in this process.

The plots are first created for an SVM model that can be approximated in a good way (model 1) and for a model that can not be approximated (model 2). Note that whether the approximation is good or bad does not depend on the training data of the SVM, but on the parameter settings. A well performing SVM model does not necessary result in a good approximation. A bad performing SVM model on the contrary can have a perfect approximation. Whether or not an approximation is good, depends on whether the approximation and the SVM model yield the same chances and output labels. More details can be found in Van Belle et al. (a).

Several diagnostic plots can be obtained using the plotperf function. Which plot is created depends on the argument type. In the code below, the two most important possibilities are used: (i) type="lp" results in a plot of the latent variable of the approximation versus the latent variable of the SVM model, (ii) type="contributions" results in a boxplot of the contributions to the approximation (all main and two-way interaction effects), the rest term (difference between latent variable of approximation and latent variable of the SVM model) and the latent variable of the SVM model. Figure 7 illustrates the results for both models. In case of a good approximation, the first plot should show points very close to the first diagonal. The second plot should show a rest term (=f(rest)) that is very small compared to the other contributions. In particular, the latent variable of the SVM model (indicated by lpmodel) should be large w.r.t. the rest term. Both plots clearly indicate that the approximation for model 1 is very good, while model 2 cannot be represented by means of the approximation. A colored nomogram for the latter model should therefore not be shown and certainly not be used to interpret the SVM model it was derived from.

```
> library(kernlab)
> data(iris)
> levels(iris$Species)[levels(iris$Species)=="setosa"] <- "other"
> levels(iris$Species)[levels(iris$Species)=="virginica"] <- "other"
> names(iris)=c("SL","SW","PL","PW","Species")
> # SVM model with a good approximation
> model1 <-ksvm(Species ~ ., data = iris,prob.model=TRUE,kpar=list(0.03),C=10)
> # SVM model with a bad approximation
> model2 <-ksvm(Species ~ ., data = iris,prob.model=TRUE,kpar=list(10),C=10)
> plotperf(model1,iris,indy=5,type=c("lp","contributions"),filename="svm")
> plotperf(model2,iris,indy=5,type=c("lp","contributions"),filename="svm2")
```

7 Summary plots for specific observations

The colored nomograms are mainly interesting to give a global overview of the model. They are able to provide insights into which predictors contribute most to the decision process and how risk estimates are obtained from more complex models. However, to summarize how risks are obtained for specific observations, it can be interesting to visualize all contributions for this observation on one plot. The *VRPM* package provides two types of plots for this purpose: contribution charts and cumulative contribution charts. Figure 8 gives an example of these plots ((a) contribution chart, (b) cumulative contribution chart), generated by means of the code below. In both charts, bars with a length equal to the contribution of each predictor are used. For contribution charts,

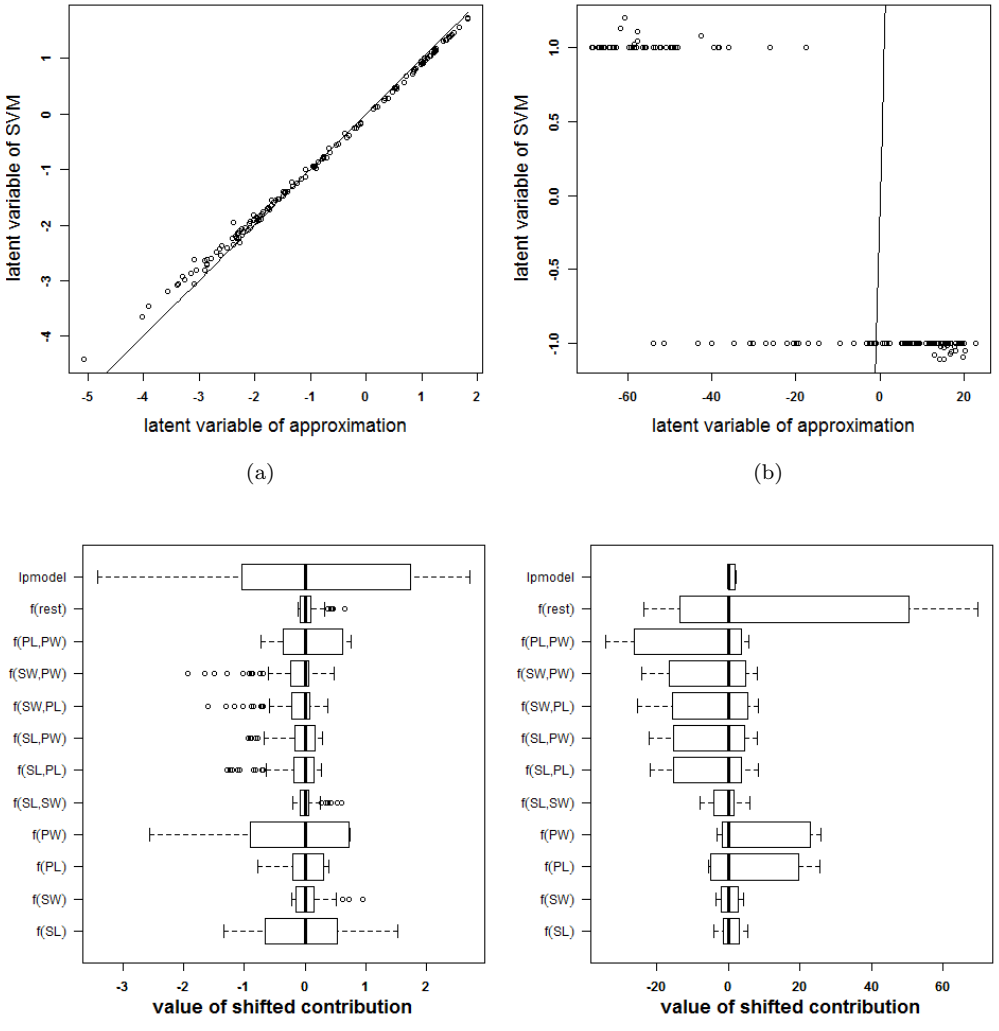

Figure 7: Diagnostic plots for an SVM model with a good approximation (a, c) and an SVM model with a bad approximation (b, d). Top: latent variable of the approximation versus latent variable of the SVM model. A good approximation results in a point cloud close to the first diagonal. Bottom: Boxplots of the different contributions in the approximation, the rest term (f(rest)) and the latent variable of the SVM model (lpmodel). A good approximation results in a small range of the rest term. All contributions are shifted to zero median to facilitate the comparison.

the bars all start at zero and black lines indicate the range of the contributions within the training set. This allows the user to observe outliers in new observations. In a cumulative contribution chart, the first bar starts at zero, whereas the other bars start where the previous bar ended. This cumulative approach directly yields the score when the last contribution is added. Both plots report the obtained score and corresponding risk.

```
> # contribution chart
> cchart(fit.lr,obs=obs1,filename="cchart")
> # cumulative contribution chart
> ccchart(fit.lr,obs=obs1,filename="ccchart")
```

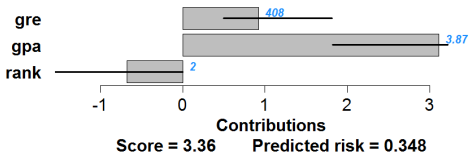

(a)

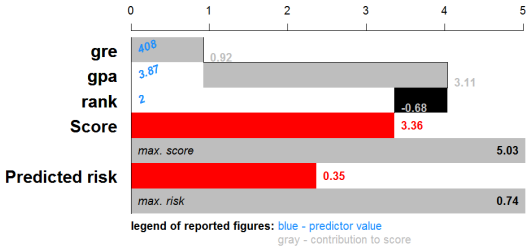

(b)

Figure 8: Summary of the contributions of all predictors to the score for one specific observation. (a) Contribution chart: the length of each bar corresponds to the contribution of a predictor to the score. The black lines indicate the range of the contributions observed in the training set. (b) Cumulative contribution chart: the length of each bar corresponds to the contribution of a predictor to the score. Thanks to the cumulation of the bars, the score is obtained automatically.

When the `HTMLsummary` function is called after a call to the `cchart` or `ccchart` function, the (cumulative) contribution plots will also be available from the created HTML page.

8 Illustrations

The *VRPM* package provides the user with several interactive applications that are created by means of the *shiny* package (Chang et al., 2015). Four applications are available. `IntermittentClaudication` to illustrate the functionality of the package on a binary logistic regression model. This application uses the model that was created by the Framingham Heart Study Group (Murabito et al., 1997). The application `xsnz` uses the *VRPM* package for multinomial risk prediction models on data from

the *VGAMdata* package (Yee, 2010). The `Iris` and `Pima` applications are available to illustrate the working for SVMs. Both applications use data from the UCI machine learning repository (Lichman, 2013; Fisher, 1936; Venables and Ripley, 2002a). To start the applications, run the following code after replacing `name` by the name of the application that needs to be launched. The application will then launch from within your browser.

```
> runVRPMexample("name")
```

9 Summary

The functionality of the *VRPM* package was illustrated for binary and multinomial logistic regression, Cox proportional hazard regression and support vector classifiers. The advantage of the color based nomograms for interactions and non-linear predictor effects was particularly useful to create nomograms for support vector machines. It was additionally shown how all nomograms for multinomial logistic regression models can be combined and presented in an appropriate way. Charts summarizing the risk prediction process for specific observations are provided as well. Finally, applications illustrating the use of the package are provided to make the interested user familiar with the possibilities of the package without the need to learn all functionality from the start.

10 Acknowledgements

V. Van Belle is a postdoctoral fellow of the Research foundation Flanders (FWO). This research was supported by: Center of Excellence (CoE): PFV/10/002 (OPTEC); iMinds Medical Information Technologies; Belgian Federal Science Policy Office: IUAP P7/19/ (DYSCO, 'Dynamical systems, control and optimization', 2012-2017); European Research Council: ERC Advanced Grant, (339804) BIOTENSORS. This paper reflects only the authors' views and the Union is not liable for any use that may be made of the contained information.

References

I. Ardoino, M. Lanzoni, G. Marano, P. Boracchi, E. Sagrini, A. Gianstefani, F. Piscaglia, and E. M. Biganzoli. Widen nomogram for multinomial logistic regression: an application to staging liver fibrosis in chronic hepatitis c patients. *Statistical Methods in Medical Research*, 2014. doi: 10.1177/0962280214560045. URL `http://smm.sagepub.com/content/early/2014/11/19/0962280214560045.abstract`.

W. Chang, J. Cheng, J. Allaire, Y. Xie, and J. McPherson. *shiny: Web Application Framework for R*, 2015. URL `http://CRAN.R-project.org/package=shiny`. R package version 0.12.2.

D. R. Cox. Regression models and life-tables (with discussion). *Journal of the Royal Statistical Society, Series B*, 34(2):187–220, 1972.

R. A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7 (Part II):179–188, 1936.

F. Harrell. *Regression Modeling Strategies. With applications to linear models, logistic regression, and survival analysis*. Springer Series in Statistics. Springer, New York, USA, 2001.

F. E. Harrell, Jr. *rms: Regression Modeling Strategies*, 2015. URL `http://CRAN.R-project.org/package=rms`. R package version 4.3-0.

A. Iasonos, D. Schrag, G. V. Raj, and K. S. Panageas. How to build and interpret a nomogram for cancer prognosis. *Journal of Clinical Oncology*, 26(8):1364–1370, 2008. doi: 10.1200/JCO.2007.12.9791. URL http://jco.ascopubs.org/content/26/8/1364.abstract.

A. Jakulin, M. Možina, J. Demšar, I. Bratko, and B. Zupan. Nomograms for visualizing support vector machines. In *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*, KDD '05, pages 108–117, New York, NY, USA, 2005. ACM. ISBN 1-59593-135-X. doi: 10.1145/1081870.1081886. URL http://doi.acm.org/10.1145/1081870.1081886.

A. Karatzoglou, A. Smola, K. Hornik, and A. Zeileis. kernlab – an S4 package for kernel methods in R. *Journal of Statistical Software*, 11(9):1–20, 2004. URL http://www.jstatsoft.org/v11/i09/.

M. Lichman. UCI machine learning repository, 2013. URL http://archive.ics.uci.edu/ml.

J. M. Murabito, R. B. D'Agostino, H. Silbershatz, and W. F. Wilson. Intermittent claudication: A risk profile from the framingham heart study. *Circulation*, 96(1):44–49, 1997.

original by Gareth Ambler and modified by Axel Benner. *mfp: Multivariable Fractional Polynomials*, 2015. URL http://CRAN.R-project.org/package=mfp. R package version 1.5.1.

J. C. Platt. Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods. In A. J. Smola, P. Bartlett, B. Schoelkopf, and D. Schuurmans, editors, *Advances in large margin classifiers*, pages 61–74, Cambridge, MA, USA, 1999. MIT Press.

P. Royston and W. Sauerbrei. *Multivariable Model - Building: A Pragmatic Approach to Regression Analysis Based on Fractional Polynomials for Modelling Continuous Variables*. Wiley Series in Probability and Statistics. Wiley, 2008. ISBN 9780470770788. URL http://books.google.be/books?id=rMzOuhz7CEcC.

V. Van Belle. *VRPM: Visualizing Risk Prediction Models*, 2015. URL http://CRAN.R-project.org/package=??? R package version 1.1.

V. Van Belle and B. Van Calster. Visualizing risk prediction models. *PLoS ONE*, 10(7):e0132614, 2015. doi: 10.1371/journal.pone.0132614.

V. Van Belle, B. Van Calster, J. A. K. Suykens, S. Van Huffel, and P. Lisboa. Explaining support vector machines: a color based nomogram. Technical report, 16-21, ESAT-Stadius, KU Leuven (Leuven, Belgium), a.

V. Van Belle, S. Van Huffel, D. Timmerman, and B. Van Calster. A color based nomogram for multinomial logistic regression. Technical report, 16-28, ESAT-Stadius, KU Leuven (Leuven, Belgium), b.

V. Vapnik. *Statistical Learning Theory*. Wiley and Sons, New York, 1998.

W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S*. Springer, New York, fourth edition, 2002a. URL http://www.stats.ox.ac.uk/pub/MASS4. ISBN 0-387-95457-0.

W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S*. Springer, New York, fourth edition, 2002b. URL http://www.stats.ox.ac.uk/pub/MASS4. ISBN 0-387-95457-0.

T. W. Yee. The vgam package for categorical data analysis. *Journal of Statistical Software*, 32 (10):1–34, 1 2010. ISSN 1548-7660. URL http://www.jstatsoft.org/v32/i10.