

5-2011

Assessment of Aortic Aneurysm Rupture Risk

Rafael Izbicki

Carnegie Mellon University, rizbicki@andrew.cmu.edu

Ann B. Lee

Carnegie Mellon University, annlee@stat.cmu.edu

Ender A. Finol

Carnegie Mellon University

Follow this and additional works at: <http://repository.cmu.edu/statistics>



Part of the [Statistics and Probability Commons](#)

This Technical Report is brought to you for free and open access by the Dietrich College of Humanities and Social Sciences at Research Showcase @ CMU. It has been accepted for inclusion in Department of Statistics by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Assessment of Aortic Aneurysm Rupture Risk

Rafael Izbicki, Ann B. Lee and Ender A. Finol

Carnegie Mellon University

May 2011

Abstract

The rupture of an abdominal aortic aneurysm (AAA) is associated with a high mortality. When an AAA ruptures, 50% of the patients die before reaching the hospital. Of the patients that are able to reach the operating room, only 50% have it successfully repaired [1]. Therefore, it is important to find good predictors for immediate risk of rupture. Clinically, the size of the aneurysm is the variable vascular surgeons usually use to evaluate this risk. Patients with large aneurysms are often sent to surgery. However, many studies have shown that even small aneurysms can rupture and deserve attention as well [2]. It is important to find good predictors of rupture that also avoid unnecessary surgery as all surgeries are associated with possible complications. Here, we use data obtained from 144 computed tomographies of patients from the Western Pennsylvania Allegheny Health System to predict the high risk of rupture of an aneurysm and also to examine which features are important for this goal.

1 Introduction

Abdominal aortic aneurysm (AAA) is the medical term used when the abdominal aorta becomes extremely large (technically, with a diameter of more than 3cm). The problem with AAAs is that their rupture is often lethal; many patients die within minutes. Patients that have AAA can undergo a surgery to have it repaired. However, this surgery is itself risky, so that only patients that are considered to have a high risk of rupture are encouraged to undergo it. Most vascular surgeons use the diameter of the aneurysm as the criteria for deciding if a patient has to undergo surgery. When the diameter is greater than 5.5cm [7], the patients are considered for repair. Patients that already have the aneurysm ruptured or have abdominal pain are also sent to surgery.

The two main goals of this study are:

1. predict the high risk of rupture of the abdominal aortic aneurysm
2. understand which covariates are important when predicting the high risk of rupture

In order to do so, we have a data base that consists of patients of two groups:

1. "elective repair" (100 patients)
2. "emergent repair" (44 patients)

Elective repair patients are individuals with large aneurysms that were suggested to undergo surgery, even though they did not display an immediate sign of rupture. On the other hand, patients from the emergent group were sent to surgery either because they had symptomatic pain or because they already had the aneurysm ruptured. All patients from this study were treated in Western Pennsylvania Allegheny Health System.

The covariates used in this study were calculated from three-dimensional models that were generated from segmented, contrast-enhanced computed tomography images [7]. There are in total 28 covariates, which include 1-dimensional indices such as diameters and heights, 2-dimensional indices such as measures of tortuosity and asymmetry, 3-dimensional indices such as the volume and the surface area as well as measurements of wall thickness at different parts of the aneurysm (a grid was created to do so). The information about thickness was condensed by using four different statistics: minimum, maximum, average and average thickness at maximum diameter. In the Appendix, one can find all variables used together with their abbreviations used here. For the formulae and more details of these quantities, see [7]. Figure 8 in the Appendix shows some of these features.

The idea of the study is to use this database to evaluate the risk of a patient having his aneurysm ruptured. We also want to understand which variables are good predictors of this event. Note that there is one major assumption here, namely that the variables which are measured do not change when the rupture occurs. In future work we plan to explore how reasonable this assumption is in depth. We also emphasize that the ruptured aneurysms collected for this work had a contained rupture, which makes this assumption less critical.

2 Methods

Most traditional methods used in machine learning [3] for classification assume that:

- The cost (loss) of classifying one patient as emergent when he is from the elective group and the cost of classifying one patient as elective when he is from the emergent group are both the same.
- The observations in a sample are independent and identically distributed (iid).

In the present problem, none of these assumptions are reasonable. The loss incurred when an emergent patient is misclassified is greater than the one when this patient is from the elective group.

Furthermore, this study was performed in a case-control scheme (retrospective study): the number of patients in each group (44 and 100) was fixed prior to data collection. Using data from past years from the Western Pennsylvania Allegheny Health System (Table 2), one can see that these proportions are quite different from what is expected in the target population. In fact, one can see that the proportions of emergent patients is about 11.4% per year for all the years presented in the table, rather than the expected percentage of 30.6% in the sample. Hence, the iid assumption is also not reasonable.

Table 1: Number of patients with abdominal aortic aneurysm in the Western Pennsylvania Allegheny Health System

	2005	2006	2007	2008	2009	Total
Elective	131	138	141	114	125	649
Emergent	20	15	16	18	15	84

The two questions that arise when dealing with this nonstandard scenario are:

- How can we evaluate if a given (fit of a) model is good, i.e., how do we define and estimate an appropriate risk function?
- How can we produce models that give good results?

As we will see, the answer to the first question is to use a modified version of the traditional cross-validation score. The second question can be resolved by adapting a traditional method in machine learning to this scenario, namely logistic regression. The use of logistic regressions was also motivated because their results are easy to interpret. In this work, the need of interpretability of the results is undeniable.

In the discussion that follows, we will make use of the following notations:

- p_0 : probability in the target population of a patient being from the Elective group
- p_1 : probability in the target population of a patient being from the Emergent group
- l_0 : cost (loss) of classifying a patient from Elective group as being from Emergent group
- l_1 : cost (loss) of classifying a patient from Emergent group as being from Elective group
- Y : group of the patient (0 for Elective and 1 for Emergent)
- X : explanatory variables (features)

Given a classifier $h : X \rightarrow Y$, the expected loss function (risk) for a new observation X is

$$R(h) = l_1 P(h(X) = 0, Y = 1) + l_0 P(h(X) = 1, Y = 0) = l_1 p_1 P(h(X) = 0|Y = 1) + l_0 p_0 P(h(X) = 1|Y = 0).$$

In order to estimate R , we randomly split the data set in 12 folds with 12 observations each. For each individual i , let

$$\hat{Y}_i = \hat{h}_{-i}(X_i),$$

where \hat{h}_{-i} is the fitted classifier using all observations but the ones in the fold for which i belongs to. In this work, we refer to \hat{Y}_i as the cross-validated label for observation i .

The estimate of the risk of a fit of classifier \hat{h} we use here is given by

$$\hat{R}(\hat{h}) = l_1 \hat{p}_1 \hat{P}(\hat{h}(X) = 0|Y = 1) + l_0 \hat{p}_0 \hat{P}(\hat{h}(X) = 1|Y = 0),$$

where

$$\hat{P}(\hat{h}(X) = 0|Y = 1) = \frac{\sum_{i:Y_i=1} I(\hat{Y}_i = 0)}{\sum_{i:Y_i=1} 1}, \quad \hat{P}(\hat{h}(X) = 1|Y = 0) = \frac{\sum_{i:Y_i=0} I(\hat{Y}_i = 1)}{\sum_{i:Y_i=0} 1},$$

and \hat{p}_0 and \hat{p}_1 a estimates of p_0 and p_1 , respectively.

Let P denote the real probability measure, and P_N denote the probability measure under the "case-control scheme". We assume that $P(x|y) = P_N(x|y)$, that is, the conditional distribution of the covariates, conditional on the group of the patients, are the same in both sample schemes. This condition is known as prior probability shift [6].

The Bayes classifier classifies a patient as Emergent if, and only if,

$$\begin{aligned} P(Y = 1|x)l_1 > P(Y = 0|x)l_0 &\iff \frac{P(Y = 1|x)}{P(Y = 0|x)} > \frac{l_0}{l_1} \iff \frac{P(x|Y = 1)P(Y = 1)}{P(x|Y = 0)P(Y = 0)} > \frac{l_0}{l_1} \iff \\ \frac{P(x|Y = 1)}{P(x|Y = 0)} > \frac{l_0 P(Y = 0)}{l_1 P(Y = 1)} &\iff \frac{P(x|Y = 1)P_N(Y = 1)}{P(x|Y = 0)P_N(Y = 0)} > \frac{l_0 P(Y = 0)P_N(Y = 1)}{l_1 P(Y = 1)P_N(Y = 0)} \iff \\ \frac{P_N(Y = 1|x)}{P_N(Y = 0|x)} > \frac{l_0 P(Y = 0)P_N(Y = 1)}{l_1 P(Y = 1)P_N(Y = 0)} &\iff P_N(Y = 1|x) > \frac{\frac{l_0}{l_1 a}}{1 + \frac{l_0}{l_1 a}}, \end{aligned}$$

where $a = \frac{P_N(Y=0)P(Y=1)}{P_N(Y=1)P(Y=0)}$. The term $P_N(Y = 1|x)$ can be estimated using a procedure such as a logistic regression, and the right side can be estimated using the data from the Hospital. In this way, we have a plug-in classifier for this classification problem. This derivation can also be found in [5].

Proceeding in an analogous way, we can estimate $P(Y = 1|x)$:

$$\frac{P(Y = 1|x)}{P(Y = 0|x)} = \frac{P_N(Y = 1|x) P_N(Y = 0)P(Y = 1)}{P_N(Y = 0|x) P_N(Y = 1)P(Y = 0)}.$$

Hence we have

$$P(Y = 1|x) = \frac{a \frac{P_N(Y=1|x)}{P_N(Y=0|x)}}{1 + a \frac{P_N(Y=1|x)}{P_N(Y=0|x)}}$$


in which the right term can be estimated by plugin of the results of a logistic regression. This derivation is also helpful in the case of other plugin classifiers, not only logistic regressions models.

Whenever necessary here, we used the value $p_1 = 84/733$. This value was chosen based on the stability shown in Table 2 between the years. We also decided to use the values $l_0 = 1$ and $l_1 = 7.72$ for the loss functions. This was motivated because it is worse to misclassify an emergent patient than an elective one, as discussed previously. Moreover, for this choice, the risk is proportional to $P(h(X) = 0|Y = 1) + P(h(X) = 1|Y = 0)$, that is, it is just the sum of the conditional probabilities of mistake. The optimal cutoff for the Bayes classifier $(\frac{\frac{l_0}{l_1 a}}{1 + \frac{l_0}{l_1 a}})$ in this case is 0.306.

Most of the models we fitted were based on logistic regressions. Rather than using the cutoff derived, we decided to use the cutoff as a tuning parameter: for each model, we chose the threshold that minimizes the empirical risk. It is worth noting that, in fact, this value was in general close to the one previously derived, as it is going to be shown in the results. Some non-logistic methods were used in order to check the robustness of the models.

3 Results

In order to describe the statistical variations of the covariates described in the Appendix, we plotted the violinplots [4] for each one of them. Violinplot are boxplots that also contain an estimate of the density of the covariate in each of the groups. Figure 1 presents the violinplots for the covariates. They are sorted by the p values obtained in a Mann-Whitney test, which are also presented.






Figure 1: Violin Plots for the variables derived from the computed tomography

In Figure 1, we see that many variables are important predictors of high risk of rupture when considered separately, especially measures related to the aneurysm's size (such as length, volume and maximum diameter). In order to understand the relationship between them, Figure 2 shows the scatter plots of the 16 most significant variables on the Mann-Whitney test that was presented.

We see that many of the covariates are very correlated to each other. In what follows we try to consider them jointly by using different models for classification. The following models were implemented (for a reference about these models, see [3]):

- (Sparse) L_1 Regularized Logistic Regression (that is, choosing the coefficients that maximize $\sum y_i(\beta_0 + \beta^t x_i) - \log(1 + \exp(\beta_0 + \beta^t x_i)) - \lambda \sum |\beta_i|$).
- (SparseL) Maximum Likelihood Logistic Regression with the variables with nonzero coefficients in the Sparse model
- (PC) Maximum Likelihood Logistic Regression with the first 5 Principal Components of the variables
- (PCGAMs) Logistic Additive Model with the first 5 Principal Components of the variables, using smoothing splines (that is, the mean is modeled as $\log \frac{\mu(X)}{1-\mu(X)} = \alpha + s(P_1) + \dots + s(P_5)$)
- (PCGAMl) Logistic Additive Model with the first 5 Principal Components of the variables, using local polynomials (1st order) (that is, the mean is modeled as $\log \frac{\mu(X)}{1-\mu(X)} = \alpha + l(P_1) + \dots + l(P_5)$)
- (AIC) Maximum Likelihood Logistic Regression with AIC criterion and stepwise selection (AIC is $2 \frac{\sum y_i(\beta_0 + \beta^t x_i) - \log(1 + \exp(\beta_0 + \beta^t x_i))}{N} - 2 \frac{d}{N}$, where p is the number of covariates of the model and N is the sample size)
- (BIC) Maximum Likelihood Logistic Regression with BIC criterion and stepwise selection (BIC is $\sum y_i(\beta_0 + \beta^t x_i) - \log(1 + \exp(\beta_0 + \beta^t x_i)) - (\log N)d$)
- (AICPC) Maximum Likelihood Logistic Regression on the Principal Components with AIC criterion and stepwise selection

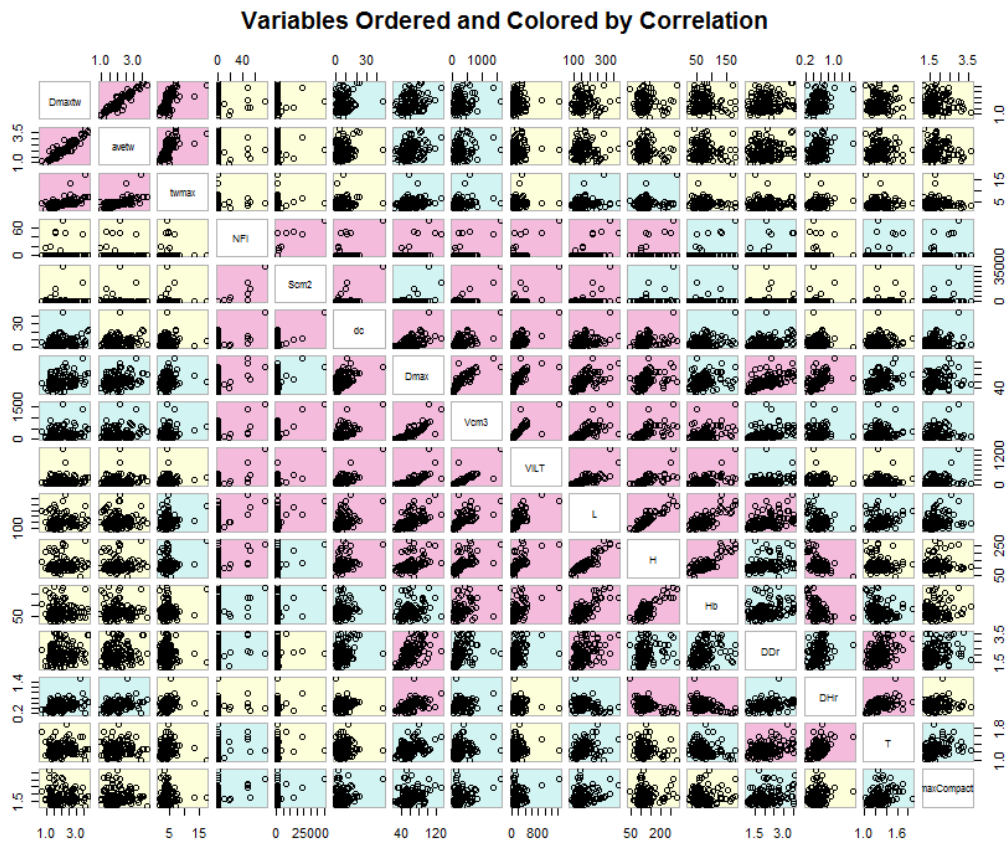


Figure 2: Scatter plots of the 16th univariately most significant variables

- (RF) Random Forest (which is an average of de-correlated trees)

Figure 3 shows the (cross-validated) estimates of the conditional misclassification probabilities as well as the estimate of the risk (expected cost) for the different models and different values of the cutoffs, as discussed in the previous section. Note that for the PCGAMl model, we have the bandwidth shown on the x axis rather than the cutoff for the logistic regression. For each value of bandwidth, we chose the optimal cutoff. Also, for the RF model, note that the same derivation from the previous section applies, once this is also a plugin estimator. For the sake of comparison, a vertical line is always displayed at the cutoff 0.306, the optimal cutoff derived on the previous section. The cross-validated labels were obtained by using a random 12-fold cross-validation [3].

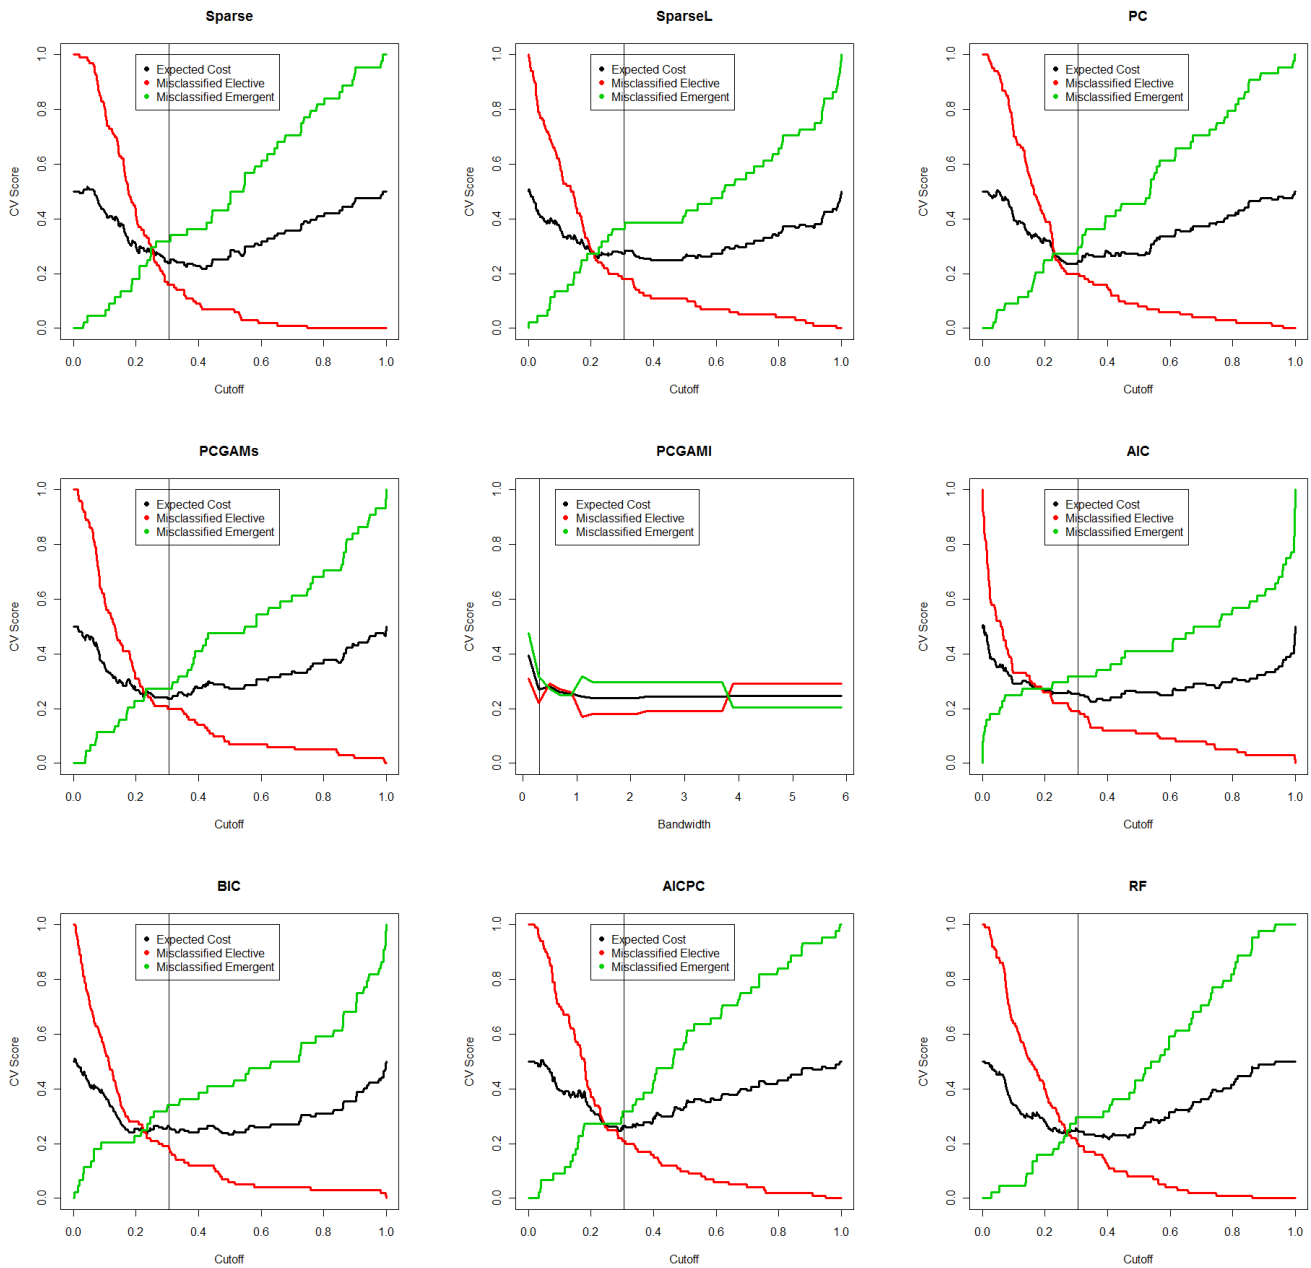


Figure 3: Results for the estimates of the risk/conditional probabilities for different models and different values of the cutoff on the logistic regression (and bandwidth for PCGAMI, see text for details). The vertical line represents the optimal cutoff according to Bayes' rule.

In order to evaluate the accuracy of the cross-validation results presented in Figure 3, we plotted the bootstrap 90% confidence intervals for each of the scores in each of the models for the optimal cutoffs. These results are presented in Figure 4. This Figure also contains another model, Support Vector Machine (SVM), which was used with the purpose of checking the robustness of our conclusions.




Figure 4: 90% bootstrap confidence intervals for the risk and conditional probabilities of mistakes in each model

We see that the performances of the models are very similar, except SVM. This is expected because no special adaptations were created to deal with different loss functions and different probabilities. We also notice that the confidence intervals are, in general, very large, indicating a certain instability of the fitted models. Table 3 shows the scores for the models.

Table 2: Summaries (conditional probabilities and risk) of the fitted models.

Model	Missclassified Elective	Missclassified Emergent	Expected Cost
Sparse	0.070	0.364	0.217
SparseL	0.110	0.386	0.248
PC	0.200	0.273	0.236
PCGAMs	0.200	0.273	0.236
PCGAMI	0.180	0.295	0.238
AIC	0.130	0.318	0.224
BIC	0.060	0.409	0.235
AICPC	0.220	0.273	0.246
RF	0.180	0.295	0.238
SVM	0.090	0.410	0.250

Figure 5 presents the ROC curves (for different cutoff values) of the fitted models which have a continuous output. Again, one can see that all of the models are very similar. In this figure, we also used cross-validated labels.




Figure 5: ROC curves for the models with continuous output (with the cross-validated labels).

Although the estimated risks of the models are very similar, we can ask ourselves if, in each of the models, we are misclassifying the same observations or not. Figure 6 shows, for each observation and each model, a dot in case this observation was misclassified on that model. The intensity of the color is proportional to the number of models in which the observation was misclassified. We also added a 2-means cluster analysis, using all covariates. Note that the observations to the right of the vertical line are patients from the emergent group, while the ones to the left are the ones from the elective group. The labels used were again the cross-validated ones.


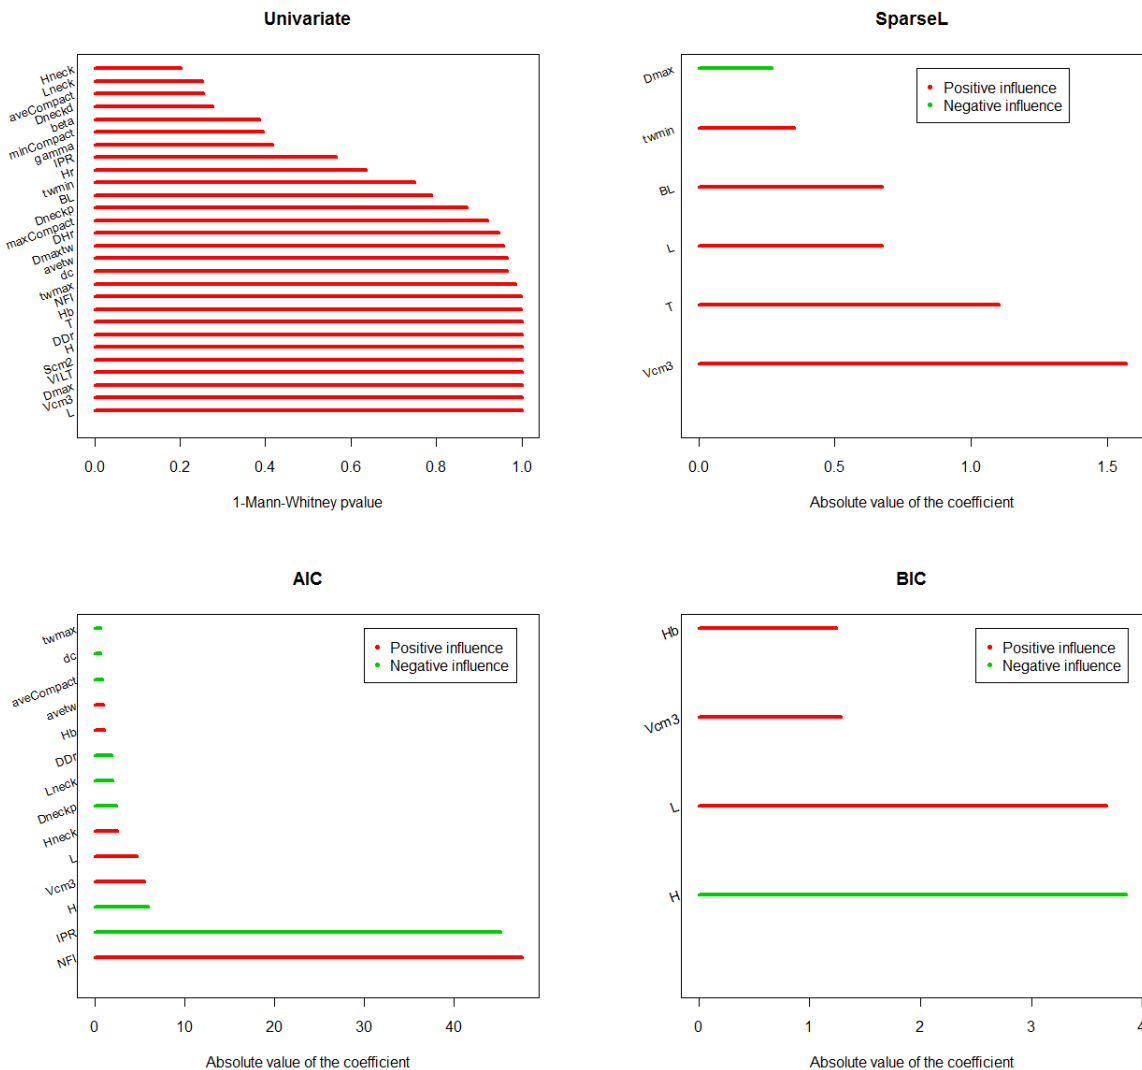


Figure 6: Cross-validated misclassifications in each models. Color intensity is proportional to the number of misclassifications for the patient. First 100 patients are elective, last 44 are emergent.

We can see that most misclassified patients have dots with dark colors. This indicates that most misclassified

patients are being misclassified in all fitted models (even on the cluster analysis, which doesn't use the labels to train the classifier). In future work, we intend to investigate which characteristics are shared by these patients.

Figure 7 shows the importance of each variable according to some of the models presented. The univariate importance was measured in terms of the Mann-Whitney pvalues, while the importance of the logistic models was measured in terms of the coefficients of the covariates after they were standardized. Also, for the Random Forest model, we calculated the importance of each variable based on the Gini splitting index [3]. The greater the Gini index, the greater the variable importance is. In the graphics associated to maximum likelihood logistic regression, a red line represents that the chance that the patient is from the emergent group increases as the covariate increases (conditional on the others). On the other hand, a green line indicates that the chance that the patient is from the emergent group decreases as the covariate increases.






Figure 7: Variable importance for univariate Mann-Whitney tests, SparseL, AIC, BIC and RF models, respectively.

It is interesting to note that some of these variables, like twmin, are not important univariately. Also, many of the important variables are related to the size of the aneurysm (as for instance L, H and Vcm3). The summaries of Wall Thickness that were considered to be important by this models were the minimum and the maximum. Tortuosity is one variable that is not directly related to the size of the aneurysm but seems to be very important.

4 Conclusions and future work

In this work, we see that the introduction of a different loss functions (that reflect that it is worse to misclassify patients from the Emergent group) and information regarding prior probabilities of each group on the target population allow us to adapt the traditional logistic regression to a new context. Moreover, it leads us to the idea of using the cutoff of the logistic regression as a tuning parameter. The same derivations work for other plugin methods.

We saw that the confidence intervals for the risks were large in the fitted models. This leads us to believe that the use of even more complex models would not be helpful for current sample sizes.

The fitted models suggest that with the covariates used we are able to do a fair work in classifying patients from the elective group, however more informative features would be desirable to classify correctly patients from emergent group. It is worth noting that the inclusion of interactions in the models did not improve the estimated risks (though these results were omitted in this paper). All fitted models had a similar performance both in terms of the final estimate of their risk and also in terms of which patients were being misclassified.

Possible further investigations for this work include:

- Among the patients from Emergent group, compare the ones that were in pain when the computed tomography was taken and those who already had the aneurysm ruptured at that time. These comparisons would be useful to try to validate the assumption that the rupture of the aneurysm does not change the covariates used in this work.
- Try to understand what common characteristics patients that were misclassified share. Sociodemographic variables may help.
- Check whether the probabilities on the logistic regressions from misclassified individuals were close to the cutoff.
- Introduce a measure of confidence for the classifications, e.g. the brier score $\sum_i (y_i - \hat{y}_i)^2$, where y_i denotes the category of observation i and \hat{y}_i denotes the fitted value (score) for observation i .

- Focus on trying to find new good predictors rather than trying to fit complex models. Combining information on wall thickness with information on wall stress at the same location may yield good results.
- Increase the number of samples, especially from the Emergent group.
- Investigate the relationship between wall stress (still to be measured) and the geometric features used here. The hypothesis is that the wall stress is highly associated to the rupture of an aneurysm.

References

- [1] M. F. Fillinger, S. P. Marra, M. L. Raghavan, and F. E. Kennedy. Prediction of rupture risk in abdominal aortic aneurysm during observation: wall stress versus diameter. *Journal of Vascular Surgery*, 37(4):724–732, 2003.
- [2] M. F. Fillinger, M. L. Raghavan, S. P. Marra, Jack L. C., and Francis E. K. In vivo analysis of mechanical wall stress and abdominal aortic aneurysm rupture risk. *Journal of Vascular Surgery*.
- [3] T. Hastie, R. Tibshirani, and J. H. Friedman. *The elements of statistical learning: data mining, inference, and prediction: with 200 full-color illustrations*. New York: Springer-Verlag, 2001.
- [4] J. L. Hintze, K.S. Narendra, and D. N. Ray. Violin plots: A box plot-density trace synergism. *The American Statistician*, 52(2):181–84, 1998.
- [5] Y. Lin, Y. L. Yoonkyung, Y. Lee, and G. Wahba. Support vector machines for classification in nonstandard situations. In *Machine Learning*, pages 191–202, 2000.
- [6] J. Quionero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence. *Dataset Shift in Machine Learning*. The MIT Press, 2009.
- [7] J. Shum, G. Martufi, E. Di Martino, C. B. Washington, J. Grisafi, S. C. Muluk, and E. Finol. Quantitative Assessment of Abdominal Aortic Aneurysm Geometry. *Annals of biomedical engineering*, October 2010.

A Covariates

1-dimensional indices:

- Dmax: Maximum diameter
- Dneckp: Distal neck diameter
- Dneckd: Proximal neck diameter
- H: Height of AAA
- L: Length of AAA centerline
- Hneck: Height of neck
- Lneck: Length of neck centerline
- Hb: Bulge Height
- dc: Centroid distance of Dmax
- maxCompact: Maximum Compactness
- minCompact: Minimum Compactness
- aveCompact: Average Compactness

2-dimensional indices:

- DHr: Diameter-Height ratio
- DDr: Diameter-Diameter ratio
- Hr: Height ratio
- BL: Bulge location
- beta: Asymmetry
- T: Tortuosity

3-dimensional indices:

- Vcm3: AAA Volume
- Scm2: AAA Surface Area
- VILT: Intraluminal thrombus volume
- gamma: AAA sac to ILT volume ratio
- IPR: Isoperimetric Ratio
- NFI: Non-fusiform Index

Summaries of wall thickness:

- twmax: Maximum wall thickness
- twmin: Minimum wall thickness
- avetw: Average wall thickness
- Dmaxtw: Maximum wall thickness at Dmax




Figure 8: Visual illustration of some of the measures. Bottomleft represents the wall stress of one patient at different locations. (Courtesy of Ender A. Finol and Judy Shum)