



Statistical evaluation of the Predictive Toxicology Challenge 2000–2001

Hannu Toivonen^{1,*}, Ashwin Srinivasan², Ross D. King³,
Stefan Kramer^{4,5} and Christoph Helma⁴

¹Department of Computer Science, PO Box 26 (Teollisuuskatu 23), FIN-00014 University of Helsinki, Finland, ²Oxford University Computing Laboratory, Wolfson Building, Parks Road, Oxford OX1 3QD, UK, ³Department of Computer Science, Penglais, Aberystwyth, Ceredigion SY23 3DB, Wales, UK, ⁴Institut Für Informatik, Albert-Ludwigs-Universität Freiburg, Georges-Köhler-Allee, Gebäude 079, D-79110 Freiburg i. Br., Germany and ⁵Institute for Computer Science, Technical University of Munich, Boltzmannstr. 3, D-85748 Garching b. Muenchen, Germany

Received on July 30, 2002; revised on November 7, 2002; accepted on November 11, 2002

ABSTRACT

Motivation: The development of *in silico* models to predict chemical carcinogenesis from molecular structure would help greatly to prevent environmentally caused cancers. The Predictive Toxicology Challenge (PTC) competition was organized to test the state-of-the-art in applying machine learning to form such predictive models.

Results: Fourteen machine learning groups generated 111 models. The use of Receiver Operating Characteristic (ROC) space allowed the models to be uniformly compared regardless of the error cost function. We developed a statistical method to test if a model performs significantly better than random in ROC space. Using this test as criteria five models performed better than random guessing at a significance level p of 0.05 (not corrected for multiple testing). Statistically the best predictor was the Viniti model for female mice, with p value below 0.002. The toxicologically most interesting models were Leuven2 for male mice, and Kwansai for female rats. These models performed well in the statistical analysis and they are in the middle of ROC space, i.e. distant from extreme cost assumptions. These predictive models were also independently judged by domain experts to be among the three most interesting, and are believed to include a small but significant amount of empirically learned toxicological knowledge.

Availability: PTC details and data can be found at: <http://www.predictive-toxicology.org/ptc/>

Contact: hannu.toivonen@cs.helsinki.fi

INTRODUCTION

The *Predictive Toxicology Challenge (PTC)* was initiated to stimulate the development of advanced techniques for

predictive toxicology models. The goal was to provide carcinogenicity predictions for a set of compounds with unknown classification, using information derived from their chemical structure alone. The challenge was taken up by 14 groups who produced 111 predictive models. The previous article (Helma and Kramer, 2003) presents the general layout of the competition, the data sets (Table 1) and the rules for participation. The focus of this paper is a statistical analysis of the significance of the submitted predictions.

METHODS

ROC curves

Ideally an *in silico* predictive carcinogenicity method should correctly identify all carcinogenic compounds while not incorrectly identifying any non-carcinogenic compounds. However, in practice, predictive results may overlap and errors be made. Two types of error are possible: errors of commission, and errors of omission. In an error of commission a compound is predicted to be carcinogenic when it is not: in an error of omission, a compound is not predicted to be carcinogenic when it is. The costs associated with these two types of error are not generally equal. Such costs are difficult to calculate and subject to considerable debate. This means that the assessment of predictive models should be robust to differing cost assumptions. This is not true for the standard approaches to measuring predictive success: use of accuracy, or maximization of coverage for a fixed false positive rate.

A better approach is to plot the predictive results in Receiver Operating Characteristic (ROC) space. ROC graphs were first developed for signal detection (e.g. Van Trees, 1971) and later extended to machine learning,

*To whom correspondence should be addressed.

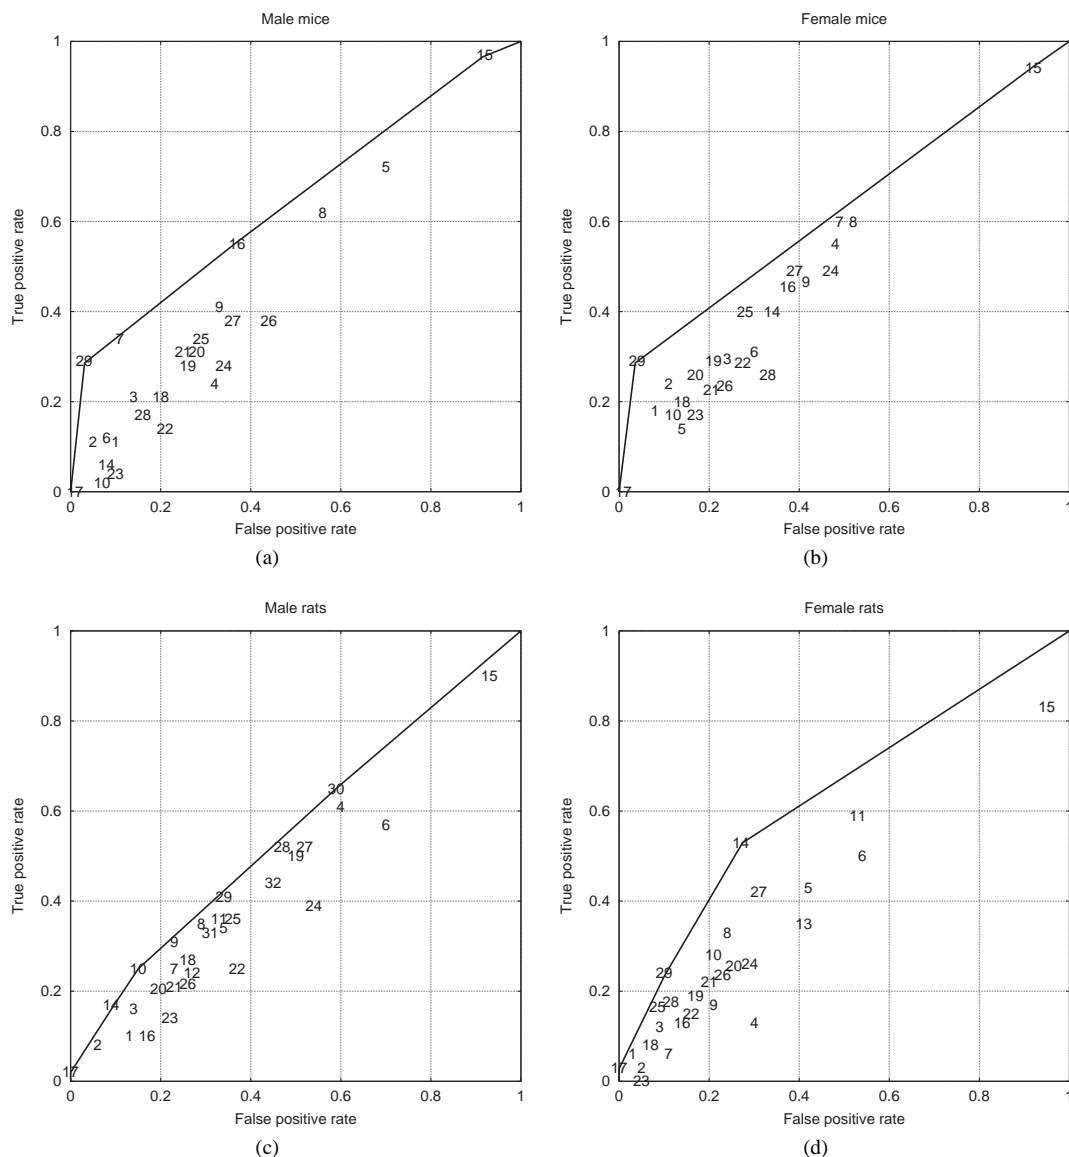


Fig. 1. ROC points and convex hulls of the submissions.

e.g. by Bradley (1995). In a ROC plot true positive rate (or sensitivity) is plotted against the false positive rate (1-specificity). Sensitivity is the probability that a compound is predicted to be carcinogenic when it is actually carcinogenic. Specificity is the probability that a compound is predicted to be non-carcinogenic when the compound is actually non-carcinogenic. Both measures are expressed in the range 0–1. This produces a square space ranging from 0 to 1 along the two axes, or unit square. This space is called the ROC space. An ideal classifier, resulting from a perfect discrimination between carcinogenic and non-carcinogenic compounds, would be presented in ROC space as a point in the top left-hand corner.

If a range of predictive results for different models are plotted in ROC space, the convex-hull made from points to the extreme top and left (see Fig. 1) defines the best available predictors. The closer the curve follows the left-hand border and the top border of the ROC space, the more accurate are the predictions made. In general a ROC curve indicates the trade-off between sensitivity and specificity, as an increase in sensitivity is accompanied by a decrease of specificity. If a predictor is not on the ROC convex-hull it is sub-optimal compared to the ones on the hull, regardless of the particular costs associated with errors of commission and omission (assuming cost is a linear function of the errors; Provost and Fawcett, 1997).

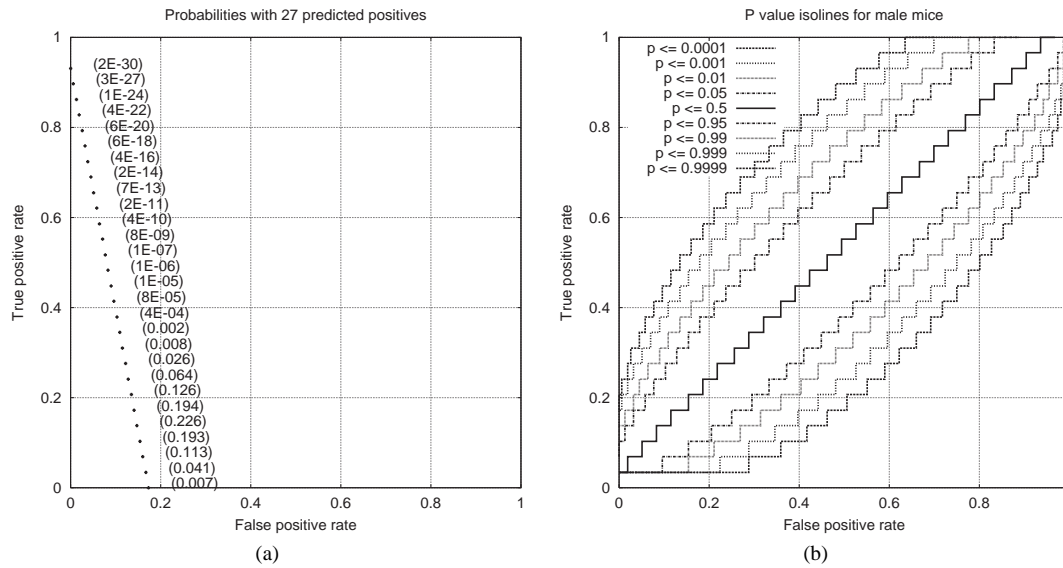


Fig. 2. (a) Probabilities of classifiers that make 27 positive predictions in the male mice test set under the null hypothesis of random predictions. (b) P value isolines for the male mice test set.

Statistical significance in ROC space

One problem with the use of convex-hulls in ROC space is that they do not include any measure of statistical significance. In this paper we adapt a standard statistical test to judge if a result in ROC space is significantly better than random predictions. The rationale for our p values in the ROC space is as follows. For a classifier C with N_C predicted positives, the null hypothesis is that the selection of the N_C examples was statistically independent of their true class. The p value of classifier C then expresses the probability with which random selection of N_C predicted positives would give at least as good a result as the one obtained by C .

Figure 2a gives, as an example, the distribution of points in the ROC space for a classifier that makes 27 positive predictions (such as the Baurin_en model) for the male mice test set (156 negatives and 29 positives). The points are along a slanted line, with the most likely points closest to the diagonal. The p value of a classifier on this line is obtained by summing the probabilities above and at the same point with the classifier. In statistical terms, this corresponds to the one-tailed statistical significance of the 2×2 contingency table (confusion matrix) describing the classifier performance in terms of the number of true and false positives and true and false negatives. (Chi-square is a well known tool for this; we obtained the exact p values using Fisher's exact test.) By computing the p values for all N_C from 0 to 185 (the number of examples in the test set), one obtains p values over the whole ROC space, i.e. for all possible classifiers.

The p value isolines can be overlaid with the ROC space, to give an overview of how the p value behaves in relation to ROC. Figure 2b shows the p value isolines for the male mice data set. Since the p values also depend on the size and the class distribution of the data set, p values are different for different data sets and also for different amounts of unclassified cases. (Some submitted predictors left some test cases unclassified.)

The p value isolines are not symmetrical since the class distribution is skewed, cf. Figure 2a. Further, the isolines have been drawn at the points where the p value is at most the one given in the label. For instance, 0.5 p value isoline is drawn above the diagonal since points in the diagonal tend to have p values larger than 0.5.

RESULTS

We plotted each of the 111 predictive models in ROC space, see Figure 1. In these figures the convex-hull of the best predictive models is displayed. We tested all the PTC submissions to see if they were significantly better than random. For this the statistical significance, expressed as a p value, was computed for the submissions in ROC space.

Figure 3 shows the predictions again in the ROC space, this time overlaid with p value isolines. All classifiers are shown, but those that did not make predictions for all cases are in parentheses. Their p values are more conservative than is apparent from the graphs, due to the smaller number of cases predicted.

Table 2 lists the statistically most significant results. Five submissions have p values below 0.05 (not adjusted

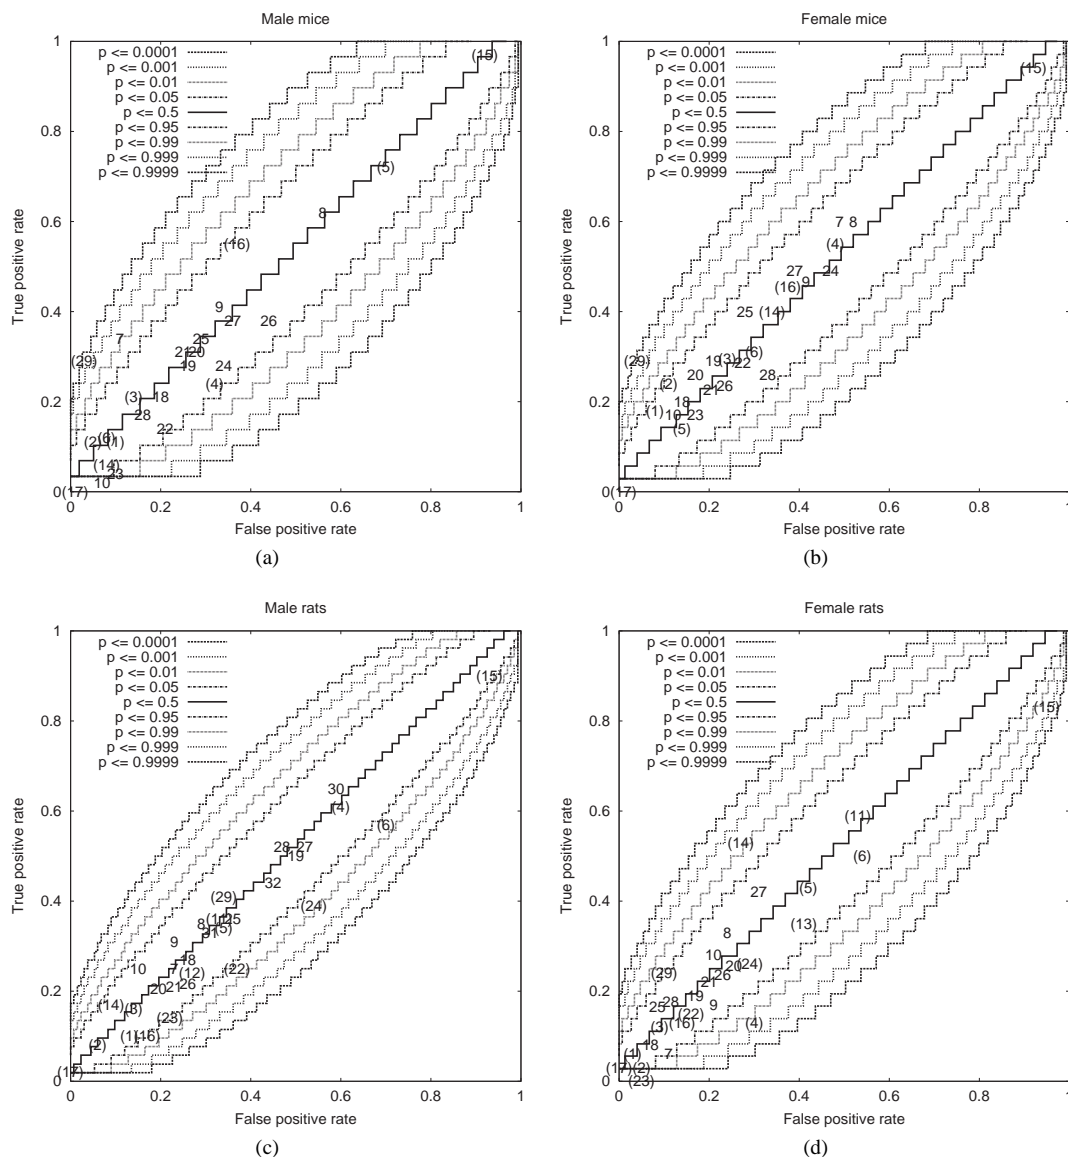


Fig. 3. *P*-value isolines and PTC submissions. Labels are the same as in Figure 1. (Parenthesis indicate classifiers that did not make predictions for all examples and for which the true *p* values are thus more conservative than appears from the figure.)

for multiple testing), three of them are below 0.005[†]. Statistically the most significant predictors are clearly Viniti on female mice, and Baurin_en and Viniti on male mice. They obviously perform significantly better than random guessing. This is arguably the first conclusive proof that it is possible to empirically learn *in silico* models which can predict chemical carcinogenicity based purely on chemical structure.

[†] In the PTE1 three from 15 models had *p*-values lower than 0.05 and one model was below 0.005, in the PTE2 three from 18 had *p*-values lower than 0.05, none of them was lower than 0.005. In both cases it was allowed to take advantage of biological information (e.g. short-term tests).

The majority of the significant models cluster near the bottom left corner. Such models only make few positive predictions, but they are accurate in their predictions. Among the five statistically best predictors, Leuven2 on male mice and Kwansai on female rats, are particularly interesting as they occur in the middle of ROC space and distant from a corner. This means that these methods are optimal for non-extreme cost assumptions; the predictors are both quite accurate and applicable to a number of examples.

Given that 111 submissions were received in total, it is interesting to look at their overall performance. How

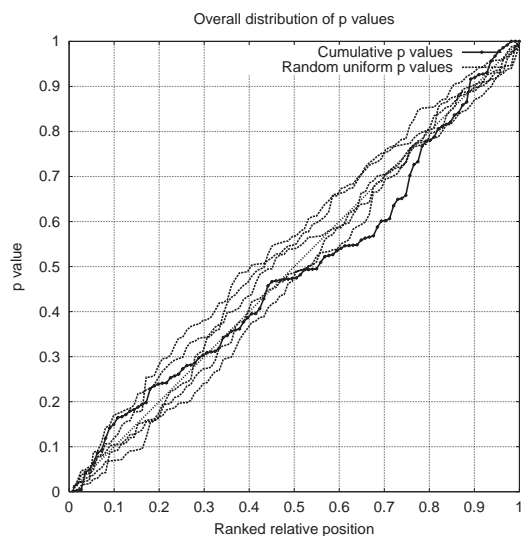


Fig. 4. Cumulative p value distribution of all 111 submissions. The dotted lines, shown for comparison, are each based on 111 i.i.d. uniformly distributed random values in $[0, 1]$.

significant, overall, are the results? Could the best results have been obtained by chance alone? The Bonferroni adjustment could be used to test this if the classifiers were independent. However, they are clearly not independent, since they had the same features and training examples available to them, some of them used similar methods, and they were tested on the same test cases. The Bonferroni adjusted p value for the statistically most significant result (Viniti, female mice) would be 0.206, but due to the contradicted independence assumption the value serves merely as a very conservative upper bound for the adjusted p value of Viniti.

A visual overview of the statistical significance of the whole set of 111 submissions is obtained by looking at the distribution of their (unadjusted) p values. By definition, under the null hypothesis p values are uniformly distributed in $[0, 1]$. Figure 4 illustrates the cumulative distribution of p values of the 111 submissions. The distribution is very close to the uniform distribution, indicating that ‘on average’ the submitted classifiers did not make informed predictions but rather collectively perform as well as random guesses. This is clearly disappointing, and requires an explanation.

DISCUSSION

Structural similarity of the training and test sets

One of the fundamental assumptions which statistical and inductive learning methods are usually based on, is that the examples which a predictive model are tested on come from the same distribution as those on which the model

Table 1. Distribution of positive and negative examples among the training and test sets

	Positive examples	Negative examples
Training set (NTP)		
Male rats	166	305
Female rats	135	371
Male mice	140	348
Female mice	155	334
Test set (FDA)		
Male rats	52	133
Female rats	36	149
Male mice	29	156
Female mice	35	150

Structure

Fragments




Fig. 5. Linear fragments of 1,4-dichlorobenzene (CAS 106-46-7).

was trained. This has two aspects:

- (1) the distribution of positive and negative examples (see Table 1);
- (2) the distribution of structural features among the training and the test set (i.e. if structural features of the test set are contained in the training set).

Although the test set contains more negative examples than the training set, the first point is of minor concern, as we have used ROC analysis for the evaluations of the results. Also note that Machine Learning algorithms outputting probabilities along with classifications are tunable for test sets with a changed class distribution (Elkan, 2001).

To estimate the structural similarity between the training and test set we have used the Molecular Feature Miner (MOLFEA; Kramer *et al.*, 2001). This technique is based on the concept of *molecular fragments*. A molecular fragment is defined as a sequence of linearly connected heavy atoms. Figure 5 shows how to decompose an example molecule into its fragments, more examples and applications can be found in Kramer *et al.* (2001).

When predicting the biological activity of an untested compound the ideal case is that all structural features of

Table 2. The 10 statistically most significant submissions. Negative and positive examples are the true class distributions of the test examples. If a model did not provide predictions for all compounds, the numbers are different from Table 1

p	Sex/species group	Model/id	Negative examples	Positive examples	False positive rate	True positive rate
0.0019	Female mice	Viniti/29	84	21	0.036	0.286
0.0027	Male mice	Baurin_en/7	156	29	0.109	0.345
0.0046	Male mice	Viniti/29	97	14	0.031	0.286
0.0433	Female rats	Kwansei/14	73	17	0.274	0.529
0.0488	Male mice	Leuven2/16	153	29	0.366	0.552
0.0643	Female mice	Animaths2v/2	148	34	0.115	0.235
0.0864	Male rats	Gonzales/10	133	52	0.150	0.250
0.0916	Female mice	Animaths1v/1	148	34	0.081	0.176
0.1186	Female mice	Smuc1/25	150	35	0.280	0.400
0.1417	Female rats	Viniti/29	69	17	0.101	0.235

this compound are already contained in the training set. Otherwise, the predictions should be labeled as unreliable, because unknown structural elements may contribute to biological activity. The Chemical Inductive Database Language of MOLFEA allows us to perform such a comparison. For a given molecule, we have to find all fragments that are present in this molecule, but not in the training set.

Of course not every unknown fragment has equal relevance. Long fragments can be decomposed into smaller ones that may provide enough information to provide a reliable estimation of biological activity. But if a short fragment, or even an element is unknown to the training set, the reliability of predictions will be much lower, if the missing fragment is relevant for a toxic mechanism. Thus there are three factors that have to be considered for the reliability of predictions:

- (1) the number of unknown fragments;
- (2) the length of unknown fragments;
- (3) the chemical ‘meaning’ of unknown fragments.

These parameters can also be used to compare the information content of two data sets. To compare the structural diversity of the PTC training set (NTP) with the test set (FDA), we have used the following procedure: For each FDA compound we determined the number and size of fragments that do not occur in the NTP learning set. To obtain a reference, the same procedure was performed on the NTP data set, using a leave-one-out procedure: Sequentially, one compound was removed from the data set and compared with the rest of the compounds to obtain the unknown fragments for this compound.

Figure 6 shows the distribution of the number of unknown fragments among the FDA test set and the NTP training set. It is obvious that the test set contains more

unknown fragments than the training set. A closer inspection reveals that the NTP data set has a high proportion of compounds with completely known fragments (211 from 417 compounds). In the FDA data there are only 23 from 285 compounds without unknown fragments. It is tempting to see this as an indication that both data sets are structurally dissimilar. Indeed the NTP data set contains a very diverse selection of compounds (industrial chemicals, pharmaceuticals, environmental contaminants,...), whereas the FDA data set is focused towards pharmaceuticals. This fact is also reflected in the distribution of the molecular sizes (measured as the number of heavy, i.e. non-hydrogen atoms per molecule) within both data sets (Fig. 7). The median size of molecules in the NTP data set is 13 heavy atoms/molecule, whereas the FDA data set contains mostly medium-sized molecules with a median of 20 heavy atoms/molecule. Since the probability to encounter an unknown fragment increases with the size of a molecule, this is a possible explanation for the high ratio of compounds with completely known fragments in the training set. Comparing the sizes of the smallest unknown fragment within a molecule (Fig. 8) gives a similar, although less pronounced result: In this case, the median size is equal in both data sets, but the FDA data is slightly skewed towards smaller (and potentially more relevant) fragments.

Summing up, we know that the test set contains structural features that are not present in the training set. But it is still unclear if the unknown features are toxicologically relevant and if they had any impact on the predictions in the challenge. To investigate this point, we have defined ‘easily predictable’ and ‘poorly predictable’ compounds for each sex/species group. Compounds that were correctly predicted by all optimal models (as determined by ROC analysis) were classified as ‘easily predictable’, those with consistently false classifications were classified as ‘poorly predictable’. Using the same procedure as




Fig. 6. Number of unknown fragments in the training and the test set.




Fig. 7. Distribution of molecule sizes in the training and the test set.

above, we determined for each molecule in these groups the length and size of fragments that are not present in the NTP training set. The results are summarized in Figures 9 and 10. These figures and the associated statistics (Wilcoxon rank sum tests were used for all comparisons) clearly indicate that the distribution of unknown fragments does not differ between ‘easily’ and ‘poorly predictable’ compounds. A similar picture is obtained by plotting the rate of correct predictions (i.e. *correct predictions/all predictions*) against the number of unknown fragments and the length of the smallest unknown fragment. Figures 11 and 12 present a summary of this data for all sex/species groups. They show clearly that there is no correlation between the number (size) of unknown fragments and the predictive accuracy.




Fig. 8. Sizes of the smallest unknown fragments in the training and the test set (compounds with completely known fragments are excluded).

We may therefore conclude that the structural differences between the training and the test set are not the main reason for the poor performance of the submitted models. This assumption is further substantiated by the cross-validation results on the NTP training set that are generally lower than 70% (Blinova *et al.*, 2003; Okada, 2003).

Toxicological knowledge and model reliability

To investigate if the predictive models involve some (known or unknown) domain knowledge, we have asked toxicological experts to judge the value (and interpretability) of the submitted models.

This judgment resulted in the following ranking:

- (1) Kwansei (model 14)
- (2) Wai (models 30–32)
- (3) Leuven (models 15–17)

It is noteworthy that two of these models are within the top five statistically strongest models. This reinforces our confidence that these models are probably the most interesting. We next review these models briefly.

The Leuven2 model for male mice considers a compound to be active if: (i) it contains bromine, or (ii) it does not contain oxygen, or (iii) it contains oxygen but no sulphur and no hydrogen bound to oxygen (i.e. if the oxygen is not alcoholic or phenolic or part of the carboxylic acid group); otherwise the compound is considered to be inactive.

This model has the benefit of being concise and simple. The first part about bromine is in general agreement with accepted toxicological knowledge. The second part has very broad coverage coupled with relatively low

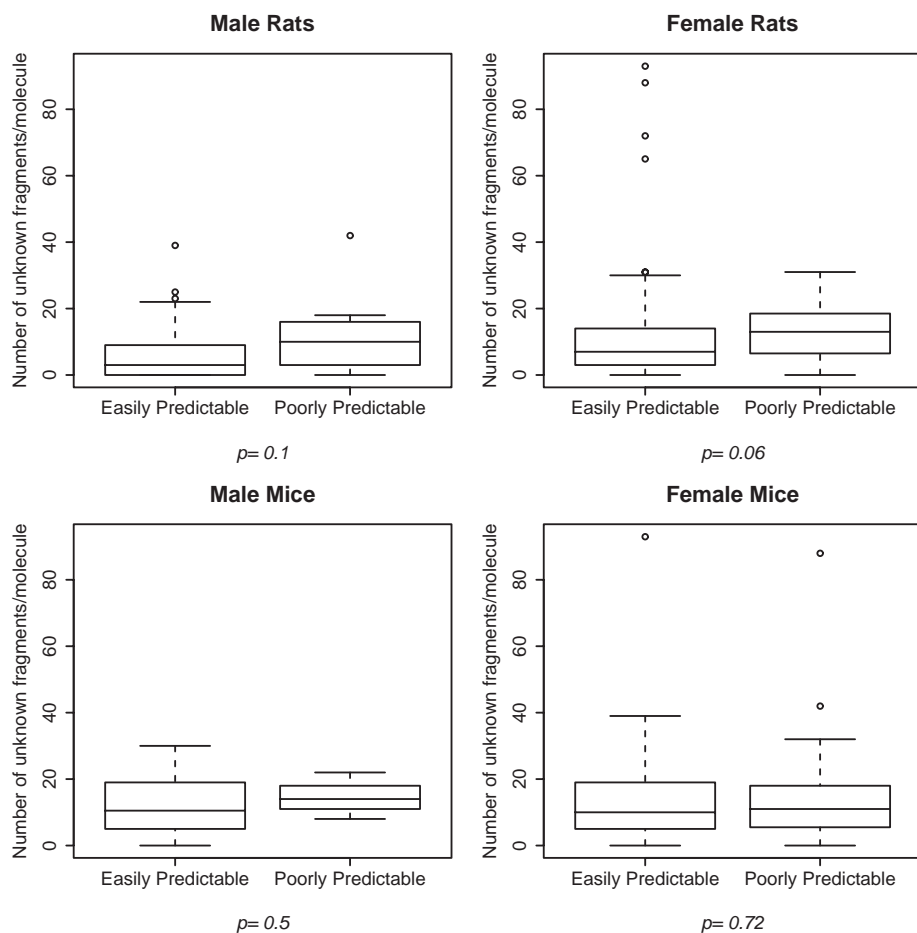


Fig. 9. Number of unknown fragments in easily and poorly predictable compounds.

accuracy. The third part of the rule covers a diverse set of compounds, including ketones, aldehydes, ethers, epoxides and nitro-organic compounds.

The Leuven rules are probably too general to be of much current value to toxicologists. We believe this reflects the bias of decision tree methods, which favors short general rules. It is interesting to note that a learner which took into account the maximum number of features would probably have produced a more toxicologically relevant set of rules.

The Kwansai model was generated by the group with probably the most experience in toxicology, and the Kwansai methodology was favored by the toxicology judge as its approach resembled that of a human toxicologist. The Kwansai model is much more complicated than the Leuven2 model. We therefore only describe a couple of its rules which are judged to contain toxicological evidence. Rule Pos1 for carcinogenicity from the Kwansai model is:

$[C-N: n] \ \& \ [C-c:c:c:c: y] \ \& \ [N: y]$

This rule requires nitrogen, but not an aliphatic carbon

nitrogen bond, and a set of aromatic carbon bonds (an aromatic ring) connected to an aliphatic carbon. It covers aromatic amines, nitroaromatics and azo compounds, all of them well known classes of carcinogens. Rule Neg2 for non-carcinogenicity is:

$[S: y] \ \& \ [HBD=0] \ \& \ [O: y]$

This rule requires sulphur and oxygen, but no hydrogen-bond donors (e.g. alcohols, amines, ...).

Despite this success in generating toxicological knowledge from data, most models generated in the PTC contained little or no toxicological knowledge, which is consistent with them performing close to random on the test data.

Why is carcinogenicity hard to predict?

From 111 models submitted to the Predictive Toxicology Challenge, only five submissions performed significantly better ($p \leq 0.05$) than random guessing. Before discussing the possible reason for this—at a first glance discouraging—result we should put it into perspective:

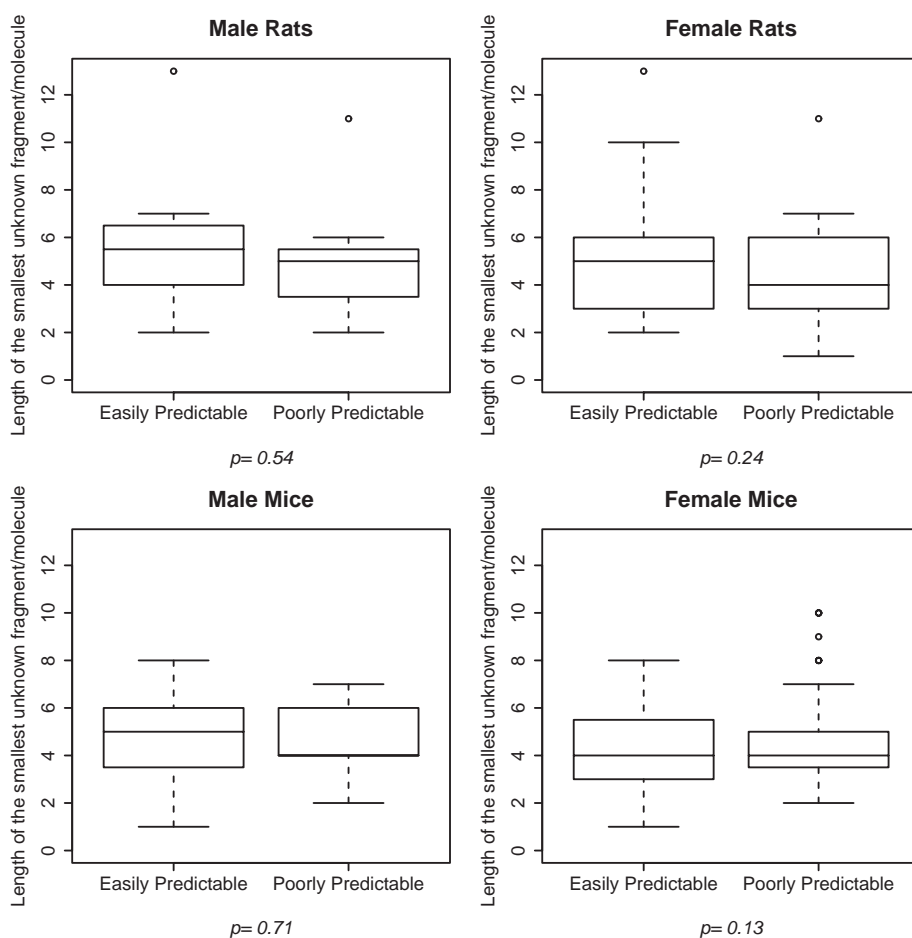


Fig. 10. Sizes of the smallest unknown fragments in easily and poorly predictable compounds.

- we are unaware of a biological alternative to rodent carcinogenicity tests that perform better than the models submitted to the PTC. (Despite decades of efforts the *Salmonella* microsome assay seems to be still the most predictive short-term test (Zeiger, 1998). The concordance with rodent carcinogenicity data is 60–70%, depending on the evaluation data.)
- the same or similar techniques as in the PTC have been successfully applied to predict other toxic effects (for a review see Helma *et al.*, 2000).

We have identified several possible causes that may lead to the poor predictability of carcinogenic effects[‡]:

[‡] Additional complications may arise from the fact that the NTP dataset is biased towards compounds whose carcinogenicity is hard to estimate. The result is that most of the carcinogenic compounds cause cancer in idiosyncratic ways, in only one species, or a few organs. This makes it hard to deduce general rules from this dataset. The FDA test set was also hard to predict because companies were careful not to submit compounds that are likely carcinogens (John Ashby, personal communication, Feb. 2003).

- (1) the biochemical mechanisms involved in chemical carcinogenicity are too complex to be modeled by machine learning or statistical techniques;
- (2) the descriptors for chemical structures and properties are inadequate for predicting carcinogenicity;
- (3) *Structure Activity Relationship (SAR)* models ignore the importance of biological variables;
- (4) rodent carcinogenicity classifications are too inaccurate to learn accurate models from them;
- (5) predictions are impossible, because compounds in the test set contain too many structural features that are unknown to the training set.

As we have discussed the last point already in the previous section, we will focus on the first arguments. As it is presently impossible to decide between the alternatives, we will present arguments for and against each point, present ideas for future exploration and indicate our opinions.




Fig. 11. Number of unknown fragments versus true prediction rate for all sex/species groups.

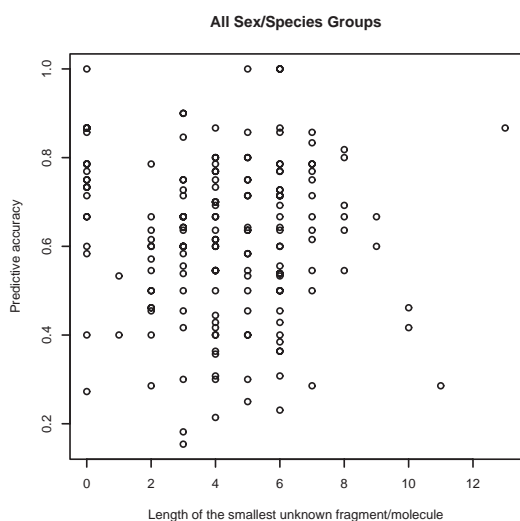


Fig. 12. Sizes of the smallest unknown fragments versus true prediction rate for all sex/species groups.

Most experimental toxicologists will favor the first argument. But we have to keep in mind that it is not necessary (and impossible) to model the whole biological system (transport, metabolism, ...), in order to make toxicity predictions. In *Structure Activity Relationship (SAR)* studies we are concerned only with the influence of chemical structure on a certain biological effect. SARs are in fact black box models that relate chemical structures to a biological outcome. (It is nevertheless possible to draw mechanistic conclusions from SAR studies.) Theoretically, machine learning techniques can

learn arbitrarily complex relationships, but they will need more data for complex models than for simple ones. With this fact in mind it is possible (but beyond the scope of this article) to perform an experiment to test the complexity hypothesis, by comparing the predictivity of models based on different numbers of training examples. [For mutagenicity, Liu *et al.* (1996) determined an optimal data base size/unit cost of ~350 compounds]. This should clarify if more than the presently publicly available ~500 compounds are needed for accurate carcinogenicity classifications.

It is theoretically possible to generate an almost unlimited number of descriptors for chemical structures and their properties. To select a subset for a SAR study is more or less a trial and error process, especially when the underlying biological processes are to a large extent unknown. One of the goals of the PTC was to provide a variety of different descriptors in Stage 1, and let participants in Stage 2 choose the best parameters. Unfortunately none of the participants reported systematic parameter selection experiments. So the optimal representation of chemical structures and the choice of parameters for predicting carcinogenicity is still an open question and has room for improvements.

It is generally desirable to include biological information in SAR studies, e.g. to improve cross-species predictions or to identify susceptible subpopulations. For the sake of clarity we have to differentiate between three types of biological data:

- (1) information about the target organism;
- (2) information about the experimental procedure;
- (3) information about other biological effects of the same compound.

In the NTP program the majority of studies were performed with genetically identical animals (F344/N rats and B6C3F1 mice), non-standardized studies were removed from the training set during the data cleaning step (Helma and Kramer, 2003). This design ensures genetically determined variations do not influence the outcome of carcinogenicity assays, therefore it does not make sense to consider biological variables apart from sex and species.

The experimental detail that has probably the highest impact on the outcome of carcinogenicity experiments is the administration route. As the majority of compounds were administered by feed or gavage it was a deliberate and debatable decision not to include this information in the PTC. Further studies are certainly needed to clarify the impact of the administration route on the outcome of carcinogenicity assays.

Another detail that was omitted in the data set is the tumorigenic dose TD_{50} . The consideration of dose information might help in distinguishing between weak and strong

carcinogens, but it adds considerable complexity to the prediction problem, because we have to combine a classification task (discrimination between carcinogens and non-carcinogens) with a regression task (prediction of TD_{50} 's for carcinogens).

Information about biological effects on other organisms or other endpoints than the one to be predicted can be useful especially if there is a mechanistic relationship (e.g. between mutagenicity and carcinogenicity). We omitted such information from the PTC for two reasons: (i) information of that kind is available only for a fraction of the compounds in the training and the test set; and (ii) we intended to test the ability to predict biological activities purely *in silico*.

It is presently impossible to estimate the reproducibility of rodent carcinogenicity assays. Due to the high costs of these assays, we have presently no replicate experiments under standardized conditions. The 2-year rodent carcinogenicity assay is a very complex experiment that involves uncountable experimental, measurement and interpretation tasks that are later condensed into a simple carcinogen/non-carcinogen classification. As errors and mistakes accumulate, it is conceivable that complex experiments like this have a lower accuracy than comparatively simple toxicity assays (e.g. the Ames test for mutagenicity). Gottmann *et al.* (2001) compared data from the general literature with NTP results and found a very low concordance of the results (<60%). The NTP and FDA data sets, however have a high concordance of >80% (Helma and Kramer, 2003), but we were unable to ascertain that the replicate experiments have been performed really independently. It should also be noted that NTP priorities have shifted towards more problematic compounds (e.g. non-genotoxic carcinogens). The interpretation of these is the subject of many expert discussions and classifications are presumably less reliable than those of strong, direct acting carcinogens. Thus, the possibility that the poor predictability of rodent carcinogenicity is due to uncertainties in the experimental results cannot be ruled out and requires further examination.

CONCLUSION

The aim of the PTC challenge was to test if it is possible for statistical/machine learning methods to learn models for the rodent carcinogenicity of chemical compounds.

Such a task is clearly scientifically challenging:

- chemical carcinogenicity is a very complex process. It involves a complex network of biochemical mechanisms (transport, metabolism, DNA damage, ...) that may differ from organism to organism;
- there are only a limited amount of training examples (a few hundred);
- the training data is not randomly distributed;
- the training data has an unknown amount of noise.

Given these difficulties it was not surprising to see that the majority of contributions did not perform better than random guessing. It was however encouraging to observe a limited number of predictive models that performed significantly better than random (see Table 2), and were judged to have empirically learned a small but significant amount of toxicological knowledge.

ACKNOWLEDGEMENTS

We would like to thank Douglas Bristol and Yin-tak Woo for their toxicological expertise.

REFERENCES

- Blinova,V., Dobrynin,D., Finn,V., Kuznetsov,S. and Pankratova,E. (2003) Toxicology analysis by means of the JSM-method. *Bioinformatics*, **19**, 1201–1207.
- Bradley,A.P. (1995) The use of area under ROC curve in the evaluation of learning algorithms. *Pattern Recognition*, **30**, 1145–1159.
- Elkan,C. (2001) The foundations of cost-sensitive learning. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI'01)*. pp. 973–978.
- Gottmann,E., Kramer,S., Pfahringer,B. and Helma,C. (2001) Data quality in predictive toxicology: reproducibility of rodent carcinogenicity experiments. *Environ. Health Perspect.*, **109**, 509–514.
- Helma,C., Gottmann,E. and Kramer,S. (2000) Knowledge discovery and data mining in toxicology. *Stat. Methods Med. Res.*, **9**, 329–358.
- Helma,C. and Kramer,S. (2003) A survey of the predictive toxicology challenge. *Bioinformatics*, **19**, 1179–1182.
- Kramer,S., De Raedt,L. and Helma,C. (2001) Molecular feature mining in HIV data. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-01)*. pp. 136–143.
- Liu,M., Sussman,N., Klopman,G. and Rosenkranz,H. (1996) Estimation of the optimal data base size for structure–activity analyses: the *Salmonella* mutagenicity data base. *Mutation Res.*, **358**, 63–72.
- Okada,T. (2003) Characteristic substructures and properties in chemical carcinogens studied by the cascade model. *Bioinformatics*, **19**, 1208–1215.
- Provost,F. and Fawcett,T. (1997) Analysis and visualization of classifier performance: comparison under imprecise class and cost distribution. In *Proceedings of the Third International Conference on Knowledge Discovery and Data Mining (KDD-97)*. pp. 43–48.
- Van Trees,H.L. (1971) *Detection, Estimation, and Modulation Theory*. Wiley, New York.
- Zeiger,E. (1998) Identification of rodent carcinogens and non-carcinogens using genetic toxicity tests: premises, promises, and performance. *Regul. Toxicol. Pharmacol.*, **28**, 85–95.