

Prediction of protein-protein interactions using one-class classification methods and integrating diverse biological data

José A. Reyes^{1,2} and David Gilbert¹

¹Bioinformatics Research Centre, Department of Computing Science, University of Glasgow.

²Facultad de Ingeniería, Universidad de Talca, Chile. Contact: jareyes@dcs.gla.ac.uk

Summary

This research addresses the problem of prediction of protein-protein interactions (PPI) when integrating diverse kinds of biological information. This task has been commonly viewed as a binary classification problem (whether any two proteins do or do not interact) and several different machine learning techniques have been employed to solve this task. However the nature of the data creates two major problems which can affect results. Firstly, imbalanced class problems due to the number of positive examples (pairs of proteins which really interact) being much smaller than the number of negative ones. Secondly the selection of negative examples can be based on some unreliable assumptions which could introduce some bias in the classification results.

Here we propose the use of one-class classification (OCC) methods to deal with the task of prediction of PPI. OCC methods use examples of just one class to generate a predictive model which consequently is independent of the kind of negative examples selected; additionally these approaches are known to cope with imbalanced class problems. We have designed and carried out a performance evaluation study of several OCC methods for this task; and have found that the Parzen density estimation approach outperforms the rest. We also undertook a comparative performance evaluation between the Parzen OCC method and several conventional learning techniques, considering different scenarios, for example varying the number of negative examples used for training purposes. We found that the Parzen OCC method in general performs competitively with traditional approaches and in many situations outperforms them. Finally we evaluated the ability of the Parzen OCC approach to predict new potential PPI targets, and validated these results searching for biological evidence in the literature.

1 Introduction

The prediction of protein-protein interactions (PPI) has emerged recently as an important problem in the fields of Bioinformatics and Systems Biology, due the fact that most of essential cellular processes are mediated by these kind of interactions. High-throughput methods for the direct identification of protein-protein interactions have been developed including yeast two-hybrid screens (Y2H) [1, 2] and mass spectrometry methods for protein complex identification [3, 4]. Even though these high-throughput techniques can increase the number of predicted PPI, in general the data obtained by these methods is often incomplete and suffers from high false-positive and false-negative rates [5]. In order to improve the accuracy and trustability of predicted protein interacting pairs, various studies have been developed in the past years based in the integrative learning analysis of diverse biological sources of information. These have Demonstrated that the combined use of direct and indirect biological insights can improve the quality of predictive PPI models.

The prediction of PPI has been commonly viewed as a classical binary classification problem where the aim is to predict whether any two proteins do or do not interact. Several traditional machine learning methods have been employed in the past for this specific task [6, 7, 8, 9, 10, 11, 12]. These methods generally use supervised learning algorithms where the final objective is to generate a classification model from a gold standard reference set of positive (really interacting protein pairs) and negative examples (non-interacting pairs). Two main drawbacks have been identified regarding these previous approaches:

- i) In general they face a highly imbalanced classification problem, where the number of positive examples is much smaller than the number of negative examples. This affects the quality of the predictive models which may be biased towards the majority class and consequently the minority class examples are poorly predicted. Under-sampling and cost sensitive strategies have been used to deal with the imbalanced problem in some of these previous works whilst others did not report any action about it.
- ii) While the selection of positive examples is based on trustable experimental techniques (i.e. small scale experiments), there is no experimental method to find pairs of proteins which do not interact (negative examples). Therefore certain assumptions have to be made for the construction of a negative gold standard set, which can introduce some bias into the learning process and consequently produces a significant effect on the performance of the classification approach [13].

As a possible solution to these issues, in this research we propose the use of one-class classification (OCC) methods. The aim of OCC is to use feature information from only one of the classes, positive examples in this case, to generate a classification model. OCC methods are known to be able to deal efficiently with highly imbalanced classification problems [14]. Additionally, unlike conventional binary classifiers, OCC methods in general produce classification models which are independent of the kind of negative gold standards traditionally employed PPI.

In this paper we present the results of the application and evaluation of various OCC methods for the prediction of PPI, based on the integrative learning analysis of diverse biological data. Additionally we have carried out a comparative evaluation between OCC methods and several common binary classification techniques, which have been previously employed for this specific task. Previous studies presented a comparative evaluation between several learning classifiers [7, 12], but did not consider OCC methods as an alternative. We also note that their results are difficult to compare because they have been generated using different reference gold standard sets and biological information. In the next section we present the main methodological aspects considered in this research, followed by a detailed description and analysis of the results obtained. Finally we give some important conclusions.

2 Methods

2.1 One-class classification methods

The common issue of one-class classification (OCC) problems is that feature information is available for only one of the classes, called the *target class*, and this is employed to generate a

classification model. The OCC model is constructed with the aim of characterizing and describing the target examples, and afterwards it is used to distinguish target examples from all other examples which have been classified into a single different category called the *outlier class*. The general task in OCC can be regarded as being similar to conventional binary classification methods, in the sense that a decision boundary or separation model is used to separate examples of the two classes (target and outliers). However OCC methods face a harder task because the decision boundary is mainly supported by examples of the target class and consequently less information is employed to build and validate it. Consequently a sufficiently representative sample of target examples is needed to generate a more accurate descriptive model in order to improve the OCC performance.

In this research we consider the task of prediction of PPI as a OCC problem, in the sense that only examples of one class (positive interaction examples) are available and/or trustable, becoming the *target class*. Because the OCC approach is mainly based on the description of examples of the target class, the resulting classification model is independent of the kind and quality of the set of negative examples employed, which could potentially solve the problem of trustability associated to this class. In order to develop a comparative performance evaluation between OCC and conventional classification methods, a set of negatives examples should be selected as the *outlier class*. This is because it is absolutely necessary to use examples of both classes for training and testing purposes when considering conventional binary classifiers. Under these conditions the performance of OCC methods can be evaluated in a manner similar to that for conventional binary classification techniques, i.e. by estimating the misclassification error on, the target class error (or false-negative rate), and when outlier examples are available, the outlier class error (or false-positive rate).

OCC methods can be classified according to the way in which they analyze, describe and generate a model for the separation of targets and outlier examples [15]. Here we consider two types, as follows. (A) *Density estimation* methods based on the estimation of the probability density distribution of the training data using some probabilistic model (i.e. Gaussian distribution); a threshold is selected and then used to compare with the density of new objects in order to classify them. (B) *Boundary* methods based in the generation of a frontier or boundary around the target objects, which is optimized to accept most of the target examples and at the same time reject most of the outliers. Four different OCC learning approaches were evaluated in this research, namely three density estimation methods (single Gaussian estimation, mixture of Gaussian and Parzen density estimation) and a boundary approach (Support vector data description SVDD). The dd_tools Matlab toolbox (http://www-ict.ewi.tudelft.nl/~davidt/dd_tools.html) was utilized to develop the experiments associated with the application and evaluation of all OCC methods. A detailed explanation of each of the OCC methods employed is presented in the supplementary material.

2.2 Reference data set

In this research we focused in the prediction of co-complexed proteins pairs (pairs of proteins which are co-members of the same complex). To evaluate different machine learning methods we need a reference data set (gold standard) containing positive and negatives examples. We used the same gold standard sets described in [7] for the study of PPI on yeast. These comprise 2,104 positive examples (true interacting protein pairs) derived from the MIPS complex catalogue [16] and 172,409 negative examples (non-interacting protein pairs) related to protein

pairs which are localized in different cell compartments and consequently are more likely to not interact. This reference data set is a subset of that used in [6], considering only examples where complete information for each one of the biological features are available.

2.3 Biological features

An important motivation for this research is that the integration of diverse kinds of biological data/information could potentially improve our ability to predict protein-protein interactions. Four different types of biological information were considered following [6] and [7]:

m-RNA expression, following the assumption that proteins which are members of the same complex are commonly expressed simultaneously. The Pearson correlation is estimated for every protein pair considering two different well known studies: the Rosetta compendium [17] and cell cycle time series analysis [18], generating two numeric values between [-1,1] which are incorporated as features.

functional similarity of protein pairs was estimated from the gene ontology (GO) [19] and the MIPS [16] functional catalog, obtaining two new numeric features. The assumption here is that proteins in the same complex tend to participate in the same biological processes.

Essentiality information [16], assuming that two proteins in the same complex are both essential or non essential for cell survival. This feature is then characterized by three possible categories (i.e. both proteins are essential or both are non-essential or only one of them is essential), and is represented by a three dimensional vector taking discrete values of +1 or -1 according to each case.

High-throughput experimental interaction data from Y2H and mass spectrometry based experiments were integrated as features. Four different experimental studies have been considered [1, 2, 3, 4]. On each case a discrete value of +1 or -1 is assigned to indicate whether the components of a protein pairs do interact or do not interact respectively.

Numerical features were normalized to obtain a distribution with a mean of zero and standard deviation of one, in order to put all data in the same range of values and to avoid possible numerical difficulties associated with imbalanced ranges. Every pair of proteins available in the reference data set was represented by a 11-dimensional vector X_i containing the information for the biological features considered here, and a label Y_i which can take two values depending on whether each of the proteins pairs do really interact ($Y_i = 1$) or not ($Y_i = -1$).

2.4 Conventional machine learning methods

A representative group of conventional or traditional machine learning techniques, which have been previously used for the task of prediction of PPI, was selected in order to undertake a comparative performance evaluation with OCC methods for this specific task. These including: Decision Trees (DT), Naive Bayes (NB), Logistic Regression (LR) and Support Vector Machines (SVM). The WEKA machine learning library [20] was used to perform the experiments related to DT, NB and LR, while the evaluation of SVM was carried out using the MATLAB interface to the SVM-light toolbox (<http://svmlight.joachims.org>).

2.5 Performance evaluation

OCC and conventional learning approaches were evaluated in different training/testing scenarios varying, for instance, the number of negative examples used to train each one of the models. A ten fold cross validation procedure was carried out for every evaluation, in order to assess the variability of the models generated. On each situation the negatives examples which were not utilized in the training step were also included in the testing evaluation. This testing strategy differs from the previous approaches used for this task, where only a fraction or sub-sample of the negative gold standard examples was considered to test the models. We think that by including all the available putative negative information each time we test our models, we are carrying out a more relevant and at the same time more challenging evaluation for the prediction of PPI.

Receiver Operator Characteristic (ROC) curves, showing the tradeoff between the false-positive rates and true-positive rates, were generated for each approach under the different scenarios evaluated. The area under the ROC curve (AUC) was calculated for each case to evaluate the overall performance of different learning algorithms. AUC scores seem to be a better evaluation measure than simple accuracy in imbalanced class problems [21].

We also calculated partial AUC scores, which are related to the normalized area under a fraction of the whole ROC curve which represents a condition of special interest. For example in severe class imbalance it seems more relevant to evaluate the performance in the region of low values of false positive-rates [22], which is the case in the prediction of PPI task. In our approach we are interested in evaluating and comparing the performance of the different classifiers under conditions of a low false-positive rate. The aim of this is to maximise the number of real interacting protein pairs predicted while minimizing the number of false-positive predicted ones. This is of especial interest for biologists working in the identification and validation of new PPI, because they can focus in the study of only the top ranked predicted PPI targets instead of evaluating many randomly selected protein pairs. We considered the area under the ROC curve up to the first 50 false-positive examples (AUC-50), which has become a commonly accepted performance measure for this specific task [11, 12].

Mean values and standard deviation for AUC and AUC-50 were calculated based on the ten fold cross-validation individual results, in order to compare the performance of different approaches. When the difference was unclear between the AUC or AUC-50 values for two methods, the Wilcoxon signed rank statistical test [23] for the median of the differences between them was computed considering a 5% significance level, in order to obtain stronger evidence that one of the methods performed better than the other.

3 Results

3.1 Evaluation of diverse OCC methods

Four different OCC methods were used for the problem of PPI prediction including: Gaussian density estimation, Mixture of Gaussian density estimation, Parzen density estimation and Support Vector Data Description (SVDD). The methods were evaluated on a balanced class set using all the positive examples available and an equal size sample of negative examples

randomly selected from the whole negative gold standard set. This was done because some of the OCC methods can take advantage of the use of a sample of negative examples to improve their performance. This procedure was repeated ten times using diverse sub-samples of negative pairs. The results of the estimation of AUC and AUC-50 scores for the OCC performance evaluation are shown in Table 1 where the mean and standard deviation are given.

Table 1: Comparison of AUC and AUC-50 values for different learning methods evaluated

Method	AUC	AUC-50
<i>OCC methods:</i>		
SVDD	0.9768 ± 0.0033	0.2455 ± 0.0325
Gaussian	0.9377 ± 0.0136	0.1224 ± 0.0136
Mixture of Gaussian	0.9853 ± 0.0096	0.2269 ± 0.0513
Parzen	0.9801 ± 0.0075	0.4010 ± 0.0282
<i>Conventional methods:</i>		
Decision trees (DT)	0.9946 ± 0.0033	0.2129 ± 0.1903
Naive Bayes (NB)	0.9908 ± 0.0017	0.2299 ± 0.0275
Logistic Regression (LR)	0.9928 ± 0.0018	0.0917 ± 0.0307
Support Vector Machines (SVM)	0.9939 ± 0.0016	0.2687 ± 0.0250

The results for global AUC scores show that there is no significant difference between most of the OCC methods evaluated, with the exception of the simple Gaussian density estimation method which exhibits the lowest overall performance. On the contrary, the analysis of the results for AUC-50 scores clearly shows that the Parzen density estimation method (AUC-50 = 0.401) by far outperforms the rest of the single OCC methods considered here. The good performance obtained by the Parzen method can be explained because this density estimation method takes into account the information of every target example available. This is different to the rest of the OCC approaches evaluated, where for example only an average density estimation from the available data is employed as in the case of Gaussian and mixture of Gaussian approaches.

The second best performance for OCC methods considering AUC-50 scores is obtained by the SVDD approach using a Gaussian kernel (AUC-50 = 0.2455). We note that recent work reported in [24] used one-class support vector machines (OCSVM) [25], which is an extension of the classical binary SVM technique, to deal with the task of prediction of PPI. In this work the authors only considered one biological feature based on protein sequence and domain information, reporting that the best results are obtained using a Gaussian kernel. In contrast, in our research we evaluated several different OCC approaches, used diverse biological features and also carried out a comparative performance evaluation with several conventional binary classification methods. Moreover, it has been shown that the SVDD and OCSVM techniques give equivalent solutions [25, 15] when using a Gaussian kernel.

3.2 Comparative evaluation between OCC and conventional classifiers

The Parzen OCC method was selected, due to its good performance, to be compared in a more exhaustive evaluation with several conventional classifiers such as Decision Trees (DT), Naive Bayes (NB), Logistic Regression (LR) and Support Vector Machines (SVM). Firstly all the

learning approaches were evaluated on the same ten different balanced class sets previously used. Estimates for AUC and AUC-50 scores for these experiments are given in Table 1.

Comparative analysis of overall AUC scores shows that conventional classifiers perform only slightly better than the Parzen OCC approach. This was expected because the task associated with OCC only uses examples of one class to generate a classification model. However in relation to the AUC-50 comparative evaluation, we found that the Parzen OCC approach clearly outperforms all conventional classification techniques. The performance of conventional classifiers in these cases is only comparable with some of the other OCC methods previously evaluated, and sometimes worse as the case of the LR approach. SVM showed the best performance for the conventional classifiers (AUC-50 = 0.2687). However the Parzen OCC approach (AUC-50 = 0.401) is able to correctly predict around the 40% of the positive examples in the AUC-50 region, while the SVM method only predicts around the 27%. It is interesting to note that DT exhibits high variability compared with the rest of the methods evaluated. The detailed analysis of AUC-50 results shows that in some of the ten fold cross validation subsets DT performs better than OCC methods but in others (the majority) it performs very poorly. The Wilcoxon signed rank test [23] was applied in this case showing that effectively the Parzen OCC method outperforms the rest of conventional classifiers.

The difference between the AUC and AUC-50 analysis can be clearly appreciated from the ROC curves of the different learning methods evaluated. Figure 1(a) shows an example of the ROC curves for the different learning techniques used in the evaluation of one cross validation subset. No important differences between these ROC curves is observed and consequently there is no significant difference in total AUC scores. When we focus in the portion of these curves related with the AUC-50 region, presented in Figure 1(b), there are clear differences in the performance of the diverse methods. In this region the Parzen OCC method outperformed the rest of the conventional learning approaches evaluated. This is still the case if we extend the partial AUC analysis up to the first 100 false-positive examples. This corroborates our assumption that analysis based on partial AUC scores (i.e. AUC-50) is more appropriate than that using overall AUC scores, for predicting PPI.

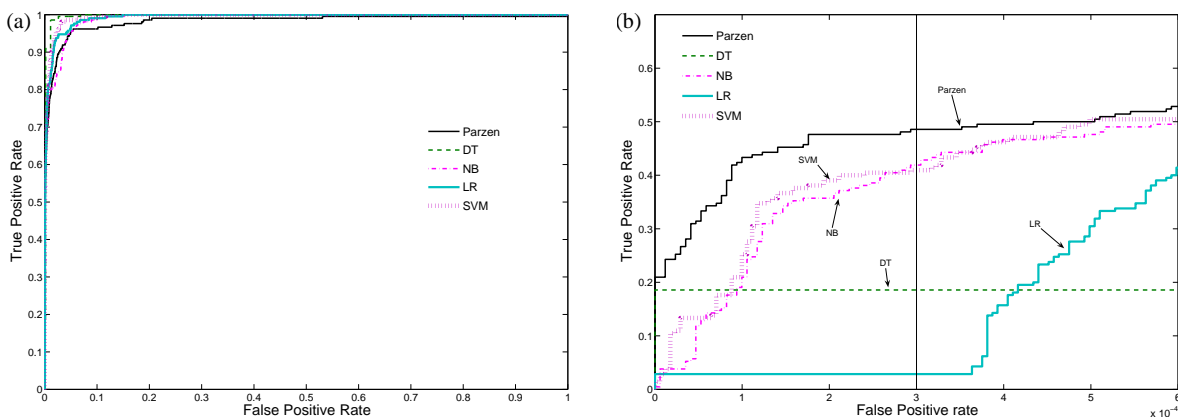


Figure 1: Example of ROC curve analysis: (a) Whole ROC curves for the different learning methods evaluated. (b) Partial ROC curves for the different learning methods evaluated. The vertical line indicates the point where approximately the first 50 false-positive examples are reached.

3.3 Comparative evaluation on different scenarios

We also evaluated and compared the effect of the use of negative examples in the performance of the diverse learning approaches. Different scenarios were generated varying the number of negative examples used for training the respective models, from none to all the negative examples available. Figure 2 shows the performance results, measured as AUC-50 scores, for all the situations considered.

Firstly we analysed the cases where less negative than positive examples were used to train the models, including the balanced class scenario when 2,104 negative examples are employed. The Parzen OCC method clearly outperforms the rest of conventional learning techniques, exhibiting a very stable performance in the different situations. This can be explained because it only uses positive examples for training purposes. On the contrary, the performance of conventional classifiers tends to decrease as less negative information is used. SVM technique shows the best performance for binary classifiers followed by the NB approach. DT and LR techniques show low performance and high variability compared with the rest of methods evaluated. Note that in the situation where no negative examples are used, only the Parzen OCC method can be employed and consequently no results for conventional classifiers are available.




Figure 2: AUC-50 comparison for different learning methods evaluated, showing the effect of reducing and incrementing the number of negative examples used to train the models. The balanced class scenario is when 2,104 negative examples are used for training. Note that no corrective action was taken for any of the imbalanced class situations.

The analysis of scenarios where more negative than positive examples are employed to train the models is quite different. The Parzen density estimation OCC technique tends to maintain its performance stability and a significant increment in the AUC-50 performance only occurs when more than 50,000 negative examples are employed. This can be explained because in these cases the models are tested on a reduced number of negative examples (most of the negative information is used to train the models). The performance of conventional classifiers tends to increase gradually as more negative examples are incorporated for the generation of their respective classification models. This was expected because these techniques can take advantage of the negative object class information.

The Parzen OCC method performs very competitively in most of the scenarios evaluated, and outperforms the other methods up to the case where 50,000 negative examples are used for training. At this point the DT technique performs as well as the Parzen OCC approach. Thereafter the DT method outperforms all the rest of the learning approaches, suggesting that DT is the traditional binary learning approach most influenced by the availability of the negative class information. Other conventional classifiers evaluated (NB, LR and SVM) do not show outstanding performance and slightly outperform the Parzen OCC method only when all available negative examples are used.

Finally we studied the effect of imbalanced classes on the performance of the different classifiers. While OCC methods are intrinsically able to cope with this situation, this is not the case for conventional classifiers. Consequently some strategy is needed to deal with the imbalanced class problem. Here we used a cost-sensitive analysis, where the misclassification cost for examples of the minority class is bigger than the misclassification cost for the majority class (note that on the different scenarios the minority class is not always the same). In situations when less negative than positive examples were used, we observe an increment in the performance of most of the conventional classifiers, reaching AUC-50 scores similar to those obtained for each approach in the balanced class scenario. The exception is the NB approach, the performance of which was almost invariant on these cases. When more negative than positive examples were used, the AUC-50 performance for all conventional classifiers tended to decrease in comparison with those obtained without cost-sensitive analysis. This can be explained because in these cases the classification model is generated considering positive and negatives examples information in a balanced way and not biased towards negative class information. Another accepted strategy to deal with the imbalanced class problem is to under-sample the majority class; we have done this when training on ten different balanced class sets (see section 3.2).

The analysis of the results presented in this section strongly suggests that the performance of conventional binary classification models is very dependant and highly affected by the presence or absence of negative examples. This can also explain the high performance (AUC-50) observed for conventional classifiers when all negative examples are employed for training. Another explanation to this observed high performance is the availability of a high-quality negative gold standard set (protein pairs located on different cell localization), which has been previously discussed in [11] and [13]. However this will not be the case when studying the prediction of PPI on other organisms when non protein cell localization is available.

3.4 Evaluation of biological feature importance

We evaluated the individual effect of the different biological features used in this research on the performance of the Parzen OCC approach. For this we removed each one of the biological attributes at time from the data set and tested the effect of this action on the AUC and AUC-50 scores, compared with those obtained when all available biological information is used. Table 2 shows the results of this procedure. The major effect on the Parzen OCC performance occurs when either functional similarity or m-RNA expression data are removed. This is consistent with the results reported in [7], [9] and [12]. It is interesting to observe that the overall AUC performance only increases when high-throughput information is removed, which can be explained due the high false-positive and false-negative rates associated with these kinds of features.

Table 2: Evaluation of the individual effect of the different biological attributes in the performance of the OCC parzen approach

Feature description	AUC	AUC-50
ALL features	0.9801 ± 0.0075	0.4010 ± 0.0282
GO removed	0.9186 ± 0.0121	0.2094 ± 0.0189
MIPS removed	0.9412 ± 0.0135	0.1983 ± 0.0225
m-RNA expression removed	0.9775 ± 0.0050	0.1883 ± 0.0238
Essentiality removed	0.9800 ± 0.0081	0.3380 ± 0.0273
High-throughput removed	0.9887 ± 0.0037	0.3463 ± 0.0261

3.5 Prediction of new potential PPI targets using Parzen OCC method

Finally we evaluated the ability of the Parzen OCC approach to predict new potential PPI, which could be used as a targets in future investigations. For this we generated a new set of randomly protein pairs which were not included on our positive and negative gold standards sets. We were able to collect a set of approximately 518,000 protein pairs examples with complete biological information from the data available in [6]. Using the Parzen OCC model trained on all positive examples available (parameters being optimized on ten fold cross validation procedure), we classified the examples in the random set, finding that 928 of them were predicted as a new potential PPI.

We focused in the analysis of the top 50 new potential PPI with the highest prediction scores generated by the Parzen OCC model. This score is the probability associated with the the positive examples class and consequently can be seen as a confidence value. To validate our predictions we employed the INTACT database (<http://www.ebi.ac.uk/intact>), which compiles molecular interactions reported in published literature, containing information of around 50,000 binary protein interactions for yeast (May 2006). We found that of the 50 top ranked examples, 36 were supported by at least one reference in INTACT. These were mostly associated with mass spectrometry experiments which are related with the identification of groups of proteins that interact to form complexes. This is statistically significant considering that if we randomly selected 50 protein pairs, not in the positive gold standard, the probability that 36 of them will be annotated in INTACT is very low ($p < 10^{-77}$) using Fisher's exact test [26]. The list of the top 50 potential new PPI targets is given in the supplementary material.

4 Conclusions

The research described in this paper has focused on the application and evaluation of one-class classification (OCC) methods for the problem of prediction of protein-protein interaction (PPI). We also considered the use of diverse biological data types in order to develop a joint integrative learning analysis.

Among various OCC methods evaluated, the Parzen OCC density estimation approach clearly exhibited the best performance. This approach was then selected to develop a comparative performance evaluation against several well known conventional machine learning methods. Different scenarios were considered varying the number of negative examples used to train the

models. We found that the Parzen OCC approach performs very competitively and outperforms the rest of conventional classifiers in most of the situations up to the case where the ratio of negative to positive examples is approximately 25 to 1.

We have demonstrated that for this specific task, the performance of conventional binary classification approaches is highly influenced by the quantity of negative examples used to train the respective models. This suggests that classification models generated from these type of methods are more reliant on negative information (in this case an untrustworthy set of negative PPI examples) than on positive information (experimentally corroborated PPI examples).

Our results suggest that the task of the prediction of PPI can indeed be formulated as an OCC problem where the predictive model is based on real (trustworthy) PPI data. In the specific case of prediction of co-complexed proteins we found that the Parzen OCC method is able to generate models which perform competitively with those generated by conventional classifiers, independently of the quality and quantity of the negative examples available. We have also carried out an initial study about the ability of the Parzen OCC approach to predict new potential PPI targets, showing that many of the highly ranked new predictions can be validated by reference to published results in the literature.

Acknowledgements

This work is supported by the Programme $Al\beta$ an, the European Union Programme of High level Scholarships for Latin America, scholarship E04D034854CL.

References

- [1] P. Uetz, L. Giot, G. Cagney, T. A. Mansfield, R. S. Judson, et al. A comprehensive analysis of protein-protein interactions in *saccharomyces cerevisiae*. *Nature*, 403(6770):623–627, 2000.
- [2] T. Ito, T. Chiba, R. Ozawa, M. Yoshida, M. Hattori, and Y. Sakaki. A comprehensive two-hybrid analysis to explore the yeast protein interactome. *Proc Natl Acad Sci*, 98(8):4569–4574, 2001.
- [3] A. C. Gavin, M. Bosche, R. Krause, P. Grandi, M. Marzioch, et al. Functional organization of the yeast proteome by systematic analysis of protein complexes. *Nature*, 415(6868):141–7, 2002.
- [4] Y. Ho, A. Gruhler, A. Heilbut, G. D. Bader, L. Moore, et al. Systematic identification of protein complexes in *saccharomyces cerevisiae* by mass spectrometry. *Nature*, 415(6868):180–183, 2002.
- [5] C. von Mering, R. Krause, B. Snel, M. Cornell, S. G. Oliver, et al. Comparative assessment of large-scale data sets of protein-protein interactions. *Nature*, 417(6887):399–403, May 2002.
- [6] R. Jansen, H. Yu, D. Greenbaum, Y. Kluger, N. J. Krogan, et al. A bayesian networks approach for predicting protein-protein interactions from genomic data. *Science*, 302(5644):449–453, 2003.
- [7] N. Lin, B. Wu, R. Jansen, M. Gerstein, and H. Zhao. Information assessment on predicting protein-protein interactions. *BMC Bioinformatics*, 5(1):154, 2004.
- [8] L. Zhang, S. Wong, O. King, and F. Roth. Predicting co-complexed protein pairs using genomic and proteomic data integration. *BMC Bioinformatics*, 5(1):38, 2004.

- [9] L. J. Lu, Y. Xia, A. Paccanaro, H. Yu, and M. Gerstein. Assessing the limits of genomic data integration for predicting protein networks. *Genome Res.*, 15(7):945–953, 2005.
- [10] Y. Qi, J. Klein-Seetharaman, and Z. Bar-Joseph. Random forest similarity for protein-protein interaction prediction from multiple sources. In *Pacific Symposium on Biocomputing*, 2005.
- [11] A. Ben-Hur and W. S. Noble. Kernel methods for predicting protein-protein interactions. *Bioinformatics*, 21(suppl 1):i38–i46, 2005.
- [12] Y. Qi, Z. Bar-Joseph, and J. Klein-Seetharaman. Evaluation of different biological data and computational classification methods for use in protein interaction prediction. *Proteins: Structure, Function, and Bioinformatics*, 63(3):490–500, 2006.
- [13] A. Ben-Hur and W. S. Noble. Choosing negative examples for the prediction of protein-protein interactions. *BMC Bioinformatics*, 7(Suppl 1):S2, 2006.
- [14] N. V. Chawla, N. Japkowicz, and A. Kotcz. Editorial: special issue on learning from imbalanced data sets. *SIGKDD Explorations*, 6(1):1–6, 2004.
- [15] D. M. J. Tax and R. P. W. Duin. Support vector data description. *Machine Learning*, 54(1):45–66, 2004.
- [16] H. W. Mewes, D. Frishman, U. Guldener, G. Mannhaupt, K. Mayer, et al. Mips: a database for genomes and protein sequences. *Nucl. Acids Res.*, 30(1):31–34, 2002.
- [17] T. R. Hughes, M. J. Marton, A. R. Jones, C. J. Roberts, R. Stoughton, et al. Functional discovery via a compendium of expression profiles. *Cell*, 102(1):109–126, July 2000.
- [18] R. J. Cho, M. J. Campbell, E. A. Winzeler, L. Steinmetz, A. Conway, et al. A genome-wide transcriptional analysis of the mitotic cell cycle. *Mol Cell*, 2(1):65–73, July 1998.
- [19] M. Ashburner, C. A. Ball, J. A. Blake, D. Botstein, H. Butler, Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. *Nat Genet*, 25(1):25–29, May 2000.
- [20] I. H. Witten and E. Frank. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, San Francisco, 2nd edition, 2005.
- [21] J. Huang and C. X. Ling. Using auc and accuracy in evaluating learning algorithms. *IEEE Trans. Knowl. Data Eng.*, 17(3):299–310, 2005.
- [22] C. Drummond and R. C. Holte. Learning to live with false alarms. In *Workshop on Data Mining Methods for Anomaly Detection*. Eleventh ACM SIGKDD, August 2005.
- [23] F. Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83, 1945.
- [24] H. Alashwal, S. Deris, and R. M. Othman. One-class support vector machines for protein-protein interactions prediction. *International Journal of Biomedical Sciences*, 1(2):120–127, 2006.
- [25] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001.
- [26] R. A. Fisher. On the interpretation of χ^2 from contingency tables and the calculation of p . *Journal of the Royal Statistical Society*, 85(1):87–94, 1922.