

Identifying novel peroxisomal proteins

John Hawkins,^{1,2*} Donna Mahony,³ Stefan Maetschke,^{1,2} Mark Wakabayashi,^{1,2}
Rohan D. Teasdale,³ and Mikael Bodén²

¹ ARC Centre for Complex Systems, The University of Queensland, St. Lucia, Queensland 4072, Australia

² School of Information Technology and Electrical Engineering, The University of Queensland, St. Lucia, Queensland 4072, Australia

³ Institute for Molecular Bioscience and ARC Centre in Bioinformatics, The University of Queensland, St. Lucia, Queensland 4072, Australia

ABSTRACT

Peroxisomes are small subcellular compartments responsible for a range of essential metabolic processes. Efforts in predicting peroxisomal protein import are challenged by species variation and sparse sequence data sets with experimentally confirmed localization.

We present a predictor of peroxisomal import based on the presence of the dominant peroxisomal targeting signal one (PTS1), a seemingly well conserved but highly unspecific motif. The signal appears to rely on subtle dependencies with the preceding residues. We evaluate prediction accuracies against two alternative predictor services, PEROXIP and the PTS1 PREDICTOR. We test the integrity of prediction on a range of prokaryotic and eukaryotic proteomes lacking peroxisomes. Similarly we test the accuracy on peroxisomal proteins known to not overlap with training data. The model identified a number of proteins within the RIKEN IPS7 mouse protein dataset as potentially novel peroxisomal proteins. Three were confirmed in vitro using immunofluorescent detection of myc-epitope-tagged proteins in transiently transfected BHK-21 cells (Dhrs2, Serhl, and Ehhadh).

The final model has a superior specificity to both alternatives, and an accuracy better than PEROXIP and on par with PTS1 PREDICTOR. Thus, the model we present should prove invaluable for labeling PTS1 targeted proteins with high confidence. We use the predictor to screen several additional eukaryotic genomes to revise previously estimated numbers of peroxisomal proteins. Available at <http://pprowler.itee.uq.edu.au>.

Proteins 2007; 69:606–616.
© 2007 Wiley-Liss, Inc.

Key words: subcellular localization; peroxisome; prediction; machine learning; peroxisomal proteins; immunofluorescent detection.

INTRODUCTION

Peroxisomes are relatively small compartments that are particularly abundant in liver cells and neurons. In plants they occur in large numbers as metabolically specialized microbodies. The peroxisome plays an important role in lipid, ethanol and glyoxylate metabolism, and detoxification of reactive oxygen species. They are believed to be essential for coping with oxidative stress. Peroxisomal disorders often involve abnormal accumulation of long fatty acids which subsequently impact membrane structure.¹ The mechanisms of protein localization are one potential avenue to be explored for disease control.²

There are numerous mechanisms of subcellular localization, depending on the organelle and the function of the protein. In general they all rely on the existence of some form of targeting signal, which is recognized by chaperone proteins that then effect the localization.^{3–5} However, the process is complicated by the dynamical aspects of cellular life. Targeting to the endoplasmic reticulum is achieved by an N-terminal signal peptide that is recognized by a signal recognition particle as it emerges from the ribosome. The localization process thus occurs in tandem with translation, and will dominate other functional localization pathways, including those for the peroxisome.

Targeting to the peroxisome only happens after the completion of translation on free ribosomes. The transport of peroxisomal matrix proteins is believed to rely on a small number of sequence motifs. The dominating targeting signal is the PTS1 which appears at the C-terminus. The PTS1 consists of a strongly conserved tripeptide, but there are constraints that range over a larger region exposed to the Pex5 that play a central role in import.^{6–9}

In recent efforts, computational methods have been used in genome-wide screens for peroxisome-like sequence features to single-out previously unknown proteins and to subsequently identify novel peroxisomal enzymes and putative regulatory proteins.¹⁰ More

Grant sponsor: Australian Research Council; Grant sponsor: Australian National Health and Medical Research Council.

*Correspondence to: John Hawkins, ARC Centre for Complex Systems, The University of Queensland, St. Lucia, Queensland 4072, Australia. E-mail: jhawkins@itee.uq.edu.au

Received 5 October 2006; Revised 18 December 2006; Accepted 11 January 2007

Published online 16 July 2007 in Wiley InterScience (www.interscience.wiley.com).

DOI: 10.1002/prot.21420

broadly, subcellular localization has been approached as a multiclass prediction problem. Models that predict localization of proteins on the basis of their amino acid sequences can be used to screen and annotate unknown proteins and to guide experimental efforts to where they are best needed. Several computational approaches have shown promises including manually programmed expert systems and data-driven machine learning models.^{3–5} Machine learning models of protein localization are often based on probabilistic techniques,^{11–13} neural networks,^{14–17} and support vector machines (SVMs).^{18–20}

Besides general subcellular localization predictors,^{20–22} there are few predictors that specialize in predicting peroxisomal targeting: PEROXIIP by Emanuelsson *et al.*²³ and PTS1 PREDICTOR by Neuberger *et al.*⁶ both address only the PTS1 signal that is located at the C-terminus. Both exclude the nonapeptide motif (PTS2) and membrane-bound peroxisomal proteins due to the lack of experimental sequence data.

PEROXIIP has three stages in its processing of candidate sequences. An initial step filters out sequences with predicted signal peptides (to eliminate those proteins cotranslationally inserted into the ER) and those sequences predicted to have a membrane spanning region. Second, a motif identification module examines the C-terminus of the sequence and filters out those proteins that do not fit a general pattern of approved tripeptide motifs. In the final step a machine learning model examines the nine amino acids preceding the PTS1 motif and provides a prediction of whether the sequence is peroxisomal or not. In the PEROXIIP model, the machine learning module consists of a neural network and a SVM that are trained independently and operate as a miniature ensemble. The final output is the union of their individual outputs, that is the sequence is predicted as peroxisomal if either the neural network or the SVM indicates so.²³

Neuberger *et al.* manually formulated a scoring function using a set of sequences experimentally confirmed to interact with the Pex5 complex responsible for PTS1 import. They also analyzed a larger sequence set extracted from public databases. The scoring function that is used by PTS1 PREDICTOR considers a “composite profile term that evaluates concordance with amino acid preferences” and a set of physical property patterns that were identified from the data sets.⁶ Owing to some species-specific variation in targeting (via PTS1), three versions were developed: Metazoa, Fungi, and General. Both sensitivity and specificity were reportedly very high.

In previously reported simulations we demonstrated that the prediction accuracy of a PEROXIIP-like architecture could be improved by including the highly conserved tripeptide in the window of residues supplied to the machine learning module. The model was making use of interdependencies between the tripeptide and the nine preceding amino acids.²⁴ The SVM model was refined by investigating the use of an amino acid composition

window and exploring the space of kernel parameters more thoroughly. The final predictor, PTS1PROWLER, utilizes a fifth order polynomial kernel with a separately trained logistic output function.²⁵

In the present paper we review the structure of our final model and perform a number of analyses to illustrate the accuracy and ability of PTS1PROWLER. We compare PTS1PROWLER against PEROXIIP and PTS1 PREDICTOR using the most recent updates of SWISS-PROT (guaranteed to not overlap with the training data of any model), and by screening proteomes of organisms known not to contain peroxisomes, both eukaryotes and prokaryotes. Furthermore, a selection of eukaryotic proteomes are screened. We report on updated estimates of the number of peroxisomal matrix proteins in each organism. Owing to the model’s high specificity, our screens should contain few false positives. Finally we screen the RIKEN IPS7 mouse protein dataset for potentially novel peroxisomal proteins and confirm three of the predictions *in vitro* using immunofluorescent detection of myc-epitope-tagged proteins in transiently transfected BHK-21 cells.

The PTS1PROWLER prediction service is integrated into the PROTEINPROWLER suite available at <http://pprowler.itee.uq.edu.au>. Related sequence data is made available at the web site.

MATERIALS AND METHODS

Replication of the PeroxiP data set

The data set of PTS1-containing peroxisomal proteins and nonperoxisomal proteins, as identified by Emanuelsson *et al.* was extracted from SWISS-PROT R39.27. The initial data set included 152 proteins identified as peroxisomal, and 308 identified as nonperoxisomal. Redundancy reduction had not been performed on this data set. The training and test data was finalized as detailed by Emanuelsson *et al.*

Highly similar proteins were removed, such that each pair of proteins differed in at least two positions in the nine residues preceding the C-terminal tripeptide.

The final stage of redundancy reduction was performed using BLASTClust. In order to reproduce a data set of the same size we found that a similarity threshold of 1.675 was required (using BLASTClust’s *-s* option). By our estimations this corresponds to between 80 and 85% similarity over the aligned region, which by default is set to 90% of the length of the sequences.

The final number of these was 90 peroxisomal proteins and 160 nonperoxisomal proteins. These numbers differ only slightly to the reported data set of 90 and 151 used as training and test data in PEROXIIP.

New data set

The new data set was collected from SWISS-PROT R45. We followed a similar process to that of PEROXIIP to

extract an initial set of peroxisomal sequences. All SWISS-PROT entries were searched for those containing a “SUBCELLULAR LOCATION” annotation in the comments field that included any of “PEROXISOM,” “GLYOXSOM,” or “GLYCOSOM” using a case-insensitive search. The proteins were also required to identify a “Microbody targeting signal” in the feature table (FT), indicating a PTS1-targeted protein. This gave an initial set of 202 proteins.

The proteins found were then filtered manually for proteins not likely to be targeted by a PTS1 and for membrane proteins. Full details of the curation process are available in the online documentation of the predictor.

The negative set was generated by extracting from SWISS-PROT R45 all eukaryotic proteins with a C-terminal tripeptide identical to that of one of the initially identified peroxisomal proteins, and a subcellular localization not specified as peroxisomal, glyoxysomal, or glycosomal. These proteins were then manually curated for potentially fallacious annotations, full details of which are given in the online documentation.

The final data set contained 206 peroxisomal proteins and 348 nonperoxisomal proteins. Redundancy reduction was performed by similar processes to those performed on the PEROXIP data set (mentioned earlier). Because the predictor was trained and tested on the complete C-terminal 12 residues, these residues were tested across both the positive and negative sets to ensure no pairs of sequences had less than two differing residues. This reduced the data set to 157 and 239 proteins for the peroxisomal and nonperoxisomal sets, respectively. Clustering of the remaining sets using BLASTClust with a similarity threshold of 1.675 (80–85% similarity over the aligned region of minimum 90% of the sequence) and extraction of representatives from each cluster (the same procedure used to replicate Emanuelsson’s data set) resulted in a final testing and training set of 124 peroxisomal proteins and 214 nonperoxisomal proteins. The ratio of positives and negatives was therefore similar to those of PEROXIP’s set, with the overall size of the data set increased by ~40%.

PTS1Prowler development

As shown in Figure 1, the overall structure of PTS1PROWLER is similar to that of PEROXIP. Processing occurs in the three distinct stages. An initial motif filter rejects sequences with a C-terminal tripeptide not occurring amongst peroxisomal proteins in SWISS-PROT R45. A machine learning module analyses the final 12 residues of the sequence and gives probabilistic prediction as to whether the protein is PTS1-targeted. Finally, a signal peptide prediction application is used to filter proteins that are likely to be secreted.

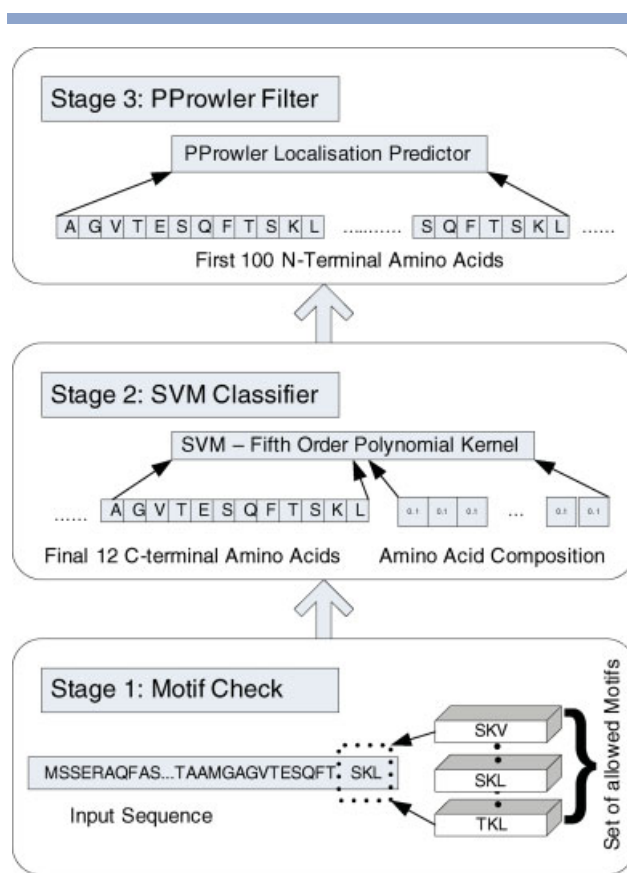


Figure 1

Overview of the three stages of the final Pts1Prowler model. The first stage filters out sequences without a PTS1 motif. The second stage runs the SVM classifier over the final 12 residues with an orthonormal encoding and an amino acid composition window. If this stage produces a positive prediction, then the third stage is executed, by which the Protein Prowler application examines the sequence for a potential signal peptide. If the Protein Prowler application gives a score greater than 0.85 then the prediction is changed to nonperoxisomal; otherwise, the output from Stage two is given. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Performance metrics

Performance is evaluated using several metrics. The Matthews correlation coefficient²⁶ (MCC) is a performance statistic that takes into account the numbers of true positives (tp), true negatives (tn), false positives (fp), and false negatives (fn). It is calculated as follows:

$$\frac{tp \times tn - fp \times fn}{\sqrt{(tp + fn)(tp + fp)(tn + fp)(tn + fn)}} \quad (1)$$

The *sensitivity* is a measure of the classifiers predilection for always identifying positive examples calculated thus:

$$\frac{tp}{tp + fn} \quad (2)$$

The *specificity* of the classifier is a measure of the classifier's tendency to avoid making false-positive predictions calculated as follows:

$$\frac{tn}{tn + fp} \quad (3)$$

The machine learning module

Perhaps the most crucial stage of the model is the machine learning module which decides whether a potential sequence with a valid PTS1 motif is genuinely a peroxisomal protein. This component has been the focus of our development efforts.

It has been shown that essential sequence dependencies occur within the last 12 residues of PTS1-targeted proteins, in particular between the PTS1 motif and the preceding nine residues.^{6,27} In our initial studies of this problem we established that the machine learning module could be improved considerably by allowing it to learn the dependencies between the PTS1 motif and the preceding nine residues.²⁴

The PEROXIP predictor was reported to have employed an amino acid composition window, as well as the 9-mer preceding the PTS1 motif as input to the machine learning module.²³ The amino acid composition encoding is formed by summing the number of occurrences of each amino acid in the entire protein, and then producing a probability vector, with each element giving the likelihood that a random position in the sequence would contain the corresponding amino acid. The amino acid composition has been shown to be at least correlated crudely with localization,²⁸ and had been used as an encoding for machine learning prediction of localization to various subcellular locations (including the peroxisome), with moderate success.²⁹ In a previous study we established that the contribution of the amino acid composition to SVM classification varied with kernel choice, however, for all but one kernel (Gaussian $\gamma = 0.01$), it improved performance, with an average improvement of 0.07 points of MCC across a range of kernels.²⁵

By default a SVM will produce a binary classification on a two-class problem. For many reasons it is desirable that a classifier service produces a posterior probability score for a given input. This can be achieved with SVMs by fitting a logistic model to the output of the trained classifier.³⁰ For this model the output probability of the target class is given by

$$P(y = 1|f(x)) = \frac{1}{1 + e^{Af(x)+B}} \quad (4)$$

where $f(x)$ is the output of the SVM for sample x , and the parameters A and B must be estimated.

In the previous study, we also established that the use of the logistic output function not only failed to impede

the performance of the model but also improved the performance for almost all kernel choices. The capacity for logistic outputs to improve performance was noted by Platt, but is deemed to be difficult to predict in advance.³⁰ Following Platt's recommendation and examples, we used an internal cross-validation within the training set (threefold) to estimate the parameters for the logistic function.³⁰ This means that two-thirds of the training set are used for training the SVM and the final third is used for estimating the parameters of the logistic function via the maximum likelihood method.

The final implementation of PTS1PROWLER is the product of all these insights into peroxisomal protein prediction. We use a fifth-order polynomial kernel without lower order terms and a logistic output function. The input is given via two windows, one containing the final 12 residues and another containing the amino acid composition. Following previous results we use the orthonormal encoding for the primary sequence window.^{23,24} The optimization procedures for adapting the SVMs and the logistic output models are drawn from Weka library of machine learning tools.³¹

In our exploration of the various free parameters available to fine-tune the SVM, we found that the complexity constant C and the round-off error ϵ had very little effect on the results. Only at extreme values did the performance of the model change, and then only in a deleterious fashion. Minor improvements were achieved by tuning the tolerance parameters, so that the average MCC increased from 0.741 to 0.749, with a final value of $T = 0.038$. The insensitivity to the setting of regularization parameters indicate that the separation in the feature space defined by the polynomial kernel is simple. However, it is worth noting that due to computation time constraints it was assumed that optimal values for these parameters could be established independently of one another.

Filtering signal peptides

The final stage of the model filters out proteins that possess a signal peptide, because they will not have an opportunity to interact with Pex5 even if they have a functional PTS1.

In PTS1PROWLER the filtering of secreted proteins is performed by the PROTEIN PROWLER subcellular localization model.¹⁷ However, in our model, this step is performed last rather than first. Its use is deferred as it is more computationally expensive than either the motif filter or the SVM-driven classification. The reordering has no effect on the classification, and only on the efficiency of the model.

Running PROTEIN PROWLER over the training sets revealed only one peroxisomal protein P24552 (D-amino-acid oxidase) with a significant prediction of containing a signal peptide (0.832). A threshold value of 0.85 was

chosen as the cutoff for the PROTEIN PROWLER filter stage, yielding no predictions from the positive set and 49 predictions from the negative training set.

Because no peroxisomal proteins are predicted as secreted by the PROTEIN PROWLER filter, the sensitivity of the model does not decline by including the filter. The PROTEIN PROWLER filter increased specificity by an average of 0.020 (an increase of 2%) and the MCC by an average of 0.017 (an increase of 2%). Decreasing the threshold could further increase the gain in accuracy, but may result in reduced sensitivity when exposed to a larger data set.

Training and testing

The entire model was trained and tested using the jackknife procedure, in which the data set is separated into as many sets as there are samples. In each set, a single example is set aside for testing. The procedure involves creating a model for each set, training each model separately, and then testing on the single test case. All of the test results are then used to produce an overall estimate of the performance of the technique on the problem. Usually this testing procedure precludes the use of multiple runs and the reporting of standard deviations, simply because the data can be split only one way, and an SVM (unlike a neural net) is not sensitive to initialization conditions and will thus perform identically on repeated exposures to the same data. However, the use of the logistic output function requires that the training data be split between that used to train the SVM and that used to fit the logistic model. Hence, our jackknife simulations are performed five times with different seeds and we report mean value of all metrics over the five runs.

In vitro peroxisome subcellular localization assay

N-Terminally tagged myc-gene of interest expression constructs were generated using a modified overlapping PCR methodology.³² Generated expression constructs comprised of a promoter fragment, the cDNA of interest, and a terminator fragment. The PCR primers, reactions, and amplification conditions used have been reported previously.³² BHK – 21 (Syrian golden hamster, kidney fibroblasts, ATCC CCL 10) cells were cultured in Dulbecco's modified essential Media (Life Technologies, Grand Island, NY) supplemented with 10% fetal bovine serum (Cambrex, NJ), and 2 mM L-glutamine (Life Technologies) and maintained in a humid, 5% CO₂ environment at 37°C. Subconfluent BHK-21 cells cultured on glass coverslips coated with 0.01% poly-L-lysine (Sigma Aldrich, St. Louis, MO) were transiently transfected with PCR expression constructs diluted with OptiMEM (Life Technologies) using Lipofectamine2000 (Life Technologies) as per the manufacturer's instructions. Twenty hours post-transfection BHK-21 cells were fixed in 4% paraformal-

dehyde (Sigma Aldrich) in phosphate-buffered saline (PBS) for 20 min prior to permeabilization in 0.1% Triton-X 100 (Sigma Aldrich) in PBS for 5 min followed by three washes in PBS. Blocking was performed in 2% bovine serum albumin (BSA, Sigma Aldrich in PBS) prior to incubation with the monoclonal anti-myc antibody (Cell Signaling Technology, Beverly, MA, 1:3000 in 2% BSA) and rabbit anti-catalase (Abcam, Cambridge, UK, 1:2000 in 2% BSA) for 60 min at room temperature. The cells were subsequently washed in PBS prior to incubation with the secondary antibodies goat anti-mouse IgG-cy3 conjugated antibody (Zymed, San Francisco, diluted 1:600 in 2% BSA) and goat anti-rabbit IgG-Alexa-488 antibody (Molecular Probes, Eugene, OR) for 45 min at room temperature and the coverslips were mounted using MO-WIOL (Calbiochem, Nottingham, UK). Representative images of the observed localization patterns were captured for each construct using an Olympus AX-70 upright fluorescence microscope. Images were prepared using Adobe Photoshop CS2 (Adobe Systems).

Latest peroxisomal proteins

To perform an independent test of the generalization of the model compared to its competitors, we extracted a set of peroxisomal proteins not present in the training set.

The protein sequences were extracted from SWISS-PROT R48 (excluding TREMBL) and filtered, such that the SUBCELLULAR LOCATION line of the CC field ended with Peroxisomal. These proteins were then filtered for the version addition identifiers "46, Created," "47, Created," and "48, Created". Proteins meeting these criteria were further filtered for a MOTIF entry in the FT containing the string Microbody targeting signal. This resulted in a final set of 17 protein sequences.

Eukaryotic genome screens

The data used for the eukaryotic genome screens was obtained from a number of different sources. In each case we used the data as it was obtained without any filtering. The URLs for each proteome file are listed as follows:

- *S. pompe*
ftp.sanger.ac.uk/pub2/yeast/pombe/Protein_data/pompep
- *S. cerevisiae*
ftpmips.gsf.de/yeast/sequences/Scerevisiae_prot_2006200
- *D. melanogaster*
ftp.ebi.ac.uk/pub/databases/edgp/sequence_sets/aa_gadfly.dros.Z
- *C. elegans*
ftp.sanger.ac.uk/pub2/wormbase/sequences/WORMPEP/wormpep

- *A. thaliana*
ftp.mips.gsf.de/cress/arabiprot/arabi_all_proteins_v020204.gz
- *M. musculus*
ftp.ensembl.org/pub/current_mouse/data/fasta/pep/Mus_musculus.NCBIM34.jul.pep.fa.gz
- *H. sapiens*
ftp.ncbi.nih.gov/genomes/H_sapiens/protein/Gnomon_prot.fsa.gz
- *O. sativa*
http://www.tigr.org/tdb/e2k1/osa1/data_download.shtml

RESULTS

PTS1Prowler

Emanuelsson *et al.* published values of 0.50, 0.78, and 0.64 for MCC, sensitivity, and specificity, respectively (values were obtained from a single run of fivefold cross-validation). When trained on the same data set, PTS1PROWLER yielded an average MCC of 0.76. With the updated data set the average MCC is 0.77, which is a 54% improvement on PEROXIIP. It is interesting to note that the sensitivity of PTS1PROWLER is identical to that of the original PEROXIIP (0.78), all of the improvement in performance has come from an increase of 45% in the specificity of the model (0.93).

An ROC analysis was performed using the average of the outputs from the five simulations (Fig. 2). The ROC

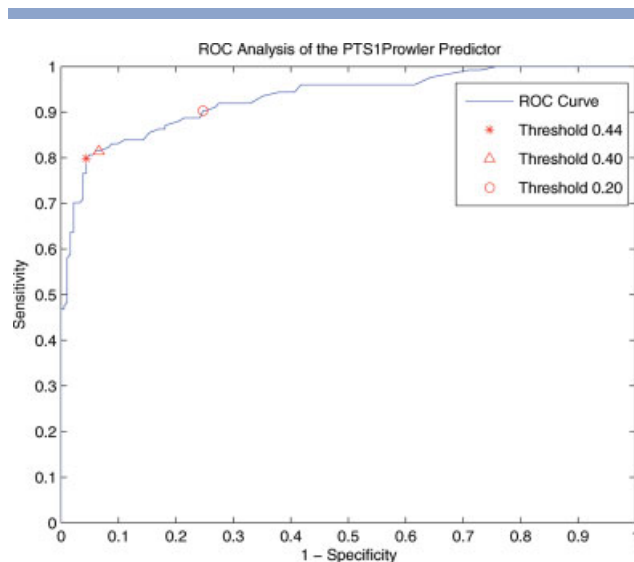


Figure 2

A receiver operating characteristic curve of the performance of Pts1Prowler. We have marked the location of three different output thresholds. At the threshold 0.44, the predictor is operating with maximal MCC. Using 0.4 gives a slight increase in sensitivity at the cost of specificity. Finally 0.2 raises the sensitivity of the predictor to 0.9. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Table I

The Prediction Results for the Three PTS1 Predictors on Experimentally Verified PTS1 Targeted Proteins Released Since R45 Until R48 of Swiss-Prot

Protein	Novel peroxisomal protein prediction					
	PeroxiP	Pts1 Predictor		Pts1Prowler		
		Method 1	Raw	$\theta = 0$	Raw	$\theta = 0.5$
NDX8_CAEEL	—	9.17	✓	0.71	✓	✓
SUOX_ARATH	—	-3.06	—	0.41	—	✓
PECR_CAVPO	—	6.63	✓	0.95	✓	✓
DECR2_MOUSE	—	4.39	✓	0.34	—	✓
PECR_HUMAN	—	6.90	✓	0.88	✓	✓
OPR3_LYCES	—	11.88	✓	0.40	—	✓
OPR3_ARATH	✓	9.46	✓	0.69	✓	✓
URIC_MACFA	✓	14.13	✓	1.00	✓	✓
NUD12_HUMAN	—	-17.77	—	0.00	—	—
URIC_AOTTR	✓	13.72	✓	1.00	✓	✓
URIC_MACMU	✓	13.72	✓	1.00	✓	✓
DECR2_RAT	—	3.53	✓	0.61	✓	✓
DCMC_RAT	✓	4.56	✓	1.00	✓	✓
ACOX_PICPA	—	3.33	✓	0.00	—	—
NUD7_MOUSE	✓	-12.22	—	0.69	✓	✓
DCMC_MOUSE	✓	2.47	✓	1.00	✓	✓
ACOX4_ARATH	✓	11.03	✓	0.63	✓	✓
Total	8		14		12	15
Accuracy (%)	47		82		71	88

When applicable, classifications are made on basis of the predictors' raw output according to the specified thresholds θ (Pts1 default $\theta = 0$ and Pts1Prowler default $\theta = 0.5$, compensated $\theta = 0.22$).

analysis revealed that the optimal threshold for a positive prediction is around 0.44 yielding an MCC of 0.78. In the name of making a conservative estimate we do not use this as the final estimate of MCC, but we recommend the threshold value for users of PTS1PROWLER.

The latest peroxisomal protein test

In order to perform an independent test of the sensitivity of PTS1PROWLER, we procured the set of peroxisomal proteins added to the SWISS-PROT database since the development of the training sets. The 17 sequences that were identified were presented to PTS1PROWLER, PEROXIIP, and PTS1 PREDICTOR. Results are shown in Table I.

For PEROXIIP we report only their Method 1, as all other methods involve relaxation of constraints seriously compromising specificity, and as such do not offer a reasonable measure of the PEROXIIP's accuracy.

On data used for model development, PTS1 PREDICTOR has been reported to have a much higher sensitivity than PEROXIIP or PTS1PROWLER. However, the specificity of PTS1 PREDICTOR was not established. The ROC analysis on PTS1PROWLER enables the determination of an adjusted threshold that the data set used for development gives the same sensitivity, 0.896, as reported for PTS1 PREDICTOR (general version, henceforth PTS1 PREDICTOR-General). With the adjustment, the accuracy of PTS1PROWLER on the novel PTS1 sequences improves from 12/17 (with

Table II

PeroxiP and Pts1Prowler Applied to All Available Proteins in Swiss-Prot Release 47.5 from Three Genera of Eukaryotes for Which There Is No Evidence of Peroxisomes

Genus	Nonperoxisomal proteome comparisons		
	No. of proteins	PeroxiP	Pts1Prowler
<i>Trichomonas</i>	15	0	0
<i>Giardia</i>	35	2	0
<i>Entamoeba</i>	55	0	0

The number of available proteins is displayed in the first column, followed by the number of proteins predicted to be peroxisomal by PeroxiP and Pts1Prowler.

0.5 as threshold) to 15/17 (with 0.22 as threshold), one better than PTS1 PREDICTOR-General.

Interestingly, PTS1 PREDICTOR and PTS1PROWLER differ in their predictions. By using the logical OR-function on the output of the two, all but one of the proteins are predicted as PTS1-targeted, the elusive protein is NUD12_HUMAN protein which was not predicted as peroxisomal by any of the three. The terminal tripeptide of this protein is PNL, which is not included in the acceptance set for either PEROXIIP or PTS1PROWLER, and hence it is automatically rejected by both classifiers.

Organisms without peroxisomes

Although peroxisomes are present in most eukaryotic cells, there are numerous species of parasitic eukaryotes which are believed not to possess a peroxisome. For protozoa of the genus *Giardia*, *Trichomonas* (agent of Vaginitis), or amoebas of the genus *Entamoeba* (agents of dysentery and ulceration of the colon and liver), there is no evidence for the presence of a peroxisome.³³ Thus, their proteomes should not contain proteins with a valid PTS1 except by mere chance.

The proteomes of peroxisome-less organisms offer an opportunity to test the integrity of classifiers. By feeding the set of known proteins from each genus through both PEROXIIP and PTS1PROWLER, we report how many proteins are fallaciously labeled peroxisomal, thus giving an independent estimate of the specificity of the models. The data sets for this study were retrieved from SWISS-PROT R47.5 (within each taxonomy, excluding protein fragments). For PEROXIIP we conservatively use only Method 1 predictions, in which the most restrictive criteria are used. For PTS1PROWLER we use the neutral cutoff of 0.5.

The results displayed in Table II reinforce the findings of the preceding simulations. PEROXIIP identified two proteins as peroxisomal, a false-positive rate of 1.9%, whereas PTS1PROWLER identified none. Owing to restrictions in the number of sequences that may be presented to the PTS1 PREDICTOR service, we were unable to easily test PTS1 PREDICTOR this way.

To provide an estimate of the accuracy of PTS1PROWLER in relation to PTS1 PREDICTOR, PTS1PROWLER was tested

on the entire proteomes of a range of prokaryotic organisms, all of which were screened previously by PTS1 PREDICTOR.⁶ This test provides an additional assessment of the predictors' relative specificity due to the fact that prokaryotes completely lack subcellular compartments and, consequently, have no selective conservation of PTSs.

The results of the PTS1PROWLER tests are shown in Table III. The collection of protein sequences were downloaded from the NCBI whole genome retrieval site.

Neuberger *et al.* report the total number of false positives for the same genomes (but older versions with a slightly lower total) for their three predictors as follows: Metazoa 70/20,544 (0.34%), Fungi 39/20,544 (0.19%), and General 152/20,544 (0.74%). Thus, with a false-positive rate on prokaryotic data of 0.08% our predictor has superior specificity to all three of the PTS1 PREDICTOR models.⁶ However, the sensitivity of PTS1 PREDICTOR is given as 0.896, significantly higher than our model using the default threshold. In order to obtain a more objective comparison of the two models, we again adjusted the threshold (0.22) to achieve a comparable sensitivity. With the adjusted threshold PTS1 PROWLER yielded an additional 35 false positives, taking the total to 52. This gives a false-positive rate of 0.24%, which is still superior to PTS1 PREDICTOR-Metazoa and PTS1 PREDICTOR-General.

Experimental validation

The model was used to screen the entire RIKEN IPS7 mouse protein dataset, a high quality proteome derived exclusively from full-length transcripts.^{34,35} This data set was chosen for the selection of candidates for experimental validation. The results of the screen are shown in Table IV. The results are broken into three categories based on the entries in the LOCATE database (as of May 5, 2006).³⁵ The label *Peroxi* is given for proteins that have some evidence of being peroxisomal, either by direct experiment or through similarity. *Nonperoxi* applies to proteins that have some evidence that they are located

Table III

Number of Prokaryotic Proteins Within Each of the Selected Organisms Predicted to Be Peroxisomal by Pts1Prowler

Species	Prokaryotic proteome screens		
	No. of proteins	Predictions	FP rate
<i>Escherichia coli</i> K12	5,379	7	0.13
<i>Pseudomonas aeruginosa</i> PA01	5,567	2	0.04
<i>Helicobacter pylori</i> -strain J99	1,491	1	0.07
<i>Staphylococcus epidermidis</i> ATCC 12228	2,419	3	0.12
<i>Staphylococcus aureus</i> subsp. <i>aureus</i> N315	2,588	2	0.08
<i>Mycobacterium tuberculosis</i> CDC1551	4,189	2	0.05
Total	21,633	17	0.08

Table IV*The Results of Screening the RIKEN IPS7 Data Set with Pts1Prowler*

RIKEN IPS7 proteome screen							Total
Positive predictions			Negative predictions				
Peroxi	Nonperoxi	Unknown	Peroxi	Nonperoxi	Unknown		
		SL			SL		
33	13	48	38	12452	20938	33,451	


elsewhere. Finally *Unknown* is given to proteins for which there is no available information.

Of the 48 proteins of no known subcellular localization, five were chosen for experimental validation. Mammalian expression constructs were engineered to encode the full-length ORF of the five proteins. The subcellular localization of the myc-epitope-tagged proteins was

determined using immunofluorescent detection of the myc epitope in transiently transfected BHK-21 cells.

Three of the proteins (PA113040.1 Dhrrs2, PA226887.1 Serhl, PA74692.8 Ehhadh) localized to distinct cytoplasmic puncta [Fig. 3(A,D,G)]. Costaining with a rabbit anti-catalase antibody [Fig. 3(B,E,H)] that detects endogenous peroxisome protein confirmed that these three proteins are localized to the peroxisome [Fig. 3(C,F,I)].

Two of the constructs computationally predicted to be peroxisomal proteins were found to have subcellular localizations outside of the peroxisome (data not shown). One protein (PA103003.1 Zfp438) localized to the nucleus and protein (PA111509.2 Hmgcl) showed cytoplasmic localization. The former was only weakly predicted as peroxisomal (0.52) and has since been identified as containing a zinc finger motif. The later protein has been generally predicted to be mitochondrial, and hence poses something of a general problem for predictors. Overall,

**Figure 3**

Experimental validation of predicted peroxisomal proteins. Subconfluent BHK-21 cells were transiently transfected with N-terminal, myc-epitope-tagged expression constructs generated by overlapping PCR. These constructs were expressed for 20 h prior to fixation, immunodetection, and subsequent visualization as described in the Materials and Methods. Three overexpressed proteins with previously unknown subcellular localization showed localization to distinct cytoplasmic puncta (A, D and G). Colabeling with a catalase antibody (B, E, and H) revealed colocalization to peroxisomal structures (merged images C, F, and I). The scale bar represents 10 μ m.

Table V

The Number of Proteins Within Each of the Eight Selected Eukaryotic Proteomes That Were Predicted to Be Peroxisomal by PeroxiP and Pts1Prowler

Species	Eukaryotic genome screens for peroxisomal proteins					
	PeroxiP			Pts1Prowler		
	Proteins	Predicted	Percentage (%)	Proteins	Predicted	Percentage (%)
<i>S. pompe</i>	4,962	10	0.20	4,990	6	0.12
<i>S. cerevisiae</i>	6,449	27	0.42	6,720	21	0.31
<i>D. melanogaster</i>	13,729	58	0.42	14,080	48	0.34
<i>C. elegans</i>	20,414	61	0.28	22,730	60	0.26
<i>A. thaliana</i>	25,826	61	0.24	26,639	62	0.23
<i>M. musculus</i>	28,097	59	0.21	36,471	65	0.18
<i>H. sapiens</i>	38,051	44	0.12	37,605	59	0.16
<i>O. sativa</i>	41,915	102	0.24	61,250	95	0.16

the PTS1PROWLER application enabled the productive identification of novel peroxisomal proteins within the mouse proteome.

Eukaryotic proteome tests

The proteomes of the eight eukaryotic species studied by Emanuelsson *et al.* were presented to PTS1PROWLER to provide an updated estimate of the number and proportion of peroxisomal proteins in these organisms (Table V and Fig. 4).

As an aside, Emanuelsson *et al.* report on an unexpected discrepancy between numbers found in human and mouse believed to stem from C-terminal annotation inaccuracies. We do not observe such a pronounced dip in the plot of predictions made by PTS1PROWLER (Fig. 4).

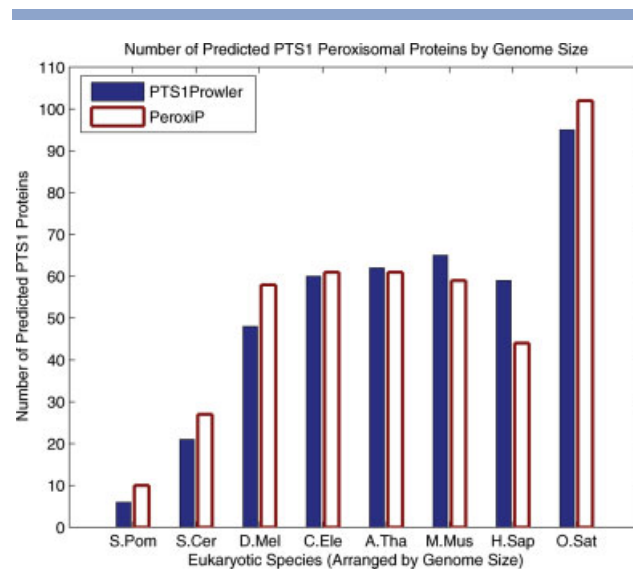
CONCLUSION

We have taken the overall three-stage classifier design of Emanuelsson *et al.* and focused on improving the machine learning stage of the process. The final model is a finely tuned SVM with a polynomial kernel of order five. Unlike Emanuelsson we include the PTS1 tripeptide in the input window and additionally we train the model to produce a probability by fitting a logistic function to the output of the SVM. The combination of all of these modifications resulted in a predictor with an average MCC of 0.77, and an improvement of 54% on the results reported for the original PEROXIIP predictor. ROC curve analysis reveals that the optimum cutoff value for accepting a protein as peroxisomal occurs at about 0.44. To achieve a comparable sensitivity to that reported for PTS1 PREDICTOR, a threshold around 0.22 should be used.

We have performed numerous analyses of the performance of our model. Notably, by testing on eukaryotic data for organisms without a peroxisome, we found that our model made no false positives, whereas PEROXIIP operated with a false positive rate of 1.9%. We used pro-

karyotic proteomes to replicate the tests in Neuberger and colleagues' published benchmarks. Using the optimal output threshold, our model operated with significantly fewer false positives than any of the three models in the PTS1 PREDICTOR suite. Both Emanuelsson *et al.* and Neuberger *et al.* compare their models against general predictors like PSORT and report that such are inferior to the specialized predictors.

We test the sensitivity of our model using the very latest peroxisomal proteins from SWISS-PROT and found that using the default threshold PTS1PROWLER is significantly better than PEROXIIP, but slightly worse than PTS1 PREDICTOR. However, with the adjusted threshold (matching PTS1 PREDICTOR's reported sensitivity of 0.896), it found

**Figure 4**

Eukaryotic proteome predictions. Graph of the number of Pts1Prowler and PeroxiP predicted peroxisomal matrix proteins for a range of eukaryotic organisms sorted by proteome size. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

one more positive than PTS1 PREDICTOR. Importantly, this adjustment did not compromise specificity to a great extent. The comparison on false-positive rates on prokaryotic sequence data showed that the lowered threshold only brought PTS1PROWLER's false positive rate to 0.24%, still below PTS1 PREDICTOR-General's 0.74% and PTS1 PREDICTOR-Metazoa's 0.34%, however, slightly worse than PTS1 PREDICTOR-Fungi's 0.19%. It should be noted that the data set of new peroxisomal proteins is very small compared to those used for the specificity tests, and hence this test has weaker significance.

We ran our model on the RIKEN IPS7 mouse protein dataset to test for novel peroxisomal proteins. Five proteins of unknown subcellular localization were identified as potentially peroxisomal and were selected for experimental validation. Their subcellular localization was determined using immunofluorescent detection of the myc-epitope-tagged proteins and three were confirmed to be novel peroxisomal proteins.

Finally we presented our classifier with the data sets of a number of eukaryotic proteomes in order to provide an update on the numbers of peroxisomal proteins found therein. The absolute numbers were much on par with that produced by the original PEROXIP study, even though the number of sequences in these data sets has increased. In general we found that our model has downgraded the number of predictions as a function of proteome size and predicts a smoother relationship between proteome size and number of peroxisomal proteins.

ACKNOWLEDGEMENTS

R.D.T. is supported by an NHMRC R. Douglas Wright Career Development Award.

REFERENCES

- Weller S, Gould SJ, Valle D. Peroxisome biogenesis disorders. *Ann Rev Genom Hum Genet* 2003;4:165–211.
- Moyersoen J, Choe J, Fan E, Hol WG, Michels PA. Biogenesis of peroxisomes and glycosomes: trypanosomatid glycosome assembly is a promising new drug target. *FEMS Microbiol Rev* 2004;28:603–643.
- Nakai K. Protein sorting signals and prediction of subcellular localization. *Adv Prot Chem* 2000;54:277–344.
- Emanuelsson O. Predicting protein subcellular localisation from amino acid sequence information. *Brief Bioinf* 2002;3:361–376.
- Schneider G, Fechner U. Advances in the prediction of protein targeting signals. *Proteomics* 2004;4:1571–1580.
- Neuberger G, Maurer-Stroh S, Eisenhaber B, Hartig A, Eisenhaber F. Motif refinement of the peroxisomal targeting signal 1 and evaluation of taxon-specific differences. *J Mol Biol* 2003;328:567–579.
- Lametschwandtner G, Brocard C, Fransen M, Van Veldhoven P, Berger J, Hartig A. The difference in recognition of terminal tripeptides as peroxisomal targeting signal 1 between yeast and human is due to different affinities of their receptor Pex5p to the cognate signal and to residues adjacent to it. *J Biol Chem* 1998;273:33635–33643.
- Baker A, Sparkes IA. Peroxisome protein import: some answers, more questions. *Curr Opin Plant Biol* 2005;8:640–647.
- Michels P, Moyersoen J, Krazy H, Galland N, Herman M, Hannaert V. Peroxisomes, glyoxysomes and glycosomes. *Mol Membr Biol* 2005;22:133–145.
- Reumann S, Ma C, Lemke S, Babujee L. Araperox. A database of putative arabidopsis proteins from plant peroxisomes. *Plant Physiol* 2004;136:2587–2608.
- Drawid A, Gerstein M. A bayesian system integrating expression data with sequence patterns for localizing proteins: comprehensive application to the yeast genome. *J Mol Biol* 2000;301:1059–1075.
- Lu Z, Szafron D, Greiner R, Lu P, Wishart D, Poulin B, Anvik J, Macdonell C, Eisner R. Predicting subcellular localization of proteins using machine-learned classifiers. *Bioinformatics* 2004;20:547–556.
- Scott MS, Thomas DY, Hallett MT. Predicting subcellular localization via protein motif co-occurrence. *Genome Res* 2004;14:1957–1966.
- Emanuelsson O, Nielsen H, Brunak S, von Heijne G. Predicting subcellular localization of proteins based on their n-terminal amino acid sequence. *J Mol Biol* 2000;300:1005–1016.
- Cai Y-D, Liu X-J, Chou K-C. Artificial neural network model for predicting protein subcellular location. *Comp Chem* 2002;26:179–182.
- Dyrlöv Bendtsen J, Nielsen H, von Heijne G, Brunak S. Improved prediction of signal peptides: signalP 3.0. *J Mol Biol* 2004;340:783–795.
- Bodén M, Hawkins J. Prediction of subcellular localisation using sequence-biased recurrent networks. *Bioinformatics* 2005;21:2279–2286.
- Hua SJ, Sun ZR. Support vector machine approach for protein subcellular localization prediction. *Bioinformatics* 2001;17:721–728.
- Chou KC, Cai YD. Using functional-domain composition and support vector machines for prediction of protein subcellular location. *J Biol Chem* 2002;29:45765–45769.
- Park K-J, Kanehisa M. Prediction of protein subcellular locations by support vector machines using compositions of amino acids and amino acid pairs. *Bioinformatics* 2003;19:1656–1663.
- Nakai K, Horton P. PSORT: a program for detecting sorting signals in proteins and predicting their subcellular localization. *Trends Biochem Sci* 1999;24:34–35.
- Chou KC, Cai YD. Prediction and classification of protein subcellular location-sequence-order effect and pseudo amino acid composition. *J Cell Biochem* 2003;90:1250–1260.
- Emanuelsson O, Eloffsson A, von Heijne G, Cristobal S. In silico prediction of the peroxisomal proteome in fungi, plants and animals. *J Mol Biol* 2003;330:443–456.
- Wakabayashi M, Hawkins J, Maetschke S, Bodén M. Exploiting targeting signal dependencies in the prediction of pts1 peroxisomal proteins. In: Gallagher M, Hogan J, Maire F, editors. *Intelligent data engineering and automated learning—IDEAL 2005: Sixth International Conference (Lecture notes in computer science, Vol. 3578)*. Brisbane, Australia: Springer; 2005. pp 454–461.
- Hawkins J, Bodén M. Predicting peroxisomal proteins. In: *Proceedings of the IEEE Symposium on Computational Intelligence in Bioinformatics and Computational Biology, Piscataway, NJ, November 2005*. IEEE, pp 469–474.
- Matthews BW. Comparison of predicted and observed secondary structure of T4 phage lysozyme. *Biochim Biophys Acta* 1975;405:442–451.
- Amery L, Fransen M, De Nys K, Mannaerts GP, Van Veldhoven PP. Mitochondrial and peroxisomal targeting of 2-methylacetyl-coa race-mase in humans. *J Lipid Res* 2000;41:1752–1759.
- Nishikawa K, Kubota Y, Ooi T. Classification of proteins into groups based on amino acid composition and other characters. I. Angular distribution. *J Biochem* 1983;94:981–995.
- Reinhardt A, Hubbard T. Using neural networks for prediction of the subcellular location of proteins. *Nucl Acids Res* 1998;26:2230–2236.

30. Platt J. Probabilistic outputs for support vector machines and comparison to regularized likelihood methods. In: Smola A, Bartlett P, Schoelkopf B, Schuurmans D, editors. *Advances in large margin classifiers*. Cambridge, MA: MIT Press; 2000. pp 61–74.
31. Witten IH, Frank E. *Data mining: practical machine learning tools with Java implementations*. San Francisco: Morgan Kaufmann; 2000.
32. Aturaliya RN, Fink JL, Davis MJ, Teasdale MS, Hanson KA, Miranda KC, Forrest ARR, Grimmond SM, Suzuki H, Kanamori M, Kai C, Kawai J, Carninci P, Hayashizaki Y, Teasdale RD. Subcellular localization of mammalian type II membrane proteins. *Traffic* 2006;7:613–625.
33. Parsons M, Furuya T, Pal S, Kessler P. Biogenesis and function of peroxisomes and glycosomes. *Mol Biochem Parasitol* 2001;115:19–28.
34. The FANTOM Consortium. The transcriptional landscape of the mammalian genome. *Science* 2005;309:1559–1563.
35. Fink JL, Aturaliya RN, Davis MJ, Zhang F, Hanson K, Teasdale MS, Kai C, Kawai J, Carninci P, Hayashizaki Y, Teasdale RD. LOCATE: a mouse protein subcellular localization database. *Nucl Acids Res* 2006;34:D213–D217.