



Mimicking Cellular Sorting Improves Prediction of Subcellular Localization

Rajesh Nair^{1,4} and Burkhard Rost^{1,2,3}

¹CUBIC, Department of Biochemistry and Molecular Biophysics, Columbia University, 650 West 168th Street BB217, New York, NY 10032, USA

²Columbia University Center for Computational Biology and Bioinformatics (C2B2), Russ Berrie Pavilion, 1150 St. Nicholas Avenue, New York NY 10032, USA

³NorthEast Structural Genomics Consortium (NESG) Department of Biochemistry and Molecular Biophysics, Columbia University, 650 West 168th Street BB217, New York, NY 10032, USA

⁴Department of Physics Columbia University, 538 West 120th Street, New York, NY 10027, USA

Predicting the native subcellular compartment of a protein is an important step toward elucidating its function. Here we introduce LOctree, a hierarchical system combining support vector machines (SVMs) and other prediction methods. LOctree predicts the subcellular compartment of a protein by mimicking the mechanism of cellular sorting and exploiting a variety of sequence and predicted structural features in its input. Currently LOctree does not predict localization for membrane proteins, since the compositional properties of membrane proteins significantly differ from those of non-membrane proteins. While any information about function can be used by the system, we present estimates of performance that are valid when only the amino acid sequence of a protein is known. When evaluated on a non-redundant test set, LOctree achieved sustained levels of 74% accuracy for non-plant eukaryotes, 70% for plants, and 84% for prokaryotes. We rigorously benchmarked LOctree in comparison to the best alternative methods for localization prediction. LOctree outperformed all other methods in nearly all benchmarks. Localization assignments using LOctree agreed quite well with data from recent large-scale experiments. Our preliminary analysis of a few entirely sequenced organisms, namely human (*Homo sapiens*), yeast (*Saccharomyces cerevisiae*), and weed (*Arabidopsis thaliana*) suggested that over 35% of all non-membrane proteins are nuclear, about 20% are retained in the cytosol, and that every fifth protein in the weed resides in the chloroplast.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: protein subcellular localization prediction; support vector machines; hierarchical ontology; sequence alignment; database search

Abbreviations used: 1D structure, one-dimensional structure (e.g. sequence or string of secondary structure, or solvent accessibility); ER, endoplasmic reticulum; GFP, green fluorescent protein; GFP-tagging, here used to refer to the large-scale experimental determination of localization through GFP tagging; GO, gene ontology; HSSP, database of protein structure-sequence alignments; IL, large scale localization of yeast proteins from high-throughput immuno-localization of epitope-tagged gene products; LOChom, predicting localization using annotation transfer through sequence homology; LOCKey, using SWISS-PROT keywords to predict subcellular localization; LOctree, hierarchical system of SVMs introduced here; NLS, nuclear localization signal; NNPSL, neural networks predicting localization; PHD, profile based neural network prediction of secondary structure, solvent accessibility and transmembrane helices; PredictNLS, prediction of nuclear proteins through nuclear localization signals; PROFphd, advanced profile-based neural network prediction of secondary structure and solvent accessibility; PSORT, knowledge-based expert system using amino acid composition and sequence motifs; SGD, *Saccharomyces cerevisiae* genome database; SignalP, neural network system predicting signal peptides; SMART, simple modular architecture research tool; SVM, support vector machine; SWISS-PROT, data base of protein sequences; SubLoc, support-vector machine-based prediction of localization; TargetP, combined method predicting chloroplast (ChloroP), extra-cellular (SignalP), and mitochondrial proteins; TMS, identification of chloroplast proteins in *A. thaliana* using tandem mass spectroscopy; TrEMBL, translation of the EMBL-nucleotide database coding DNA to protein sequences.

Assignment and prediction of subcellular localization indispensable

The sequencing of the genomes, i.e. all DNA sequences, of over 260 organisms (February 2005), including the human genome^{1,2} has been completed. For over 200 of the entirely sequenced organisms, the protein sequences are publicly available; 105 have been analyzed in the PEP database†, and contribute about 413,000 protein sequences, i.e. about one-fourth of all currently known protein sequences.^{3–5} With this explosion of genome sequences, the major challenge in modern biology is to follow suit in advancing the knowledge of the expression, regulation, and function of the entire set of proteins encoded by an organism, i.e. its proteome. This information will be invaluable for understanding how complex biological processes occur at a molecular level, how they differ in various cell types, and how they are altered in disease states. Proteins must be localized in the same subcellular compartment to cooperate towards a common function. Therefore, experimentally unraveling the native compartment of a protein constitutes one step on the long way to determining its role. Using experimental high-throughput methods for epitope and green fluorescent protein (GFP) tagging, two groups have recently reported localization data for most proteins in *Saccharomyces cerevisiae* (baker's yeast).^{6,7} So far, the majority of large-scale experimental efforts to predict localization have been restricted to yeast, or to particular compartments, such as a recent analysis of chloroplast proteins in *Arabidopsis thaliana* (weed).⁸ As of now, these large-scale experiments cannot be repeated for mammalian or other higher eukaryotic proteomes. One major obstacle is that large scale production of a collection of cell lines each with a defined gene chromosomally tagged at the 3' end is not yet possible.⁹ In contrast, computational tools can provide fast and accurate localization predictions for any organism.^{10–13} Attempts to predict subcellular localization have increasingly become one of the central problems in bioinformatics/computational biology.^{14–20}

Most reliable predictions cover less than 50% of all proteins

A number of methods predict localization by identifying short sequence motifs, such as signal peptides^{21–26} or nuclear localization signals (NLS)^{15,27–29} that are responsible for protein targeting. Most proteins destined for the secretory pathway, the mitochondria and the chloroplast contain N-terminal peptides that are recognized by the translocation machinery.^{30,31} The term “signal peptide” is used to describe the peptides in secreted

proteins that are cleaved in the endoplasmic reticulum (ER) by signal peptidases; the peptides responsible for targeting proteins to the mitochondria and chloroplast are referred to as “transit peptides”. Signal and transit peptides can be recognized by generic prediction methods, that by the detection of these peptides also predict subcellular localization.^{17,24,32,33} Many proteins destined for the nucleus contain NLS motifs that may occur anywhere in the sequence. Recently, we have collected a data set of experimental and potential NLS motifs as an aid to predicting nuclear localization.²⁹ However, the vast majority of nuclear proteins have no known motif. For mitochondria and chloroplast, a number of alternative targeting pathways have also been discovered recently.^{34–37} Additionally, proteins such as fibroblast growth factors are targeted to the extra-cellular space *via* non-classical secretory pathways, i.e. they do not possess N-terminal signal peptides.^{38,39} Furthermore, a particular problem for methods detecting N-terminal signals is that start codons are predicted with less than 70% accuracy by genome projects.^{1,2,40} Overall, known and predicted sequence motifs enable annotating about 30% of the proteins in six entirely sequenced eukaryotic proteomes.^{4,41,42} Other methods that can be reliably used to annotate localization but are not always applicable are annotation transfer from sequence homologues⁴³ and text analysis.^{44–47} A particular variant of homology-based predictions is the domain projection method that is based on similarity to SMART domains of known subcellular localization.⁴¹ Despite recent high-throughput experiments, the most reliable prediction methods together cover less than 50% of entirely sequenced multi-cellular proteomes.

De novo predictions of localization restricted by limited “biophysical reality”

In the near future, the only hope of assigning compartments to the remaining half of all multi-cellular proteins is using methods that predict localization from features other than known import/export motifs. The most promising approach is to exploit the correlation between localization and amino acid composition of a protein,^{48,49} which is mostly due to the altering of the protein surface in response to changing environmental conditions.⁵⁰ Methods using only amino acid composition to predict localization are *de novo* methods; they predict localization without any explicit experimental knowledge of the protein under investigation. In particular, they are as accurate if any information about function is available for a target as when the target is merely a “hypothetical protein”. Higher-order residue correlations (between residues i and $i+n$, for $n=2,3,4$) have been accounted for by using pseudo-amino acid composition.^{51–53} Recently, we showed

† <http://cubic.bioc.columbia.edu/db/PEP/>

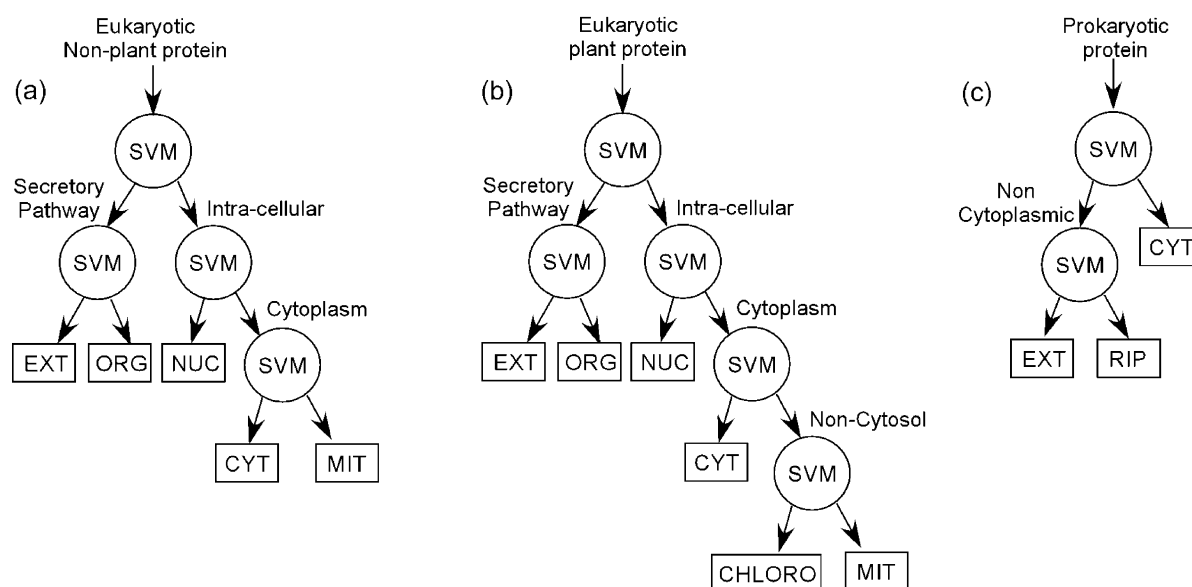


Figure 1. Hierarchical architecture of LOctree. LOctree uses specialized architecture to predict subcellular localization of proteins from different organisms: (a) architecture for eukaryotic non-plant proteins; (b) architecture for plant proteins; and (c) the architecture for prokaryotic proteins. At each branch point a support vector machine (SVM) is used to accomplish a binary classification (either protein belongs to localization class L or does not belong to L). The hierarchical architecture has been designed to mimic the biological protein sorting mechanism as closely as possible. The branches of the tree represent intermediate stages in the sorting machinery while the nodes represent the decision points in the sorting machinery. The different levels of SVMs in the hierarchical tree are labeled Level 0, Level 1, etc. For example, Level 0 represents the top node SVM which discriminates between secretory pathway proteins and other intra-cellular proteins ((a) and (b)) or proteins which remain in the cytoplasm from the rest (c). The intermediate node SVMs in the next level are represented as Level 1, and are responsible for separating extra-cellular proteins from proteins sorted to the organelles and nuclear proteins from cytoplasmic proteins ((a) and (b)). For the prokaryotic architecture (c), Level 1 is the terminal level for Gram-negative bacteria and separates extra-cellular proteins from periplasmic proteins. In addition, Level 1 also contains the cytoplasmic leaf which is propagated without branching from Level 0. For Gram-positive bacteria, Level 0 is the terminal level and separates cytoplasmic proteins from extra-cellular proteins (non-cytoplasmic branch). The leaves of the tree, represented by rectangular boxes represent the final localization classes for which prediction is made. If a leaf has a depth smaller than the overall depth of the tree it is propagated without branching for the remainder of the tree. Level 2 is the terminal level for the eukaryotic non-plant architecture (a) and is responsible for sorting proteins into one of five subcellular classes (mitochondria and cytosol plus the three leaves from Level 1), while Level 3 is the terminal level for the plant architecture (c) and separates proteins into one of six classes (mitochondria and chloroplast plus the four leaves from Level 2). The prediction accuracy of the parent nodes is higher than the child nodes leading to a significantly improved prediction accuracy for the intermediate localization states. Abbreviations: EXT, extra-cellular; NUC, nucleus; CYT, cytosol; MIT, mitochondria; CHLORO, chloroplast; RIP, periplasm; and ORG, organelle. Organelles are the endoplasmic reticulum, Golgi apparatus, peroxysomes, lysosomes, and vacuolar compartments.

that incorporating structural and evolutionary information significantly improves prediction accuracy.⁵⁴ With the availability of many completely sequenced genomes, phylogenetic profiles have been employed to identify subcellular localization.⁵⁵ So far, this approach has been much less accurate than methods based solely on composition. Drawid & Gerstein have proposed a Bayesian system, based on a diverse range of 30 different features, to predict the localization of yeast proteins.⁵⁶ The problem with all these methods is that they are based on sequence features that may reveal localization but are not the reason why proteins are transported, such as signal and transit peptides and nuclear localization signals. Furthermore, all general methods, with the exception of PSORT,^{57,58} implicitly assume that all localizations are equidistant, i.e. if a method predicts a nuclear protein to

be cytoplasmic it makes the same mistake as another method which predicts this protein to be extra-cellular. In reality, however, some compartments are more similar to each other than others, e.g. ER is closer to extra-cellular than to nuclear due to the proximity in the space of the biological sorting machinery.

Here, we describe a novel system of support vector machines (SVMs) that predict subcellular localization by incorporating a hierarchical ontology of localization classes modeled onto biological processing pathways. By construction, the system penalizes confusions of classes along the same pathway (e.g. ER instead of extra-cellular) less than confusions between classes from different pathways (e.g. ER instead of nuclear). The biological similarities are incorporated from the description of cellular components in the gene

Table 1. LOctree on non-redundant test set of eukaryotic non-plant proteins

Hierarchy level	Class	Nprot	Acc	Cov	GAv	Q (StdDev)	MCC	MI	Nstates
Level 0	Secretory pathway	415	81	80	81	89 (2)	0.73	0.44	2
	Intra-cellular	1090	92	93	93				
Level 1	Extra-cellular	363	83	81	82	78 (4)	0.55	0.40	4
	Organelles	52	51	52	52				
	Nuclear	562	78	78	78				
Level 2	Cytoplasm	528	76	78	77	74 (6)	0.55	0.39	5
	Cytosol	330	63	66	64				
	Mitochondria	198	70	67	68				

Abbreviations used: hierarchy level and class are as illustrated in Figure 1; *Nprot*, number of proteins in sequence-unique test set with a given localization; *Nstates*, number of effective states predicted at given level (note that Level 1 contains four states while Level 2 contains five states, namely the Level 1 leaves (extra-cellular, organelles and nuclear) + cytosol + mitochondria). Performance measures: Acc, accuracy or specificity (equation (2)); Cov, coverage or selectivity (equation (3)); gAv, geometric average between Acc and Cov (equation (4)); Q, overall prediction accuracy for a given level in the hierarchy (equation (5)); note depending on the level this is a two-state, four-state, or five-state value); MCC, Mathews correlation coefficient (equations (6) and (8)); MI, mutual information (equations (7) and (10)). Note 1: Q=74% at Level 2 is the overall accuracy for classification into one of five localization classes (extra-cellular, organelles, nuclear, cytosol or mitochondria).

ontology (GO).^{59,60} We simplified and tailored the GO definitions to the problem of protein sorting. For example, in GO both the ER and the Golgi apparatus are subcategories of the cytoplasm. However, proteins destined for the extra-cellular space, the ER, the Golgi, endosomes and lysosomes are targeted *via* the same secretory pathway. By this criterion, proteins from the secretory pathway are more similar to each other than they are to other intra-cellular proteins.³⁰ Hence, in our classification scheme these compartments are grouped together and are designated as belonging to the secretory pathway. Technically, we incorporated the ontology through a decision tree with SVMs as the nodes (Figure 1). We favored SVMs over neural networks due their improved performance (data not shown). The final system, LOctree, was extremely successful at learning evolutionary similarities among subcellular localization classes and was significantly more accurate than other traditional networks at predicting subcellular localization.

We have applied LOctree to analyze the subcellular localization of complete genomes of a number of eukaryotic and prokaryotic organisms. The LOctree subcellular localization prediction server and the results of our localization annotations for entire proteomes are available.†

Results

Data sets and cross-validation results

More data with noise better than less data with less noise

Proteins with experimentally annotated subcellular localization were extracted from SWISS-PROT⁶¹ (Methods and Materials). For this study, we excluded membrane proteins, i.e. all our results are valid for a subset of 75–80% of all proteins.^{42,62–65}

In total, we had 8980 eukaryotic and 13,186 prokaryotic non-membrane proteins with explicit experimental annotations (Methods and Materials). Training and test sets were constructed by partitioning the data such that test sequences had less than 25% sequence identity to any sequence in the training set over an alignment length of 250 residues (HVAL=5; equation (1) Methods and Materials). To avoid overestimating performance, we reduced redundancy such that our final sequence-unique test set contained 1505 non-redundant eukaryotic non-plant sequences, 304 plant and 672 prokaryotic test sequences. All results of the methods described here were based on sixfold cross-validation experiments, i.e. we cycled six times through the entire sequence-unique data such that each protein was used for testing once. To increase the size of the training set, we included homology-based (LOChom⁴³) and keyword-based (LOCKey⁴⁵) predictions in the training data. While adding these noisy predictions, we ascertained that no homologues to any of the test proteins were included. This procedure almost quadrupled the training data; it increased prediction accuracy by nearly seven percentage points (data not shown). The major improvement resulted from the addition of keyword-based annotations using LOCKey.⁴⁵ Plants were treated separately since their compositional features differed significantly from non-plant eukaryotes (not shown). Using the SVM-light package,⁶⁶ we found the radial basis function (RBF) kernel to perform better than linear and polynomial kernels. This result was obtained on a small subset of all proteins without cross-validation. In particular, we did not optimize this solution for the final test set.

Very accurate distinction between secretory pathway proteins and all others

To predict the localization of an unknown eukaryotic protein, LOctree first determines if it is sorted using the secretory pathway. The SVM that

† <http://www.rostlab.org/services/LOctree/>




Figure 2. Reliability of LOctree. The curves show prediction accuracy of LOctree for eukaryotic animal sequences. (a) Overall performance: the prediction accuracy decreases as we descent the hierarchical tree (Figure 1(a)). The Level 2 accuracy shown includes the accuracy of all Level 1 leaves like the extra-cellular, organelle and nuclear classes (Figure 1(a)), and represents the accuracy of classifying the protein into one of five subcellular classes. At 75% coverage the prediction accuracy is around 94% for Level 0, dropping to 84% for Level 1 and 77% for Level 2. The ability of the hierarchical system to predict intermediate localization states at a significantly higher accuracy is evident from the 17% difference in prediction accuracy between Level 0 and Level 2. Level 1 separates proteins into one of four subcellular classes and is over 7% more accurate than Level 2, which separates proteins into one of five classes. (b) Class-wise performance: LOctree is best at discriminating secretory pathway proteins from all other proteins (91% accuracy at 50% coverage). Prediction of nuclear and extra-cellular proteins was only slightly less accurate (84% accuracy at 50% coverage) while performance was significantly worse for cytosolic proteins with only 64% correctly predicted. The standard deviation in the prediction accuracy for each of the localization classes was roughly 7%.

makes this distinction achieved an overall prediction accuracy around 90% for both eukaryotic non-plant (Table 1; Figure 2(a)) and plant proteins (Table 2). Using the signal peptide prediction of SignalP²³ as one input to the SVM improved accuracy by over one percentage point (data not shown). We also confirmed our previous observation that using overall composition in conjunction with N-terminal composition improved performance.⁵⁴ Our methods also distinguished intra-cellular proteins very accurately from those entering the secretory pathway (>90% accuracy; Tables 1 and 2).

Overall accuracy of 74% for non-plants

If a protein is predicted as belonging to the secretory pathway it is further sub-classified into extra-cellular or not (Figure 1(a)). The non extra-cellular proteins belong to either of the following organelles: endoplasmic reticulum (ER), Golgi apparatus, peroxisome, lysosome, or vacuole. Proteins native in one of these organelles were predicted at levels around 50% accuracy and 52% coverage (values were higher for plant proteins; Table 2); these values were much lower than the averages for all other classes. The sub-classification

Table 2. LOctree on non-redundant test set of plant proteins

Hierarchy level	Class	Nprot	Acc	Cov	gAv	Q (StdDev)	MCC	MI	Nstates
Level 0	Secretory pathway	42	77	79	78	94 (4)	0.74	0.49	2
	Intra-cellular	262	97	96	97				
Level 1	Extra-cellular	22	68	68	68	88 (5)	0.58	0.49	4
	Organelles	20	57	60	58				
	Nuclear	32	70	81	75				
	Cytoplasm	230	95	93	94				
Level 2	Cytosol	77	73	74	74	77 (7)	0.59	0.44	5
	Non-cytosol	153	84	80	82				
Level 3	Mitochondria	50	61	74	67	70 (3)	0.58	0.42	6
	Chloroplast	103	77	63	70				

Abbreviations used as for Table 1.

Table 3. LOctree on non-redundant test set of prokaryotic proteins

Hierarchy level	Class	Nprot	Acc	Cov	gAv	Q (StdDev)	MCC	MI	Nstates
Level 0	Cytoplasm	426	89	97	93	90 (4)	0.79	0.52	2
	Non-cytoplasm	246	93	80	86				
Level 1	Periplasmic	125	86	62	73	83 (2)	0.55	0.45	3
	Extra-cellular	42	59	74	66				

Abbreviations used as for Table 1. Note 1: Level 1 is applicable to Gram-negative bacteria only. For Gram-positive bacteria, the system performs a two-state classification with non-cytoplasmic proteins being classified as extra-cellular. For Gram-negative bacteria, non-cytoplasmic proteins are further separated into periplasmic and extra-cellular proteins. Note 2: the Level 0 prediction accuracy did not differ significantly between Gram-positive and Gram-negative bacteria. The prediction accuracy reported above is the combined prediction accuracy for Gram-positive and Gram-negative bacteria. The overall two class prediction accuracy for Gram-positive bacteria was $Q_2=90\%$ while the three class prediction accuracy for Gram-negative bacteria was $Q_3=83\%$.

of intra-cellular proteins into nucleus and cytoplasm was less accurate; however, levels of accuracy and coverage were still above 76% (Tables 1 and 2). Of the final localization classes (the leaves in Figure 1(a)), the extra-cellular class was predicted most accurately while the nuclear class was predicted only slightly less accurately (Figure 2(b)). Cytosolic proteins were, as expected, predicted with the lowest accuracy by the non-plant method (Table 1 and Figure 2(b)), while mitochondrial proteins were the least accurately predicted by the plant specialist (Table 2). Predictions for all classes of the non-plant system were extremely balanced between accuracy and coverage (Table 1). For reporting prediction accuracy, if a leaf (terminal node) has a smaller depth than the overall tree (the extra-cellular leaf has a depth of 1 while the overall tree depth is 2), the predictions for this leaf are propagated without branching for the depth of the tree (Figure 1(a)). Hence, the quoted Level 2 accuracy of $Q_5=74\%$ (Table 1) is the overall accuracy for classification into one of five localization classes (extra-cellular, organelles, nuclear, cytosol or mitochondrial). This was over seven percentage points more accurate than another system that used the same input with the traditional pair-wise SVMs (data not shown). The unique feature of our method is that it predicts “intermediate” localizations such as intra-cellular and cytoplasm (Tables 1 and 2). These intermediate localizations are predicted with a much higher accuracy as is evident from the progressive decrease in prediction accuracy as we descent the hierarchical tree (Figure 2(a)).

Accurate distinction of three prokaryotic classes

For prokaryotic proteins, LOctree first determines if the protein is cytoplasmic or not. The SVM discriminating between these localizations reached an overall accuracy of 90% (Table 3). Prediction accuracy did not differ significantly between Gram-positive and Gram-negative bacteria. For Gram-negative bacteria, the non-cytoplasmic proteins are further classified into periplasmic and extra-cellular. The overall three class (cytoplasmic, periplasmic or extra-cellular) prediction accuracy for Gram-negative bacteria was $Q_3=83\%$; the two class

(cytoplasmic or extra-cellular) accuracy for Gram-positive bacteria was $Q_2=90\%$. For Gram-negative bacteria, the distinction between periplasmic and extra-cellular proteins was at a much lower accuracy than cytoplasmic proteins.

Comparison with other methods using additional test sets

Other methods tested on new data set

We compared our method to the following publicly available methods: TargetP,³³ SubLoc,⁶⁷ NNPSL,⁴⁰ and PSORT II.²⁵ In contrast to all other methods, TargetP focuses exclusively on N-terminal sorting signals (secreted, chloroplast, mitochondria); it does not predict proteins targeted using other mechanisms or in any other compartment, such as in the nucleus or cytoplasm. Of the other servers that predict at least four classes, SubLoc⁶⁷ is also based on SVMs; NNPSL⁴⁰ is based on neural networks. SubLoc and NNPSL rely solely on amino acid composition while PSORT²⁵ combines information from local sequence motifs and a neural network based method. All publicly available methods were tested on smaller data sets than LOctree, and on data sets with little mutual overlap. We could run all servers on our non-redundant test set from the cross-validation experiments (Tables 1–3). However, most of the proteins in our data set had been used to develop those servers and we could not cross-validate any method other than our own. A benchmark with our data set would, therefore, have very limited value. For completeness, we reported the results of this test which as expected, over-estimated the public servers significantly (Supplementary Data, Table 1). The most meaningful comparison of prediction methods is based on a significantly sized, sequence-unique data set of proteins that have neither been used for the development of any of the methods tested, nor have significant sequence similarity to any of the methods tested. Unfortunately, such sets are often difficult to get. If we ignored the most recent improvement of one of the components of TargetP, namely SignalP 3.0,¹⁷ we could find such a data set in proteins added between SWISS-PROT version 40 and 41

Table 4. Comparison on identical sequence-unique set of new SWISS-PROT non-plant eukaryotic proteins

Class ↓	Method →		<i>LOCtree</i> here	<i>TargetP</i> 23, 24, 33	<i>SubLoc</i> 67	<i>PSORT</i> 25, 57, 58, 97	<i>NNPSL</i> 40
	Score						
<i>Secretory Pathway</i>	<i>Acc</i>	87	93				
	<i>Cov</i>	90	73				
	<i>gAv</i>	88	82				
<i>Ext</i>	<i>Acc</i>	86		73	91	62	
	<i>Cov</i>	93		53	32	63	
	<i>gAv</i>	89		62	54	63	
<i>Nuc</i>	<i>Acc</i>	77		64	56	67	
	<i>Cov</i>	85		71	75	59	
	<i>gAv</i>	81		67	65	63	
<i>Cyt</i>	<i>Acc</i>	82		43	47	42	
	<i>Cov</i>	64		56	47	38	
	<i>gAv</i>	72		49	47	40	
<i>Mit</i>	<i>Acc</i>	73	54	48	46	30	
	<i>Cov</i>	78	75	59	59	67	
	<i>gAv</i>	75	64	53	52	45	
<i>Overall accuracy</i> Q_4			78		57	51	52

Abbreviations used as for Table 1, with the following exceptions. Data set: all sequence-unique eukaryotic non-plant proteins added between release 41 and 40 of SWISS-PROT (*Non-plant new unique* in Table 4 of the Supplementary Data). Note that none of the proteins in this set had significant sequence similarity to any of the proteins that had annotations about localization in SWISS-PROT at the time of development of the prediction methods for which results are shown. In this sense, our test set could also provide an independent and likely more accurate estimate for the sustained performance than some of the original publications for some of the methods. Localization: Ext, extra-cellular; Nuc, nuclear; Cyt, cytosolic; Mit, mitochondria; Chloro, chloroplast. Methods: Predictions from methods other than LOCtree, introduced here, were taken from their public Internet servers (Methods and Materials), except for PSORT II that was run locally; numbers in square brackets under methods refer to the original publication (References). Numbers in bold: in each row, the best method(s) is (are) marked in bold letters; methods are grouped according to significant differences (below), i.e. all values that are statistically indistinguishable are marked as one best group. Significant differences: For LOCtree, the standard deviation in the five-state accuracy was roughly six percentage points. The following estimates for standard deviations were published: TargetP,³³ about one percentage point; NNPSL,⁴⁰ about 2.5 percentage points; PSORT II,²⁵ about 3.5 percentage points. Since no error estimates were published for SubLoc,⁶⁷ we used 2.5 percentage points as the mean over the other three.

(Methods and Materials). While SignalP 3.0 has been developed after release 41, all the methods that we compared have been developed before release 41. Note that we deliberately restricted our development of LOCtree to proteins available in version 40 so that we could carry out this comparison. In many ways, our evaluation was also informative of the sustained performance of the methods tested, some of which had fallen prey to severe over-estimates of performance in their original publications. Note that PSORT, TargetP and the different versions of SignalP stood out in that their authors had correctly estimated the sustained performance all along.

LOCtree over 20 percentage points more accurate than other general servers

In the benchmark of proteins that had not been used for the development of any method, LOCtree outperformed all other servers (Table 4). TargetP was more accurate at predicting proteins targeted *via* the secretory pathway but its coverage was lower than that of LOCtree. The reason was that TargetP slightly under-predicted the secretory pathway (imbalance in *gAv* (equation (4)), i.e. the geometric average over *Acc* and *Cov*; Table 4). On our data set, we found the accuracy of SignalP 3.0¹⁷ to be slightly lower than that of TargetP, since the

Table 5. Performance of LOctree based on large scale yeast localization data

Method	Nuclear			Mitochondrial		
	Obs	Acc	Cov	Obs	Acc	Cov
GFP ⁷	586	82	68	418	83	51
IL ⁶	124	88	64	60	76	62

Abbreviations used: Methods: Experimental subcellular localization data for proteins in yeast were obtained from two methods: GFP, large scale localization using green fluorescent protein tagging;⁷ IL, large scale localization using high-throughput immunolocalization of epitope-tagged proteins.⁶ Data: Obs, number of non-membrane proteins for which localization was predicted in this compartment by the respective large-scale experiment. All proteins observed to be in multiple compartments by the large-scale methods were excluded from our analysis. LOctree was used to predict localization of the remaining proteins. The prediction accuracy (Acc) and coverage (Cov) of LOctree was calculated by assuming that the localization observed in the large-scale experiment represents the true localization of the protein. Note 1: cytoplasmic proteins were excluded from our analysis, since a large fraction (45–70%) of all proteins was observed to be in the cytoplasm in the two large-scale experiments. Nearly half of all cytoplasmic proteins were observed to be associated with more than one compartment and many are likely to be further sorted to other compartments.

difference was not significant, SignalP 3.0 was not shown separately in order to simplify. PSORT II was the most accurate server for predicting extracellular proteins; however, this was achieved at the cost of an extremely low level of coverage; in the geometric average between accuracy and coverage, PSORT II was more than 30 percentage points lower than LOctree. As shown in our cross-validation experiments (Tables 1–3), LOctree was very balanced in its compromise between accuracy and coverage, i.e. between under and over-prediction, for all classes, and it was much more balanced than any other server. In terms of the overall four-state accuracy (Q_4 ; equation (5)), LOctree scored 21 percentage points higher than its best competitor SubLoc (Table 4).

Performance better than existing methods even for incorrect sequences

LOctree explicitly used information from the first 50 residues (N termini) and compositions from the entire protein. Both these values are likely to be wrong for many proteins taken from large-scale sequencing projects.^{68–71} We tried to estimate the effect of such mistakes through two different “models”: (1) we cleaved off 30 N-terminal residues for all proteins; and (2) we randomly picked positions to remove one-third of the sequence for each protein. These tests constituted worst-case scenarios in the sense that they all over-estimated sequencing errors substantially. We found that the overall prediction accuracy of LOctree on the randomly cleaved fragments was 68% (Supplementary Data, Table 2), 10% less than what was obtained using the full protein sequence. For the N-term cleaved sequences, the accuracy further dropped to 55% due to the explicit dependence of LOctree on N-terminal sequence information. This is still

accurate enough to provide reliable first estimates of localization for genomic sequences.

About 80% agreement between predictions and large-scale experiments in yeast

Over the last years the large-scale experimental determination of subcellular localization for a substantial fraction of all yeast proteins has become increasingly accurate. Using high-throughput immuno-localization (IL) of epitope-tagged gene products, the Snyder group⁶ determined the localization for about 60%, while the O’Shea group⁷ exploited high-throughput GFP tagging to cover about 66% of all yeast proteins. Both studies did not distinguish between membrane and non-membrane proteins, and both did not capture secreted proteins. Many proteins were experimentally associated to more than one single compartment: 35% for Snyder *et al.* and 31% for O’Shea *et al.* We compared the LOctree predictions of all proteins predicted to not contain membrane helices, and observed to be nuclear or mitochondrial in the two large-scale experiments. Proteins observed to be in the cytoplasm in the two large-scale studies were excluded from our analysis, since a large fraction (43% for Snyder *et al.* and 55% for O’Shea *et al.*) of cytosolic proteins

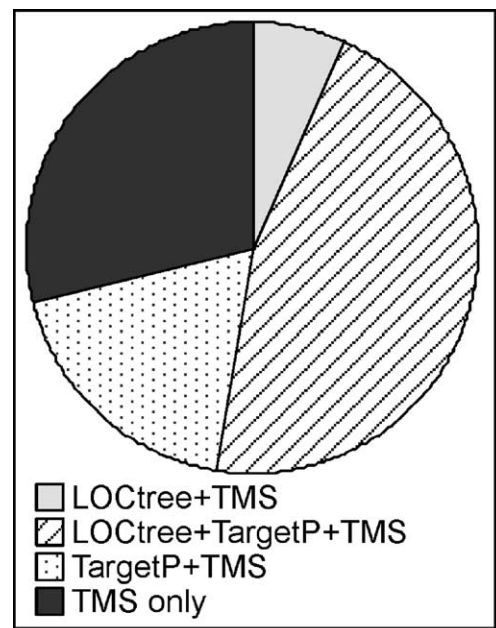


Figure 3. Benchmarking LOctree using large-scale experimental data. Both LOctree and TargetP³³ were used to predict the localization of nearly 500 chloroplast proteins from *A. thaliana* which were identified using tandem mass spectroscopy (TMS) by Kleffmann *et al.*⁸ LOctree and TargetP both predicted over 45% of these proteins to be localized in the chloroplast lending strong support to the large-scale experimental data using TMS. Over 70% of the proteins were predicted to be in the chloroplast by at least one server. TargetP showed a high degree of agreement with LOctree, agreeing with over 87% of the predictions using LOctree.




Figure 4. Composition of compartments in three representative proteomes. Note that 100% of the pie charts represents the number of proteins without transmembrane helices (predicted by PHDhtm^{98–100} and taken from PEP³). The final estimates were corrected in order to account for our compartment-specific estimates of accuracy and coverage (equation (11)). For all three proteomes the nucleus appeared to take the lion's share of all proteins, only the chloroplast came near this value for the plant representative weed (*A. thaliana*). Human (*H. sapiens*) has significantly more secreted proteins than do weed and yeast (*S. cerevisiae*). Yeast appeared to have the highest fraction of mitochondrial proteins. For the proteomes the percent fractional error for the estimates of the different compartments are given by: extra-cellular ($\pm 9\%$), organelle ($\pm 34\%$), nuclear ($\pm 10\%$), cytosol ($\pm 20\%$), mitochondria ($\pm 25\%$) and for chloroplast ($\pm 17\%$).

were also observed in alternative compartments. Similarly, we also excluded all other proteins experimentally associated with more than one compartment. This filtering left over 1000 proteins from the GFP data and about 200 proteins from the IL data. For both these data sets, about 80% of the predictions from LOctree were identical with the experimental results (Table 5). This is comparable to the 80% agreement between the GFP data and traditional non-high-throughput results previously annotated in SGD.^{72,73} The agreement between GFP and IL is about 75% for the subset of 146 proteins found to be nuclear or mitochondrial by GFP that were also found in the IL data set. The agreement between GFP annotations and yeast proteins annotated using homology to SWISS-PROT proteins was 79%. This is for five subcellular classes and using an HVAL > 10 (equation (1)) for homology annotations. For the IL data the agreement with SWISS-PROT was 72% (note due to the small data set this number is a very inaccurate estimate).

Chloroplasts: experimental data supported by predictions

Using tandem mass spectroscopy (TMS), Kleffmann *et al.*⁸ recently identified 690 proteins localized in the chloroplast of *A. thaliana* (weed). Of these we predicted 190 to contain membrane helices. We compared LOctree and TargetP³³ for

the remaining 500 proteins (Figure 3). The following results stood out: (1) less than half of these proteins were identified by all three methods; (2) a considerable fraction (29%) of the 500 proteins was only identified by TMS; and (3) when comparing the chloroplast predictions for all weed proteins, we found that LOctree and TargetP agreed in about 87% of their predictions. TargetP, however, predicts more chloroplast proteins than LOctree. This could be due to the over-prediction of chloroplast proteins by TargetP which has been reported by an independent group.⁷⁴

Application to representative proteomes

We used LOctree to annotate the subcellular localization for all non-membrane proteins in the entire proteomes of *Homo sapiens* (human),^{1,2} *A. thaliana* (weed),⁷⁵ and *S. cerevisiae* (yeast)¹⁹ (Figure 4). The results of our proteome annotations can be queried (downloaded) from the LOctree website[†]. We estimated that over 60% of all non-plant and over 50% of all plant proteins are nuclear or remain in the cytosol (Figure 4 and Table 3 of Supplementary Data). While over 75% of the non-membrane proteins in all genomes appeared intracellular, the fraction of secreted proteins varied substantially between 8% and 20%, with plants

[†] <http://www.rostlab.org/cgi-bin/var/nair/loctree/query-genome.pl>

having fewer than 10% extra-cellular proteins and the number exceeding 20% in human. Nuclear proteins were overabundant in yeast. In general, the unicellular yeast was somewhere in between human and weed in its composition of compartments.

Discussion

Tree-based system provided additional advantages to boosting performance

Our results demonstrated how the prediction of subcellular localization can be substantially improved by mimicking the biological protein trafficking mechanism as closely as possible through, LOCTree, a hierarchical tree of SVMs (Figure 1 and Tables 1–3). PSORT^{25,57,58,76} is based on an implementation of a reasoning tree that is conceptually most similar to LOCTree. However, unlike LOCTree, PSORT is not based on an explicit “ontology” of subcellular localization. Instead, the nodes of the PSORT reasoning tree assign a probabilistic value to the presence/absence of a single feature and have no intrinsic meaning. In contrast, the nodes of LOCTree separate proteins belonging to different cellular sorting pathways. In addition to being more accurate, our machine learning system provided two added benefits. The first was the prediction of “intermediate stages”. The prediction of intermediate stages, such as secretory pathway, was achieved at much higher levels of accuracy than the native compartments. This is not too surprising given that large-scale experiments using IL and GFP-tagging^{6,7} suggest that 30% of all proteins are ambiguous, i.e. spend a considerable portion of their life-time in more than one native compartment. These experimental data might strike many biologists as “expected,” since the vast majority of proteins travel through the cell, e.g. most extra-cellular proteins “visit” at least three other compartments (ER, Golgi, vesicles) before they eventually are secreted. However, the very fact that our system reaches levels above 74% accuracy for the distinction of proteins into one of five compartments (extra-cellular, cytoplasmic, mitochondrial, nuclear, and organellar) suggested that most proteins have very strong preferences for one single compartment imprinted onto their sequences. The importance of post-translational modifications in altering these sequence signals as a means to increase the “fitness” for other compartments might explain the difference between these two extreme opposite perceptions of proteins as native to “one native compartment” and as “frequent travelers between compartments”.^{12,50,77} The second advantage of our hierarchical system might appear to be of more technical nature, namely that our modular architecture allows the addition of more fine-grained modules at later stages (e.g. the split of nuclear into nuclear-matrix and other⁷⁸). However, this seemingly technical detail actually

once again was borne out of the advantage of mimicking the actual sorting system. We observed that as one descends the hierarchical tree the prediction accuracy progressively decreases, since the classification task becomes increasingly complex and the SVM has to discriminate between increasingly similar proteins. One problem with our decision tree-like implementation was that a prediction mistake at a top node could not be corrected at nodes lower in the hierarchy. The appropriate choice of the evolutionary hierarchy was, therefore, crucial. The fact that we cannot correct mistakes from higher levels was by no means a feature of the design: we tried to recover from higher-level sorting mistakes by predicting localization of a protein at all nodes and averaging over the prediction strengths over all higher level nodes. Sometimes this worked, however, most of the time such an alteration introduced new mistakes.

Over 20 percentage points improvement over existing generalized methods

We also showed that the increase in the size of the training set through the addition of noisy predictions was more relevant than the noise added from the mistakes in these annotations obtained through text-analysis⁴⁵ and homology-transfer.⁴³ This surprising finding suggested that prediction methods might improve even more through the continued addition of large-scale experimental tackling of localization. Finally, we confirmed our previous findings⁵⁴ that predicted structure and evolutionary profiles contain information relevant for the prediction of localization. All these data combined with our hierarchical tree-based system improved the overall accuracy over 20 percentage points over the best competitor that generically predicted localization in four states (extra-cellular, nuclear, cytoplasmic, mitochondrial; Table 4). The only method that performed at a similarly high level as LOCTree was TargetP³³ (Table 4) that focuses on particular classes (secreted, mitochondria, chloroplast). TargetP³³ was also the only method that appeared significantly better at predicting one particular compartment, namely, chloroplast proteins (Table 1 of Supplementary Data; note, however, that in this test we did compare our method in cross-validation mode to TargetP in not-cross-validation mode, i.e. were likely to have over-estimated the performance of TargetP). Predictions using LOCTree had the added advantage of being extremely balanced between accuracy and coverage (Tables 1–4). In contrast, methods such as SignalP,²³ the secreted/not-secreted component of TargetP, are either very prone to over (high coverage, low accuracy) or under-prediction (low coverage, high accuracy; e.g. TargetP, PSORT II). The only other general method that was as well balanced as LOCTree was NNPSL⁴⁰ that had an overall performance of 27 percentage points below LOCTree (Table 4).

Many estimates for performance had a rather short-life span

Another problem that we noticed was that only TargetP and PSORT had published estimates for performance that were close to our results on a “never-seen-before” set of sequence-unique proteins. In particular, SubLoc⁶⁷ was estimated to achieve an overall accuracy of $Q_4=79\%$, while it reached only 57% on our data (Table 4). The differences may be explained by the fact that up to 90% pair-wise sequence identity was allowed between testing and training set for the original publication of SubLoc and NNPSL.⁴⁰ Cai *et al.*⁷⁹ also claim very high level of accuracy (73%). That value was more difficult to compare because their methods are not available as servers and because their publications did not rigorously describe protocols for removing redundancy. In fact, it appears that only proteins identical between training and testing set were excluded. Furthermore, the accuracy was compiled on a different partition of the prediction goal. More recently this group published even higher estimates using similar data sets with unspecified sequence similarity between testing and training.^{51,52} In general, the problem of correctly estimating performance is a very difficult one as illustrated by the bi-annual meetings for the critical assessment of structure prediction (CASP^{80–84}) and by servers that evaluate the performance of servers such as EVA†.^{85,86} The task is particularly difficult in a field in which we have too few and no continuous resource of experimental data.

As accurate as large-scale experiments?

Numerically, our predictions from LOCTree agreed as much with traditional “small-scale” biochemical determinations of subcellular localization as did the recent large-scale experiments^{6–8} (Tables 4 and 5). Interestingly, our predictions reached a similar level of performance as large-scale experiments (GFP: 79%, IL: 72%, and LOCTree: 74–78%) if analyzed against more careful traditional approaches as the standard-of-truth. This by no means implies that we aimed at the replacement of experiments. Rather, we see predictions from LOCTree as a reasonable, cheap starting point for careful experiments and as a complement for the interpretation of large-scale results. Furthermore, our prediction method had slightly different potential than the large-scale experiments, e.g. while we could identify secreted proteins, we currently could not clearly distinguish between Golgi and vesicles, nor did we include membrane proteins.

Open tasks

Our current system marked in some ways the end

of a very long series of methods addressed at predicting localization. While methods using homology-transfer (LOChom⁴³) and text-analysis (LOCKey⁴⁵) were crucial for achieving our new state-of-the-art level of performance we will have to tie some loose ends, in particular, we currently exclude membrane proteins, treat each protein as one without any experimental annotations (LOChom and LOCKey are used for training and for our prediction server; however, they are not generically integrated into LOCTree), and do not distinguish between proteins that are generically native to more than one compartment and those which are not. Furthermore, the task of annotating more than a few representative proteomes remains.

Conclusion

Previous attempts at predicting subcellular localization have implemented machine-learning algorithms using the standard parallel architecture as is common practice in computer science and have focused on improving prediction by incorporating additional sequence features that are correlated with localization. Here we have shown that prediction accuracy can be significantly improved by using a hierarchical architecture of support vector machines to mimic the protein sorting mechanism. This result is likely to hold for other aspects of protein function and can significantly aid the development of more accurate predictors of protein function. The ability of many proteins to function in more than one native subcellular compartment makes the prediction task especially difficult. In fact, over 30% of the more than 4000 yeast proteins for which localization has been determined using high-throughput experiments⁷ are associated with more than one compartment. The hierarchical architecture of LOCTree can better incorporate proteins which spend time in more than one native compartment by predicting “intermediate” localization states, which span multiple subcellular classes, at a much higher accuracy. The fact that the system achieved an overall five-state prediction accuracy of 74% seems to indicate that the native subcellular localization is imprinted somehow onto the protein sequence and that a majority of proteins carry only one strong sequence signal for one particular compartment. In future, it should be possible to further extend the abilities of LOCTree by adding modules that can make fine-grained distinctions such as discriminating among the different organelles and various substructures like the nucleolus.

Methods and Materials

Data sets used for development and evaluation

We selected all eukaryotic and prokaryotic proteins with explicit annotations about subcellular localization in

† <http://www.rostlab.org/EVA>

SWISS-PROT release 40.⁸⁷ We excluded proteins annotated as MEMBRANE, POSSIBLE, PROBABLE, SPECIFIC PERIODS or BY SIMILARITY. We also excluded proteins annotated with multiple localizations. This left about 9000 eukaryotic proteins and 13,000 prokaryotic proteins in our trusted set of experimentally annotated localization (“SWISS-PROT annotated” set; Table 4 of Supplementary Data). Training and test sets were constructed from this set such that no pair of proteins from any two sets had sequence similarity levels corresponding to $HVAL > 5$ (equation (1)). We picked this value, since below this threshold assigning subcellular localization based solely on homology leads to significant errors.⁴³ Furthermore, the test set was redundancy reduced at $HVAL < 10$ using a simple greedy search.⁸⁸ This ensured that no two proteins in the test set had greater than 25% sequence identity over more than 250 residues (number of sequence unique proteins given in Table 4 of Supplementary Data). The reason for this reduction was to find a balance between biased data known to yield over-estimates^{89,90} and between too small data sets likely to yield incorrect estimates. We did not have to define thresholds for significant sequence similarity between motifs such as signal peptides,⁸⁹ since we never explicitly used this information, rather we used the entire protein information†.

SWISS-PROT new set used for testing, only

After we completed the development of all our methods, we used an additional data set to re-examine performance, namely, we collected all proteins that had been added to SWISS-PROT between release 40 and 41 (results presented in Table 4). We excluded all new proteins that had $HVAL > 5$ (equation (1)) to any previously used protein and found the sequence-unique subset of the new proteins (Table 4 of Supplementary Data). We never used any of these proteins for development, and it is rather unlikely that any of the other methods tested (Table 4) used any of these, since all methods were developed based on SWISS-PROT releases < 41 .

HSSP-value to measure pair-wise sequence similarity

The simplest way to measure sequence similarity is percentage pair-wise sequence identity (PIDE), i.e. the percentage of residues identical between two proteins (not counting gaps). Another measure is the statistical expectation values as reported by BLAST. Here, we used a third measure, namely the HSSP-value (HVAL) because it more accurately allowed the separation between proteins pairs for which similarity in localization is recognizable from sequence than the other two.⁴³ The HVAL^{91,92} is given by:

$$HVAL = PID - \begin{cases} 100 & \text{for } L \leq 11 \\ 480L^{-0.32\{1+\exp^{-L/1000}\}} & \text{for } L \leq 450 \\ 19.5 & \text{for } L > 450 \end{cases} \quad (1)$$

where L was the number of residues aligned between two proteins, PID the percentage of pair-wise identical residues.

Increasing size of training set

Preliminary results suggested that a larger training set improved SVM performance through increased coverage of the sequence space. Another source of improvement was using a sequence redundant set of proteins to train the SVM. Two strategies were used to increase the size of the training set. (1) SWISS-PROT keyword-based annotations: using LOCKey,⁴⁵ we first annotated localization for all proteins in the SWISS-PROT database for which adequate keyword functional information was present in the database. Next, proteins with $HVAL > 5$ to proteins in the test set were excluded. The remaining proteins were added to the training set. (2) Homology-based annotations: using LOChom,⁴³ we annotated localization for all sequence homologues in the SWISS-PROT database of proteins in the training set. Using both procedures increased the size of the training set by almost a factor of 4.

Building evolutionary profiles

We have shown⁵⁴ that using evolutionary information in the form of sequence profiles significantly improves prediction accuracy. Profiles were built by aligning the sequences against the SWISS-PROT + TrEMBL database using the MaxHom dynamic programming algorithm.⁹³ The aligned sequences were filtered for redundancy at 95% pair-wise sequence identity, i.e. pairs exceeding this limit were removed. Finally, we included only those proteins that had $HVAL > 5$ and $PID > 50\%$ with respect to the guide sequence. These thresholds were previously found to be optimal for a rather different prediction method.⁵⁴ Finally, the profile composition was calculated by replacing each amino acid residue in the protein by the residue frequencies in the profile. All composition information was input to the SVM in the form of profile-based composition.

Hierarchical architecture and support vector machine training

Each decision node in the hierarchical architecture of LOctree (Figure 1) was implemented using a support vector machine (SVM). The SVMs were implemented using the SVM-light package.⁶⁶ We used the following input information: amino acid composition (20 units), composition of the 50 N-terminal residues (20 units), and amino acid composition in the three secondary structure states (60 units). For the eukaryotic plant and non-plant systems, raw output from the SignalP server²³ was used as additional input to the SVM at the top node which determined whether a protein is sorted through the secretory pathway or not. Altogether 100 variables (104 for the top-level SVM in the eukaryotic architecture) were used as input to the SVMs. Each SVM was trained using the radial basis function (RBF) kernel. The γ parameter for the RBF-kernel and the C parameter for the trade-off between training error and margin were determined by optimization on small subset of the training data. We used a constant value of $\gamma = 15$ and $C = 500$ for all SVMs. The predictions from the SVMs were observed to be quite resilient to changes in kernel parameters.

Final decision through simple winner-takes-all

We experimented with the following two methods for determining the localization of a protein. (1) Decision tree: at each node in the hierarchical architecture (Figure 1), a simple yes/no decision was made based only on the SVM

† www.rostlab.org/results/2004/LOctree/

output at that node to determine which branch of the localization tree the protein belongs to. (2) Summing over prediction strengths: the branch of the localization tree the protein is sorted through at each node was determined by summing the prediction strengths over all previous nodes. The simple decision tree architecture was finally used, since it outperformed the summing architecture by over one percentage point.

Cross-validation

The “SWISS-PROT annotated” data was partitioned into six sets of equal size using the HVAL criteria described above: five of these sets were combined to give the training set and the sixth one was used for testing. Finally, we rotated through the test set such that each protein was used for testing exactly once. We never used any information from the test set to optimize parameters.

Evaluating performance

As a simple measure for performance we used the percentage accuracy (equation (5)). The accuracy/specificity and coverage/sensitivity of the two-state networks were measured using four ratios derived from TP (true positives, i.e. the number of proteins predicted to be in localization L and experimentally observed to be in localization L), FP (false positives, i.e. the number of proteins predicted to be in localization L and observed in *not-L*) and FN (false negatives, i.e. the number of proteins predicted not to be in *not-L* and observed to be in L). We used:

$$\text{Acc}(L) = 100 \frac{TP}{TP + FP} \quad (2)$$

$$\text{Cov}(L) = 100 \frac{TP}{TP + FN} \quad (3)$$

In other words, $\text{Acc}(L)$ is the accuracy/specificity in predicting localization L , and $\text{Cov}(L)$ is the corresponding coverage/selectivity. We combined these two numbers through the geometric average:

$$\text{gAv}(L) = 1/100 \cdot \sqrt{\text{Acc}(L) \cdot \text{Cov}(L)} \quad (4)$$

We omitted the specific qualifier L in text and Tables whenever the localization that we referred to was obvious. The overall accuracy was measured by Q :

$$\begin{aligned} Q &= 100 \frac{TP}{TP + FN} \\ &= 100 \frac{\text{number correctly predicted}}{\text{number of proteins in data set}} \end{aligned} \quad (5)$$

Note that the values of Q we presented in text and Tables referred to different number of states, depending on which level in the hierarchy (Figure 1) we monitored the overall accuracy. Where needed, we indicated the number of states through subscripts, e.g. Q_5 is the overall five-state accuracy that measures the accuracy in predicting one of five classes of localization for non-plant eukaryotes.

For the two class predictions we also reported the Mathews correlation coefficient (MCC)⁹⁴ and the normalized mutual information (MI).⁹⁵

$$\text{MCC} = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FN)(TP + FP)(TN + FP)(TN + FN)}} \quad (6)$$

$$\begin{aligned} \text{MI} &= -\frac{TP + FN}{N} \cdot \log\left(\frac{TP + FN}{N}\right) \\ &\quad - \frac{TN + FP}{N} \log\left(\frac{TN + FP}{N}\right) \end{aligned} \quad (7)$$

The Mathews correlation coefficient and the mutual information as defined above are only applicable for two classes. For more than two classes, the Mathews correlation coefficient was modified to the generalized correlation coefficient:⁹⁶

$$\text{GC}^2 = \frac{\sum_{ij} \frac{(z_{ij} - e_{ij})^2}{e_{ij}}}{N(K - 1)} \quad (8)$$

where N was the number of proteins, K the number of localization classes, z_{ij} the confusion matrix and e_{ij} was given by:

$$e_{ij} = \frac{\text{Acc}_i \times \text{Cov}_j}{N} \quad (9)$$

For many classes the mutual information was modified to the information coefficient:⁹⁶

$$\text{MI} = \sum_i \left\{ -\frac{\text{Cov}_i}{N} \log\left(\frac{\text{Cov}_i}{N}\right) + \sum_j \frac{z_{ij}}{N} \log\left(\frac{z_{ij}}{\text{Acc}_j}\right) \right\} \quad (10)$$

Prediction methods

The prediction accuracy of four publicly available methods was evaluated using the “Non-plant unique”, “Plant unique” and the “Non-plant new unique” sets (Table 4 of Supplementary Data). The four methods were: (1) TargetP: neural network based tool predicting localization based on N-terminal sequence information,^{23,24,33} (2) SubLoc: support vector machine prediction of localization from amino acid composition,⁶⁷ (3) PSORT II: integrated method based on detecting sorting signals and predictions from other methods like NNPSL and SignalP,^{25,58,97} and (4) NNPSL: neural network based tool predicting localization from amino acid composition.⁴⁰ All methods were run with default parameter settings.

Estimate for composition in entire proteomes

Three different values determined our final estimates for the percentage of proteins in localization L in entirely sequenced proteomes (Figure 3), namely the number of proteins actually predicted by LOctree to be in localization L , our estimate for the accuracy (equation (2)) and coverage (equation (3)) of that prediction. In detail:

$$N_{\text{prd}}(L) = \frac{N_{\text{prot_notTMH}}}{N_{\text{prd_corrected_sum}}} \cdot N_{\text{prd_corrected}}(L) \quad (11)$$

with

$$\begin{aligned} N_{\text{prd_corrected}}(L) &= \frac{\text{Acc}(L) \cdot N_{\text{prd_raw}}(L)}{\text{Cov}(L)} \text{ and } N_{\text{prd_corrected_sum}} \\ &= \sum_{\forall L} N_{\text{prd_corrected}}(L) \end{aligned}$$

where $N_{\text{prd_raw}}(L)$ was the number of proteins predicted directly by LOctree to be in localization L , $\text{Acc}(L)$ and $\text{Cov}(L)$ the estimates for prediction accuracy and

coverage, respectively, and $N_{\text{prot_nofTMH}}$ the number of proteins without membrane helices. The error in the number of proteins predicted to be in localization L was estimated by:

$$\frac{\sigma_{N(L)}}{N(L)} = \sqrt{\left(\frac{\sigma_{\text{Acc}(L)}}{\text{Acc}(L)}\right)^2 + \left(\frac{\sigma_{\text{Cov}(L)}}{\text{Cov}(L)}\right)^2} \quad (12)$$

Acknowledgements

Thanks to Jinfeng Liu (Columbia University) for computer assistance and the collection of genome data sets and to Kazimierz Wrzeszczynski (Columbia University) for proof reading the manuscript. Thanks to Astrid Reinhardt (Baylor College of Medicine, Texas), Tim Hubbard (Sanger Centre, Hinxton), Sujun Hua (Tsinghua University), Zhirong Hun (Tsinghua University), Olof Emanuelsson (Stockholm University), Henrik Nielsen (CBS, Copenhagen), Søren Brunak (CBS, Copenhagen), Gunnar von Heijne (Stockholm University), Kenta Nakai (Tokyo University) and Paul Horton (UCB) for access to their prediction methods. Thanks to both anonymous referees for important corrections. Last, not but least, thanks to Amos Bairoch (SIB, Geneva), Rolf Apweiler (EBI, Hinxton), Phil Bourne (UCSD), and their crews for maintaining excellent databases and to all experimentalists who enabled this analysis by making their data publicly available.

Supplementary Data

Supplementary data associated with this article can be found, in the online version, at [doi:10.1016/j.jmb.2005.02.025](https://doi.org/10.1016/j.jmb.2005.02.025)

References

1. Venter, J. C., Adams, M. D., Myers, E. W., Li, P. W., Mural, R. J., Sutton, G. G. *et al.* (2001). The sequence of the human genome. *Science*, **291**, 1304–1351.
2. Lander, E. S., Linton, L. M., Birren, B., Nusbaum, C., Zody, M. C., Baldwin, J. *et al.* (2001). Initial sequencing and analysis of the human genome. *Nature*, **409**, 860–921.
3. Carter, P., Liu, J. & Rost, B. (2003). PEP: predictions of entire proteomes. *NAR*, **31**, 410–413.
4. Liu, J. & Rost, B. (2002). Target space for structural genomics revisited. *Bioinformatics*, **18**, 922–933.
5. Pruess, M., Fleischmann, W., Kanapin, A., Karavidopoulou, Y., Kersey, P., Kriventseva, E. *et al.* (2003). The Proteome Analysis database: a tool for the *in silico* analysis of whole proteomes. *Nucl. Acids Res.* **31**, 414–417.
6. Kumar, A., Agarwal, S., Heyman, J. A., Matson, S., Heidtman, M., Piccirillo, S. *et al.* (2002). Subcellular localization of the yeast proteome. *Genes Dev.* **16**, 707–719.
7. Huh, W. K., Falvo, J. V., Gerke, L. C., Carroll, A. S.,

- Howson, R. W., Weissman, J. S. & O'Shea, E. K. (2003). Global analysis of protein localization in budding yeast. *Nature*, **425**, 686–691.
8. Kleffmann, T., Russenberger, D., von Zychlinski, A., Christopher, W., Sjolander, K., Gruissem, W. & Baginsky, S. (2004). The *Arabidopsis thaliana* chloroplast proteome reveals pathway abundance and novel protein functions. *Curr. Biol.* **14**, 354–362.
9. Davis, T. N. (2004). Protein localization in proteomics. *Curr. Opin. Chem. Biol.* **8**, 49–53.
10. Bork, P. & Koonin, E. V. (1998). Predicting functions from protein sequences—where are the bottlenecks? *Nature Genet.* **18**, 313–318.
11. Koonin, E. V. (2000). Bridging the gap between sequence and function. *Trends Genet.* **16**, 16.
12. Nakai, K. (2001). Prediction of *in vivo* fates of proteins in the era of genomics and proteomics. *J. Struct. Biol.* **134**, 103–116.
13. Rost, B., Liu, J., Nair, R., Wrzeszczynski, K. O. & Ofran, Y. (2003). Automatic prediction of protein function. *Cell. Mol. Life Sci.* **60**, 2637–2650.
14. Eisenhaber, F. & Bork, P. (1998). Wanted: subcellular localization of proteins based on sequence. *Trends Cell Biol.* **8**, 169–170.
15. Nakai, K. (2000). Protein sorting signals and prediction of subcellular localization. *Advan. Protein Chem.* **54**, 277–344.
16. Schneider, G. & Fechner, U. (2004). Advances in the prediction of protein targeting signals. *Proteomics*, **4**, 1571–1580.
17. Bendtsen, J. D., Nielsen, H., von Heijne, G. & Brunak, S. (2004). Improved prediction of signal peptides: SignalP 3.0. *J. Mol. Biol.* **340**, 783–795.
18. Rapp, M., Drew, D., Daley, D. O., Nilsson, J., Carvalho, T., Melen, K. *et al.* (2004). Experimentally based topology models for *E. coli* inner membrane proteins. *Protein Sci.* **13**, 937–945.
19. Peltier, J. B., Emanuelsson, O., Kalume, D. E., Ytterberg, J., Friso, G., Rudella, A. *et al.* (2002). Central functions of the luminal and peripheral thylakoid proteome of *Arabidopsis* determined by experimentation and genome-wide prediction. *Plant Cell*, **14**, 211–236.
20. Jensen, L. J., Gupta, R., Blom, N., Devos, D., Tamames, J., Kesmir, C. *et al.* (2002). Prediction of human protein function from post-translational modifications and localization features. *J. Mol. Biol.* **319**, 1257–1265.
21. von Heijne, G. (1995). Protein sorting signals: simple peptides with complex functions. *Exs*, **73**, 67–76.
22. Nielsen, K. J., Hill, J. M., Anderson, M. A. & Craik, D. J. (1996). Synthesis and structure determination by NMR of a putative vacuolar targeting peptide and model of a proteinase inhibitor from *Nicotiana glauca*. *Biochemistry*, **35**, 369–378.
23. Nielsen, H., Engelbrecht, J., Brunak, S. & von Heijne, G. (1997). A neural network method for identification of prokaryotic and eukaryotic signal peptides and prediction of their cleavage sites. *Int. J. Neural Syst.* **8**, 581–599.
24. Emanuelsson, O., Nielsen, H. & von Heijne, G. (1999). ChloroP, a neural network-based method for predicting chloroplast transit peptides and their cleavage sites. *Protein Sci.* **8**, 978–984.
25. Nakai, K. & Horton, P. (1999). PSORT: a program for detecting sorting signals in proteins and predicting their subcellular localization. *Trends Biochem. Sci.* **24**, 34–36.

26. Emanuelsson, O. & von Heijne, G. (2001). Prediction of organellar targeting signals. *Biochim. Biophys. Acta*, **1541**, 114–119.
27. Cokol, M., Nair, R. & Rost, B. (2000). Finding nuclear localisation signals. *EMBO Rep.* **1**, 411–415.
28. La Cour, T., Gupta, R., Rapacki, K., Skriver, K., Poulsen, F. M. & Brunak, S. (2003). NESbase version 1.0: a database of nuclear export signals. *Nucl. Acids Res.* **31**, 393–396.
29. Nair, R., Carter, P. & Rost, B. (2003). NLSdb: database of nuclear localization signals. *Nucl. Acids Res.* **31**, 397–399.
30. Rusch, S. L. & Kendall, D. A. (1995). Protein transport *via* amino-terminal targeting sequences: common themes in diverse systems. *Mol. Membr. Biol.* **12**, 295–307.
31. Schatz, G. & Dobberstein, B. (1996). Common principles of protein translocation across membranes. *Science*, **271**, 1519–1526.
32. Nielsen, H., Engelbrecht, J., Brunak, S. & von Heijne, G. (1997). Identification of prokaryotic and eukaryotic signal peptides and prediction of their cleavage sites. *Protein Eng.* **10**, 1–6.
33. Emanuelsson, O., Nielsen, H., Brunak, S. & von Heijne, G. (2000). Predicting subcellular localization of proteins based on their N-terminal amino acid sequence. *J. Mol. Biol.* **300**, 1005–1016.
34. Rapaport, D. (2003). Finding the right organelle. Targeting signals in mitochondrial outer-membrane proteins. *EMBO Rep.* **4**, 948–952.
35. Hoogenraad, N. J. & Ryan, M. T. (2001). Translocation of proteins into mitochondria. *IUBMB Life*, **51**, 345–350.
36. Paschen, S. A. & Neupert, W. (2001). Protein import into mitochondria. *IUBMB Life*, **52**, 101–112.
37. Schnell, D. J. (2000). Functions and origins of the chloroplast protein-import machinery. *Essays Biochem.* **36**, 47–59.
38. Nickel, W. (2003). The mystery of non-classical protein secretion. A current view on cargo proteins and potential export routes. *Eur. J. Biochem.* **270**, 2109–2119.
39. Bendtsen, J. D., Jensen, L. J., Blom, N., Von Heijne, G. & Brunak, S. (2004). Feature-based prediction of non-classical and leaderless protein secretion. *Protein Eng. Des. Sel.* **17**, 349–356.
40. Reinhardt, A. & Hubbard, T. (1998). Using neural networks for prediction of the subcellular location of proteins. *Nucl. Acids Res.* **26**, 2230–2236.
41. Mott, R., Schultz, J., Bork, P. & Ponting, C. P. (2002). Predicting protein cellular localization using a domain projection method. *Genome Res.* **12**, 1168–1174.
42. Liu, J. & Rost, B. (2001). Comparing function and structure between entire proteomes. *Protein Sci.* **10**, 1970–1979.
43. Nair, R. & Rost, B. (2002). Sequence conserved for subcellular localization. *Protein Sci.* **11**, 2836–2847.
44. Fleischmann, W., Moller, S., Gateau, A. & Apweiler, R. (1999). A novel method for automatic functional annotation of proteins. *Bioinformatics*, **15**, 228–233.
45. Nair, R. & Rost, B. (2002). Inferring sub-cellular localization through automated lexical analysis. *Bioinformatics*, **18**, S78–S86.
46. Kretschmann, E., Fleischmann, W. & Apweiler, R. (2001). Automatic rule generation for protein annotation with the C4.5 data mining algorithm applied on SWISS-PROT. *Bioinformatics*, **17**, 920–926.
47. Lu, Z., Szafron, D., Greiner, R., Lu, P., Wishart, D. S., Poulin, B. *et al.* (2004). Predicting subcellular localization of proteins using machine-learned classifiers. *Bioinformatics*, **20**, 547–556.
48. Nishikawa, K., Kubota, Y. & Ooi, T. (1983). Classification of proteins into groups based on amino acid composition and other characters: I. Angular distribution. *J. Biochem.* **94**, 981–995.
49. Nakashima, H. & Nishikawa, K. (1994). Discrimination of intracellular and extracellular proteins using amino acid composition and residue-pair frequencies. *J. Mol. Biol.* **238**, 54–61.
50. Andrade, M. A., O'Donoghue, S. I. & Rost, B. (1998). Adaptation of protein surfaces to subcellular location. *J. Mol. Biol.* **276**, 517–525.
51. Cai, Y. D., Liu, X. J., Xu, X. B. & Chou, K. C. (2002). Support vector machines for prediction of protein subcellular location by incorporating quasi-sequence-order effect. *J. Cell Biochem.* **84**, 343–348.
52. Chou, K. C. & Cai, Y. D. (2003). Prediction and classification of protein subcellular location—sequence-order effect and pseudo amino acid composition. *J. Cell Biochem.* **90**, 1250–1260.
53. Pan, Y. X., Zhang, Z. Z., Guo, Z. M., Feng, G. Y., Huang, Z. D. & He, L. (2003). Application of pseudo amino acid composition for predicting protein subcellular location: stochastic signal processing approach. *J. Protein Chem.* **22**, 395–402.
54. Nair, R. & Rost, B. (2003). Better prediction of subcellular localization by combining evolutionary and structural information. *Proteins*, **53**, 917–930.
55. Marcotte, E. M., Xenarios, I., van Der Bliek, A. M. & Eisenberg, D. (2000). Localizing proteins in the cell from their phylogenetic profiles. *Proc. Natl Acad Sci. USA*, **97**, 12115–12120.
56. Drawid, A. & Gerstein, M. (2000). A Bayesian system integrating expression data with sequence patterns for localizing proteins: comprehensive application to the yeast genome. *J. Mol. Biol.* **301**, 1059–1075.
57. Nakai, K. & Kanehisa, M. (1992). A knowledge base for predicting protein localization sites in eukaryotic cells. *Genomics*, **14**, 897–911.
58. Nakai, K. & Kanehisa, M. (1991). Expert system for predicting protein localization sites in Gram-negative bacteria. *Proteins: Struct. Funct. Genet.* **11**, 95–110.
59. Ashburner, M., Ball, C. A., Blake, J. A., Botstein, D., Butler, H., Cherry, J. M. *et al.* (2000). Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. *Nature Genet.* **25**, 25–29.
60. Lewis, S., Ashburner, M. & Reese, M. G. (2000). Annotating eukaryote genomes. *Curr. Opin. Struct. Biol.* **10**, 349–354.
61. Boeckmann, B., Bairoch, A., Apweiler, R., Blatter, M. C., Estreicher, A., Gasteiger, E. *et al.* (2003). The SWISS-PROT protein knowledge base and its supplement TrEMBL in 2003. *Nucl. Acids Res.* **31**, 365–370.
62. Rost, B., Casadio, R. & Fariselli, P. (1996). Topology prediction for helical transmembrane proteins at 86% accuracy. *Protein Sci.* **5**, 1704–1718.
63. Jones, D. T. (1998). Do transmembrane protein superfolds exist? *FEBS Letters*, **423**, 281–285.
64. Wallin, E. & von Heijne, G. (1998). Genome-wide analysis of integral membrane proteins from eubacterial, archaean, and eukaryotic organisms. *Protein Sci.* **7**, 1029–1038.
65. Melen, K., Krogh, A. & von Heijne, G. (2003). Reliability measures for membrane protein topology prediction algorithms. *J. Mol. Biol.* **327**, 735–744.
66. Joachims, T. (2000). *Estimating the Generalization*

- Performance of a SVM Efficiently Proceedings of the Seventeenth International Conference on Machine Learning 2000*, Morgan Kaufmann Publishers, Inc., San Francisco, CA.
67. Hua, S. & Sun, Z. (2001). Support vector machine approach for protein subcellular localization prediction. *Bioinformatics*, **17**, 721–728.
 68. Brenner, S. E. (1999). Errors in genome annotation. *Trends Genet.* **15**, 132–133.
 69. Koonin, E. V. (2001). Computational genomics. *Curr. Biol.* **11**, R155–R158.
 70. Kyrpides, N. C. & Ouzounis, C. A. (1998). Errors in genome reviews. *Science*, **281**, 1457.
 71. Hubbard, T. J. P. (1997). New horizons in sequence analysis. *Curr. Opin. Struct. Biol.* **7**, 190–193.
 72. Chervitz, S. A., Hester, E. T., Ball, C. A., Dolinski, K., Dwight, S. S., Harris, M. A. *et al.* (1999). Using the *Saccharomyces* Genome Database (SGD) for analysis of protein similarities and structure. *Nucl. Acids Res.* **27**, 74–78.
 73. Christie, K. R., Weng, S., Balakrishnan, R., Costanzo, M. C., Dolinski, K., Dwight, S. S. *et al.* (2004). *Saccharomyces* Genome Database (SGD) provides tools to identify and analyze sequences from *Saccharomyces cerevisiae* and related sequences from other organisms. *Nucl. Acids Res.* **32**, D311–D314.
 74. Richly, E. & Leister, D. (2004). An improved prediction of chloroplast proteins reveals diversities and commonalities in the chloroplast proteomes of *Arabidopsis* and rice. *Gene*, **329**, 11–16.
 75. Arabidopsis Genome Initiative. (2000). Analysis of the genome sequence of the flowering plant *Arabidopsis thaliana*. *Nature*, **408**, 796–815.
 76. Horton, P. & Nakai, K. (1996). A probabilistic classification system for predicting the cellular localization sites of proteins. In *Fourth International Conference on Intelligent Systems for Molecular Biology* (States, D., Agarwal, P., Gaasterland, T., Hunter, L. & Smith, R. F., eds) *Fourth International Conference on Intelligent Systems for Molecular Biology, St Louis, MO, USA*, vol. 1, AAAI Press, St Louis, MO, USA.
 77. Mann, M., Hendrickson, R. C. & Pandey, A. (2001). Analysis of proteins and proteomes by mass spectrometry. *Annu. Rev. Biochem.* **70**, 437–473.
 78. Mika, S. & Rost, B. (2005). NMPdb: database of nuclear matrix proteins. *Nucl. Acids Res.*, **33**, D160–D163.
 79. Cai, Y. D. & Chou, K. C. (2000). Using neural networks for prediction of subcellular location of prokaryotic and eukaryotic proteins. *Mol. Cell Biol. Res. Commun.* **4**, 172–173.
 80. Moulton, J., Fidelis, K., Zemla, A. & Hubbard, T. (2003). Critical assessment of methods of protein structure prediction (CASP)-round V. *Proteins: Struct. Funct. Genet.* **53**, 334–339.
 81. Moulton, J., Fidelis, K., Zemla, A. & Hubbard, T. (2001). Critical assessment of methods of protein structure prediction (CASP): round IV. *Proteins: Struct. Funct. Genet.*, 2–7.
 82. Moulton, J., Hubbard, T., Bryant, S. H., Fidelis, K. & Pedersen, J. T. (1999). Critical assessment of methods of protein structure prediction (CASP): round III. *Proteins: Struct. Funct. Genet.*, 2–6.
 83. Moulton, J., Hubbard, T., Bryant, S. H., Fidelis, K. & Pedersen, J. T. (1997). Critical assessment of methods of protein structure prediction (CASP): round II. *Proteins: Struct. Funct. Genet.* **1**, 2–6.
 84. Moulton, J., Pedersen, J. T., Judson, R. & Fidelis, K. (1995). A large-scale experiment to assess protein structure prediction methods. *Proteins: Struct. Funct. Genet.* **23**, ii–iv.
 85. Eyrich, V., Marti-Renom, M. A., Przybylski, D., Fiser, A., Pazos, F., Valencia, A. *et al.* (2001). EVA: continuous automatic evaluation of protein structure prediction servers. *Bioinformatics*, **17**, 1242–1243.
 86. Koh, I. Y. Y., Eyrich, V. A., Marti-Renom, M. A., Przybylski, D., Madhusudhan, M. S., Narayanan, E. *et al.* (2003). EVA: evaluation of protein structure prediction servers. *Nucl. Acids Res.* **31**, 3311–3315.
 87. Bairoch, A. & Apweiler, R. (2000). The SWISS-PROT protein sequence database and its supplement TrEMBL in 2000. *Nucl. Acids Res.* **28**, 45–48.
 88. Hobohm, U., Scharf, M., Schneider, R. & Sander, C. (1992). Selection of representative protein data sets. *Protein Sci.* **1**, 409–417.
 89. Nielsen, H., Engelbrecht, J., von Heijne, G. & Brunak, S. (1996). Defining a similarity threshold for a functional protein sequence pattern: the signal peptide cleavage site. *Proteins: Struct. Funct. Genet.* **24**, 165–177.
 90. Rost, B. (2002). Enzyme function less conserved than anticipated. *J. Mol. Biol.* **318**, 595–608.
 91. Rost, B. (1999). Twilight zone of protein sequence alignments. *Protein Eng.* **12**, 85–94.
 92. Sander, C. & Schneider, R. (1991). Database of homology-derived protein structures and the structural meaning of sequence alignment. *Proteins*, **9**, 56–68.
 93. Sander, C. & Schneider, R. (1991). Database of homology-derived structures and the structural meaning of sequence alignment. *Proteins: Struct. Funct. Genet.* **9**, 56–68.
 94. Mathews, F. S. (1985). The structure, function and evolution of cytochromes. *Prog. Biophys. Mol. Biol.* **45**, 1–56.
 95. Rost, B. & Sander, C. (1993). Prediction of protein secondary structure at better than 70% accuracy. *J. Mol. Biol.* **232**, 584–599.
 96. Baldi, P., Brunak, S., Chauvin, Y., Andersen, C. A. & Nielsen, H. (2000). Assessing the accuracy of prediction algorithms for classification: an overview. *Bioinformatics*, **16**, 412–424.
 97. Horton, P. & Nakai, K. (1997). Better prediction of protein cellular localization sites with the k nearest neighbors classifier. *Ismb*, **5**, 147–152.
 98. Rost, B. (2005). How to use protein 1D structure predicted by PROFphd. In *The Proteomics Protocols Handbook* (Walker, J. E., ed.), pp. 879–908, Humana, Totowa, NJ.
 99. Rost, B. (1996). PHD: predicting one-dimensional protein structure by profile based neural networks. *Methods Enzymol.* **266**, 525–539.
 100. Chen, C. P., Kernytsky, A. & Rost, B. (2002). Transmembrane helix predictions revisited. *Protein Sci.* **11**, 2774–2791.

Edited by G. von Heijne

(Received 27 December 2004; received in revised form 8 February 2005; accepted 9 February 2005)