Multidimensional visualization to support analysis of image collections

Danilo M. Eler   Marcel Y. Nakazaki   Fernando V. Paulovich   Davi P. Santos
M. Cristina F. Oliveira   João E. S. Batista Neto
Rosane Minghim

Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo
CP 668, 13560-970, São Carlos, SP, Brazil
{eler,myn,paulovic,davi,cristina,jbatista,rminghim}@icmc.usp.br

Abstract

Multidimensional Visualization techniques are invaluable tools for analysis of structured and unstructured data with variable dimensionality. This paper introduces a methodology and a software tool called PEx-Image – Projection Explorer for Image *for analysis and exploration of image collections employing visualizations. The visual mappings proposed here are similarity-based multidimensional projections and point placements, which layout the data on a plane for visual exploration. The proposed approach supports various image analysis tasks such as feature selection and classification, improving data exploration capabilities. We also illustrate how it can be successfully employed in simultaneous analysis of different data types, such as text and images, offering a common visual representation for data expressed in different modalities.*

1. Introduction

Image analysis and image processing applications typically compute feature vectors from images, so that they can be compared based on content (dis)similarity. Pattern recognition algorithms typically inspect the m-dimensional space defined by the extracted features. Techniques such as Neural networks, Support Vector Machines and Clustering are widely employed for image comparison and classification. A common difficulty is the large number of features, which define a high-dimensional space that strongly affects performance of classification and clustering. Handling this problem typically requires dimension reduction and feature selection techniques. Reduction impacts the behavior of classification algorithms, which must be tuned for optimal performance on data sets almost on an individual basis.

Multidimensional Projections are commonly applied to generate graphical representations of multidimensional data. They work by projecting data originally defined in an m-dimensional space into a p-dimensional space where $p << m$ (typically $p = 2, 3$). Techniques vary in their approaches, but a common goal is that data representation in the projected space should preserve relevant data relationships defined in the original space.

We employ projection techniques to support unsupervised classifications of image data sets aimed at interactive user-directed exploratory analysis. Such techniques have been successfully employed before in *Projection Explorer* (PEx) [17] to map document sets [16, 5] based on their content similarity. We adapt and extend this underlying framework to support user-driven exploration of image data and associated textual information, now called *Projection Explorer for Image* (PEx-Image). We illustrate how integration of both input types into the same exploratory environment enhances the capability of the visualizations in finding new interesting cases within the data.

We also investigate how multidimensional projections can benefit from an embedded feature selection scheme and how a classifier can aid in the visual data exploration process. The former yields projections with more uniform groups and the latter enhances data exploration potentialities by adding information to that already revealed by multidimensional projections alone.

Textual data sets are often expressed as high-dimensional vector spaces for classification, retrieval and visualization tasks. Feature vectors are also extracted from images to define an m-dimensional vector image space. Such high-dimensional data points may be input to projection techniques, which generate a representation of the points in R^2 that can be displayed and interacted with. Figure 1 illustrates a projection from an image data set generated with *PEx-Image*, employing a point-placement technique called Neighbor-Joining (NJ) similarity tree [5]. This particular data set contains 6 different classes of medical images, each

described by 45 samples. If a good similarity metric is employed the NJ-tree effectively separates groups of similar images (or classes) in branches, as observed in Figure 1. It is therefore a useful representation, e.g., to validate the metric itself for further processing of comparable data sets.
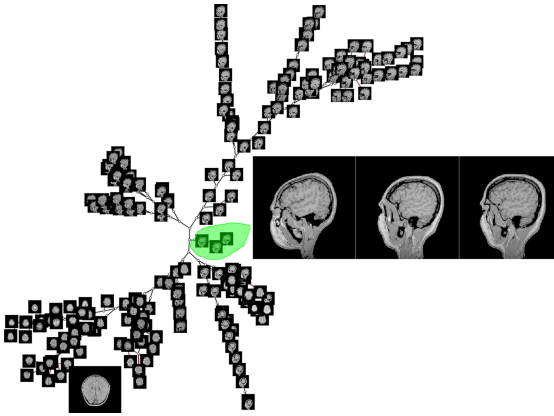
Figure 1. 2D projection of a medical image data set and details. Layout generated with the Neighbor-Joining technique.

Views generated by different techniques from the same data, or from distinct but somehow related data sets, may be coordinated in the exploration tool. We illustrate how such coordinated multiple views (CMV) benefit analysis of image data sets, favoring identification of relevant associations.

2 Related Work

Organizing large image collections for effective automated search and retrieval is a major research issue in Computer Vision and Computer Databases. A recent trend investigates how visualization and mining techniques may assist user exploration of such databases, where the quest is for unexpected albeit potentially useful information. Chen et al. [3], for example, employ visualization techniques to compare images based on three distinct feature extraction schemes that use color, layout and texture information, respectively. Images are classified by a system that employs associative networks obtained with Pathfinder Network Scaling (PFnets) [21]. Nodes in the networks are connected to reflect proximity and similarity between image pairs, preserving only the most important connections.

Unlike Chen et al., who employ an unsupervised approach to establish image similarities, the method by Fan et al. [7] classifies images based on automatic annotations. An ontology is constructed from manual annotations associated to some images. Objects are extracted from the images and receive annotations that allow assigning an image

to a specific domain. New images added to the database are classified based on the ontology previously constructed. For that, the main objects in the image are automatically extracted and their features computed. Resulting features are compared with those from the ontology, and a classification is performed with Support Vector Machines. A visual representation is created using Principal Component Analysis and a hyperbolic tree, which groups images related to similar concepts.

Several applications that employ visual representations of image data to support user exploration employ dimensionality reduction to generate two-dimensional displays. Dimension reduction can be applied to the image as a whole [19, 22] or to image feature vectors [14]. Similar to the work proposed here, Yang [24] shows a classification of similar images by projecting the data from a high-dimensional space to a lower-dimensional one. He describes experiments with data sets containing the same underlying image with variations in shape, brightness and orientation. No features are extracted, rather all pixels are employed to describe the multidimensional image object.

Our work contrasts with previous ones in that it offers a visual framework to support the overall process of analyzing image data sets, from simple exploration tasks to evaluating image manipulation and mining algorithms. Moreover, it relies on effective visual maps capable of unifying different data types into a common visualization framework, an integration that favors identification of associations between different views of the same data or of different data types related to the same phenomenon. Also, to the best of our knowledge, there are currently no attempts to explore the benefits of combining multidimensional projections, automatic feature selection and classification into a unified framework.

3 Background

Image analysis and visualization techniques related to this work are now briefly reviewed.

3.1 Multidimensional Point Placement and Projections

Distance-based projections strive to preserve, in the projected space, the distance relationships established by the similarity measurements in the original space. A good projection groups highly related 'points' close together (proximity measured by their Euclidean distance in R^2). Figure 2 presents a sample projection of the same data set depicted in Figure 1, generated with the *PEx-Image* tool described in Section 4. *PEx-Image* implements recent projection techniques known to perform well on multidimensional data

[16, 13] and also many classical projection and dimension reduction techniques.
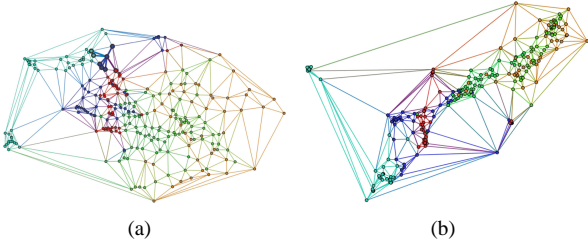
Figure 2. Alternative projection of the medical image data shown in Figure 1. Colors indicate image classes. a) IDMAP projection [13]; b) LSP projection [16].

Similarity trees (Figure 1), on the other hand, rely on data object similarity to derive a hierarchical organization. Particularly, the NJ-tree technique employs the same rationale adopted by algorithms that reconstruct phylogenic trees, replacing the concept of ancestry by a virtual object node with combined dissimilarity. The tree is created based on the property that similar nodes – that share common properties – are assigned to the same branch, in a bottom-up strategy. Therefore, a node is assigned as an ancestor of another one when they share similar content. The tree is un-rooted and its reduced depth yields rational use of screen space, and a useful hierarchical interpretation that reveals both local and global similarity relationships.

3.2 Feature Selection by Salience Measures

Feature selection is generally employed to reduce measure acquisition costs and improve precision of image classification systems [10]. We employ a feature selection scheme based on salience measures originally proposed by Garson [8] and further extended by Santos [20].

The fact that neural networks make no prior assumption on the nature of the data distribution has been explored by many researchers to avoid the restrictive conditions of statistical methods [2]. Garson [8] proposed a salience measure that indicates the importance of each MLP (Multi-Layer Perceptron) input node. His heuristic approach, which is an estimate based on the cost of all connection paths from the input to the output layer, was shown to be computationally viable when compared with traditional selection methods [15].

First of all, a normalization factor N_h is computed for each h hidden neuron. It is the sum of the absolute values of the neuron weights $w_{i,h}$ (index i indicates the input node number): $N_h = \sum_i |w_{i,h}|$

Each possible path (where variable o stands for output neurons) returns a value $w*_{i,h,o}$ based on N_h and output weights: $w*_{i,h,o} = \frac{|w_{i,h}||w_{h,o}|}{N_h}$

Finally, salience $S_{i,o}$ of input i, with respect to output neuron o, is given by: $S_{i,o} = \sum_h w*_{i,h,o}$

The salience measure enables to eliminate a feature with no knowledge on its related contribution to classification rate. Hence, rather than training the neural network several times to evaluate each feature contribution, a single training suffices to rank the relevance of features based on their salience values. For improved accuracy intermediate trainings can be performed until the end of selection process.

Salience is usually applied to reduce selection complexity. For instance, for each feature removed, a training could be performed; for n features, the complexity would be $O(n)$ (more specifically $O(n - m)$, because the algorithm ends when an m-sized subset is reached). This linear complexity algorithm is preferable to an exhaustive search, although there is no guarantee of optimality.

We adopt in our studies a feature selection by salience that overcomes two major drawbacks of Garson's original proposal. First, the approach is extended to handle an n-class problem, as opposed to a 2-class problem only (single MLP output). Second, when faced with irregular features (uneven levels of mean values and/or within different numerical ranges), the MLP tends to provide higher weights for features with lower mean values. This has a direct influence on salience measures for they are computed over such weights. The normalization scheme based on *centering* and *scaling* techniques [20] overcomes this problem, so that the salience measure computed actually reflects the discriminant power of each feature.

Another advantage of the feature selection process by salience with MLP neural nets is that classification can be attained at virtually no extra computational effort, as the weights computed during selection can be employed during classification without another training process.

## 4 *PEx-Image*: A Tool for Exploring Image Data Sets

PEx-Image tailors and extends previous projection techniques to handle image databases, offering several positioning techniques that emphasize image content similarity plus complementary functionalities to assist focused data exploration.

Input data for *PEx-Image* may be a set of images, from which characteristics may be extracted and combined into image feature vectors. Each feature vector represents an image as an m-dimensional point, where m is the feature vector size. Alternatively, *PEx-Image* may take as input any set of m-dimensional points – e.g., feature vectors previously extracted, possibly by another tool, or even other data types

represented as vectors, such as textual or record data. A user may then choose the type of projection to generate a 2-D layout, as discussed in Section 3.1. Various object similarity metrics are supported, such as Euclidean Distance, Cosine Similarity and City Block.

Once the high-dimensional data points (images, in this case) are projected they are displayed on the screen using proper graphical marks. Typical marks are circles, which may be connected with lines to represent relationships between data objects, such as similarity, proximity or a triangulation of some type. In situations where the data set has one or more associated scalar fields – e.g., an assigned class or type – scalar values can be mapped to circle color or size. Alternatively, miniatures of the images themselves may be employed as graphical markers, as illustrated in Figure 1.

As long as the distance metric employed is effective in capturing content similarity, the resulting layout provides a visual map that approximates and groups similar images, providing a powerful starting point for further exploring the data set. Several functionalities assist users in interaction tasks (illustrated in the video provided). The analyst may resort to capabilities such as zooming, moving points to resolve overlapping, moving points automatically according to their connections, highlighting the nearest neighbors of a selected image (Figure 3(a)) and recovering the correspond images (Figure 3(b)), or select a region and recover the corresponding images.

(a)   (b)
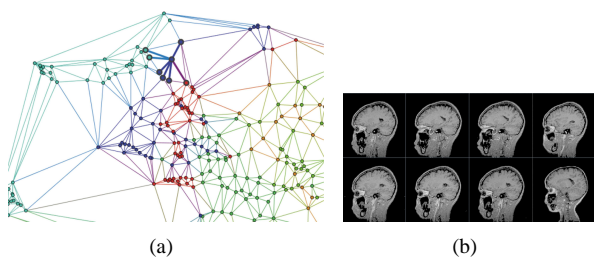
Figure 3. Visualizing the nearest-neighbors of a user-selected image: (a) Selected image and its neighbors (highlighted circles); and (b) Image and its neighbors (most similar images).

PEx-Image also supports coordinated interaction with multiple views of the image data – e.g., distinct layouts generated by multiple projection techniques – as well as with views of related data, for example, textual information associated to images. Coordination implies that user selection of an image or a group of images in a view highlights the corresponding images (or their related objects) in the other coordinated views. Coordination features are illustrated in the following section.

5   Results

The following case studies employ five distinct families of image features, four of which (a–c and e) conveying textural information: a) Fourier Descriptors from image histogram and Energy computed from 2D image Fourier descriptors [9]; b) a set of 16 Gabor Features, taken from 4 frequency levels (0.02-0.3) and 4 distinct orientations [11]; c) 72 co-ocurrence matrices features computed from 6 Haralick's functions over 12 combinations of distances and directions [23]; d) mean intensity and standard deviation computed over the entire image; e) a set of 1,024 wavelet features [6]. The combination of features (a) to (d) is referred to as *group one* features.

The first study describes how projections can help assessing the efficiency of a particular feature set in grouping similar images, for a particular data set. Figure 4 illustrates four distinct NJ-tree projections (with Euclidean distance) of a set of 512 MRI medical images distributed into 12 classes, each indicated by a different color. Projections (a), (b), (c) and (d), respectively show the results for: 16 Gabor features; 20 Fourier descriptors from the image histogram; 6 Energy Fourier descriptors, mean and standard deviation; 72 Co-ocurrence matrices features; all features combined.

View coordination is an important aid to compare different feature combinations and here it helps identifying which feature set shows the best grouping results. By coordinating the four different data views (node color indicates image class), and selecting in the view shown in Figure 4(d) those images that do not belong to head images (dark blue), one observes the corresponding images highlighted in the four projections. Inspection reveals that the fourth choice of features (Figure 4(d)) is the most effective in grouping the different image classes.

The same data set was employed in a second experiment which evaluates the grouping capabilities of different distance metrics under NJ-tree projections. Figures 5(a), 5(b) and 5(c) show, respectively, the projections obtained with the Euclidean, Cosine Similarity and the City Block distances. *Group one* features were employed in all cases. Projections in Figure 5 show that the NJ-tree successfully captures the groups of similar images. Images belonging to a given class have been either mapped into branches close to each other (in most of the cases) or into a single branch (in some cases). This behavior is reinforced by the projection in Figure 6, which shows the same NJ-tree illustrated in Figure 5(b), now displaying miniature images as graphical markers. Closer inspection of the projections reveals that better class separation is achieved with the Euclidean distance (Figure 5(a)) and Cosine Similarity (Figure 5(b)).

The next study uses a public X-ray image data set called ImageCLEF 2006 (available at http://ir.shef.ac.uk/imageclef/2006/). Figure 7 shows
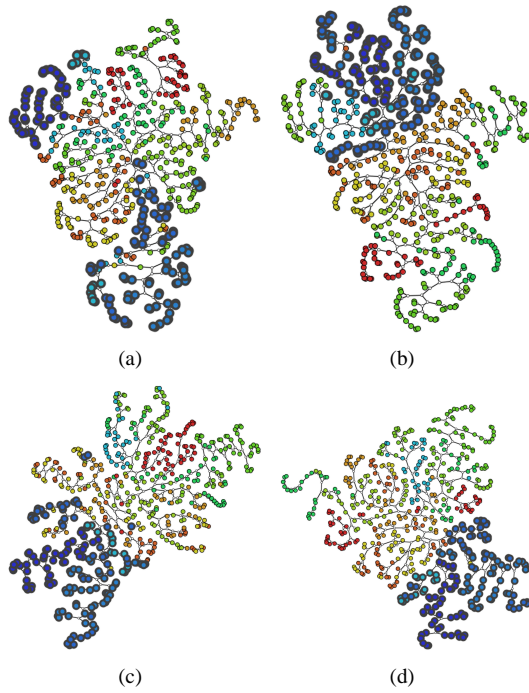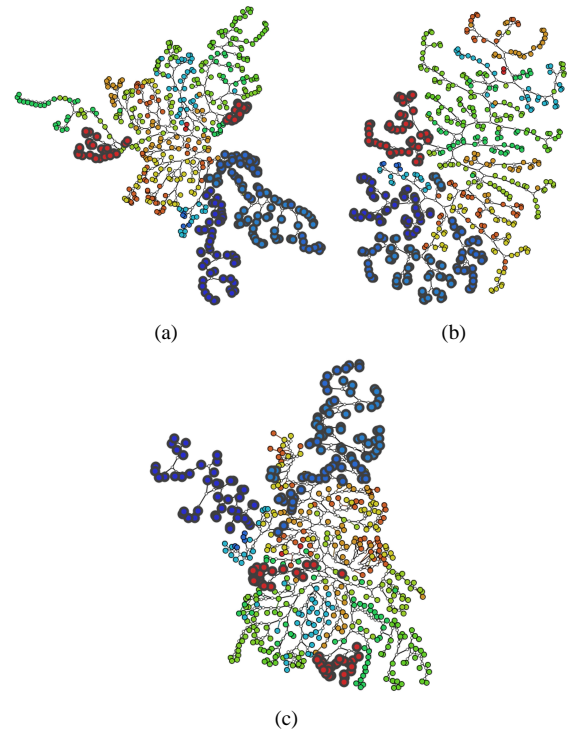
Figure 4. Projections of the medical image data set with Euclidean distance for different feature sets. (a) Gabor; (b) Fourier descriptors, mean and standard deviation; (c) Co-occurrence matrices and (d) all features combined (*group one*).

Figure 5. NJ-tree projections of the medical image data set for *group one* features and 3 different distance metrics: (a) Euclidean; (b) Cosine Similarity; and (c) City Block.

two NJ-trees with 1,000 test samples from such data set. Figure 7(a) shows the NJ-tree obtained from the same feature set employed in the projection shown in Figure 4(d). On the other hand, Figure 7(b) shows the NJ-tree for the 1,024 wavelet features. Coordinated examination of both views, by selecting and highlighting image clusters, reveals that the wavelet features outperform the other set.
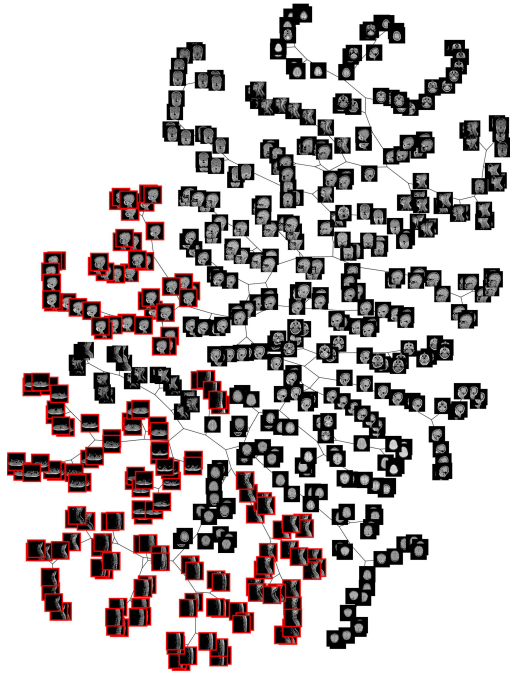
In real life applications class information is usually not available for testing purposes, and classifiers can aid in the exploration process. In the ImageCLEF training data, we sought for classes with no less than 50 and no more than 200 samples. As a result, a total of 3,345 image samples (within 35 distinct classes) have been selected as a training set, from which both the *wavelet* and *group one* feature sets were computed. The MLP-based classification process was then carried out on both feature sets. The primary purpose here was not to reduce the number of features but, instead, classify all samples and analyze how this could contribute to a projection-based exploration of the test data. Figure 8 shows 3 NJ-tree projections for a set of 537 test images which includes the same classes found in the training data set.

The MLP-based classifier derives a classification hypothesis, and combining this information and the NJ-tree projection can enhance the user-driven exploration task. Figure 8(a) shows an NJ-tree colored by the given image classes previously labeled by physicians, whereas Figures 8(b) and 8(c) show NJ-trees whose node colors indicate the class assigned by the classifier to each test sample. The former was obtained with the *group one* feature set, whereas the latter employed the *wavelet* features. The neural network training was interrupted (converged) with a 63.67% error in the *group one* features training, and with a 23.28% error in the wavelet features training – the latter produced the best results.

The next case study illustrates how the MLP-based feature selection can contribute to the exploration process. Figure 9 shows a comparison between two NJ-trees computed from the wavelet feature vectors with different dimensions, i.e., distinct feature spaces. Figure 9(a) shows an NJ-tree created with the complete (1024) feature set, whereas an NJ-tree with the best 37 wavelet features is shown in Figure 9(b) – this subset includes only those features with salience values greater than 0.5. A feature ranking was generated based on a neural network training with an error rate equal
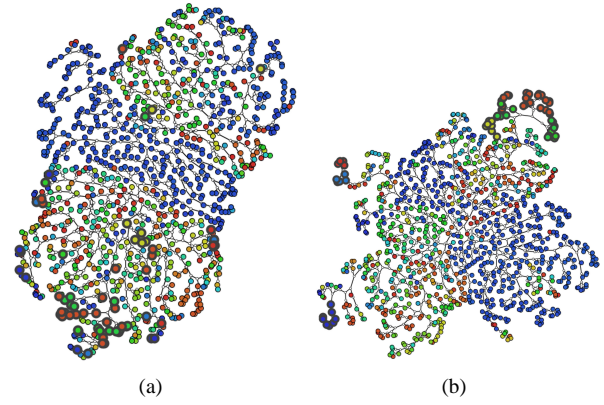
Figure 6. NJ-tree projection of Figure 5(b), with miniature images instead.

(a)   (b)

Figure 7. NJ-tree projections with Euclidean distance metric for the X-ray ImageCLEF data set. (a) *group one* **features; (b) wavelet features.**

29.62%. Coordinated explorations of the projections reveal that the smaller feature set (Figure 9(b)) actually yields better groupings than the complete one (Figure 9(a)), resulting in similar images placed in the same or in nearby tree branches.

Visual analysis strategies as presented here are also useful to explore biological data sets such as protein sequences. Such sequences may be expressed through the construction of particularly elaborated images. Protein structures can be represented as a 2D matrix, where element $[i, j]$ carries information about the interactions of residues i and j [18]. Using a tool such as PROTMAP2D [12] one can generate 2D maps from such information, named *contact maps*, and obtain images that carry relevant information about the genetic sequences. Figure 11 shows four sample contact maps from a protein sequence data set.

In the following study we use sequences from the RSCB *Protein Data Bank* (www.rcsb.org/pdb), which contains information on thousands of sequences and their respective structures. We started with 11 sequences with different enzyme classification, randomly selected from RSCB-PDB. These sequences were compared, using BLAST [1], against the PDB repository containing 63,762 sequences. By selecting the most similar sequences to each source sequence, using a threshold of e^{-100} for the e-value attribute, we obtained 138 sequences belonging to 11 different classes.

We then generated the contact maps of those sequences, and independently projected the resulting images and the amino-acid sequences themselves (see Figure 10). Coordinated exploration of the views enables comparing the information that can be gathered from each visualization. Figure 10(a) shows the NJ-tree built from the contact map images, using the same procedure employed for the other image sets. A total of 28 features have been used: 26 Fourier descriptors from image histogram and whole image, mean and stardard deviation. Figure 10(b) is the NJ-tree built from the sequences, using as similarity measurement an approach known as Normalized Compression Distance - NCD [4], which is computed directly from the sequences. Similarity is computed by approximating the Kolmogorov complexity of the data strings using compression algorithms. We have used gzip and bzip2 compressions, with the equivalent results. Node color in the projection indicates the original enzyme class of the sequence.

One observes that both trees have extremely similar structures. Same sequences are placed next to each other on the projections from either input data. The highlighted nodes, all belonging to the same class, appear shifted from one projection to another but still remain in the same neighborhood. A few groups are circled and labeled in the picture, to stress their positioning on the visual map. Both trees yield very similar overall interpretations regarding this data set, suggesting that either means of expression could lead to similar organization of the data based on similarity. By focusing on individual groups, one observes that the contact maps disclose information that differs from the one revealed by the similarity obtained from the sequences themselves, as reflected by the changes in branch topology. This is bound to be so, since contact maps do contain informa-
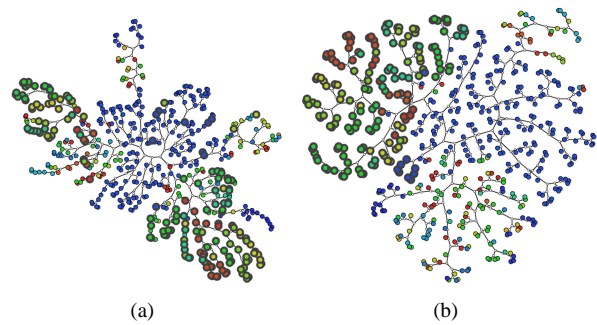
(a)

(b)   (c)
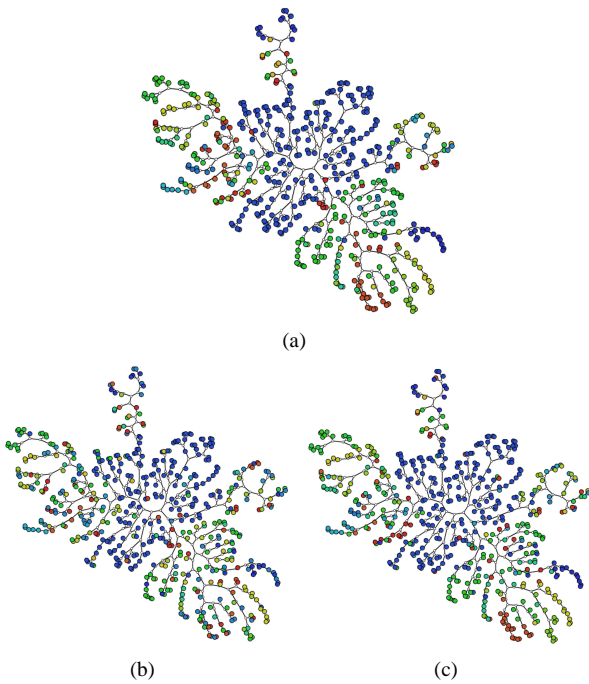
Figure 8. NJ-tree exploration (with 1,024 wavelet features and cosine similarity) aided by an MLP classifier (a) colors from image classes. Colors assigned by classifier: (b) *group one* **features and (c) 1,024 wavelet features.**

(a)   (b)

Figure 9. NJ-tree projection with cosine similarity. Feature selection for saliences larger than 0.5**. (a) 1,024 wavelet features; (b) 37 wavelet features.**

(a)   (b)
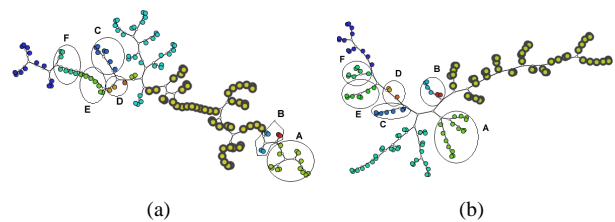
Figure 10. NJ-trees of 138 protein sequences in 11 enzyme classes. (a) Contact Map images; and (b) Protein sequences, using NCD as similarity measurement. Color maps class, and areas labeled correspond to the same sequences in both maps. Highlighted nodes belong to a particular class.

tion that is not directly expressed by sequence similarity. This example highlights the advantages of having a similar framework to examine information of different natures regarding the same basic data. Similarities and differences in information content can be located and explored immediately.

A discrepancy immediately noticeable is observed in Figure 10(a), where two 'green' nodes (next to group D) seem to be misplaced, well away from their original class (group A). In the map built from the sequences, they are placed in the same branch as the remaining nodes A. Close examination (see Figure 11) reveals that, although the sequences are similar, the contact maps of the two isolated nodes are actually quite different from the others, revealing particularities of the residual information for those sequences. Such information would have been very difficult to find without this type of visualization.

6   Conclusions

We introduced *PEx-Image*, a set of tools and a novel approach to map an image data set onto a two-dimensional

space, and demonstrated how they make data analysis and exploration more effective. Both feature selection strategies and classification algorithms can be evaluated using the visual framework presented here. The visualizations provided can facilitate and accelerate comprehension of results from image analysis algorithms. These algorithms, in turn, improve the process of visual mapping and mining.

Since the same visual analysis framework can be applied to multiple distinct data sources, e.g., images and texts, as illustrated in the biological data set example, the approach supports comparing and associating different types of information on the same phenomenon. It enables complementing information through coordination of multiple visual maps, in order to enhance analysis tasks.

7   Acknowledgements
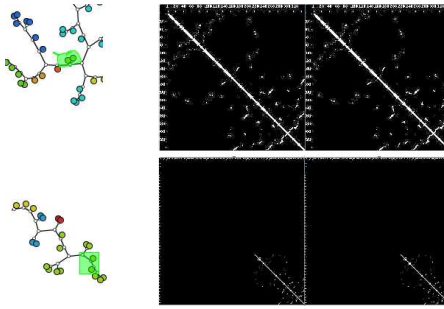
Figure 11. Slices of the map shown in Figure 10(a), showing the points seemingly misplaced. Two sequences (top) placed away from the others (bottom) in their class are actually considerably different (in image type) from the others in the group.

References

[1] S. F. Altschul, W. M. W. Gish, E. W. Myers, and D. J. Lipman. A basic local alignment search tool. *J. of Molec. Bio.*, (215):403–410, 1990.

[2] G. Castellano and A. M. Fanelli. Variable selection using neural-network models. *Neurocomp.*, 31(1-4):1–13, 2000.

[3] C. Chen, G. Gagaudakis, and P. Rosin. Similarity-based image browsing. In *XVI IFIP World Comp. Cong., Int. Conf. on Intell. Inf. Proc.*, pages 206–213, Beijing, China, 2000.

[4] P. Cilibrasi, R. Vitnyi. Clustering by compression. *IEEE Trans. Information Theory*, 51(4):1546–1555, 2005.

[5] A. M. Cuadros, F. V. Paulovich, R. Minhgim, and G. P. Telles. Point placement by phylogenetic trees and its application for visual analysis of document collections. In *IEEE Symp. Visual Analytics Sci. and Tech. 2007*, pages 99–106, Sacramento, California, USA, 2007.

[6] L. A. da Silva, R. A. Moreno, S. S. Furuie, and E. D. M. Hernandez. Medical image categorization based on wavelet transform and self-organizing map. In *Proc. VII Int. Conf. on Intel. Syst. Design and Applications*, pages 353–356, Washington, DC, USA, 2007. IEEE CS Press.

[7] J. Fan, Y. Gao, and H. Luo. Hierarchical classification for automatic image annotation. In *Proc. XXX ACM Int. Conf. on Research and Develop. in Inf. Retriev.*, pages 111–118, New York, NY, USA, 2007. ACM Press.

[8] G. D. Garson. Interpreting neural net connection weights. *AI Expert*, 6(4):46–51, 1991.

[9] K. Huang and S. Aviyente. Rotation invariant texture classification with ridgelet transform and fourier transform. In *ICIP*, pages 2141–2144. IEEE, 2006.

[10] A. K. Jain and D. Zongker. Feature selection: Evaluation, application and a small sample performance. *IEEE Trans. Pattern Anal. Mach. Intell*, 19:153–158, 1997.

[11] B. S. Manjunath and W.-Y. Ma. Texture features for browsing and retrieval of image data. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 18(8):837–842, 1996.

[12] I. T. Michal J. Pietal and J. M. Bujnicki. Protmap2d: visualization, comparison and analysis of 2d maps of protein structure. *Bioinformatics*, 23(11):1429–1430, March 2007.

[13] R. Minghim, F. V. Paulovich, and A. A. Lopes. Content-based text mapping using multi-dimensional projections for exploration of document collections. In *Visualization and Data Analysis 2006, Proc. SPIE-IS&T Electronic Imaging*, 2006.

[14] B. Moghaddam, Q. Tian, N. Lesh, C. Shen, and T. S. Huang. Visualization and user-modeling for browsing personal photo libraries. *Int. J. Comput. Vision*, 56(1-2):109–130, 2004.

[15] R. Nath, B. Rajagopalan, and R. Ryker. Determining the saliency of input variables in neural network classifiers. *Computers & OR*, 24(8):767–773, 1997.

[16] F. V. Paulovich, L. G. Nonato, R. Minghim, and H. Levkowitz. Least square projection: a fast high precision multidimensional projection technique and its application to document mapping. *IEEE Trans. on Vis. and Comp. Graph.*, 14(3):564–575, 2008.

[17] F. V. Paulovich, M. C. F. Oliveira, and R. Minghim. The projection explorer: A flexible tool for projection-based multidimensional visualization. In *Proc. XX Braz. Symp. on Comp. Grap. and Im. Proc.*, pages 27–36, Washington, DC, USA, 2007. IEEE CS Press.

[18] D. C. Philips. The development of crystallographic enzymology. *asf*, 30:11–28, 1970.

[19] S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, December 2000.

[20] D. P. Santos and J. B. Neto. Feature selection with equalized salience measures and its application to segmentation. In *Proc. XX Brazilian Symp. on Comp. Grap. and Im. Proc.*, pages 253–262. IEEE CS Press, 2007.

[21] R. W. Schvaneveldt, editor. *Pathfinder associative networks: studies in knowledge organization*. Ablex Publishing Corp., Norwood, NJ, USA, 1990.

[22] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, December 2000.

[23] M. Tuceryan and A. K. Jain. Texture analysis. In *The Handbook of Pattern Recognition and Computer Vision (2nd Edition)*, pages 235–276, 1998.

[24] L. Yang. Distance-preserving projection of high-dimensional data for nonlinear dimensionality reduction. *IEEE Trans. Patt. Anal. Mach. Intell.*, 26(9):1243–1246, 2004.