

Coastal typology development with heterogeneous data sets

B. A. Mather · Robert W. Bilbrey

Abstract This paper presents a data-driven expert-guided method of coastal typology development assigning colors to different classes of data based on class similarity in a high-dimensional space. This combination of techniques results in a methodology through which one or more experts can easily develop a useful coastline typology with results that are similar to preexisting expert typologies, but which makes the process more quantitative, objective, consistent, and applicable across space and time.

Keywords Coastal zone · Typology · Clustering · Visualization · Distance measure

Introduction

The Land–Ocean Interaction in the Coastal Zone project (LOICZ) is a component of the International Geosphere–Biosphere Programme (IGBP) that focuses on the area of the earth’s surface where land, ocean, and atmosphere meet and interact. The overall goals of this project are (1) to determine at regional and global scales the nature of that dynamic interaction; (2) to determine how changes in various compartments of the Earth system are affecting coastal zones and altering their role in global cycles; (3) to assess how future changes in these areas will affect their use by people; and (4) to provide a sound scientific basis for future integrated management of coastal areas on a sustainable basis (Pernetta and Milliman 1995).

A primary LOICZ objective is developing global-scale estimates of biogeochemical fluxes of carbon, nitrogen, and phosphorus (C, N, and P) in and through the coastal zone (CZ) (Gordon et al. 1996). The strategy adopted is to identify “type-specimen” CNP budgets for well-characterized coastal regions, to further identify the coastal regions around the world of which such functional observations might be typical, and to use this typology relationship to upscale the limited local data to an estimate of global coastal zone function. Within this context, a typology is defined as a classification system that divides coastal zones into a set of classes according to one or more physical, geological, atmospheric, or human-related variables.

The development of an inventory of standard-format CZ budgets is in progress in the Biogeochemical Budgets task of LOICZ (<http://data.ecology.su.se/mnode/>). The typology project (<http://www.nioz.nl/loicz/typo.htm>) is responsible for developing the coastal classification approach needed for budget upscaling. One of the major strategies adopted is the development of clustering and visualization techniques suitable for classifying coastal areas in terms of their similarity with respect to environmental variables relevant to biogeochemical function.

The task is challenging because of the need to rely on globally available data, and to incorporate many different types of variables – marine, terrestrial, climatic, biotic, and “human dimension” (i.e., socioeconomic and environmental alteration). Although a variety of data are available, data sets differ in format, resolution, classes, and completeness, and the data themselves are typically not normally distributed or amenable to standard statistical analyses.

Received: 10 April 2000 / Accepted: 14 June 2001
 Published online: 2 March 2002
 © Springer-Verlag 2002

B.A. Mather (✉) · R.W. Bilbrey
 Strategic Center,
 Coastal Geoscience
 Strategic Center, PA 19081, Lancaster,
 Pennsylvania 19081, USA
 e-mail: mather@strategiccenter.com
 Tel.: +1-610-3288081
 Fax: +1-610-3288082

Traditional approaches to typology development for geospatial data take either a top-down or a bottom-up approach. In a top-down approach, experts design a decision tree based on different variables and variable ranges that seem appropriate for the environment being considered (Lankford 1977; Smith et al. 1997; Smith and Crossland 1999). The experts then apply this scheme to a data set and iteratively refine the classifications. A variation on this approach is to have experts classify a training set for a pattern classifier – either symbolic or subsymbolic like an artificial neural network – and then have the pattern classifier learn the classes from the training set and generalize the classification strategy to unseen data. In the bottom-up approach, a clustering method is used to determine groups of similar data points which then form standard classes. Traditional clustering methods include agglomerative clustering and the K-means clustering algorithm, also known as vector quantization (VQ) (MacQueen 1965; Anderberg 1973; Hartigan and Wong 1979; Rabiner and Juang 1993). In a geographic/geological context, researchers have used a variation on bottom-up clustering termed regionalization, which locates spatially contiguous class members after applying a general agglomerative clustering to the data set that ignores spatial location (Harff and Davis 1990). Researchers have also used K-means clustering on Landsat-4 data to examine the geographic differences between coastal areas of large lakes (Maktav 1985). In both of these cases the data sets were fully populated and the number of variables small and statistically well behaved.

In the top-down typology approach, the result is dependent upon expert decisions. In a bottom-up approach, the resulting typology is affected by two major issues, both of which can be guided by expert input. First, how many classes should there be in the typology? Second, how do we measure similarity between data points? The second is especially important when we consider multi-dimensional heterogeneous vectors – data points that have multiple variables with different ranges, variances, and meanings.

The answer we propose to the first question – how many classes? – is that expert opinion guided by data analysis is most appropriate. The data analysis we propose using in the section “Using description length to determine the optimal number of clusters” is an information-theoretic criterion that balances the costs and benefits of using more or fewer classes. If we used agglomerative clustering – which generates a tree or dendrogram of classification schema for all possible numbers of clusters – then the expert user would generally select the appropriate level. As we note below, our information-theoretic criterion would also apply in this situation.

The traditional answer to the second question – how do you measure similarity of heterogeneous data? – is to use a statistical measure that incorporates variances and covariances of the variables. However, when dealing with potentially incomplete heterogeneous global data sets at different scales, traditional techniques begin to break down. The first casualty is that the covariance matrix

becomes noninvertible, making it impossible to use the Mahalanobis distance that handles covariances. The second casualty is that with missing variables in some locations we cannot use a simple “sum-of” technique because different pairs of data points will sum different numbers of values. In sections “Traditional distance measures for heterogeneous data” and “An alternative distance measure based on extremes”, we present our solutions, which include two measures that degrade gracefully in the face of missing data and still incorporate variance information.

In the section “Experiments and results”, we present the results of typology development on Australasia, which is a good example location because of the existence of both expert typologies for the region and a large number of budget sites which we can use for flux estimation (Smith and Crossland 1999). In the section “Discussion and future directions”, we discuss the results and present directions for future work. Finally, we conclude with a summary of the typology development process in “Conclusion and summary”.

Theory and methodology

Before we can begin to develop classifications and typologies, we must first define what we mean when we say two data points are similar. Only then can we think about grouping similar points and building conceptual structures. With a mathematical definition of similarity, we can bring to bear numerous useful concepts and algorithms from statistics and pattern recognition. This section defines two reasonable definitions of similarity and then presents a suite of algorithms that use these definitions for typology development.

Traditional distance measures for heterogeneous data

A useful way to think about similarity is as the distance between two data points. If the two points are similar, the distance between them is small. As their similarity decreases, the distance between them gets larger. In a homogeneous multivariable space – such as the 3D space we inhabit – a useful measure of distance is Euclidean distance. For the N-dimensional data points x and y , we can define Euclidean distance D_E as:

$$D_E = \sqrt{\sum_{i=1}^N (x_i - y_i)^2} \quad (1)$$

This measure of similarity is appropriate for measuring the distance between points with homogeneous dimensions. Homogeneous in this context can be defined as each dimension having similar ranges and variances. For the LOICZ data set, however, and in fact any heterogeneous data set – defined as a data point whose dimensions do not have similar ranges or variances – Euclidean distance is not a useful measure of similarity. Put simply, it

does not take into account the fact that dimensions with larger scales and variances will dominate the distance measure. Instead, we need to normalize the ranges and variances of the different dimensions so that, for example, an elevation variable measured in meters does not dominate one measured in kilometers. One way to do this is to use a scaled Euclidean distance, D_S , which divides the squared difference in each dimension by the variance of that dimension as shown below:

$$D_S = \sqrt{\sum_{i=1}^N \frac{(x_i - y_i)^2}{\sigma^2}} \quad (2)$$

Intuitively, this distance is a statistical measure that weights the difference in a given dimension according to how significant that difference is based on the estimated variance of the data in that dimension. For a heterogeneous data set whose data points are fully populated – data exist in every dimension of every point – this definition of distance is reasonable. For a heterogeneous data set whose data points are not fully populated in every dimension, such as the LOICZ data set, we need to deal with the missing data.

One method of dealing with the missing data problem is to use the average scaled Euclidean (ASE) distance between two points (Backer 1995). This distance measure, D_A , is defined as:

$$D_A = \frac{\sum_{i \in \text{Valid}} \frac{(x_i - y_i)^2}{\sigma^2}}{\text{card}(\text{valid})} \quad (3)$$

where *Valid* is the set of dimensions that have valid data in both x and y , and $\text{card}(\text{Valid})$ is the number of valid dimensions.

The distance measure D_A can be interpreted in the following intuitive manner. If the value is less than one, then the average difference between x and y in any one dimension is less than a standard deviation. If the value is greater than one, then the average difference is greater than a standard deviation. Taking the square root of D_A would provide an exact measure in terms of standard deviations.

When the covariance matrix of the data set – which specifies the correlation between variables – is invertible, then we can bring the Mahalanobis distance measure to bear on the problem (Backer 1995). The Mahalanobis distance takes into account not only the scaling and variance of a variable, but also the variation of other variables based on the covariances. Thus, if there are multiple variables in the data set that are all providing similar information, for example, if there were multiple variables dependent upon temperature, then their contribution is weighted appropriately so that the single factor, i.e., temperature, does not dominate the distance measure. Unfortunately, with high-dimensional heterogeneous data sets the covariance matrix can become singular, i.e., noninvertible, which means that the Mahalanobis distance is not possible to compute.

An alternative distance measure based on extremes

An alternative distance measure for geographic classification is to use the maximum scaled difference (MSD) between corresponding variables rather than the average scaled distance. In other words, two vectors that are identical except for a single variable x_i will have the scaled difference in x_i as their distance. Compare this with a traditional measure, where the fact that most of the differences are zero drives the Euclidean or scaled Euclidean distance towards zero as the number of dimensions increases. A formal definition of the distance is given by:

$$MSD(A, B) = \max_{i \in \text{Valid}} \left(\frac{(A_i - B_i)^2}{\sigma_i^2} \right) \quad (4)$$

The MSD is a well-behaved distance measure since it obeys the properties of identity, symmetry, and the triangle inequality. In other words, two vectors that contain all variables can only have a distance of zero if they are equal to one another (identity property). Two vectors have the same distance no matter the order in which they are considered (symmetry property); and if $MSD(A, B) \neq 0$ and $MSD(A, C) \neq 0$, then $MSD(B, C) \neq 0$ (triangle inequality), which just states that if two points are not equal, they cannot both be equal to some third point. The MSD also behaves nicely both with respect to missing variables – it just considers variables that exist in both data points – and multiple variables that carry the same information – it considers only the maximum difference.

Another way of thinking about the MSD is that it lets the extremes rule judgements of similarity; two vectors cannot be similar if they have a single variable that is very different. In our implementation of MSD distance, we use the maximum normalized squared difference, where the normalization constants are the variances of the specific variables. In some ways, this distance measure may better capture what we think of as similarity in coastlines. Two habitats that are very much the same except for one variable – such as temperature or precipitation – may end up being very different. Conversely, we would think of two locations that have small differences in all variables as being fairly similar. The average scaled Euclidean distance could rate both of these cases as being equally similar, but the MSD distance would say the latter case – lots of small variations – should be more similar. Thus, the MSD distance starts to capture some of our intuition on the problem.

Other researchers have also attempted to use alternatives to an Euclidean-based distance for environmental classifications. One example is to use a multidimensional scaling approach where the rank of a data point's distance to another data point is weighted more than the actual distance (Dyer et al. 2000).

The MSD distance is inspired by the Hausdorff distance, which is a measure of similarity between *sets* that has been used successfully in image comparisons and object recognition tasks in the field of computer vision (Huttenlocker et al. 1993). It has also recently been used in data mining applications to select variables and build decision trees

(Piramuthu 1999). The Hausdorff distance says that the distance between two sets A and B is the maximum of the minimum distances between all points in A and all points in B.

Unsupervised K-means clustering

Given a definition of similarity, we can now start to look for natural groupings of similar points that may indicate the existence of a meaningful class. A standard method for clustering similar points is unsupervised K-means clustering, also called vector quantization (VQ) (MacQueen 1965; Hartigan and Wong 1979; Rabiner and Juang 1993). Overall, the algorithm takes as input a distance measure, a data set, and a desired number of clusters. It then attempts to find a set of vectors that best represents the data set. Each of these vectors is the mean vector of a unique subset of the data points. The output of the VQ algorithm is the set of mean cluster vectors and a tag for each data point, indicating its cluster membership.

The algorithm is briefly defined below. The inputs to the algorithm are the distance measure $D(P_1, P_2)$, the number of clusters K , and the data points $Q[1...N]$. The output is a set of mean cluster vectors $V[1...K]$.

- Assign randomly selected data points to $V[1...K]$
- Loop
 - Calculate a tag value for each data point $Q[1...N]$. The tag is the index of the closest $V[i]$ according to $D()$
 - Calculate a new set of mean cluster vectors $V'[1...K]$
 - If V' is the same as V , then terminate
 - Otherwise V gets V' and the loop continues
- Return V

Since there is a random element to the VQ algorithm, it is important to run it multiple times with the same inputs. The best set of cluster vectors V is the set that minimizes the overall representation error, which can be defined as the sum of the distances between each point and its nearest mean cluster vector.

Using description length to determine the optimal number of clusters

One problem with the VQ algorithm for typology development is that the user must specify the number of clusters beforehand. If an expert has some idea of the number of desired clusters, this is not a problem. However, the expert may not know a priori how many natural clusters there are.

One way to approach this problem is to look for a natural breakpoint in the error as the number of clusters increases. As the number of clusters gets larger, the representation error tends to zero – which it will become when there are as many clusters as there are data points. Figure 1 shows a plot of error versus the number of clusters that demonstrates this tendency. However, the utility of increasing the number of clusters is not a constant function. At some point, the reduction in the representation error is not worth adding another cluster.

To generate the error information, we first specify a range of K values and run the VQ algorithm multiple times for each K . A plot of the resulting error values relative to the number of clusters provides a graphical description of the

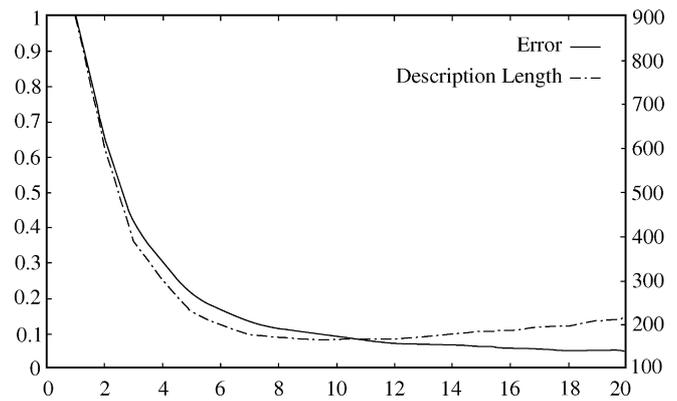


Fig. 1 Representational error versus number of clusters for a three-variable data set for the Australian coast

benefit of increasing the number of clusters, as shown in Fig. 1. In this plot there appear to be two natural break-points in this plot, one around 5 clusters and another around 10. This indicates that the first five clusters are critical, the next five are still significantly decreasing the error, and beyond 10 clusters the benefit of using more clusters is minimal. For this example, which is a 3-variable data set constituting the Australian, Oceania, and New Zealand coastlines, therefore, a 10 class typology would be appropriate.

Note that we can take a more rigorous approach to determining how many clusters is appropriate. The Minimum Description Length Principle, as defined by Rissanen (1989), gives us a mathematical way of defining when we have enough clusters. Rissanen provides an information-theoretic definition of description length that is a combination of the number of parameters in a model plus the representational error for that model. The best model balances these two factors so that their sum is minimized. In the context of clustering, the model is the set of average cluster vectors that represent the data set; and the representational error for a given model is the sum of the squared distances between each point and its associated average cluster vector. The description length equation is given as:

$$DL = -\log P(x^n | \Theta) + \frac{k}{2} \log n \quad (5)$$

This says that the description length (DL) is equal to the log of the probability of the data (x^n) given the model (Θ) plus the number of model parameters k multiplied by the log of the number of data points. For the LOICZ data set and our definition of distance, the probability of the data given a model, e.g., the set of mean cluster vectors, is the sum of the squared average scaled distances from each point to its associated mean cluster vector multiplied by the number of variables per point. The number of parameters is the number of variables per point multiplied by the number of clusters. Note that since the number of variables per point is in both terms of the equation, we leave it out when calculating the description length.

A plot of the description length for each cluster is given in Fig. 1. Note that the minimum description length is reached between 9 and 12 clusters and begins to get larger again beyond that. Therefore, by this measure – a bit more rigorous than the eyeball – we should be using 9–12 clusters to represent this set of coastline data.

Segmentation through region growing

Whereas clustering is, in a sense, a global algorithm, we can also take a local approach to determining groups of similar data points. Since the LOICZ data have a geographic context, it makes sense to identify contiguous sections of coastline consisting of similar data points. The clustering approach described above does not necessarily take into account geographic considerations when deciding what data points are alike, although this has been done by other researchers (Prakash et al. 1996).

Region growing is a commonly used technique in computer vision, where local context is extremely important. The basic idea is to begin with a seed point and then add neighboring points to that region as long as they are (1) similar enough to their neighbor and (2) similar enough to the seed region. Note that the requirement for neighboring points to be similar – which we can define as a local threshold – is usually tighter than the requirement for points to be similar to the seed region – which we can define as a global threshold. When one region stops growing – because its neighbors are too different – then we can select a second seed point and grow another region. This process continues until all data points are labeled. The result of this process is a set of connected regions consisting of similar pixels. How many regions there are is dependent upon the local and global thresholds that control the growing process. If the thresholds are rigorous there will be more regions; if the thresholds are loose there will be fewer. Note that this approach removes the need to specify the number of clusters, but replaces it with the specification of the local and global similarity thresholds. What a region-growing algorithm provides is a starting point for building a hierarchy based on variable-sized contiguous building blocks.

Methods for merging regions

Once we have a set of regions – whether found through clustering or region growing – we may want to merge similar regions together regardless of their spatial location. Especially in the case of region growing, where all regions are spatially contiguous, it is important to begin matching up discontinuous but similar stretches of coastline. We can use a step-wise optimal approach to merging – also called agglomerative clustering (Anderberg 1973) – which iteratively merges the two regions with the closest mean cluster vectors. Membership in a cluster is strictly maintained with the hierarchy that develops. The algorithm for selecting and merging two regions is as follows:

- Find the pair of mean cluster vectors with minimum distance
- Give all the data points in both clusters the same label
- Calculate a mean cluster vector for the new cluster

Strict membership in a cluster hierarchy is maintained because points are not relabeled based on their distance to a mean cluster vector. At the end of the process, any given mean cluster vector represents an archetype point for its cluster and is not necessarily the closest mean cluster vector for all points in the cluster. This method of merging is appropriate for merging regions found through the segmentation/region-growing method.

The combination of region growing or VQ followed by region merging provides a method for automatically developing a hierarchical typology, if one is desired. Note that we could start the merging process from the initial set of data points, rather than the output of a K-means or segmentation algorithm. However, since these algorithms are grouping similar points in an optimal or near-optimal fashion, to start at the individual data points is probably unnecessary and may not give as good results – although this is definitely a good future comparison to make. The section “Experiments and results” presents the results of using the segmentation and merging algorithms on subsets of the LOICZ data set.

Iterative refinement for visualization of cluster relationships

Throughout the process of cluster or region development and merging it is important to be able to visualize the process and the results. The LoiczView program provides an intuitive graphical user interface to the set of tools that implement the methods described above. In particular, it allows the user to visualize both the spatial distribution of clusters and, through color relationships, the similarity of clusters in the data space. Other researchers have used spatial location to represent similarity on a 2-D plane with carefully selected dimensions (Ankerst et al. 1998). The latter style of visualization can be useful for validating clusters, but does not connect the data points to their geographic location.

The LoiczView program uses a novel iterative refinement technique for selecting the display colors to represent distances between color vectors. This is a hard problem because the distances calculated between clusters reside in a high-dimensional space – up to 100 dimensions – while color resides in a 3-D space. Therefore, in most cases we cannot select a set of colors whose distances exactly mirror the true distances between the mean cluster vectors.

As a simple example of this, consider five points in a 5-D space that are all equidistant from one another. One set of points that meets this criteria is the set $[(1, 0, 0, 0, 0), (0, 1, 0, 0, 0), (0, 0, 1, 0, 0), (0, 0, 0, 1, 0), (0, 0, 0, 0, 1)]$. In this case, each 5-D point is $\sqrt{2}$ away from every other point. In a 3-D space, it is only possible to have four points equidistant from one another – a tetragon. It is not possible to generate five points that are equidistant from one another in a 3-D space. Therefore, the best we can do when selecting colors is to approximate the true distances in color space.

The problem can be set up as follows. First, calculate the matrix of distances between each cluster vector. Normalize this matrix by dividing each element by the largest element of the matrix. Now all of the distances are in the range $[0, 1]$. Second, generate a set of random colors

and assign one color to each cluster. Now calculate the matrix of distances between the colors in color space. In this development of the technique, we will use the RGB color space, where each axis ranges from [0, 1]. Now we have two matrices whose elements are in the range [0, 1]. The following algorithm will iteratively modify the cluster colors so that it reduces the difference between the two matrices:

- Calculate the normalized cluster distance matrix D
- Assign a random color to each cluster
- Set the adjustment rate A (e.g., 20%)
- Loop
 - Calculate the color distance matrix C
 - Let E_{ij} be the largest magnitude element of $D-C$
 - Let I and J be the clusters whose error is E_{ij}
 - Let C_{ij} be the color vector from color _{j} to color _{i}
 - Adjust the color values of I and J to reduce E
- Until the matrices are close enough or we have looped enough

The update rule for the cluster colors is given as:

$$\begin{aligned} color_i &= color_i + C_{ij}AE_{ij} \\ color_j &= color_j - C_{ij}AE_{ij} \end{aligned} \quad (6)$$

The number of iterations required to produce a good result is dependent upon the size of the matrix and the number of clusters. For a 10×10 matrix, 200 iterations achieves a result that no longer changes significantly in terms of the largest error between the two matrices. For a much larger matrix, more iterations may be required. The adjustment rate is an important parameter of the problem. The adjustment rate needs to be fast enough to allow improvement, but not so large that the system overshoots good solutions. Unless otherwise specified, all visualizations involving color were developed using this algorithm.

Experiments and results

The methods described above allow us to analyze and visualize large heterogeneous data sets such as the LOICZ data set. To test and refine these methods we have applied them to a subset of the LOICZ data set and compared the results with expert judgements. Our process for developing and validating a horizontal typology (not hierarchical) is as follows:

1. Select the variables to use
2. Select how many classes (clusters) to create
3. Apply the VQ algorithm using an appropriate distance measure
4. Apply semantic labels to each cluster
5. Compare with expert judgement or pre-existing typologies

For our prototype typology development we use a subset of the LOICZ data set corresponding to the Australia/New Zealand coastline. This data set has a spatial resolution of 1°.

Variable selection

In this experiment the variable selection was based on two factors. First, did the variable provide good coverage of the area (<10% missing data). Second, did the variable actually provide useful information (vary in a reasonable way over the data set). Beyond these two considerations, the primary concern was not to give too strong a weight to any one aspect of the environment. The end result was a set of 17 variables.

The variables we selected included: seasonal precipitation (max and min), seasonal air temperature (max and min), seasonal sea-surface temperature (max and min), seasonal soil moisture (max and min), seasonal salinity (max and min), seasonal Coastal Zone Color Scanner (CZCS; max and min), average annual runoff, an annual evaporation proxy, average wave height, standard deviation of elevation, and a tidal mixing proxy. Precipitation and air temperature information are from the Intergovernmental Panel on Climate Change (1999); the remaining variables are from the LOICZ typology data set (<http://www.nioz.nl/loicz/typo.htm>). For the Australasia coast we modified the LOICZ typology data by interpolating it to cover locations with no data. For the most part this meant taking land cell variables and interpolating them onto adjacent coastal cells, and taking sea cell variables and interpolating them onto adjacent coastal cells (a coastal cell is defined as a cell that contains both land and sea).

The evaporation proxy is a combination of wind speed and vapor pressure. The proxy variable is the product of the two multiplied by 10 (vapor pressure is water vapor pressure multiplied by 10). The vapor pressure variable came from the Intergovernmental Panel on Climate Change (1999) and the wind speed from the LOICZ typology data set. The tidal mixing proxy is a combination of a tidal form variable (semidiurnal, mixed, diurnal) and tidal range. The tidal mixing proxy is tidal range multiplied by tidal frequency, where tidal frequency is (semidiurnal=2, mixed=1.5, and diurnal=1). The two base variables came from the LOICZ typology data set.

Determine an appropriate number of classes

We used the minimum description length principle, described above, to determine the appropriate number of clusters for the data set. Figure 2 shows the plot of error and description length versus number of clusters. From this graph, the appropriate number of clusters is between 10 and 15. We selected 12 classes in this example.

Cluster the data

We used the VQ algorithm using the average scaled Euclidean distance measure to generate a set of representative classes. To get a good set of classes we ran it ten times and took the lowest error result. This provided us with a reasonable set of representative classes for the data. Figure 3a shows a visualization of the resulting classes by mapping them into an image using latitude and longitude, and using color to identify the class of each data point. Figure 3b shows a visualization of the same clustering result, but with the class colors selected using the algorithm from the section "Iterative refinement for visualization of

cluster relationships” to show the relationships between classes. Note that three distinct classes exist, while the others merge into more of a continuum in the color-similarity presentation (see Table 1 for class identifications).

Apply semantic labels to each class

The next step in the typology development process involves looking at the variable averages in the class file generated by the previous step. The class files contain the average values for each variable in each class. By looking at these average values we can see their typical geophysical characteristics. Based on these characteristics we can not only give them semantic labels, but also see what differentiates two classes for which many variables have similar values. A summary of the clusters and their labels is given in Table 1. The labels were assigned based upon the average values of each variable for each class, using terminology consistent with the earlier, independent expert typology (Smith and Crossland 1999).

We can compare the classes identified in the unsupervised clustering of Australasia with a pre-existing expert typology to see how well the process compares. Figure 4 visu-

ally compares our 12-class typology with a typology developed in Smith and Crossland (1999). Despite the difference in variables used for the two typologies, the general form of the different classes is similar, with our data-driven typology showing more detail in terms of local phenomena such as the Alpine region of New Zealand and the special characteristics of bays.

Table 1 shows a direct comparison of the semantic labels for the data-driven and expert typologies, showing a good match between the two. The biggest difference is that a number of the localized classes in the data-driven typology do not show up in the coarser expert typology. Clearly this “error analysis” is based upon visual matching of two images: one produced by a cluster analysis, one by a top-down expert typology. We could obtain a quantitative estimate of the match between the two typologies by labeling each geographic cell according to the expert typology and looking at the percent overlap between the clusters. It is important to realize, however, that typologies in the absence of ground truth or an application are difficult to evaluate quantitatively. Thus, our analysis in this paper is based on qualitative expert judgements and qualitative comparison with expert typologies. This is appropriate in the case of the Australia example since we are comparing typologies that were generated using different data sets and approaches.

We can undertake the same process of typology development using the alternative MSD distance measure.

Figure 5 compares 12-class clusterings using the average scaled Euclidean distance and the Hausdorff distance. Note the similarities and differences between the two results. The biggest differences occur on the southern and northern coasts of Australia where the southern coast apparently has fewer extreme differences (but higher average differences) than the northern coast. Thus, the MSD distance does not divide the southern coast into two sections in a 12-class clustering, but the average scaled Euclidean distance does.

Before making conclusions about the utility of the MSD distance based on this figure, however, it is important to note that using the different distance measure also causes the representational error to change. This, in turn, causes the description length to have a different minimum – in this case one that is much larger. The MDL analysis (see section “Using description length to determine the optimal number of clusters”) says that instead of 12 classes, there

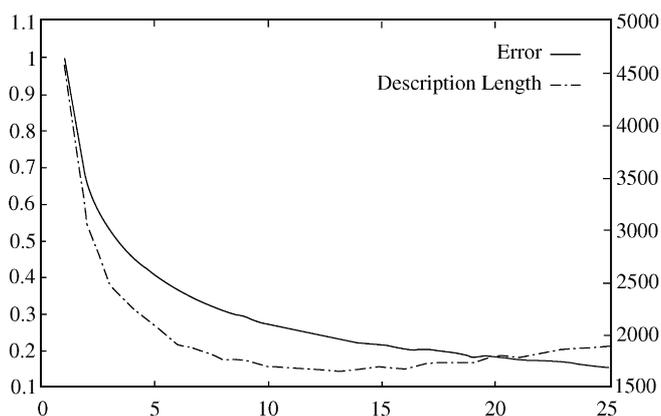


Fig. 2

Representational error and description length versus number of clusters for the 17-variable data set of Australasia

Fig. 3

a Twelve-class clustering result for Australasia using average scaled Euclidean distance and randomly selected colors; b same clustering result but with colors selected to reflect the similarity of the classes

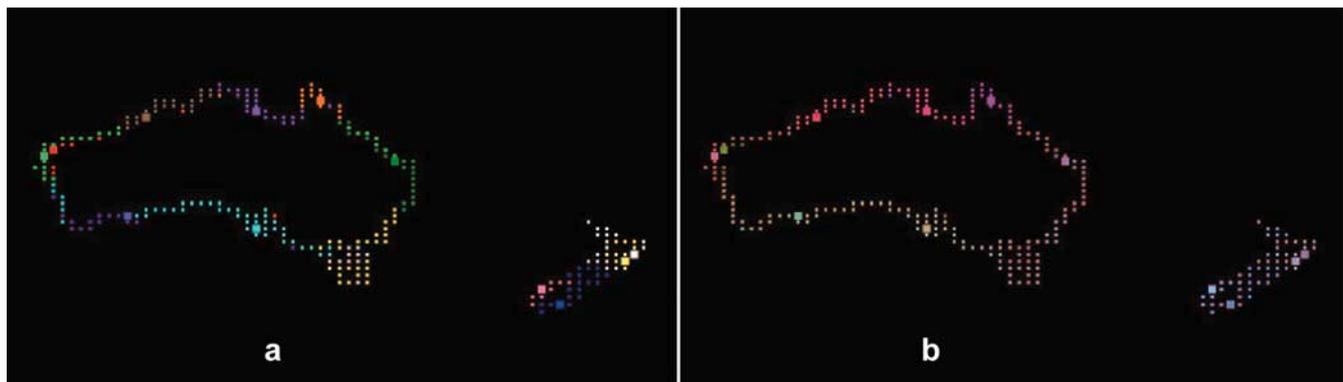
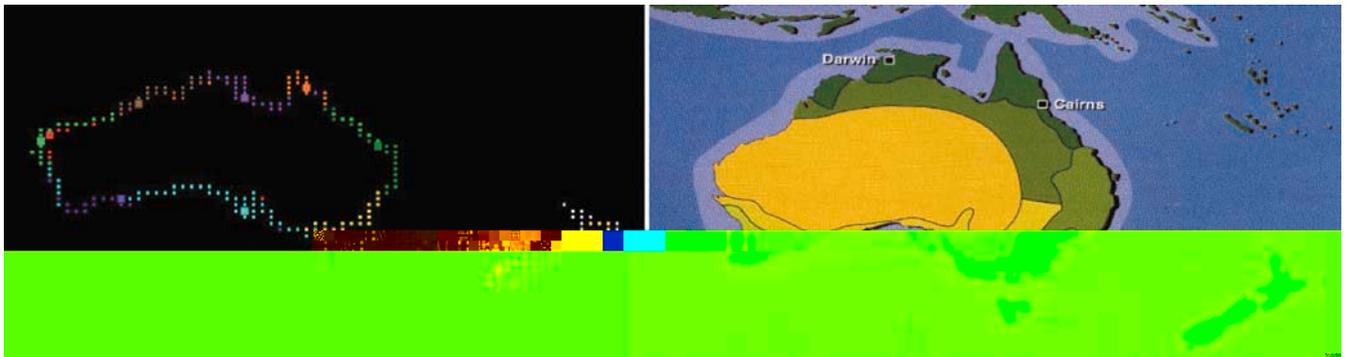


Table 1

Semantic labels, key variables, and comparison with expert typology of Australasia using the average scaled Euclidean distance measure (Fig. 3a)

Color	Class label	Expert class	Key variables
White	Wet temperate I	Wet temperate	Soil moisture/runoff (low)
Red	Grassy bay	No match	CZCS (lots of green)
Lighter green	Dry tropical	Dry tropical/subtropical	Precipitation (low)
Blue	Cool wet temperate	Wet temperate	Temperature (high)
Yellow	Wet temperate II	Wet temperate	Temperature (low)
Purple	Moist temperate	Mediterranean	Precipitation (medium)
Orange	Wet/dry tropical I	Wet/dry tropical (border on wet trop.)	Soil moisture/runoff (high)
Cyan	Dry temperate	Dry temperate	Precipitation (medium)
Pink	Alpine	No match	High waves
Dark green	Wet/dry	Wet/dry	Minimum moisture variables (higher)
Light purple	Tropical/subtropical	Tropical/subtropical	Precipitation/runoff (very low)
Brown	Wet/dry tropical II	Wet/dry tropical	Runoff (very high)
			Elevation variance (very high)
			Precipitation (high)
			Minimum precipitation (medium)
			Temperatures (medium)
			Elevation variance (low)
			Maximum precipitation (high)
			Minimum moisture variables (very low)

**Fig. 4**

a Twelve-class clustering using average scaled Euclidean distance; b expert typology of Australasia (Smith and Crossland 1999)

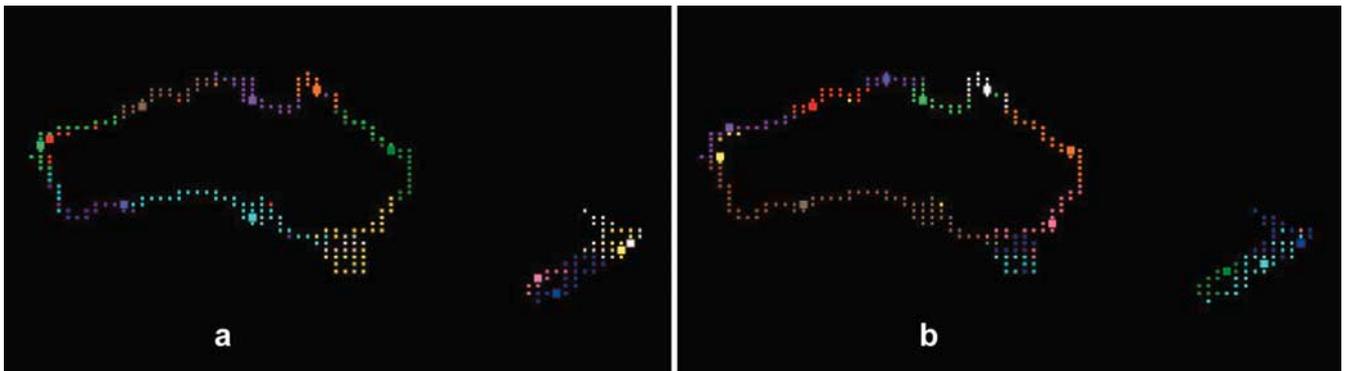
should be more like 24–40. In other words, when you are looking at extremes rather than averages, there are more extremes to be considered.

Figure 6a shows an example of a 24-class clustering result using the ASE distance measure. Figure 6b shows a 24-

class clustering result using the MSD distance measure. In the MSD plot, the places that significantly increased in complexity were the southern coast (one class became four), New Zealand (three classes became six), and the grassy bay cluster (one became three). These three regions

Fig. 5

Twelve-class clustering using a average scaled Euclidean distance and b MSD distance



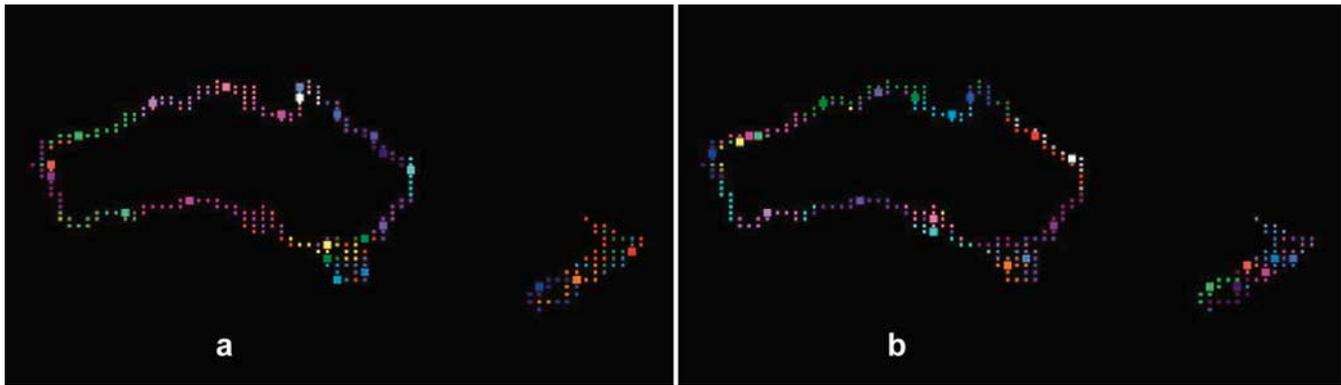


Fig. 6 Clustering result for 24 classes using **a** the ASE distance and **b** the MSD distance

account for 8 of the 12 new classes, and highlight where significant localized changes in geographic variables are taking place. The other subtle difference between the two plots is that the MSD appears to generate more contiguous regions and pick up on more details than the ASE.

As a final comparison, we applied the region-growing and merging technique to the same data set. For the local parameter threshold – how similar a neighboring point must be – we used 1.0 standard deviation. For the global parameter threshold – how similar a point must be to the original seed point – we used 1.22 standard deviations. Points had to be within two degrees of one another to be considered adjacent, and the distance measure was the average scaled Euclidean distance.

The resulting segmentation contained 170 different contiguous regions: three significant regions and 167 one to three point regions. Applying the merge technique to this set of regions, the graph of description length versus number of clusters gives us a guide as to when to stop merging. Figure 7 shows this graph, which bottoms out between 16 and 28 classes. In Fig. 8 we show the 28 class result, which appears to highlight a number of localized phenomena, similar to the MSD distance. Note that the localized phenomena each tend to occupy a different class, however, since the segmentation process requires data points to be contiguous.

Discussion and future directions

The first question we need to answer is whether the typology development process outlined above gives us something useful. From the Australasia example, the answer seems to be that it does produce a reasonable set of classes. The results show broad agreement with the previous expert typology. Furthermore, they highlight localized phenomena that do not show up in the expert version, but nevertheless exist in the data. Note that we obtained these results despite heterogeneous variables with some missing data, indicating that the distance measures we used are appropriate for the task.

The primary real benefit of the data-driven methodology is that it gives us a quantitative, consistent, and objective way

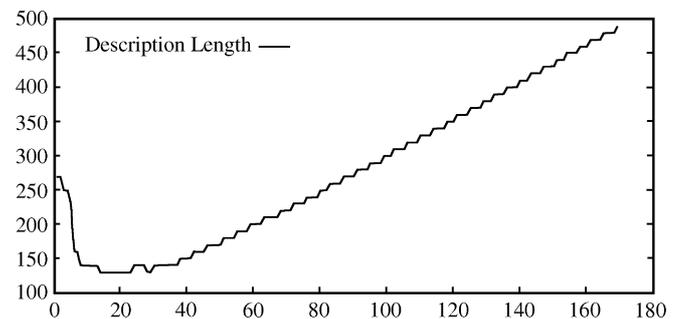


Fig. 7 Description length versus number of clusters during merging (merge process goes from *right to left* from 170 down to 1)

to compare classes across both space and time. Thus, this approach can be used not only to compare coastlines across the world within a temporally fixed data set, but also to compare how coastlines change based on actual or predicted climate change. A second real benefit of using the bottom-up expert-guided approach is in the time-saving aspect of the process. A group of experts that included the authors was able to develop the complete Australasia typology with the span of an hour. This was largely because of the tools we developed for automating the process – MDL analysis, clustering, and visualization. A third real benefit of this approach is its utility for integration of data and communication about results across disciplinary boundaries. Human dimension variables and physical variables that are both environmental forcing functions can be effectively combined, even though their mechanisms of operation are very different. The visual presentation of results, and the specification of relationships among the clusters, make it a tool that can be used for calibration by a variety of experts – and therefore a tool for integrating their individual judgments.

Looking at the distance measures in detail, it appears that the MSD distance measure gives us slightly different classes than the ASE distance measure when applied to the same problem. The MSD appears to more readily identify local phenomena, while overall similarity drives the ASE. It will be up to the experts driving the process to decide which approach they want to emphasize.

In comparing both the average scaled Euclidean and MSD clustering results with the local segmentation, we have to ask whether the latter gives us any more or different

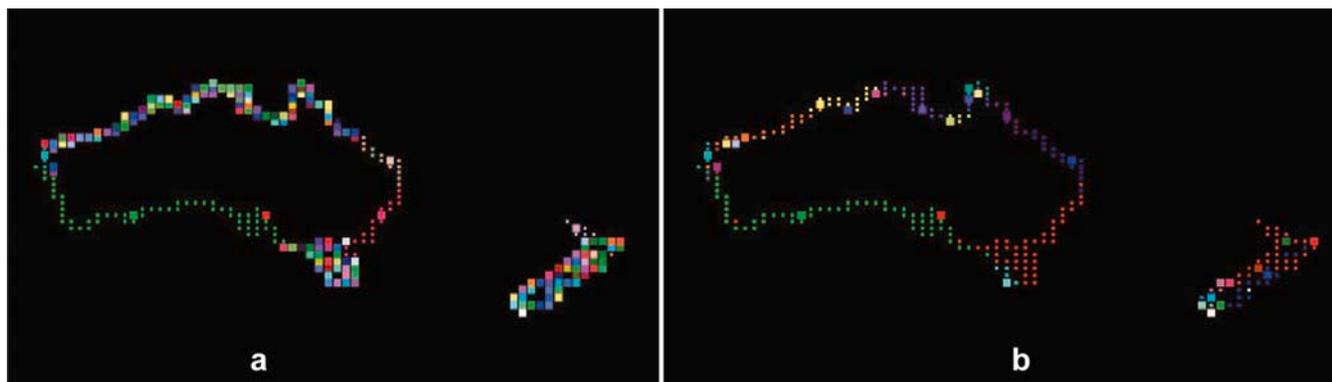


Fig. 8
a 170-region result of segmentation process; b 28-region result of segmentation followed by merging

information. When compared with the 24-class ASE or MSD results, the primary difference is the existence of longer stretches of similar coastline in the segmentation/merging result. This occurred because the long southern coastline was a single group to start with, which meant it never split during the process.

Whether the segmentation/merging offers any new insights is unclear from the single example. We need to undertake more experimentation before deciding whether to pursue it any further as a potential typology development tool.

Conclusion and summary

This paper presents a set of methods that permit clustering, classification, and comparison of environments at regional and global scales. Clustering of high-dimensionality data sets can be based on scaled Euclidean distances in ways that permit the use of data sets that are incomplete, not normally distributed, or otherwise unsuitable for more traditional statistical analysis. Two different distance criteria – the average scaled Euclidean distance and the maximum scaled distance – provide alternative ways to explore the nature of environmental similarities and differences. The paper also presents ancillary techniques that expand the applicability and ease of use of these methods. Use of a minimum description length algorithm makes it possible to estimate objectively the optimum number of clusters for a given data set, and a novel color-similarity approach permits visualization of the similarities of spatially distributed clusters. Methods of cluster merging provide a means of establishing hierarchical relationships among the clusters, and of aggregating smaller clusters into larger groups. We have demonstrated these techniques on a 17-variable coastal data set for Australia and neighboring regions. The results are highly consistent with an independent expert-judgement coastal typology, and the differences and similarities between the various approaches to cluster definition are intuitively understandable in terms of the variables and techniques used.

The methods provide a novel and potentially powerful set of tools for classifying environments and for upscaling environmental functions, either on a purely data-driven basis or as a component in an expert-driven examination of selected and weighted variables. Our initial application of this methodology will be to regionalization and globalization of coastal CNP budgets as part of the LOICZ projects. However, the techniques are further applicable to issues of global and regional change by comparison of clusters based on present data, historical records, and future scenarios or model predictions. This will permit space-for-time trade-offs in analysis by identifying a region's potential trajectory between clusters over time.

Acknowledgments The authors would like to thank the LOICZ project for its support of the typology project and the work described herein. We would also like to thank the American Association for the Advancement of Science (AAAS) for the Earth Systems Science Conference in 1997 in South Dakota which brought together scientists from a variety of disciplines – including the authors – and launched this approach to coastline typology development.

References

- ANDERBERG MR (1973) Cluster analysis for applications. Academic Press, New York
- ANKERST M, BERCHTOLD S, KEIM D (1998) 'Similarity clustering of dimensions for an enhanced visualization of multidimensional data.' In: Proc Conf Information Visualization '98, Oct, pp 52–60
- BACKER E (1995) Computer-assisted reasoning in cluster analysis. Prentice Hall, Englewood Cliffs
- DYER KR, CHRISTIE MC, WRIGHT EW (2000) The classification of intertidal mudflats. *Cont Shelf Res* 20(10–11):1039–1060
- GORDON DC JR, BOUDREAU PR, MANN KH, ONG J-E, SILVERT WL, SMITH SV, WATTAYAKORN G, WULFF F, YANAGI Y (1996) LOICZ biogeochemical modelling guidelines, 2nd edn. LOICZ Reports and Studies no 5. LOICZ, Texel, The Netherlands, pp vi, 96
- HARFF J, DAVIS JC (1990) Regionalization in geology by multivariate classification. *Math Geol* 22(5):573–588
- HARTIGAN JA, WONG MA (1979) A k-means clustering algorithm. *Appl Stat* 28:100–108
- HUTTENLOCHER DP, KLANDERMAN GA, RUCKLIDGE WJ (1993) Comparing images using the Hausdorff distance. *Pattern Anal Mach Intelligence* 15(9):850–863.
- INTERGOVERNMENTAL PANEL ON CLIMATE CHANGE (1999) IPCC data distribution centre for climate change and related scenarios for impacts assessment. CD-ROM, Version 1.0, April. IPCC-Data

- Distribution Centre, Climatic Research Unit, University of East Anglia, Norwich
- LANKFORD R (1977) Coastal lagoons of Mexico: their origin and classification. In: Wiley M (ed) *Estuarine processes*. Academic Press, New York, pp 182–215
- MACQUEEN J (1965) On convergence of k-means and partitions with minimum average variance. *Ann Math Stat* 36:1084
- MAKTAV D (1985) The study of the natural geographic differences in the coastal areas of water covered parts of Marmara region in Turkey with the help of Landsat-4 MSS data using an unsupervised classification algorithm with Euclidean distance. In: *Proc 11th Int Symp on Machine Processing of Remotely Sensed Data*, West Lafayette, Indiana, USA, pp 122–127
- PERNETTA JC, MILLIMAN JD (eds) (1995) *Land–Ocean Interactions in the Coastal Zone: Implementation Plan*. Rep 33. International Geosphere–Biosphere Programme, Stockholm, 215 pp
- PIRAMUTHU S (1999) The Hausdorff distance measure for feature selection in learning applications. In: *Proc 32nd Annu Int Conf on Systems Sciences*, Hawaii, IEEE Computing Society, New York
- PRAKASH HNS, KUMAR SR, NAGABHUSHAN P, GOWDA KC (1996) Modified divisive clustering useful for quantitative analysis of remotely sensed data. In: *Proc Int Symp on Geoscience and Remote Sensing, IGARSS '96*, IEEE Computing Society, New York, vol 3, pp 1858–1860
- RABINER L, JUANG B-H (1993) *Fundamentals of speech recognition*. Prentice Hall, Englewood Cliffs
- RISSANEN J (1989) *Stochastic complexity in statistical inquiry*. World Scientific, Singapore
- SMITH SV, CROSSLAND C J (1999) Australasian estuarine systems: carbon, nitrogen and phosphorus fluxes. *LOICZ Reports and Studies no 12*, LOICZ, Texel, The Netherlands, pp ii, 182
- SMITH SV, IBARRA-OBANDO S, BOUDREAU PR, CAMACHO-IBAR VF (1997) Comparison of carbon, nitrogen and phosphorous fluxes in Mexican coastal lagoons. *LOICZ Core Project of IGBP*, Texel, The Netherlands