ELSEVIER

A general framework for determining cutoff values to select pollen analogs with dissimilarity metrics in the modern analog technique

Eugene R. Wahl [*]

Quaternary Paleoecology and Conservation Biology Graduate Programs, University of Minnesota, St. Paul, MN, USA

Abstract

The modern analog technique (MAT) is a quantitative calibration tool for using modern pollen assemblages to interpret fossil pollen assemblages for vegetation and climate reconstruction. When the MAT is applied using multivariate distance metrics, a cutoff value for the metric is often used to determine the presence/absence of analogs in a modern pollen reference set. Two kinds of error arise when a cutoff value is used: (1) *false positive error*, which occurs when analogy is falsely determined to exist between the vegetation (or other parameter) of a sample of interest and that of a sample in the reference set; and (2) *false negative error*, which occurs when analogy is falsely determined not to exist. The existing literature focuses primarily on examining cutoff thresholds from the perspective of reducing false positive error, with relatively little attention paid to false negative error and to the inherent trade-off between the two errors. This paper sets forth a general analytical framework for determining cutoff thresholds that minimize the joint occurrence of the two errors, and employs the squared chord distance metric with a newly developed reference set of modern pollen surface samples from southern California, USA, as a demonstration case. It also examines the nature of the tradeoffs that occur if an analyst decides to accept increased risk (beyond the joint minimum) of one of the kinds of error for additional reduction of the other. An asymmetric tradeoff in these risks above and below the joint error minimizing cutoff(s) is described (a more rapid proportionate increase of false negatives at cutoffs below the joint minimum in relation to the proportionate increase of false positives at cutoffs above it), which is controlled by the relative variances of the distributions of like- and non-like-vegetation sample comparisons in terms of the distance metric. This asymmetry is found to be general among sample sets reported using the squared chord distance, but is not general across other distance metrics.
© 2003 Elsevier B.V. All rights reserved.

Keywords: modern analog technique; squared chord distance; cutoff values; false positive error; false negative error; Type I error; Type II error; paleoecology; pollen analysis; ROC analysis; southern California

1. Introduction

One of the most widely-used methods in the modern analog technique (MAT) of reconstructing paleovegetation and paleoclimate from fossil

* Present address: National Center for Atmospheric Research, Environmental and Societal Impacts Group, Boulder, CO, USA. Fax: +1-303-497-8125.

E-mail address: wahl@ucar.edu (E.R. Wahl).

pollen assemblages employs a multivariate distance (or dissimilarity) metric as a measure of closeness for quantitatively determining analogy between modern and the fossil pollen samples. In particular, the squared chord distance metric (SCD) has been used for a variety of pollen-based reconstructions of forest and climate dynamics in North America (e.g. Overpeck et al., 1985; Anderson et al., 1989; Bartlein and Whitlock, 1993; Davis, 1995; Calcote, 1998; Davis et al., 1998; Davis, 1999; Davis et al., 2000). As a 'signal-to-noise' enhancing measure (Prentice, 1980; Overpeck et al., 1985), SCD dampens the effect of dominant pollen taxa, heightening somewhat the importance of less abundant taxa; and has shown optimizing characteristics in relationship to other dissimilarity measures when used with pollen assemblages from a wide variety of vegetation types (Overpeck et al., 1985; Gavin et al., 2003). Other distance metrics, such as the Canberra metric (Prentice, 1980; Overpeck et al., 1985; cf. Juggins, 2003), are now being explored for use in situations in which the information content of the pollen assemblages is concentrated in poorly represented taxa (e.g. Oswald et al., 2003).

A key issue in using any distance metric to select analogs, SCD included, is the cutoff value chosen, i.e. the quantitative level of dissimilarity below which two samples are considered reasonable matches. Overpeck et al. (1985) and Bartlein and Whitlock (1993) examined this issue for vegetation formations and forest types in the semitropical, temperate, and boreal forests of eastern North America, and Calcote (1998) examined it for specific stand types in the forests of northern Wisconsin and the Upper Peninsula of Michigan. In western North America, Anderson et al. (1989) examined it for grasslands, parklands, boreal forests, and tundra in interior Alaska and northwestern North America, and Davis (1995) examined it for a variety of vegetation types in significant portions of the western and southwestern USA. The focus of these examinations has been mainly on reducing the chance of falsely identifying two pollen spectra as being from similar vegetation when they are not (*false positive error* – analogous to Type I error in statistical hypothesis testing). The SCD cutoffs determined by Overpeck et al. (1985)

and Calcote (1998) are 0.15, 0.12, and 0.05 for vegetation formations, forest types, and forest stands, respectively. The cutoff used by Bartlein and Whitlock (1993) is 0.205. The cutoff determined by Anderson et al. (1989) for 'good analogs' or 'strong analogs' is 0.095, and the 'conservative' cutoff for separating 'analog' from 'no analog' situations is 0.185. The cutoff determined by Davis (1995) is 0.15. The range of values possible for the SCD statistic is 0 to 2 when the pollen data are characterized in terms of proportions.

This paper presents a new, general analytical framework for evaluating the issue of appropriate cutoff levels when using multivariate distance metrics for analog selection. It employs the SCD with a newly developed reference set of 41 modern pollen surface samples from southern California, USA (Wahl, 2003a) as a demonstration case. This reevaluation is motivated by two considerations.

First, and most generally, motivation for reexamining cutoff value determination in the MAT arises from a gap in the analytical structure of previous work. Although the discussions by Overpeck et al. (1985), Anderson et al. (1989), Davis (1995), and Calcote (1998) rightly focus on reducing the chance of generating false positives when choosing analogs, they do not rigorously consider the concomitant risk of failing to correctly identify pollen spectra from similar vegetation (*false negative error* – analogous to Type II error in statistical hypothesis testing). These two types of error are inversely related; reduction of false positives by lowering cutoff values carries a price in terms of generating false negatives, and vice versa. The nature of the tradeoff between these co-occurring types of error needs to be addressed if maximal information is to be recovered in analog selection, and to guide decisions to accept elevated levels of one kind of error in order to reduce the other. The analysis of cutoff values in this study achieves these goals; it uses the comparative relationships between all possible pairings of the surface samples in the new reference set to determine an analog selection threshold that minimizes the joint occurrence of both errors. It also explicitly examines the nature of the tradeoffs that occur when it is considered appropriate to
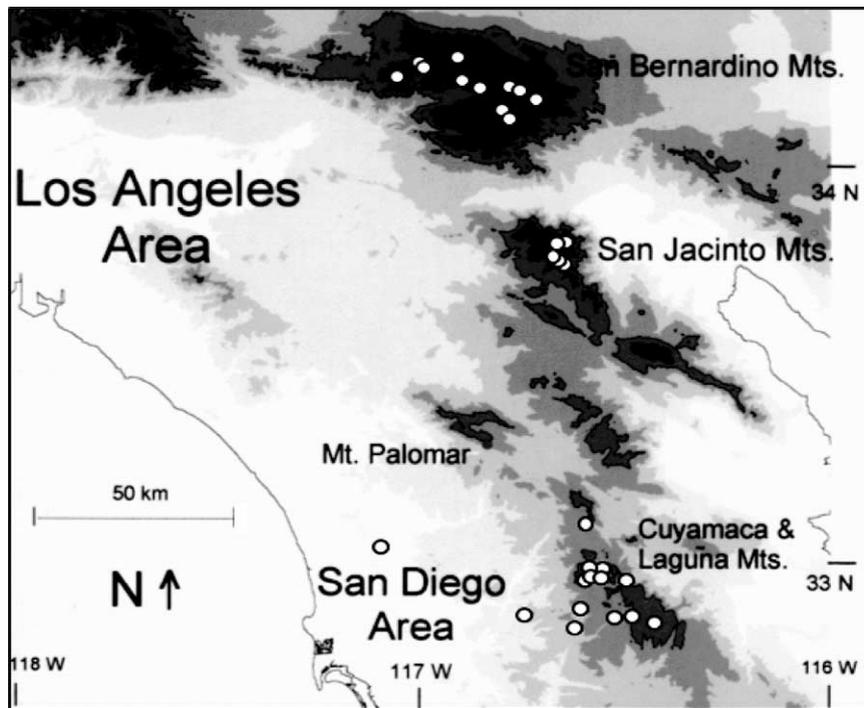
Fig. 1. Digital elevation map of southwestern California, USA. The 41 surface sample sites used in this study are shown as dots. In some cases multiple surface sample sites are indicated by one dot at the scale shown. Contour lines are shown at 0-m and 1500-m elevations. Darkness of color is associated with increasing elevation; dark gray–black shaded polygons approximately represent regional montane areas (> 1500 m; max. elevation 3474 m). (Source data for DEM base map from US Geological Survey).

accept increased risk of one of the types of error (beyond the joint minimum) for additional reduction of the other type of error.[1] Numerically, the approach presented here is identical to 'Receiver Operating Characteristic' (ROC) analysis, first developed for examination of signal detection characteristics of radar systems and widely employed in a number of other fields: e.g., in the evaluation of medical and psychological laboratory tests, and

the evaluation of weather forecast success (Henderson, 1993; Zweig and Campbell, 1993; cf. Green and Swets, 1988, for a canonical exposition of ROC methods; cf. Mason and Graham, 2002, for a review of ROC development). The examination of data and the presentation of results in this paper are tailored to considerations specific to the use of the MAT with multi-taxon microbiological data sets. The approach presented here, along with analyses developed by D. Gavin and W. Oswald that emphasize other issues in the use of the MAT, represent the first application of ROC methods for pollen-based paleoenvironmental reconstruction (cf. Gavin et al., 2003; Oswald et al., 2003).

Second, the forests of cismontane southern California (west of the eastern, desert crest of the coastal mountains) have unique vegetation and species-specific physiological characteristics that distinguish them from other montane forests in

---

[1] Precursors to the analysis presented here include: Anderson et al. (1989), who provide quantitative measurement of the relative risks of generating false positives and false negatives at three different cutoff levels; Davis (1995), who provides proportionate relationships for the occurrences of same-type and different-type samples over the range of SCDs between 0 and 1; and Calcote (1998), who examines the issue of creating false negatives by setting a very low cutoff to strongly reduce the chance of getting false positives. However, none of these examinations develops an analytical framework in which the joint minimization of the two kinds of error can be evaluated.

Table 1
Surface sample site information

Sample	MAIN VEGETATION CATEGORY		Latitude	min	sec	Longitude	min	sec	Elevation (m)
1	Lodgepole/Jeffrey/Fir Forest		34	12	10	116	46	55	2774
2	Western Juniper/Lodgepole Forest/Woodland		34	12	0	116	45	53	2746
3	Fir/Lodgepole Forest		34	12	4	116	45	57	2707
4	Fir/Limber/Juniper Woodland		34	12	30	116	51	19	2597
5	Fir/Jeffrey/Lodgepole Forest		34	7	21	116	46	51	2499
6	Jeffrey/Fir Forest		34	7	29	116	46	56	2481
7	Pine/Fir/Oak Forest		34	14	6	117	3	7	2243
8	Pine/Fir Forest		34	15	58	117	0	21	2228
9	Pine/Fir Forest		34	16	5	117	0	22	2228
10	Pine/Western Juniper/Fir Forest	UPPER- & MID-MONTANE CONIFER-DOMINATED FORESTS	34	10	20	116	43	4	2438
11	Lodgepole/Limber Forest		33	48	2	116	40	21	2926
12	Pine/Fir Forest		34	13	19	116	53	58	2402
13	Pine/Fir Forest		34	16	53	116	54	51	2286
14	Lodgepole/Fir Forest		33	48	24	116	39	12	2682
15	Fir/Pine Forest--Aspen Grove (> dense with Aspen)		34	8	54	116	47	59	2170
16	Fir/Pine Forest--Aspen Grove (< dense with Aspen)		34	8	54	116	47	59	2170
17	Pine/Fir/Oak/Cedar Forest		32	56	56	116	36	16	1890
18	Pine/Alder/Fir Forest		33	46	9	116	39	44	2405
19	Pine/Fir Forest		33	46	9	116	39	44	2405
20	Pine/Fir Forest		33	46	9	116	39	44	2402
21	Manzanita Chaparral		32	56	38	116	29	28	1692
22	Pine/Oak Forest		32	50	15	116	25	24	1829
23	Oak/Conifer Forest		34	8	5	116	58	56	1682
24	Cedar/Pine/Oak/Fir Forest		32	57	15	116	35	16	1573
25	Pine/Oak Forest		32	51	9	116	28	56	1554
26	Oak/Cedar/Pine Forest	LOWER-MONTANE CONIFER-OAK FORESTS, MEADOWS, CHAPARRAL	32	58	2	116	35	8	1451
27	Oak/Cedar Stand--in Pine/Oak/Cedar/Fir Forest		32	58	10	116	35	1	1448
28	Pine/Oak Open Forest Clearing		32	58	15	116	34	26	1439
29	Oak/Pine Forest		33	5	22	116	35	40	1228
30	Cedar/Oak/Fir Forest--Burn Site		33	20	13	116	54	27	1573
31	Meadow--in Pine/Oak/Cedar/Fir Forest		32	58	8	116	35	4	1439
32	Meadow--in Pine/Oak/Cedar/Fir Forest		32	58	6	116	35	5	1434
33	Mixed Chaparral--Mt. Mahogany dominated		32	56	58	116	33	35	1384
34	Sagebrush Steppe (> open phase)		32	51	2	116	31	17	1132
35	Sagebrush Steppe (> closed phase)		32	51	5	116	31	22	1125
36	Oak Woodland (> closed phase)	STEPPE, MIXED CHAPARRAL, OAK WOODLAND	33	17	28	116	50	2	1414
37	Oak Woodland (> open phase)		32	49	9	116	37	7	1073
38	Mixed Chaparral		32	52	9	116	36	29	1091
39	Mixed Chaparral		32	49	21	116	37	17	1061
40	Chamise Chaparral	CHAMISE CHAPARRAL	32	51	21	116	44	30	817
41	Coastal Sage Scrub	SAGE SCRUB	33	1	49	117	5	34	244

Vegetation characterizations follow the usage of Munz (1974), Thorne (1988), and Hickman (1996). 'Jeffrey', 'Limber', and 'Lodgepole' in the conifer-dominated forest group are common names of important regional pine species.
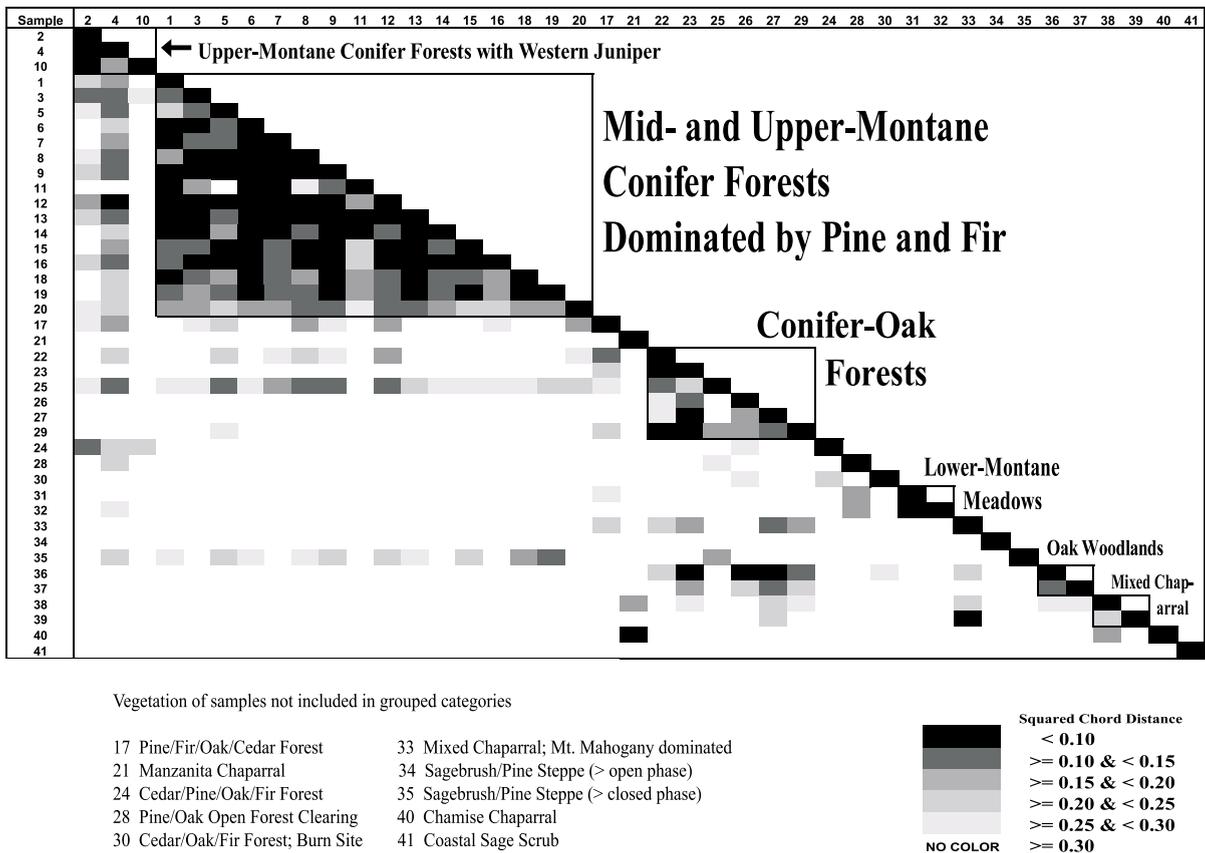
Fig. 2. Pair-wise comparison of 41 southern California surface samples in terms of SCD dissimilarity statistic. Boxes with vegetation names outline groups of samples with similar vegetation (Wahl, 2003a). Sample numbers are from Table 1.

the western USA. At the level of overall vegetation, these forests are often characterized as having a depauperate Sierran flora (Munz, 1974; Thorne, 1988), because they lack important tree taxa – particularly the entire hemlock (*Tsuga*) genus and important species of fir (*Abies*) and pine (*Pinus*) – along with numerous herbaceous plants that characterize the forests of the Sierra Nevada Mountains of northern California. At the level of individual species, nearly all of the regional montane species of pine, fir, cedar (*Calocedrus*), and arboreal oak (*Quercus*) are at their modern range limits in southern California and adjacent Baja California del Norte, Mexico (Little, 1971; Griffin and Critchfield, 1976; Burns and Honkala, 1990; Roberts, 1995). Trees of these taxa in the southern California region are often shorter and of less girth at maturity than is typical nearer the center

of their ranges (Zedler and Krofta, 1995), which suggests that they are either exhibiting phenotypic responses to abiotic stress or showing genetic adaptation in situ to climate differentiation (Davis and Shaw, 2001). The oak woodland, chaparral, and coastal sage scrub vegetation just west and below the montane zone are also unique to coastal California and northern Baja California del Norte, Mexico (Griffin and Critchfield, 1976; Hickman, 1996). These characteristics of the regional vegetation call for reevaluation of cutoff values that were developed for vegetation in other regions of North America, including other parts of California and the West (Davis, 1995). This evaluation follows the result of Overpeck et al. (1985) that there is need to reexamine previously determined cutoff values depending on the scale of application and the pollen registration charac-
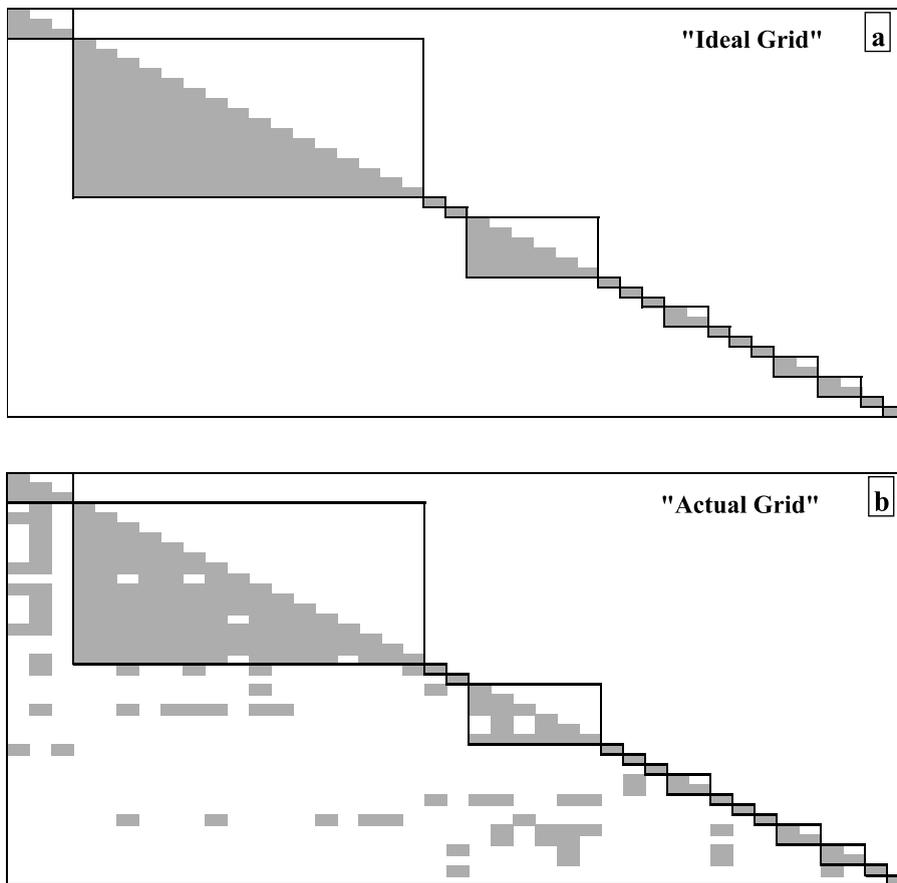
Fig. 3. Discrimination of vegetation groups in surface sample comparisons. (a) 'Ideal Grid' matrix representing perfect discrimination of surface samples: within-group and same-sample comparisons from Fig. 2 are assigned the value 1 (shaded); outside-of-group and non-same-sample comparisons from Fig. 2 are assigned the value 0 (not shaded). (b) 'Actual Grid' matrix showing SCDs between pairs of surface samples in Fig. 2 in relation to a given cutoff value: SCDs < cutoff are assigned the value 1 (shaded); SCDs ≥ cutoff are assigned the value 0 (not shaded). Example cutoff value is 0.225.

teristics of the surface and fossil samples to be used in the MAT.

2. Methods

2.1. Surface sample set

The locations of the 41 sample sites in the study region are shown in Fig. 1 and are listed with vegetation and elevation information in Table 1. At each site 5–10 surface-soil sub-samples were collected and mixed together, according to the collection protocol of Adam and Mehringer (1975). The pollen registration of the samples is

generally able to distinguish vegetation differences (e.g. an oak–cedar stand in a matrix of pine–oak–cedar–fir forest) at distances on the order of 50–150 m, and there are very low levels of background pollen of abundant producers (e.g. background *Quercus* pollen at conifer-dominated forest sites or background *Pinus* pollen at non-forested sites) (Wahl, 2003a). Vegetation characterizations are based on surveys taken around the sample sites; the vegetation categories conform to those identified in the botanical and vegetation literature of the region (Munz, 1974; Thorne, 1988). The collection and laboratory methods used and the vegetation/pollen relationships of the samples are described in detail in Wahl
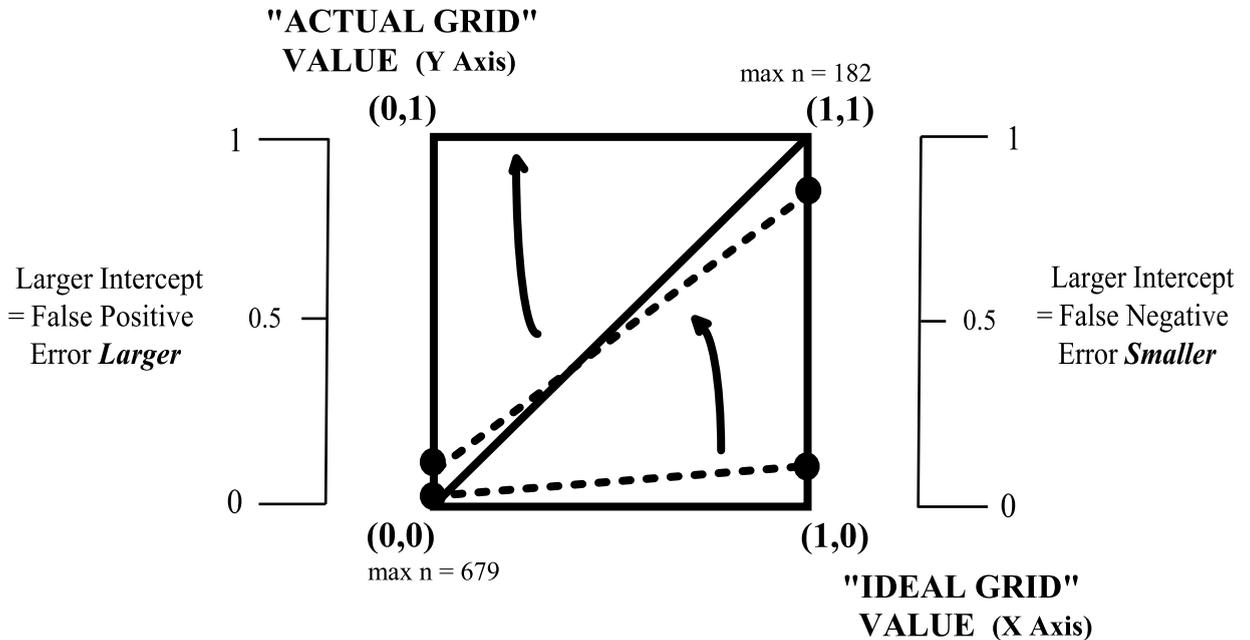
Fig. 4. Geometric representation of joint minimization of false positive and false negative errors. This graph shows three levels of generality. (a) The bold unit square with (1,1), (0,0), (0,1), and (1,0) vertices shows the four possible states that *a comparison of two corresponding grid cells* from the 'Ideal Grid' and the 'Actual Grid' (Fig. 3) can have for a given SCD cutoff. Matches between the 'Ideal Grid' and 'Actual Grid' are at (1,1) for true positives and (0,0) for true negatives. Discrepancies are at (0,1) for false positives and (1,0) for false negatives. (b) The cumulative numbers of the (1,1), (0,0), (0,1), and (1,0) values generated for *all 861 intercomparisons among the 41 surface samples at a given cutoff value* can be pictured as summations of the counts for each state at their appropriate vertices. The maximum number of (1,1) and (0,0) values possible is indicated next to these vertices. The ratio $\Sigma(1,1)/182$ (called the *true positive fraction*) gives the proportion of surface samples that should be chosen and *are* [$\Sigma(1,1)$] to the total that should be chosen [182]; *false negative errors* are *smaller* as this ratio is closer to *one*. The ratio $\Sigma(0,0)/679$ (called the *true negative fraction*) gives the proportion of surface samples that should not be chosen and *are not* [$\Sigma(0,0)$] to the total that should not be chosen [679]; *false positive errors* are *smaller* as this ratio is closer to one, or as [$(1-\Sigma(0,0)/679)$] is closer to zero. The values $\Sigma(1,1)/182$ and [$(1-\Sigma(0,0)/679)$] are represented geometrically by pairs of right and left vertical intercepts, respectively – shown as dots on the sides of the unit square, and scaled by the separate scale bars flanking the sides of the unit square. The *difference* between these values (in the order, right-to-left) determines the *slope* of the line segment connecting the two intercepts. (c) The situation of *varying the cutoff value over its entire range of 0–2* can be represented by calculating the slopes determined by the changing left- and right-side intercept values as the cutoff value in varied (explained in detail in text). An ideal situation is represented by the diagonal in the unit square with a slope of one; in this situation all sample intercomparisons for a given cutoff are either true positives [$\Sigma(1,1)=182$] or true negatives [$\Sigma(0,0)=679$]. The cutoff value(s) that jointly minimize false positive and false negative error are those that generate the *maximum slope* value for a given data set, which in actual situations will be <1.

(2003a). The primary pollen count data are archived with the North American Pollen Data Base and are accessible on the World Wide Web at http://www.ngdc.noaa.gov/paleo/pollen.html.

2.2. Squared chord distance formula and included pollen types

SCD is defined by the formula:

$$SCD_{jk} = \Sigma_i \ (p_{ij}^{0.5} - p_{ik}^{0.5})^2$$

where p is the pollen proportion (expressed in the range 0–1) of a type included in the comparison, $i = 1...n$ are the included pollen types, and j and k represent the two samples being compared (Overpeck et al., 1985).

The included pollen types and groupings of pollen categories used in the SCD calculations are

described in the Appendix. The proportions used in the calculations were determined in relation to the sum of included pollen (Calcote, 1998). The SCD relationships between all possible pairings of the surface samples, grouped by similarity of vegetation – with a few samples removed from the groups and treated individually because of unique pollen representation characteristics (Wahl, 2003a; cf. Calcote, 1998) – are shown in Fig. 2.

2.3. Analytical determination of cutoff values in the modern analog technique that jointly minimize false positive and false negative errors

In order to determine a cutoff value for the analog selection threshold that jointly minimizes both false positive and false negative errors, the following method was used. Two grids of 1's (ones) and 0's (zeros) were established, conforming to the structure of the SCD comparison grid in Fig. 2 (Fig. 3). The 'Ideal Grid' (IG) in Fig. 3a represents the results of an ideal cutoff that would sort all the like-vegetation samples together (represented by the 1's) and would distinguish these from all the non-like-vegetation samples (represented by the 0's). Such perfect discrimination between vegetation types does not occur in the real pollen samples, which overlap across the groupings to some degree to form the 'Actual Grid' (AG) (Figs. 2 and 3b; cf. Calcote, 1998). At the particular cutoff value of 0.225, which is the best for this data set in terms of jointly minimizing false positive and negative errors (reported below, Fig. 5), most of the 'within-group' comparisons receive 1's and most of the 'outside-of-group' comparisons receive 0's, representing good, although not perfect, discrimination (Fig. 3b). Most of the 'outside-of-group' comparisons receiving 1's in Fig. 3b represent comparisons between samples from vegetation types that have significant similarity in plant composition, e.g. the montane conifer forests dominated by pine and fir compared with conifer forests that have significant cover of western juniper (*Juniperus occidentalis*, var. *australis*; common name *sensu* Munz, 1974), and the mixed conifer–oak forests compared with the conifer-dominated forests and the oak woodlands.

Quantitative and geometric comparison of the patterns of 1's and 0's in the two grids allows rigorous determination of how well a given cutoff value compares to the ideal cutoff (Fig. 4). The two sets of values in the comparison are the 861 [$n(n+1)/2$, where $n = 41$] 1's and 0's from the cells along and below the diagonal in each grid, both sets having the same order of cells. The vertices of the unit square in Fig. 4 represent the four possible combinations of 1's and 0's that each of the 861 pairs of cells can take. Matches between the IG and AG are at (1,1) for true positives (correct identification of like-vegetation sample comparisons) and (0,0) for true negatives (correct identification of non-like-vegetation sample comparisons). Discrepancies are at (0,1) for false positives (failure to correctly identify non-like-vegetation sample comparisons) and (1,0) for false negatives (failure to correctly identify like-vegetation sample comparisons).

Perfect discrimination of 'within-group' and same-sample comparisons from 'outside-of-group' and non-same-sample comparisons at a particular cutoff value is represented geometrically by the solid line that connects (0,0) and (1,1) in the figure. If all the AG values were identical to the IG values, the overall comparison of the AG and IG would be this line, with a *slope* of one. This situation represents absolute minimization of both false positive and false negative errors, and helps demonstrate the analytical definition of these errors, their corresponding true negative and true positive fractions, and the individual and joint minimization of the two kinds of error.

False negative error is measured using the ratio of actual (1,1) values to the maximum number possible, 182 – shown in Fig. 4 by the reduction of the height of the right-side intercept below one. This error equals $1 - \Sigma(1,1)/182$, and is individually minimized when $1 - \Sigma(1,1)/182$ is minimized, or when the corresponding true positive fraction [$\Sigma(1,1)/182$] is maximized.

False positive error is measured using the ratio of actual (0,0) values to the maximum number possible, 679 – shown in Fig. 4 by the increase in the height of the left-side intercept above zero. This error equals $1 - \Sigma(0,0)/679$, and is individually minimized when $1 - \Sigma(0,0)/679$ is mini-

mized, or when the corresponding true negative fraction [$\Sigma(0,0)/679$] is maximized.

Joint minimization of both errors can be achieved using any one of four equivalent criteria: (1) minimizing the sum of false negative and false positive errors; (2) maximizing the sum of the true positive and true negative fractions; (3) maximizing the difference between the true positive fraction and false positive error; and (4) maximizing the difference between the true negative fraction and false negative error. ROC analysis conventionally uses criterion (3), which numerically means maximizing the difference [$\Sigma(1,1)/182-(1-\Sigma(0,0)/679)$] – represented in Fig. 4 as maximization of the slope of the line connecting the left- and right-side intercepts.

At a cutoff value of 0, all the values in the comparison are true negatives (0,0) or false negatives (1,0). No SCD comparisons are selected as matches. This situation represents *maximum* false negative error coupled with *minimum* false positive error. In this case, the line connecting the intercepts lies on the 'X axis', with a slope of 0. As the cutoff value is increased to any arbitrarily small number greater than 0 (e.g. 1×10^{-10}), the (0,0) values are all maintained (all the 'outside-of-group' and non-same-sample comparisons are rejected when they should be), but 41 of the (1,0) values change to true positives (1,1) (the same-sample comparisons along the diagonal are 'picked up' because their SCDs are identically zero). This situation still has minimum false positive error and a small amount less than maximum false negative error, represented by the zero left intercept and small positive right intercept of the lower dashed line in Fig. 4. As the cutoff value is increased, more and more of the 'within-group' comparisons are correctly selected along with the same-sample comparisons (increasing the number of (1,1) values and decreasing the number of (1,0) values), and eventually some of the 'outside-of-group' and non-same-sample comparisons are erroneously selected because of their relatively low SCDs (increasing the number of false positive (0,1) values and decreasing the number of (0,0) values).

Up to a point, increasing the cutoff value generally correctly selects more 'within-group' comparisons than it erroneously selects 'outside-of-group' and non-same-sample comparisons, which makes the false negative error decrease more rapidly than the false positive error increases, indicated graphically by the right intercept rising faster than the left intercept (lower curved arrow rising towards the upper dashed line). Even though the slope of the line is rising in this case, it can never be as large as one because some (0,1) values are being added while the number of (1,1) values is quickly rising. With further increases in the cutoff value this process reverses; there are fewer and fewer 'within-group' comparisons that are additionally selected and more and more of the 'outside-of-group' and non-same-sample comparisons are selected – the number of (0,1) values now increases faster than the number of (1,1) values, indicated graphically by the upper curved arrow rising from the upper dashed line. Eventually a cutoff value is reached below which *all* the SCDs fall, which means that all the compared values are (1,1) (correct selection of like samples) and (0,1) (incorrect selection of unlike samples). At this point the slope is again 0, but the left intercept is 1. Empirically, this geometry means that the slope of the line connecting the two intercepts is ≥ 0, and has a maximum value < 1. When the slope is at its maximum for a particular data set, the best possible combination of low false positive and low false negative errors is achieved.[2]

To examine the relationship between cutoff value and the joint minimization of both errors for the new surface sample data set, 400 cutoffs were input into the analysis described, covering the range of possible SCD values between 0 and 2 in increments of 0.005 (starting at 0.005). The

---

[2] Numerically, the slope value determined by the formula [$\Sigma(1,1)/182-(1-\Sigma(0,0)/679)$] is identical to the slope determined by a simple linear regression of the actual vs. ideal sets of 1's and 0's. In this case, the geometry of Fig. 4 can be interpreted as follows: the paired values of the elements in the regression vectors all lie on the four vertices in Fig. 4; the slope of the regression (dashed) line comes closest to the diagonal between (0,0) and (1,1) when the 'pull' on the regression line from (0,0) and (1,1) values is the greatest, i.e. when there are the greatest combined cases of true negatives and true positives.
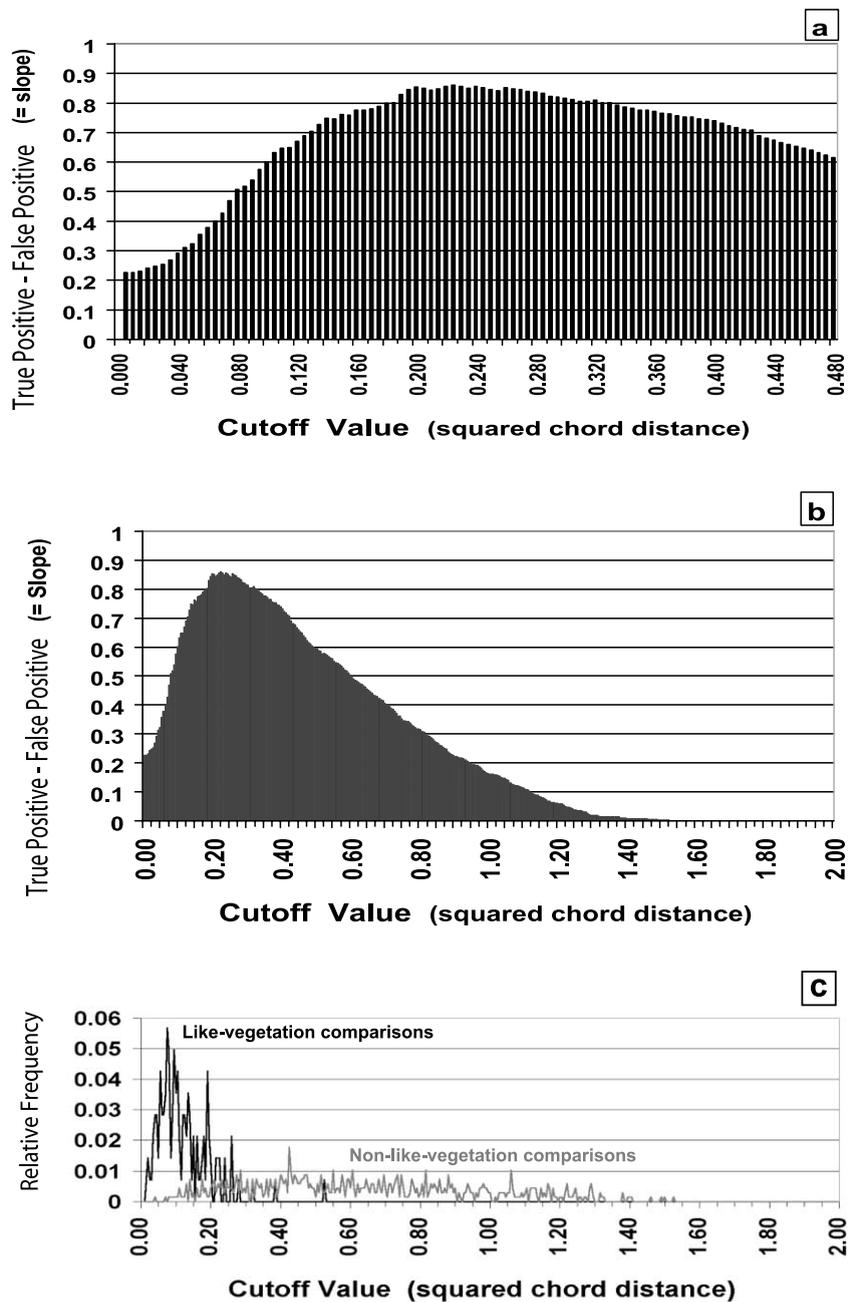
Fig. 5. Relationship between SCD cutoff value and joint occurrence of false positive and false negative errors for surface sample data set. Difference between true positive fraction and false positive error (slope value), as described in text and Fig. 4, for: (a) range of SCD values between 0–0.480, at 0.005 increments; (b) entire range of possible SCD values (0–2), at 0.005 increments. Joint minimization of the two kinds of error (= maximum discrimination between vegetation types in terms of the SCD metric) occurs at highest values. For comparative purposes, panel (c) shows the normalized distributions of like-vegetation and non-like-vegetation sample comparisons in relation to SCD.

slopes of the intercept lines (i.e. the difference $[\Sigma(1,1)/182-(1-\Sigma(0,0)/679)])$ determined by each of these cutoff values were then compared.

3. Results

The slope values of the intercept lines for the examined cutoffs are shown in Fig. 5a,b. The single lowest joint occurrence of false positive and false negative errors (highest slope) occurs at a cutoff of 0.225. Around this value, the cutoffs in the range between ∼0.19 and ∼0.285 all perform similarly well in simultaneously reducing the two errors, which implies that within this range there is latitude in selecting a specific cutoff for use with little diminution of minimization characteristics. For comparative purposes, Fig. 5c shows the normalized distributions of like-vegetation and non-like-vegetation sample comparisons in relation to SCD.

The data from Fig. 5a,b are represented in traditional 'ROC curve' format (true positive fraction vs. false positive error) in Fig. 6. In contrast to Fig. 5a,b, which focuses on the relationship between cutoff value and the joint occurrence of false positive and false negative errors (the specific goal of this paper), the ROC curve format focuses on the overall power of the data to discriminate between like and non-like cases, represented as the area under the curve (AUC). The closer AUC is to one, the greater is the capacity of the data to distinguish like- and non-like comparisons (in this case, in terms of vegetation). The AUC for the new surface sample set is 0.978, which represents excellent overall discrimination between the vegetation types in terms of their pollen representation (Henderson, 1993; Zweig and Campbell, 1993).

In order to test the extent to which including unknown and unidentifiable pollen in the SCD calculations (described in the Appendix) affects the determination of cutoff levels that jointly minimize the two kinds of error, the entire analysis was repeated excluding the grouped unknown, unidentifiable, *Ceanothus*, *Chenopodium/Amaranthus*, and 'Other Fern/Ally' categories. The results of this experiment are close to those reported above. In this case, the lowest joint occurrences
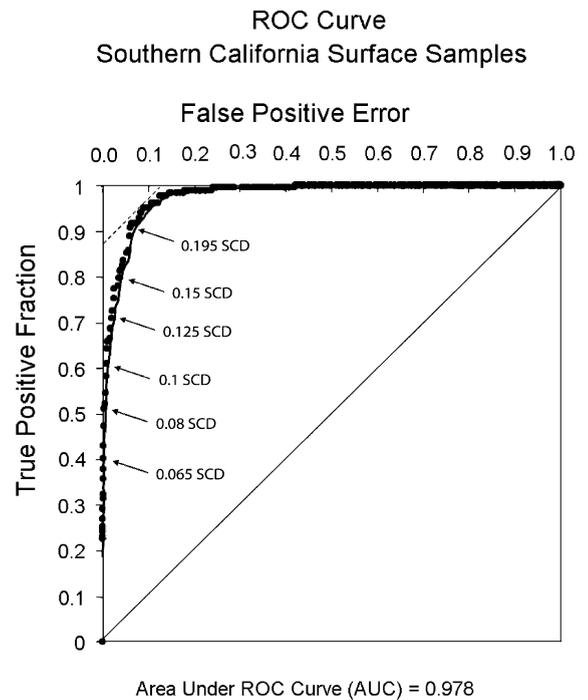
Fig. 6. ROC curve for surface sample data set. Maximal joint reduction of false positive and false negative errors occurs at the point of intercept between the ROC data and a 1:1 line that is nearest to the upper-left corner of the graph, shown by the dashed line. The SCD associated with this combination of true positive fraction and false positive error is, as in Fig. 5, 0.225. The SCDs associated with true positive fractions nearest 0.4, 0.5, 0.6, 0.7, 0.8, and 0.9 are shown for comparative purposes. Geometrically, the ideal case of perfect discrimination between samples from like and non-like vegetation would be represented by a kinked ROC curve that follows the left and upper boundaries of the graph; indicating that a cutoff can be found that identifies *all* true positive comparisons before *any* false positives are generated. The AUC in this case is equal to 1. This situation is equivalent to the diagonal in Fig. 4 that connects the (0,0) and (1,1) vertices. [The solid line tracking the data points is a five-point moving average. The AUC calculation sums the areas of adjoining histogram rectangles determined from the individual data points.]

of false positive and false negative errors are at slightly lower cutoff values (0.205–0.22), and the surrounding range of similarly well-performing cutoffs is slightly wider and slightly lower (∼0.16 to ∼0.275) than it is with the grouped pollen category included. Because the grouped pollen category carries useful information in terms of the distinction between conifer-dominat-

ed forests and other types of vegetation that is the most salient feature of the surface sample set (cf. Appendix), and because including or not including this category has little effect on the cutoff values that best reduce the joint occurrence of false positive and false negative errors, I retained the grouped pollen category in developing the primary results reported here and for comparing these surface samples with fossil pollen assemblages (Wahl, in preparation).
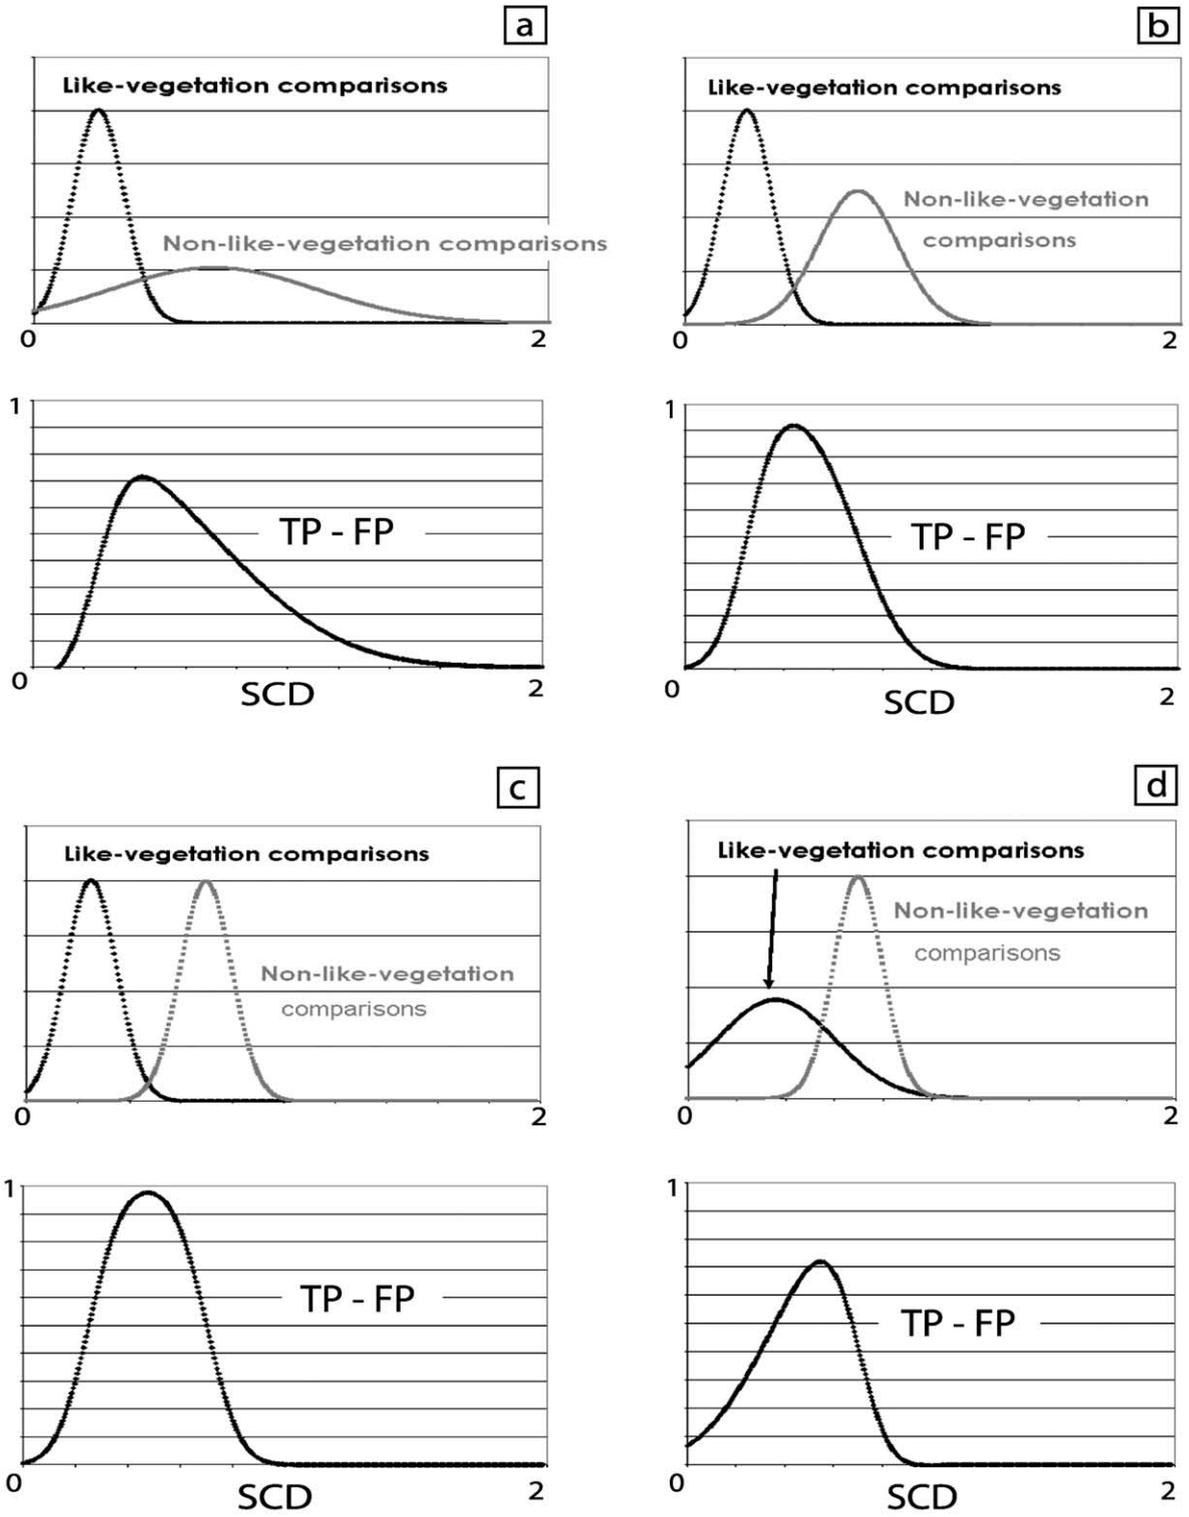
A second check for robustness was done by eliminating selected samples from the data set, and recalculating Fig. 5a,b for each case. The samples were chosen to represent three distinctive cases in the data set: (1) sample 6, a conifer-dominated forest site from the vegetation group with the most samples, which has strong similarities to the other conifer-dominated forests, but few close similarities with samples outside its group; (2) sample 25, a pine–oak forest site that has relatively few close similarities to the other conifer–oak forests, but a number of close similarities with the conifer-dominated forests; and (3) sample 17, a pine–fir–oak–cedar forest site with unusual pollen representation characteristics (cf. Appendix) that cause it to be categorized separately from the other forest samples. In each of these cases the results are virtually identical to those shown in Fig. 5a,b, which indicates that the analysis of joint error minimization is highly robust in relation to the characteristics and inclusion/non-inclusion of individual samples, even from a relatively small data set.

4. Discussion

The joint minimization of false positive and false negative errors establishes a benchmark criterion against which the analog selection performance of different cutoff values can be objectively and quantitatively evaluated. I have avoided characterizing cutoff(s) that achieve this criterion as representing an 'optimal' situation because there may be cases in which it is considered appropriate to deviate from joint error minimization (JEM) for particular analytical purposes. Using the criterion of JEM implies that equal weight is being given to avoiding occurrences of false positives and false negatives. In the absence of a clear rationale for assessing the impact of the errors differentially, an equal weighting scheme represents a rational default criterion (much as it does in establishing Bayesian prior probabilities in the absence of information suggesting specific weighting of 'priors'; cf. Gavin et al., 2003), but it is not the only weighting scheme that could be reasonably applied.

Following the use of ROC methods in the medical literature, the analytical possibilities that arise from differential weighting of the two kinds of error can be explored by characterizing low numbers of false positives as high *specificity*, and low numbers of false negatives as high *sensitivity* (Henderson, 1993; Zweig and Campbell, 1993). In the ROC curve format (Fig. 6) increasing specificity is measured from right-to-left along the horizontal axis and increasing sensitivity from bottom-to-top along the vertical axis (keeping in mind that one *minus* true positive fraction equals false negative error). It is easy to imagine cases in which setting a decision threshold above or below the JEM value(s) could be considered appropriate, for the purpose of boosting either sensitivity (threshold above JEM) or specificity (threshold below JEM). In medicine, for example, if a test

Fig. 7. Idealized distributions of normalized like-vegetation and non-like-vegetation sample comparisons in relation to SCD as in Fig. 5c (top part of each panel), with companion diagram showing the difference between the true positive fraction (TP) and false positive error (FP) for each scenario as in Fig. 5b (bottom part of each panel). (a) Like-vegetation comparisons with relatively narrow variance and non-like-vegetation comparisons with relatively wide variance. (b) Like-vegetation comparisons with relatively narrow variance and non-like-vegetation comparisons with relatively wide variance, but less wide than in (a). (c) Like-vegetation and non-like-vegetation comparisons with same variance. (d) Like-vegetation comparisons with relatively wide variance and non-like-vegetation comparisons with relatively narrow variance. The smooth continuum across (a)–(d) of changing relative steepness on either side of the highest values for the TP–FP curves was confirmed with other examples (not shown) for intermediate and more extreme cases of the relative variances of the like-vegetation and non-like-vegetation comparisons.

is being used to identify cases of a highly contagious disease (with the diseased state being indicated by low test statistic values), it might be considered best to set a threshold for positive identification above the JEM value(s), boosting the test's sensitivity (reducing false negative results) although sacrificing some of its specificity (allowing relatively high amounts of false positives). In the opposite case, decreased false positive identifications might be sought by lowering the threshold for positive identification of a disease below the JEM value(s), boosting the test's specificity although sacrificing some of its sensitivity (allowing relatively high amounts of false negative results).

For the purpose of choosing analogs for fossil pollen samples from a modern pollen data set, using a cutoff above the JEM value(s) could be worth consideration when the modern data set contains a relatively small number of samples – i.e. deciding that the ability to identify a larger number of analogs (increased sensitivity) is worth extra uncertainty in terms of some loss of vegetation precision in analog selection (decreased specificity). The new southern California data set is a potential example of this situation. Using a cutoff below the JEM value(s) could be worth consideration when employing a very sample-rich modern data set – i.e. deciding that some likely analogs can be sacrificed (decreased sensitivity) for the purpose of reducing uncertainty in the vegetation precision of the analogs selected (increased specificity). The modern pollen data available in eastern North America are a potential example of this latter situation (cf. North American Pollen Data Base). As a quantitative illustration of the latter situation, the true positive fractions associated with six cutoff values below the JEM value are highlighted in Fig. 6, which illustrates how much false negative error (one *minus* true positive fraction) is associated with increasingly strict cutoff values in the case of the southern California data set.

The format for presentation of results in Fig. 5 provides additional information on the nature of the sensitivity/specificity tradeoff that is relevant in examining whether to use a cutoff above or below the JEM value(s). The *relative steepness of*

the envelope of the histograms on either side of the JEM range in Fig. 5a,b represents the relative tradeoff between: (1) adding false positives as the cost of decreasing false negatives (increasing sensitivity/decreasing specificity) as cutoff value is increased above JEM; and (2) adding false negatives as the cost of decreasing false positives (increasing specificity/decreasing sensitivity) as cutoff value is decreased below JEM. For the data presented here, the steepness of this envelope is generally less for cutoffs above JEM than for cutoffs below it, indicating that the additional cost of adding false positives in order to reduce false negatives at cutoffs above JEM generally increases *less quickly* than the additional cost of adding false negatives in order to reduce false positives at cutoffs below JEM. (The positive or negative values of the slopes are not relevant in this evaluation; rather, they indicate the direction of the false positive/false negative tradeoffs as described.) The asymmetric tradeoff of these costs implies that decisions to use cutoffs below the JEM value(s) should be considered with particular care, both conceptually and quantitatively; e.g. in the southern California data set the increase of false negatives rises particularly fast as cutoffs are lowered in the range between 0.140 and 0.030. (These considerations hold true for the alternative definition of the pollen types included in the SCD calculation, described in Section 3.)

In general, asymmetric cost tradeoffs between the two types of error, favoring 'too lenient' cutoffs over 'too stringent' cutoffs, can be seen to result from the relative variances of the distributions of like-vegetation and non-like-vegetation sample comparisons that underlie the JEM process (cf. Fig. 5c). This characteristic is demonstrated geometrically in Fig. 7, which shows idealized distributions of like-vegetation and non-like-vegetation comparisons along with their corresponding distributions of true positive fraction *minus* false positive error, as in Fig. 5 (cf. Green and Swets, 1988, p. 95). Fig. 7 shows that a relatively narrow variance of the like-vegetation comparisons coupled with a relatively wide variance of the non-like-vegetation comparisons leads to an asymmetric cost tradeoff of the kind reported here. In contrast, a relatively wide variance of the

like-vegetation comparisons coupled with a relatively narrow variance of the non-like-vegetation comparisons leads to the opposite kind of asymmetry: favoring 'too stringent' cutoffs over 'too lenient' cutoffs when a value above or below JEM is being considered.

The combination of a relatively narrow variance of the like-vegetation comparisons with a relatively wide variance of the non-like-vegetation comparisons noted for the southern California data set (Fig. 5c) is consistently observed in other modern pollen surface sample sets for which SCD histograms of like-vegetation vs. non-like-vegetation sample comparisons are reported (Anderson et al., 1989; Davis, 1995; Gavin et al., 2003; Oswald et al., 2003). This characteristic is not general across all distance metrics, however, as demonstrated for use of the SCD in relation to the 'equal weight' Canberra Metric (Prentice, 1980; Overpeck et al., 1985) by Oswald et al. (2003). The SCD histograms of like- and non-like-vegetation comparisons reported by Oswald et al. result in a true positive *minus* false positive distribution (reported in the equivalent form of true positive fraction *plus* true negative fraction) with the same asymmetry noted here, whereas the true positive *minus* false positive distribution that results from the corresponding Canberra Metric histograms of like- and non-like-vegetation comparisons is symmetric (cf. Fig. 7c). This counter-example is sufficient to demonstrate that, although the asymmetric cost tradeoffs noted for cutoffs above and below JEM appear to be consistent across modern pollen surface sample sets when sample-to-sample comparisons employ the SCD, there is no general asymmetry for these tradeoffs across the multivariate metrics that have been examined for use with pollen data (Prentice, 1980; Overpeck et al., 1985; Juggins, 2003; Gavin et al., 2003). A general examination of this issue across metrics is beyond the scope of this paper, and constitutes an important area for further research.

The methods, results, and discussion in this paper have focused on the analytical capacity of multivariate distance metrics (especially the SCD) to distinguish pollen samples from like and non-like vegetation. A parallel analysis is possible in terms of determining cutoff values to use

in reconstructing climate with modern pollen analogs for fossil pollen assemblages. In an independent analysis of the southern California data set that examines how different SCD cutoff levels affect the ability of the samples to act as estimators of modern temperature (when used as analogs for each other), degradation of the temperature reconstructions on either side of the best-performing cutoff values (0.20–0.25) shows the same tradeoff asymmetry as the results reported here for discriminating like from non-like vegetation (Wahl, 2003b). Other work by Lytle and Wahl, which employs a large ($n = 884$) regional surface sample reference set in conjunction with Monte Carlo-based subsampling of fossil pollen spectra to generate replicated reconstructions of paleotemperature for a site in northern Michigan, USA, also shows a related asymmetry (Lytle and Wahl, in review). In this analysis, the degradation of precision of reconstructed climate is much more pronounced as the cutoff value is lowered from a SCD of 0.15 to 0.05, in comparison to degradations above the best reconstruction range of 0.20–0.35.

An additional dimension of the cost of using *highly* conservative cutoff values in order to strongly reduce false positive errors is demonstrated by Lytle and Wahl in terms of the interactions among pollen count size, cutoff value, and reconstruction precision. Count sizes of ≤ 150 grains (included in the SCD sum) at the best cutoffs yield reconstruction precisions as good or better than those that can be achieved with counts of 1000 grains at very low cutoffs of 0.05 and 0.10. This result underscores the cost of employing very low cutoffs, which (at least with the SCD) leads to large numbers of false negative identifications that would otherwise add useful information. In effect, highly conservative cutoffs lead to under-utilization of the vegetation and climate information in the pollen assemblages used in the reconstruction process. From the standpoint of temperature reconstruction, this information loss cannot be removed even by increasing count size several-fold. This cost (in terms of both reconstruction precision and compensating analytical effort) cannot be properly appreciated when selection of analog threshold values is examined primarily in terms of

reducing false positive errors. When analysis of potential cutoffs is done in terms of discriminating like- and non-like-vegetation samples, the ROC-based analytical methods described here provide powerful tools that can lead toward full utilization of the information content in pollen assemblages, and can guide decisions to select cutoffs that are either above or below JEM value(s) for particular analytical purposes.

The methods of analysis presented are general and can be used with any modern surface sample reference set. In contrast, the reported JEM value(s) and rates of error–cost tradeoffs for cutoffs above and below the JEM value(s) are specific to the data set examined. This consideration is apparent from the comparison between the southern California data set and the regional northeastern North American data discussed above, and can be seen in relation to other regional data sets (Gavin et al., 2003; Oswald et al., 2003). In addition to differences in regional vegetation, specificity of results to particular data sets and kinds of analyses also arises from: (1) the number of taxa that are included in the SCD pollen sum (Sawada et al., 2001); (2) the use of a nearest-neighbor criterion for analog selection (Gavin et al., 2003) instead of examination of all the potential paired comparisons done here; and (3) the scale of vegetation differences among vegetation types when ROC analysis is used sequentially within a given data set to compare each vegetation type in turn with all the other types (Gavin et al., 2003). These considerations indicate that the analysis of appropriate cutoff values needs to be considered as a good-practice portion of research design for every paleoecological research project that uses multivariate distance metrics in the MAT, unless the case can be made that an existing analysis can be reasonably applied for a particular project.

Finally, the methods described in this paper are not unique to MAT applications with pollen data. They should be equally applicable to any technique of paleoenvironmental reconstruction that uses multivariate comparison to modern multi-taxon reference sets. Examples of these other kinds of techniques include reconstructions based on diatom, phytolith, and foraminifera assemblages.

Appendix. Pollen types and groupings used in SCD calculations

The included pollen types are listed below, following Wahl (2003a):

Pinus	*Adenostoma fasciculatum*
Abies	*Rhus/Toxicodendron* type
Cupressaceae	Ericaceae
Quercus	*Ceanothus* type
Artemisia	*Pteridium aquilinum*
Other Asteraceae	Other fern/ally spores
Poaceae	Unidentifiable pollen
Chenopodium/Amaranthus type	Unknown pollen

Pine pollen was included at the genus level of taxonomic resolution because most pine grains counted in the surface samples were either of the mixed *Pinus ponderosa/jeffreyi/coulteri* type or indeterminate to sub-generic categories, and because few pine grains can be distinguished to sub-generic categories in the fossil pollen record for which these surface samples are used as a potential analog set (Wahl, in preparation).

The *Ceanothus*, *Chenopodium/Amaranthus*, and 'Other Fern/Ally' pollen taxa were summed together with the unknown and unidentifiable categories in the SCD calculations. This grouping adds together the least represented non-arboreal

pollen types (with the exception of the Ericaceae-undifferentiated taxon, which in this region is largely associated with manzanita (*Arctostaphylos* spp.) chaparral) with the unidentifiables and unknowns to form a generic category of largely non-arboreal pollen. (*Ceanothus* occurs as non-arboreal shrubs in the study region.) Categorizing unknown and unidentifiable pollen as largely non-arboreal is justifiable based on experience in pollen counting, in which the pollen types of all the regional trees are known and the only kind of arboreal pollen that may have been included in the unidentifiable category due to poor preservation is oak. (Poorly preserved material with some resemblance to pine, fir, alder (*Alnus*), aspen/cottonwood (*Populus*), or Cupressaceae pollen either was positively identifiable as one of these types or could not confidently be considered to be pollen.) The possible oak grains that were included in the unidentifiable category are generally much less than 30% of the total unidentifiable pollen; the exceptions are in samples with otherwise high levels of oak pollen. At the least, the combined, largely non-arboreal category represents pollen that is not from coniferous trees. Since the distinction between the conifer-dominated forests and other kinds of vegetation in the region is the single most important characteristic of the pollen record utilized here (Wahl, 2003a), including the unidentifiables and unknowns preserves useful information.

Because of idiosyncratic biases in the pollen representation characteristics of sample 17 (montane conifer forest), caused by orographic and vegetation factors, *Adenostoma fasciculatum* and *Ceanothus* were excluded from the pollen sum of this sample in calculating the SCDs (Wahl, 2003a).

All taxonomic references conform with Hickman (1996).

References

Adam, D.P., Mehringer, P.J., 1975. Modern pollen surface samples – an analysis of subsamples. J. Res. US Geol. Soc. 3, 733–736.

Anderson, P.M., Bartlein, P.J., Brubaker, L.B., Gajewski, K., Ritchie, J.C., 1989. Modern analogues of late-Quaternary pollen spectra from the western interior of North America. J. Biogeogr. 16, 573–596.

Bartlein, P.J., Whitlock, C., 1993. Paleoclimatic interpretation of the Elk Lake pollen record. In: Bradbury, J.P., Dean, W.E. (Eds.), Elk Lake, Minnesota: Evidence for Rapid Climate Change in the North-Central United States. Geological Society of America, Boulder, CO, pp. 275–293.

Burns, R.M., Honkala, B.H., 1990. Silvics of North America, vol. 1, Conifers. United States Department of Agriculture, Washington, DC.

Calcote, R.R., 1998. Identifying forest stand types using pollen from forest hollows. Holocene 8, 423–432.

Davis, M.B., Shaw, R.G., 2001. Range shifts and adaptive responses to Quaternary climate change. Science 292, 673–679.

Davis, M.B., Calcote, R.R., Sugita, S., Takahara, H., 1998. Patchy invasion and the origin of a hemlock–hardwoods forest mosaic. Ecology 79, 2641–2659.

Davis, M.B., Douglas, C., Calcote, R.R., Cole, K.L., Winkler, M.G., Flakne, R., 2000. Holocene climate in the western Great Lakes national parks and lakeshores: Implications for future climate change. Conserv. Biol. 14, 1–16.

Davis, O.K., 1995. Climate and vegetation patterns in surface samples from arid western U.S.A.: Application to Holocene climatic reconstructions. Palynology 19, 95–117.

Davis, O.K., 1999. Pollen analysis of a late-glacial and Holocene sediment core from Mono Lake, Mono County, California. Quat. Res. 52, 243–249.

Gavin, D.G., Oswald, W.W., Wahl, E.R., Williams, J.W., 2003. A statistical approach to evaluating distance metrics and analog assignments for pollen records. Quat. Res. 60, in press.

Green, D.M., Swets, J.A., 1988. Signal Detection Theory and Psychophysics. Peninsula Publishing, Los Altos, CA.

Griffin, J.R., Critchfield, W.B., 1976. The Distribution of Forest Trees in California. United States Department of Agriculture, Berkeley, CA.

Henderson, A.R., 1993. Assessing test accuracy and its clinical consequences: A primer for receiver operating characteristic curve analysis. Ann. Clin. Biochem. 30, 521–539.

Hickman, J.C. (Ed.), 1996. The Jepson Manual: Higher Plants of California. University of California Press, Berkeley, CA.

Juggins, S., 2003. C2 1.3: Win95/98/NT2000/XP Program for Analysing and Visualising Paleoenvironmental Data (http://www.campus.ncl.ac.uk/staff/Stephen.Juggins/software/c2home.htm).

Little, E.L. Jr., 1971. Atlas of United States Trees, vol. 1, Conifers and Important Hardwoods. United States Government Printing Office, Washington, DC.

Lytle, D., Wahl, E., in review. The effect of pollen count size and analog threshold levels on the quality of paleoclimate reconstructions using the modern analog technique.

Mason, S.J., Graham, N.E., 2002. Areas beneath the relative operating characteristics (ROC) and relative operating levels (ROL) curves: statistical significance and interpretation. Q.J.R. Meteorol. Soc. 128, 2145–2166.

Munz, P.A., 1974. A Flora of Southern California. University of California Press, Berkeley, CA.

Oswald, W.W., Brubaker, L.B., Hu, F.S., Gavin, D.G., 2003. Pollen-vegetation calibration for tundra communities in the Arctic Foothills, northern Alaska. J. Ecol. 91, in press.

Overpeck, J.T., Webb, T., III, Prentice, I.C., 1985. Quantitative interpretation of fossil pollen spectra: Dissimilarity coefficients and the method of modern analogs. Quat. Res. 23, 87–108.

Prentice, I.C., 1980. Multidimensional scaling as a research tool in Quaternary palynology: A review of theory and methods. Rev. Palaeobot. Palynol. 31, 71–104.

Roberts, F.M. Jr., 1995. Illustrated Guide to the Oaks of the Southern California Floristic Province: The Oaks of Coastal Southern California and Northwestern Baja California, Mexico. F.M. Roberts Publications, Encinitas, CA.

Sawada, M., Viau, A., Gajewski, K., 2001. Critical thresholds of dissimilarity in the modern analog technique (MAT) for quantitative paleoclimate reconstruction. In: Chylek, P., Lesins, G., (Eds.), First Annual Conference on Global Warming and the Next Ice Age, Halifax, NS, Canada, Dalhousie University, Halifax, NS, pp. 149–152.

Thorne, R.F., 1988. Montane and subalpine coniferous forests of the Transverse and Peninsular Ranges. In: Barbour, M.G., Major, J. (Eds.), Terrestrial Vegetation of California. Wiley, New York, pp. 537–557.

Wahl, E.R., 2003a. Pollen surface samples for paleoenvironmental reconstruction from the Coast and Transverse Ranges of southern California. Madroño 50, in press.

Wahl, E.R., 2003b. Assigning climate values to modern pollen surface sample sites and validating modern analog climate reconstructions in the southern California region. Madroño. 50, in press.

Wahl, E.R., in preparation. Holocene paleoenvironmental reconstruction in the southern California Peninsular and Transverse Ranges.

Zedler, P.H., Krofta, D.A., 1995. Cuyamaca Forests: Summary Discussion and Management Recommendations. Report to the California Department of Parks and Recreation.

Zweig, M.H., Campbell, G., 1993. Receiver-operating characteristic (ROC) plots: A fundamental evaluation tool in clinical medicine. Clin. Chem. 39, 561–577.