



QUANTITATIVE BIOSTRATIGRAPHY—ACHIEVING FINER RESOLUTION IN GLOBAL CORRELATION

Peter M. Sadler

*Department of Earth Sciences, University of California, Riverside, California 92521;
email: peter.sadler@ucr.edu*

Key Words timescale, constrained optimization, graphic correlation, slotting, simulated annealing

■ **Abstract** The fossil record preserves a wide range of events that might be used to build timescales and correlate strata from place to place. The events include the originations and extinctions of species, the occurrence of distinctive faunal assemblages, magnetic field reversals, changes in ocean chemistry, and volcanic ash falls. The fundamental task is to determine the global sequence of all these events. Modern computer algorithms can build high-resolution timescales by sequencing and calibrating thousands of such events from hundreds of localities. Each successful sequencing algorithm can be understood in terms of a simple two-dimensional visual device. Graphic correlation, permutable sequences, permutable matrices, and slotting devices are each suited to a different data problem, such as contradictory evidence of sequence, a lack of information about sequence, or uncertain correspondence of events. Algorithms based upon permutable sequences and evolutionary programming heuristics have the flexibility to optimize sequences with a wide variety of event types and data problems; they are slower than the more narrowly tailored methods. All methods will be challenged to keep pace with the amount of biostratigraphic information that is accumulating in the latest generation of shared databases. Currently, it is possible to sequence sufficiently large numbers of events to imagine a potential resolution of 10,000 to 50,000 years over time spans on the order of 50 to 100 million years. The actual resolving power is less because the solutions to these sequencing problems are not unique. Multiple equally good solutions typically emerge and, to extract a consensus sequence with which all solutions agree, some runs of events must be collapsed into unresolvable clusters. Nevertheless, quantitative methods have been shown to improve resolution up to tenfold over traditional methods that discard many potentially useful events.

SEQUENCING THE FOSSIL RECORD

The raw data of sedimentary geology are mostly local in scope. Locally exposed successions of sedimentary strata, so-called stratigraphic sections, record local geologic history. Each sediment layer represents one small and short-lived patch in the dynamic mosaic of Earth surface environments; its fossils record some of the biota in that patch; and local relative age is evident because younger layers must be

deposited on top of older. But many critical questions concerning the interactions between environment and biota involve climate and have global scope. Global patterns of cause and effect do not emerge until geologic strata of the same age can be identified across widely dispersed locations. Furthermore, to resolve abrupt changes and coupled effects, this interregional correlation must resolve the briefest possible time intervals.

Beginning more than two centuries ago, solutions to the stratigraphic correlation problem have been sought primarily in the fossils that the strata contain. A fundamental challenge for time correlation, therefore, coincides with one for paleobiology in general: to determine the correct sequence of as many as possible of the biologic events, such as speciation, extinction, and immigration, preserved in the fossil record. The mid-twentieth century saw the beginnings of two significant improvements in the attempts to meet this challenge. First, several other forms of reliable stratigraphic information became routinely available to augment the paleontologic data, especially for Cenozoic strata. Radioisotopic dating, for example, can reveal the relative age of strata from widely distant locations; geochemical fingerprints improve the ability to trace unique sedimentary horizons; paleomagnetic reversals provide globally isochronous events that may be recorded in sediments; and climate-related fluctuations in the isotopic composition of ocean water provide an independent set of globally recognizable time intervals. Second, the use of paleontologic data in correlation was given more rigorous mathematical footings by a few pioneering stratigraphers who developed algorithms to solve the sequencing task on mainframe computers (see Tipper 1988 for a review). After desktop computers became commonplace tools in the 1980s, a new generation of correlation programs was developed that have enjoyed wider usage. Some of the new programs capitalize on advances in operations research, which have added to the arsenal of optimizing techniques that can be adapted to stratigraphic sequencing problems. The computer algorithms enable the sequencing task to be attempted for large numbers of biologic events and approach correspondingly fine resolution.

How short are the time intervals that these methods can resolve? The current state of the art might be assessed from two examples of the largest sequencing problems that have been shown to be manageable by individual researchers and desktop computers (Table 1). For land-mammal faunas on a single large continent, for example, data sets on the order of 6500 events and 5000 localities are evidently tractable, even if very few of the localities are arrayed in local stratigraphic sections that reveal their relative age by stratal superposition. Richly fossiliferous sections are more common in the marine invertebrate record. For the best of them, detailed range charts have been developed that resolve the locally observed appearance and disappearance horizons for tens of taxa. For global correlation of the graptolite clade (early Paleozoic marine macroplankton), for example, data sets have been sequenced that contain on the order of 3000 events from 250 local range charts that summarize in excess of 4000 individual fossiliferous horizons. The potential resolving power of these two examples, estimated by dividing the total time span

TABLE 1 Comparison of two large instances of the stratigraphic correlation problem

	North American Cenozoic land mammals	Global Ordovician to early Devonian graptolites
Instance size	6475 sequenced events for 3243 taxa from 4978 faunal lists	2856 sequenced events for 1410 taxa, 36 marker-beds and 22 bentonites from 256 range charts
Observed coexistences	217,673 coexistences 289,141 event orderings	16,986 range overlaps 26,101 coexistences
Calibration	186 dated events	22 dated events
Implied pairwise coexistences	514,088 (max likelihood) 509,359 (parsimony)	105,387
Time interval	>65 million years	~100 million years
Potential resolution	~10,000 years	~37,000 years
Mean resolution		~50,000 years
Methods	FAD-LAD matrix permutation	Sequence permutation by simulated annealing
Sources	Alroy 1992, 1994, 2000	Sadler & Cooper 2003 Cooper & Sadler 2004 Cooper et al. 2004

by the number of biological events, is on the order of 10,000 to 50,000 years. Typically, however, the sequence will not be uniquely resolvable for every event and the actual resolving power of the timescale is lower. Also, relatively few datable events of known age, such as ash fall layers, are available to calibrate the sequences. For these two examples, the average spacing of the suitable radioisotopically dated events is only 0.5 to 5.0 million years.

Neither of these case histories enjoys ideal data. Relative to marine Cenozoic correlation problems, nonmarine instances suffer from a lack of continuous sections, and Paleozoic instances include less radiometric, paleomagnetic, and stable isotopic data. But neither example will have exhausted the resolving power of the fossil record until data from contemporaneous clades have been combined. Mammal events could be combined with floral events; the graptolite clade coexisted with conodonts, brachiopods, trilobites, etc. Proprietary solutions to such larger multiclade sequencing problems have undoubtedly been developed by research teams in petroleum companies, notably Amoco (Mann & Lane 1995a).

Most recently, the paleobiological sciences have begun to invest in huge shared databases (The Paleobiology Database, Chronos, FaunMap, etc.). These data sources will surely outgrow those of the petroleum industry and rapidly increase the size of readily available data sets to which correlation algorithms should be applied.

This information explosion presents three serious challenges. One is to develop algorithms that can efficiently solve much larger instances of the correlation problem than are currently attempted; the second is to compile these databases with full regard to their usefulness in time correlation and the minimum demands of the correlation algorithms; the third is to increase the number of radioisotopically dated events.

Rather than attempt a consumer report of all existing algorithms for stratigraphic correlation, it seems appropriate at this time to review the range of mathematical approaches that have enjoyed success and the varied nature of the available data, which determines the suitability of different methods for any given correlation problem. Tipper (1988) has compiled a more complete catalog of earlier methods.

DIFFERENT METHODS FOR DIFFERENT DATA

Geologic time correlation proceeds by constructing a global calendar of past events in which the appearances and extinctions of fossil species dominate the entries. Other events include changes in ocean chemistry, reversals of Earth's magnetic field, and the deposition of volcanic ash beds, some of them dated by radiometric methods. The challenge is to merge incomplete inventories of physical events and partly contradictory faunal successions from many local thickness scales (measured stratigraphic sections) onto a single calendar that correctly sequences all the events and scales the time intervals between them. Because correctly sequenced events serve the purpose of correlation, with or without knowledge of their numerical ages, sequencing is the fundamental task and the focus of this review. Numerical estimates of age are available for very few events, especially in the older periods of the Phanerozoic. Furthermore, estimates of the relative size of time intervals between events rest largely upon questionable assumptions about rates of sediment accumulation and biological turnover. Consequently, scaling and calibration tasks are best attempted after the optimal sequence of events has been determined.

The way to improve the resolving power of the geologic calendar is obvious but not easy—increase the number of events and thus reduce the average time intervals between them. There is no shortage of species to add. The real problem is to keep all the appearance and extinction events in their correct sequence. The difficulty increases dramatically with the number of species for three reasons: First, the number of possible sequences of appearance and extinction events grows faster than exponentially as a function of the number of species (Figure 1). Also, events that are separated by smaller time intervals are more likely to be preserved in contradictory order from place to place. Finally, as the list of species grows it must include more provincial organisms that will be missing from many locations. Traditionally, the solution was to sacrifice resolving power and limit the number of events to the few whose order is easy to determine. But the rules for determining the best order are very straightforward and, with computer assistance, can readily be applied to otherwise dauntingly large numbers of events.

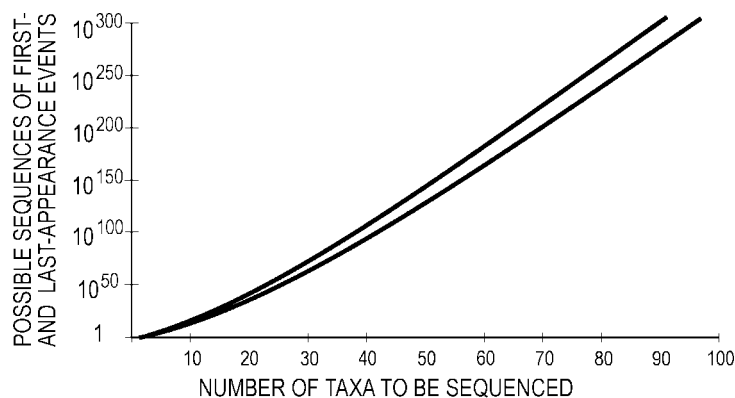


Figure 1 The problem of sequencing paleobiological first and last appearances is “NP complete” (Dell et al. 1992); i.e., the number of possible sequences grows faster than exponentially as the number of species increases. The upper bound assumes ~~that~~ no species coexist; the lower bound applies the extreme constraint that all species coexist.

The aim of this review is to enable readers to choose intelligently between current computer-assisted methods of correlation. Space does not allow full description of any method, but we can distinguish their different strategies, the questions that they can answer, and the data to which they are properly suited. Logically, the methods must come last. We follow Edwards’ (1982) premise that “the methods should suit the data.” The data make the correlation problem difficult and they vary from instance to instance. They limit the suitable methods, which, in turn, determine precisely which questions can be answered. So, we first examine the data from a rather abstract, mathematical view. This enables us to formulate the problem more rigorously. Then we may usefully turn to the methods.

Some of the methods that were implemented on institutional main-frame computers in the 1970s and 1980s, most notably those based upon principal components analysis, failed to generate much interest after the personal computer revolution. The surviving methods seem to be those that can readily be illustrated using simple two-dimensional visual devices. Intuitive heuristic analogies and easy adaptation to graphical user interfaces seem to have been the keys to popularity. We examine the popular methods in terms of four visual devices: piecewise linear regression on an x-y graph, sequence permutation, pattern management in permutable matrices, and least-cost traverses across the cells of fixed matrices. The rules uncovered in these simple games suffice to understand the algorithms that solve large instances of the problem. We complete the progression from data to method for paleobiological data and then consider how they may be augmented by other kinds of information.

PALEOBIOLOGICAL DATA

The local appearances and disappearances of fossil taxa are the most readily available data for correlation. Biological species are very numerous and conveniently unique. For the purpose of global correlation, many species are usefully widespread and short-lived. But it is crucial to acknowledge that local first and last appearances are also uniquely troublesome as recorders of calendar events: The local stratigraphic horizons at which they are observed do not reliably reproduce the true global sequence of origination and extinction events. Discrepancies must be expected because local appearances and disappearances are likely to be migration events and probably displaced by lapses in fossilization. We explore the implications below.

The true global origination and extinction events (first and last appearance datums or FADs and LADs) of a single taxon are discrete points on the geologic timescale—event horizons. It is useful to our broader purposes to recognize that there are closely related interval data and continuous time series. The continuous variable is the abundance of the taxon (Figure 2). The FAD and LAD are endpoints of the time interval that spans all the nonzero abundances in the fossil record—the taxon range. Ideally, the data for correlation are charts of the ranges of several taxa, constructed from an ordered suite of samples in a single stratigraphic section.

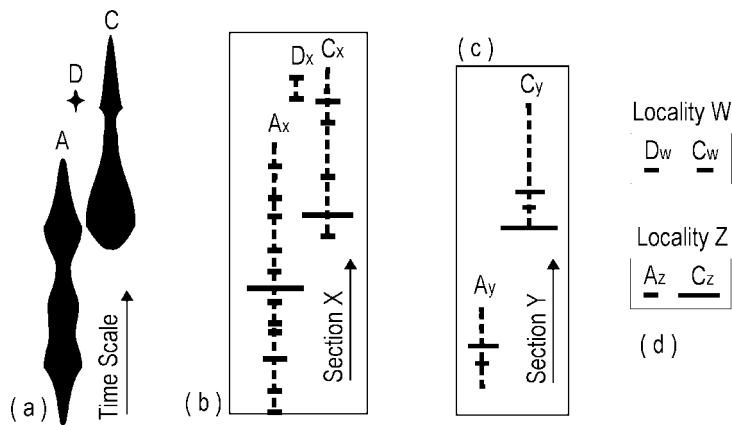


Figure 2 Paleobiological data. (a) Spindle width shows changing abundance of taxa A, C, and D through time. (b) Range chart for section X; black bars indicate relative abundance at each fossiliferous horizon; overlapping ranges prove coexistence of taxon pairs A:C and C:D. (c) Range chart for section Y. (d) Isolated localities W and Z prove coexistence C:D and A:C, respectively. Section X preserves true sequence of range ends; section Y does not. Section Y, locality W, and locality Z together reproduce most of the relationships recorded in section X. Neither the locally preserved range ends nor the preserved abundance peaks are correlative.

In less favorable instances, isolated faunas for which the stratal separation and relative age are not known dominate the data.

Biology demands that the global abundance of a species cannot fall to zero within its temporal range. Unfortunately, species distributions are patchy, the patches may shift, few individuals are fossilized, and fossils may be overlooked. Consequently, the local taxon ranges observed in single stratigraphic sections reflect local conditions and do include gaps. More critically, and for the same reasons, gaps of unknown extent occur at the ends of observed ranges. Thus, local horizons of highest and lowest finds of a species do not correspond to the global FAD and LAD. The discrepancies vary from place to place, and locally observed taxon range charts contradict one another in detail concerning the sequence of range-end events. Fortunately, our stratigraphic freedom to adjust the observed endpoints of a range is limited. Assuming that the fossils were recovered *in situ*, the global FAD lies at or below the stratigraphically lowest find in each section, and the global LAD lies at or above the horizons of the locally highest finds. In other words, each locally observed range is likely to be shorter than the global range, but cannot be longer. It follows that all locally observed coexistences are trustworthy and must appear in the true calendar.

The temporal ranges defined by a suite of taxa are, in general, not contiguous; they may overlap or be disjunct. The overlap intervals are coexistences that have considerable value in correlation. We may also manufacture arbitrary intervals of high and low abundance and recognize points of peak abundance (so-called *epiboles* or *acme horizons*). Their value in correlation is limited. Even when variations in the abundance of living organisms are synchronized over a wide area, the process of fossil preservation will surely distort the abundances.

The various kinds of observed paleobiological data, their stratigraphic freedom, and the corresponding temporal events are summarized below. Unless otherwise stated, the associated logic assumes that none of the fossils has been reworked into younger strata or washed down a well bore into older strata. Burrowing organisms can displace tiny fossils. Waves and currents can exhume and redeposit larger fossils if they are not destroyed in the reworking process. Let us assume, for now, that reworked fossils can be recognized and discarded from the data. Some micropaleontologists exclude very low abundance occurrences at the ends of the observed range. This is a ploy to avoid reworked fossils. We return to this problem later.

1. Lowest find of a taxon in a range chart: a unique local observation that may fall at or above the horizon of the global FAD event and is free to be adjusted down-section only (Shaw 1964).
2. Highest find of a taxon in a range chart: a unique local observation that may fall at or below the horizon of the global LAD event and is free to be adjusted up-section only (Shaw 1964).
3. Nonoverlapping ranges in a range chart: proof that the FAD of the taxon observed lower in the section must precede the LAD of the taxon observed above. This follows from the asymmetrical degrees of freedom given above.

The observed sequences of all other combinations (LAD-below-FAD, two FADs, or two LADs) may have been preserved in the wrong order (Alroy 1994).

4. Overlapping ranges in a range chart: proof of coexistence and a constraint on acceptable sequences of events; the FADs of both taxa must occur before either LAD (Guex 1977, 1991; Alroy 1992; Kemple et al. 1995).
5. Reversal of order of two observed range ends from one range chart to another: proof of overlapping ranges on global timescale; the observed contradiction is an artifact of preservation or a record of migration (Sadler 2003). Alternatively, if reworking and caving cannot be excluded and numerous range charts are available, the relative frequency of the two possible pairwise sequences may be used as the basis for a probabilistic estimate of the true order (Agterberg & Nel 1982, Agterberg & Gradstein 1999).
6. Coexisting taxa in an isolated fauna: proof of coexistence and a constraint on acceptable sequences of events; the FADs of both taxa must occur before either LAD (Guex 1977, Alroy 1992), unless the fauna is a very condensed or time-averaged accumulation (Sadler 2003). In exceptional circumstances, the total assemblage and relative abundances may be a unique and traceable event horizon.
7. Peak abundance events: nonunique observations whose value in time correlation depends upon ad-hoc rationales for a globally synchronous acme; the observed level is free to be adjusted up- or down-section within the observed taxon range.
8. Arbitrary abundance intervals: nonunique partial ranges of questionable value for correlation except, perhaps, for estimating probabilistic ranges or eliminating the potential effects of reworking. (Gradstein & Agterberg 1998).

THE PALEONTOLOGICAL CORRELATION PROBLEM

The true global sequence of FADs and LADs is not knowable in detail and the locally preserved sequences of highest and lowest finds are incomplete and contradictory. The practical and tractable problem is to find a hypothetical sequence of FADs and LADs that enjoys the lowest net misfit with all observations in local range charts and isolated faunas, or requires the smallest net adjustment of all observed ranges. It is an optimization problem. The objective function sums the misfit increments or adjustments across all taxa, sections, and faunas; it must be minimized. The problem becomes a constrained optimization (Kemple et al. 1995) because there are some necessary features of feasible solutions that must be enforced: The FAD of every taxon must precede its LAD, and every proven coexistence must be honored. It is usually desirable to add the constraints that the adjustments may lengthen, but not shorten, the

locally observed ranges and that all observed FAD below LAD sequence must be honored.

Notice that contradictions in the sequence of FADs and LADs are only possible (ignoring taxonomic blunders) for taxa whose ranges do overlap and only likely for events that are more closely spaced in time than the gaps in their local ranges. Thus, the paleontologic correlation problem may be greatly simplified, at the expense of resolving power, by considering only a few, widely separated events. Although this traditional solution seems rather crude in the computer-assisted age, it is the reason why biostratigraphy has always been able to solve the correlation problem, at least at low resolution.

The various computer-assisted methods differ in two critical regards: how they measure misfit and the way they search for the optimal ordering of events. Although the search strategy is primarily a matter of programming that influences the speed of the optimization, there is one geologic consideration: In general, the fastest search strategies result from restricting the range of data or event types for which an algorithm is suited. The choice of the measure of misfit is always a critical geological decision: Different measures lead to different optimal sequences. Thus, our discussion of the methods must consider the measures of misfit to which they are suited.

METHODS

The logic of several methods is illustrated with simple two-dimensional graphics. These visual analogs readily indicate the data to which a method is suited and the logic of the underlying algorithms. They are discussed according to the following data progression: solutions for two range charts, multiple range charts, isolated faunas, and finally sequential faunas in a time interval that is shorter than the taxon ranges. This last category provides a logical and operational transition to nonpaleobiological data.

Graphic Lines of Correlation

A rigorous mathematical basis for stratigraphic correlation was firmly established by Shaw's (1964) development of graphic correlation. In this method, the sequence and spacing of range-end events in two range charts are plotted as orthogonal X and Y axes as in a linear regression exercise (Figure 3). The best-fit regression line is termed the line of correlation (LOC). The LOC may not have a negative slope, but may be curved or be piecewise linear (Macleod & Sadler 1995). Changes in slope correspond to changes in the ratio of accumulation rates in the two sections. Vertical and horizontal segments point to surfaces of hiatus in one section. Time intervals that are recorded as a hiatus in both sections collapse to a point on the LOC. Only hiatuses of different age or duration are expressed as a distinct LOC segment.

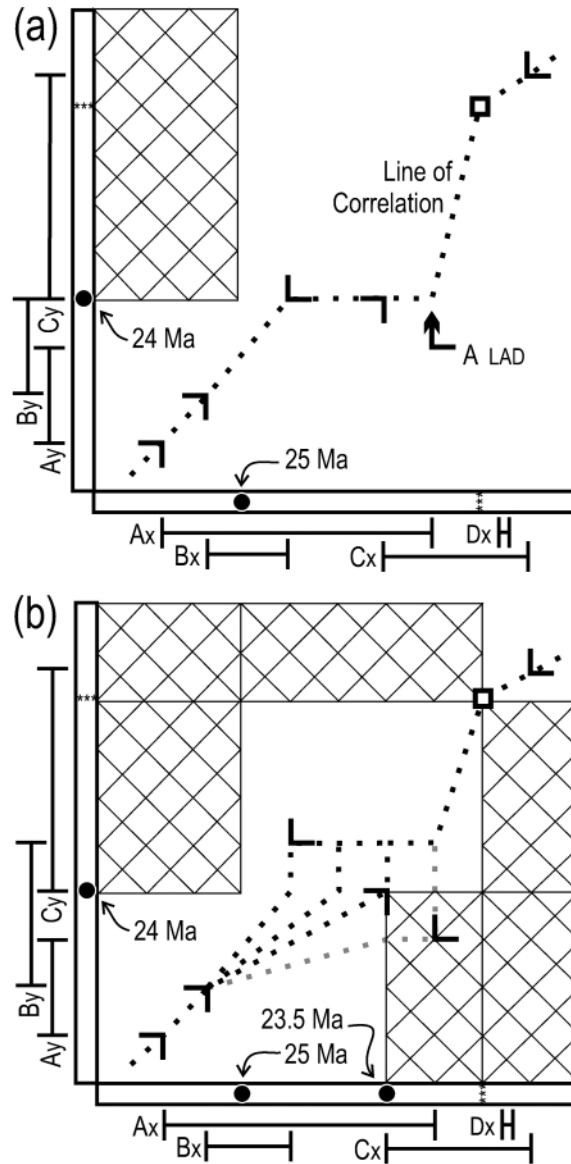


Figure 3 Graphic correlation of four taxon ranges (*A, B, C, D*), a marker bed (*little stars*), and two radiometric dates (*filled circles*) in two sections. Right-angle symbols: coordinates of shared range ends, showing freedom to adjust range ends in direction of ends of symbol. Open square: coordinates of marker bed with no freedom to adjust. (a) Cross-ruled area is out of bounds for all LOCs that honor the sequence of the two radiometric ages. (b) A different range for taxon B in section Y allows several feasible LOCs, one of which is eliminated by an additional radiometric date.

In graphic correlation, the coordinates of first- and last-appearance events are adjusted onto the LOC parallel to one axis or the other, so that the adjustment represents a range extension in one section. The best-fit LOC is the one that requires the smallest net range extension, summed across all taxa and sections. This method solves simultaneously for the sequence and the spacing of events. In the process, it introduces into the sequencing task some unnecessary assumptions about accumulation rate (Kemple et al. 1995). Procedures for fitting the best LOC include deterministic regression techniques (Macleod & Sadler 1995) and heuristic search algorithms from evolutionary programming (Zhang & Plotnick 2001).

Suppose that instead of two local range charts, the problem involves numerous measured sections. It is still possible to work forward from the data to a solution via a long series of two-dimensional steps using graphic correlation (Shaw 1964). The next paragraph explains a multidimensional alternative.

Permutable Sequences

Faced with numerous range charts, it is preferable to invert the problem and search for the best-fit sequence by comparing numerous trial sequences in turn with the raw data (Kemple et al. 1995). Starting from a random sequence, the method proceeds by trial and error. Each trial sequence is generated by a mutation of the previous trial (Figure 4). The search can be steered toward the best-fit sequence by one of several heuristic search procedures developed in the discipline of operations research to solve NP-complete problems. The simulated annealing heuristic is favored because it combines simplicity with effectiveness. Tabu-search, and greedy algorithms with random restarts have also been implemented (Dell et al. 1992).

Mutations are rejected automatically if they violate constraints. Otherwise, under simulated annealing, their acceptance or rejection depends upon how they change the misfit between the trial sequence and the empirical data. A rejected mutation is removed before the next mutation is attempted. The results of accepted mutations are cumulative. The search heuristic accepts some temporarily deleterious mutations in order to avoid getting trapped in a suboptimal sequence. The mutation method must be chosen with regard to the efficiency of the search. Figure 4 illustrates three mutation sizes that have been used with simulated annealing (Sadler 2003). The big mutation (Figure 4*b*) is most efficient because the mean size of the mutation grows with the size of the data set. The double mutation is too big: It moves the search away from bad sequences very rapidly but may fail to converge on the optimal solution. The smallest mutation (Figure 4*a*) enables the fit to be recalculated most rapidly. Unfortunately, it becomes increasingly small as the size of the data set grows and eventually requires far too many trials to be competitive.

Large correlation problems typically yield more than one equally well-fit solution. The relative position of some blocks of events simply cannot be resolved by the measure of misfit. The resolving power of each event is then expressible as a best-fit interval, which is the range of positions in sequence that it can occupy in the full set of equally well-fit solutions (Sadler & Cooper 2003). The closely related

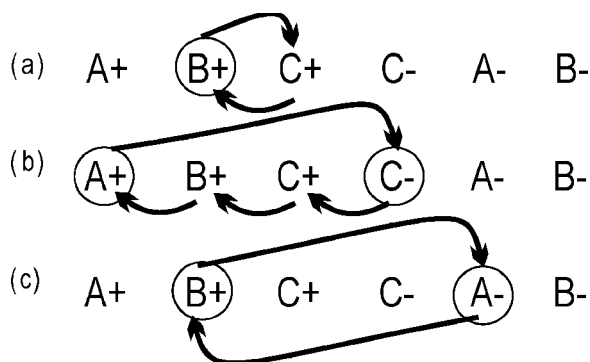


Figure 4 Permutable event sequences. Time runs from left to right. A+ and A- are, respectively, the first- and last-appearance events of taxon A. Small mutation (a) switches a random event with its younger neighbor. Big mutation (b) moves one random event to a new position at random; all intervening events shift one position in opposite direction. Double mutation (c) switches two randomly chosen events. All mutations are subject to constraints.

relaxed-fit curve tracks the increasing range of each event as the set of acceptable solutions is expanded to include increasingly suboptimal fit. The resolving power of the whole data set is given by the maximum number of events that may be selected, subject to the constraint that their best-fit intervals do not overlap (Sadler 2003).

MEASURES OF MISFIT As in graphic correlation, the measure of misfit used to optimize a permutable sequence may be the sum of all range extensions needed to fit the data to the trial sequence. It is an advantage of inversion, however, that a wide variety of measures can be implemented within the same search algorithm. In this way, permutable sequences can be suited to the data and misfit measures of all the other methods. Thus, the permutable sequence approach can readily assess the relative resolving power of different measures and compare the best-fit sequences that they generate. This flexibility results from the fact that inversion works from possible answers to the data; the answer is always a sequence, regardless of the data, the assumptions, or the method. Other methods may be specific to certain formats of raw data, particularly the relative abundance of range charts and isolated faunas. In the permutable sequence method, it is the measure of misfit that must adjust to suit the nature of the observed data. The choice of the measure of misfit is a critical stratigraphic decision that determines the outcome of the sequencing task. In effect, different measures of misfit entail different assumptions or answer different questions about the sequence of events. The measure of misfit does also influence the speed with which a solution can be found. Sequence permutation is likely to be relatively slow when mimicking methods other than graphic correlation. That is the cost of flexibility.

To measure misfit in terms of range extensions, the algorithms must be able to interrogate measured sections. The length of the extensions can be expressed as stratigraphic thickness or the number of event levels they encompass. The former measure tends to favor the sequences preserved in the thickest sections; the latter favors the most fossiliferous sections and is better behaved when the sections have different rock types and preservation potential. Observed coexistences, including those observed in isolated faunas (Webster et al. 2003), serve as constraints on the acceptable mutations of trial sequences. The threshold for evidence of coexistence may be set to include any co-occurrence in a single sample or raised to require the observation of overlapping ranges. The stiffer of the two criteria is less susceptible to problems caused by organisms that churn the near-surface sediments or sampling protocols that aggregate the fossils collections from adjacent layers.

Other possible measures of misfit do not require such detailed interrogation of the stratigraphic sections. They are ordinal in the sense that they use only the observed sequence of events, not their separation. They may be used to mimic any of the permutable matrix methods discussed below or applied in concert with range extensions. One measure totals the number of pairwise coexistences that are implied by a trial sequence but not observed in any range chart or fauna. A closely related measure totals all the pairwise sequences of FADs-before-LADs that are implied but not observed. Yet another measure counts the number of times that an observed sequence of any two range-end events contradicts the order of the same pair in the trial sequence. The contradictions of this kind can be expressed as a fraction of the total number of observations of the pair to give the procedure a probabilistic basis. Further implications of these essentially ordinal measures are discussed under their implementation in permutable-matrix methods.

Permutable Matrices

Now suppose that most of the raw data are isolated faunas, rather than range charts from stratigraphic sections. In other words, the observations include many instances of coexistence but little direct evidence of sequence. This situation may arise in nonmarine strata because the rocks are sparsely fossiliferous and patchily exposed or because the fossils are found in a complex of lenticular channel deposits rather than simple layers. Even in fossiliferous marine strata, such as radiolarian cherts, complex folding and faulting may obscure stratal superposition.

Now the method of choice is to display the observed relationships in the cells of a matrix in which the rows and columns are assigned to taxa or events, in arbitrary order. The order of rows and columns is then permuted until the contents of the matrix match the expectation for a good solution. Permutable matrix applications can be recast for solution using graph theory (Guex 1991). Three forms of permutable matrix are discussed here. The first assigns a row and a column to each taxon; the second assigns the rows to FADs and the columns to LADs; the third assigns a row and a column to every range-end event.

COEXISTENCE MATRIX The first application uses the cells of a permutable matrix to record observed coexistences between two taxa: one indexed to the row in which the cell is located, the other to the column (Figure 5). If rows and columns array the taxa in order of decreasing age, then the observed coexistences will be more likely between near neighbors than between pairs of taxa whose rows and columns are widely separated in the matrix. In the correctly ordered matrix, therefore, cells marked as observed coexistences must be concentrated about the main diagonal of the matrix (Guex 1991, Alroy 1992, Blackham 1998) and away from the top right and lower left corners. Accordingly, good solutions for the sequence of taxa are found by permuting the order of taxa in the rows and columns until a maximum concentration of observed coexistences is achieved along the main diagonal. Some cells without observed coexistences may remain embedded in the diagonal concentration of observed coexistences, regardless of

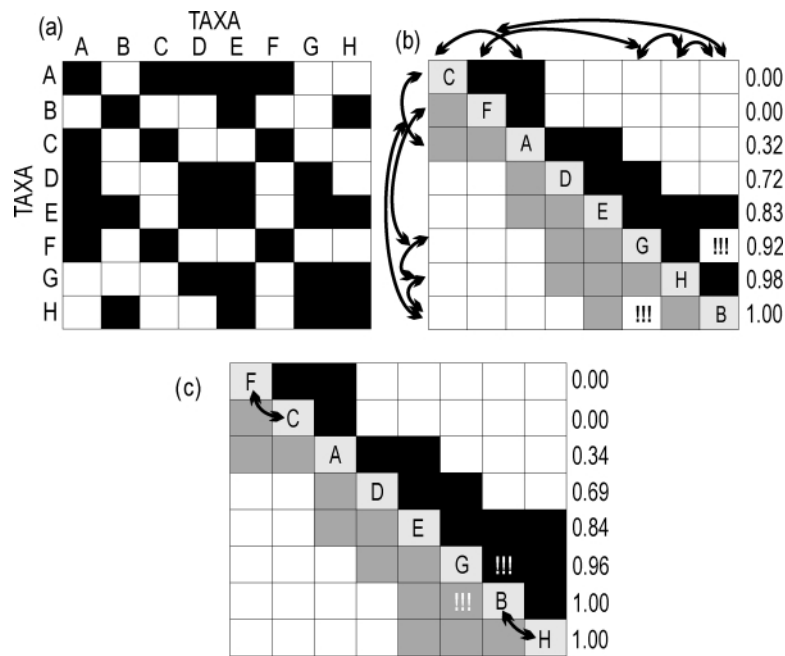


Figure 5 Permutable coexistence matrix. (a) Black cells record observed pairwise coexistences for eight taxa, A to H (extended from Alroy 1992). (b) Permuted matrix achieves diagonal concentration of filled cells, placing taxa in feasible sequence. Following Guex's (1991) convention, black cells are shown only to right of main diagonal. Exclamation marks highlight one coexistence (B,G) that is implied (b) but not observed (a). (c) Repermuted matrix after addition of the implied coexistence. The internal order of the pairs F:C and B:H cannot be resolved. Numbers to the right of matrices are reciprocal averaging scores (explained in Alroy 1992).

permutation. These cells correspond to coexistences that are implied by the solution but have not been observed, either owing to failures of preservation or because the taxa occupy different provinces. The total number of these unobserved, but implied, coexistences is a measure of misfit between the solution and the field observations.

If the permutation process is quantified, as in the method of reciprocal averaging (Alroy 1992), the position of each taxon in the optimal sequence can be assigned a score between 0.0 and 1.0. The scores may be used to compile a composite sequence of first and last appearance events as follows: return to the isolated faunas and assign each a relative "age," which is the average of the scores of all the taxa that it contains. (When faced with data in the form of range charts, we found solution strategies that proceed directly to a composite sequence of events. Now that the available data are predominantly isolated faunas, the preferred methods return estimates of relative age for those faunas.) Use their average scores to assemble the faunas into a composite section that uses average scores in place of rock thickness. From this composite section, develop a final range chart and a full sequence of events. Alternatively, the optimally permuted matrix may be used to identify a series of assemblage zones (Guex 1991).

FAD-LAD MATRIX Alroy (1994) realized that the information in a coexistence matrix can also be cast as a second type of permutable matrix in which rows are first-appearance events and columns are last-appearance events (Figure 6). The observation that a pair of taxa coexists is proof that both FADs occur in sequence before either LAD. Two corresponding cells are marked to indicate the FAD-before-LAD constraints. Other FAD-before-LAD constraints can be added from disjunct ranges in range charts. Thus, the FAD-LAD matrix can contain more information than the coexistence matrix. It differs from the coexistence matrix in two other important respects. The order of rows and columns may be permuted independently. When the rows and columns of this matrix are arrayed in order of decreasing age, the marked cells should fill the upper right half. Of course, finding this target arrangement requires a different permutation protocol. Alroy (1994, 2000) discusses the options. His preference is one that assigns weights according to the likelihood that each event pair will be observed. Because this likelihood varies with the true sequence of events, an iterative improvement process must seek the optimal permutation. Again the number of embedded unobserved cells emerges as a natural measure of misfit, which can also be implemented in sequence permutation methods.

The matrix permutation process provides numerical scores, which can be used to construct a master sequence of events, estimate the relative ages of the isolated faunas, and draw up a composite range chart. But the process begins a little differently from that already described for coexistence matrices. Notice that the FAD scores (right-hand end of the rows in Figure 6) and LAD scores (bottom of the columns) are independent scales that cannot be merged directly into one numerical series to represent relative age: Some short-lived taxa receive

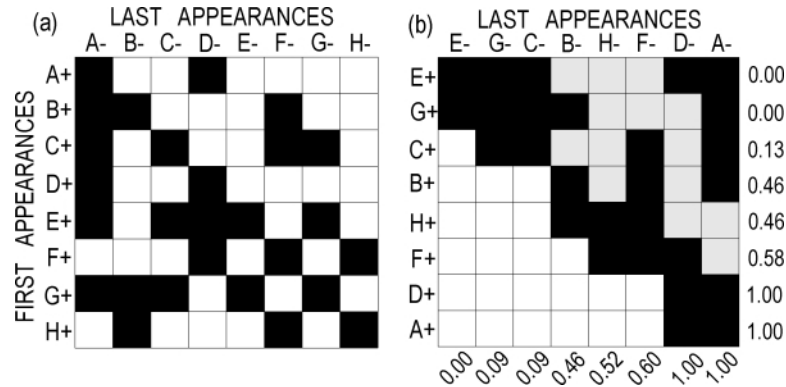


Figure 6 Permutable FAD-LAD matrix or event ordination matrix (example from Alroy 1994). (a) Matrix showing order of observed range ends for eight taxa (A to H). Black cells indicate observations that the first-appearance event (+) precedes the last-appearance event (-). (b) Matrix rows and columns rearranged to concentrate filled cells to the upper right. Row and column scores increase from oldest to youngest events. Gray cells are implied but unobserved pairwise sequences. Numbers are Alroy's appearance event ordination scores from the permutation process.

FAD scores that are higher than their LAD scores. The FAD and LAD sequences are merged into one series that preserves the order of scores for each and respects all the FAD before LAD implications. Then the merged sequence is provided a new ordinal age scale by numbering events from first to last. Alroy (2000) recommends numbering by consecutive runs of like events (FADs or LADs), not by every individual event. The score for each fauna is then the span of numbered positions in the merged sequence through which all of its taxa are present.

ALL-EVENT MATRIX A third and larger event matrix is imaginable. Its rows and columns include all events, both the first and the last appearances. The cells of this matrix record all the locally observed pairwise orderings of events. Cells contain the number or frequency of observations of the row event before the column event and can be filled with information from range charts. Permutation would then seek to place as many of the high frequencies as possible above and to the right of the main diagonal. This is a probabilistic ranking strategy. It powered one of the earliest correlation programs (Agterberg & Nel 1982) and remains in use in the RASC, or ranking-and-scaling, program (Agterberg & Gradstein 1999). This approach differs in one crucial respect from all that we have considered so far: It seeks the average or probable range of each taxon.

We can now return to the question of fossils that are reworked into younger strata or washed down well bores into older strata. All the methods except the last seek extreme taxon ranges that extend from the earliest of the first appearances to

the youngest of the last appearances. They all assume (implicitly or explicitly) that reworking and borehole caving have not modified the observed ranges (or that displaced fossil occurrences have been removed from the data); otherwise, observed ranges could be longer than the true ranges and the underlying logic of the methods would fail. As a result, these methods find the true extreme ranges plus additional increments owing to unrecognized reworking or caving. The overestimation of the true global range is small if the cryptic reworking distances are correctly assumed to be short relative to the pervasive underestimation of true ranges by preserved ranges (Edwards 1982). Because the underestimation errors should resemble the size distribution of the observed gaps within fossil ranges (Sadler 2001), there is some as yet unexploited scope for testing this assumption.

The probabilistic ranking strategy replaces these assumptions with a model in which the observed range ends differ from the true range ends by a symmetrically distributed random error. It accepts the alternative risk of underestimating the true taxon ranges and effectively assumes that the likelihood of reworking and borehole caving are equal to the likelihood of gaps in preservation. Gradstein & Agterberg (1998) have suggested applying probabilistic ranking algorithms to shortened taxon ranges drawn only from the first to the last common occurrence, rather than the rarer extreme finds in each section or from the first to the last consistent occurrence. The latter mark the ends of the longest gap-free portion of the observed range. These somewhat subjective approaches probably impart more symmetry to the range-end errors; thus they improve the assumption behind the permutation of all-event matrices and, perhaps, explain why the resulting correlations seem more robust.

One purported advantage of probabilistic ranking is the ability to generate statistical confidence statements about the optimal sequences. In contrast, methods that seek maximal ranges determine best-fit intervals from the set of equally well-fit sequences (Cooper & Sadler 2003). Each strategy for quantifying the quality of the solution is suited to its underlying assumptions. Similarly, comparison with the logic of Alroy's (1994) FAD-LAD matrices might leave the impression that the larger matrices used in probabilistic ranking have missed the critical information that FAD-before-LAD observations are always correct. But this assertion depends upon the assumption that reworking and caving are insignificant. Without good information concerning the impact of reworking, it is not worthwhile to argue the merits of the two matrices. The choice of method must suit the assumptions and the data. Not surprisingly, average or probabilistic ranges were developed for marine microfossils in oil wells; these data are prone to caving and reworking. Severe caving may require abandoning FADs altogether (Cooper et al. 2001). The FAD-LAD matrix was developed for land mammals; these microfossils have a lower propensity for reworking than marine microfossils. Cooper et al. (2001) observe that the different methods should be recognized to have answered different questions. Logically, the average ranges predict the local sequence of events that is most likely to be found in the next measured section. Maximum ranges are a conservative estimate of the total time span of strata in which a taxon may be




Figure 7 Immutable horizon-slotting matrix. (a) Matrix rows and columns show event horizons in their observed, immutable, order in sections *X* and *Y*. Cell numbers are measures of the dissimilarity between the corresponding row and column horizons. The lowest-net-dissimilarity path across the matrix (*gray cells*) may take only horizontal or vertical steps. Notice that the perimeter values are repeated to allow runs of events from one section at the ends of the optimal sequence. (b) Events are interleaved in the order in which they are encountered along the path.

found worldwide, including unrecognized cases of reworking. Maximum ranges are needed for studies of global biodiversity.

Pathways Across Immutable Matrices

Graphic correlation and matrix permutation are appropriate methods when the correspondence of events from locality to locality is not in doubt, but their order is uncertain. Consider instead the correlation of sequences of faunas in which the events are changes in the proportions of taxa, not first and last appearances. Immutable matrices now become appropriate because the local order of faunal assemblages is known but not the correspondence of assemblages from one section to the next. Assemblages from two sections are arrayed in their immutable order as rows and columns in a matrix (Figure 7). Each cell now holds a numerical estimate of the dissimilarity between the corresponding row and column assemblages. The cell-by-cell path across the matrix that accumulates the smallest sum of dissimilarity values represents the best interleaving or slotting of events from the two sections into a composite sequence (Hawkins & Merriam 1974, Gordon & Reyment 1979).

Many advances in slotting algorithms have been developed in response to the properties of pollen data from lake cores (Clark 1985, 1995; Pels et al. 1996). Slotting becomes the most suitable method for data like these in which the stratigraphic sections are shorter than the ranges of the taxa in the floras that they preserve. Now,

the taxon ranges extend beyond the limits of the section; there are no meaningful observed range-end events. Any attempt to correlate must rest on the changing mix of species and their abundances, from horizon to horizon. In effect, the faunal assemblages are used like any physical characteristics of the strata.

Slotting by least-dissimilarity pathways assumes that in the true composite sequence, floras will resemble their near neighbors more closely than their distant neighbors. This potentially weak assumption is strengthened by the immutable matrix, which forces the solution to keep faunas from the same section in their observed superpositional order. In other words, we cannot just select cells with the lowest difference measures as evidence of correlative faunas; we need a path of contiguous cells. Thus, for example, it is possible to interleave faunas from two sections with assemblages of pollen or foraminifera that change in response to climate fluctuations on a shorter timescale than the duration of individual species.

The method can accommodate almost any measure of dissimilarity, including a wide range of multivariate options that can be tailored to the data. The more sophisticated options (e.g., Pels et al. 1996) include the abundance of multiple species, the abruptness of the changes in abundance between adjacent faunas (peak and trough features in the abundance curves), and weights based upon the thickness intervals between faunas. Thus, the method can match the shape of two curves of relative abundance. Slotting is easily constrained by correlative horizons known from both sections; the least-dissimilarity pathway is simply forced to pass through the corresponding cells.

Conventional two-dimensional slotting algorithms pioneer their route across the immutable matrices in one pass by scouting alternative choices just one or two cells ahead. Ryan et al. (1999) propose an iterative improvement approach that examines random mutations of the path with the simulated annealing heuristic—the same procedure by which Kemple et al. (1995) permute sequences of range-end events. A single permutable sequence or a single two-dimensional permutable matrix can contain all the taxa or events in a correlation problem. Immutable slotting matrices, however, need one dimension for every measured section. Multidimensional slotting algorithms are reputedly rather unwieldy. The simulated annealing approach ought to be more readily extended to multiple sections, however, because it works backward from trial solutions to the data. Several users of immutable matrices (Gordon et al. 1988, Thompson & Clark 1993, Pels et al. 1996) derive a measure of the confidence of their solutions by mapping the full set of equally optimized paths or the larger set of next-best paths. This is the same strategy as the best-fit intervals and relaxed-fit curves derived from permutable sequences (Sadler & Cooper 2003).

CORRELATION AND SERIATION

Graphic correlation and probabilistic ranking methods were developed for stratigraphic sections that do correlate, in the sense that their time spans overlap considerably. The large Paleozoic correlation problem in Table 1 includes many pairs of

sections that do not overlap in age. They must be stacked in the correct order and impart to the problem a significant component of seriation. Seriation is the essence of the problem when the data are isolated faunas. The mixed correlation-seriation problem is most efficiently solved by a combination of two measures of misfit. Correlation is achieved primarily by minimizing range extensions as in graphic correlation. Proper seriation is enforced by simultaneously minimizing the number of implied but unobserved coexistences as in matrix permutation. In effect, the second measure uses the negative evidence that nonoverlapping sections may have no taxa in common. It also introduces the need to assign weights to the two penalties to determine their relative importance. Although the relative weights do influence the solution, there appear to be no objective criteria to determine the appropriate size of the weights. Subjectively better solutions appear to result from relatively low weights for the implied coexistences.

NONPALEOBIOLOGICAL EVENTS

Unlike the first and last appearances of fossil taxa, physical stratigraphic events do not usually allow any freedom to adjust their observed levels. Consider a chemically fingerprinted ash-fall horizon. If it is preserved in a stratigraphic section, it must be preserved at the appropriate level. Where more than one such event is preserved in the same section, the observed sequence is immutable. Therefore, the value of nonpaleobiological events to correlation is best realized in immutable matrices or as powerful constraints in methods that are chosen to suit the paleobiological data. Immobile, uniquely identified events seen in more than one section provide cells in immutable matrices through which least-dissimilarity paths must pass, points through which LOCs must pass in graphic correlation (Figure 3*b*), and severe restrictions on the permissible mutations of the order of events in a permutable sequence.

Unique immovable events may include ash falls, distinctive turbidite beds, isotopic spike events, and identified paleomagnetic reversals. Determinations of isotopic composition or petrophysical rock properties may yield nonunique but nearly continuous series of values. They are ideally suited to slotting matrices. Thompson & Clark (1993) illustrate the process by slotting the isotopic composition of individual cores against a global standard isotope curve. They experiment with a constraint on the least-cost pathway that limits the length of runs of consecutive events from the row or column sequences. Because they were dealing with closely spaced isotope analyses, runs of events from the global standard can be interpreted as hiatuses in the core. Thus, the constraint on runs represents hypotheses concerning the likely scales of stratigraphic incompleteness.

Many other physical and chemical events with potential value for correlation have only a binary or uniform intrinsic identity. Paleomagnetic reversals, for example, have binary identity, switching either from normal to reversed or from reversed to normal polarity. Isotopic stage boundaries are also binary; the isotopic

ratio increases or decreases. Stratigraphic sequence boundaries are intrinsically uniform, each one perhaps representing the same kind of sea-level event. Such events correlate to one another, sometimes on a one-to-many basis, but the correlative events must be recognized by extrinsic evidence from the rocks in the intervals that separate them. Such events might be slotted. The extrinsic evidence from the adjacent intervals provides the dissimilarity measures. For binary events, half the cells in the slotting matrix will be out of bounds except when sandwiched within horizontal or vertical runs that represent one-to-many correlations, such as those that arise when a pair of polarity reversals is missing in one section. Logically, the slotting matrix for binary events should permit diagonal steps; it becomes an aligning matrix (Figure 8) because the similarities are determined by the intervals themselves rather than the boundary events.

The inclusion of binary or uniform events in graphic correlation is easy to imagine, but rarely undertaken. All the possible pairwise correlations of these events generate a grid of possible points on the LOC. Simply imagine the corresponding immutable matrix, rotated 90° counter-clockwise, scaled to the thickness of the stratigraphic sections, and superimposed on the graphic plot. Other matters of fit being resolved, the LOC should then choose among accessible points in the grid according to the lowest dissimilarity values.

Similarly, it is possible to imagine a blend of the immutable matrix information into permutable sequences. Every local sequence of events is constrained to be included in the grand sequence in its locally observed order. Each locally ordered set of events may be arbitrarily interleaved with other events. Retention of the local immutable sequences negates many pairwise mutations of the trial sequences.




Figure 8 Immutable interval-aligning matrix. In contrast to Figure 7, the matrix rows and columns now represent contiguous intervals in two sections (a). The lowest-cost path is permitted to make diagonal steps (a). The two sequences are aligned in a fashion that matches intervals and allows one-to-many alignments (b). The two are correlated, not interleaved.

When the misfit is determined for each trial sequence, every binary or uniform event will add a misfit increment determined by its dissimilarity with the nearest events of the same kind from a different section. After a solution is found in this way, it may be desirable to recode the most robust adjacencies as correlative marker beds and repeat the optimization process. Notice that this suggests an alternative strategy in which a whole series of possible correlations of binary and uniform events is coded in turn and run separately to determine which yields the lowest misfit for the paleobiological data.

The second strategy is tedious, but manageable with existing programs. The first strategy has not yet been implemented. It complicates the computational bookkeeping and introduces a vexing question: How to assign weights to the importance of dissimilarity values relative to the measures of paleontological misfit. A good answer is hard to find because the two measures have no obvious dimensional equivalence. But the answer surely determines the relative influence of paleontological and physical events and, thus, the outcome of the search. This nuisance reemphasizes a critical point. The selection of different properties to define a good fit between the solution and the data must be expected to change the solution to a correlation problem. The possible measures of fit are limited by the nature of the data and in turn limit the possible methods of solution. Methods that permit different combinations of measures of fit must be expected to yield a range of different answers.

RADIOMETRIC CALIBRATION

Pairs of radiometrically dated horizons have the special property that their relative age is known independently of any superpositional evidence; they do not need to be found in the same section. In the best circumstances, they are tied to uniquely identifiable marker horizons that need not be dated everywhere. Confidently calibrated paleomagnetic and isotopic boundaries may also be treated as dated marker horizons. This strategy very effectively constrains the sequencing task for late Cenozoic marine correlations. For very favorable marine Neogene instances of the correlation problem, calibrated physical events and measured sections may be sufficiently numerous that the ages of many FADs and LADs can be independently estimated prior to correlation.

For Mesozoic and Paleozoic problems, however, relatively few radiometric dates must be used to interpolate the ages of a much larger number of paleontologic events. Again, the data determine how this is done. If the data are range charts from measured sections, then the dated events are used from the outset in the sequencing process. In graphic correlation (Figure 3*a,b*) and in immutable matrices, pairs of dated events delineate areas that are out of bounds for the LOC or the least-dissimilarity pathway. Dated events are also included from the outset by inverse methods that permute trial sequences (Cooper & Sadler 2003); mutations of the trial sequences that would reverse the order of a pair of dated

events are disallowed. Because first- and last-appearance events may only be adjusted down- or up-section, respectively, their age may be fixed relative to some dated events above or below them in the same section. Thus, the addition of dated events from fossiliferous sections may reduce the number of permissible sequence mutations for FADs and LADs and significantly accelerate the trial and error process.

If the data are isolated faunas, then the dated events relate to individual faunas. Either the fossiliferous stratum is dated directly or dated events above or below the fauna provide lower or upper bounds on the age of the fauna. The sequence and spacing of the faunas is estimated using scores derived from the matrix permutation process, without recourse to estimates of age. This proxy timescale of scores can then be calibrated by regression against a true timescale using the faunas for which age estimates are available (Alroy 1994, 2000).

SOME CURRENT PROGRAMS

It is now appropriate to offer a few remarks about the different strengths of some currently available software for stratigraphic correlation. All provide reproducibility and an explicit, rigorous numerical basis to the otherwise rather subjective exercise of correlation. All implement simple rules of optimization, but solve correlation problems that are larger than is practical to attempt manually. All offer the opportunity to perform sensitivity analyses by omitting different observations and examining the impact on the solution. Any of these algorithms may be expected to improve resolving power up to tenfold relative to traditional methods that discard many events (Cooper et al. 2001, Sadler & Cooper 2003). All are likely to produce sequences of events that are at odds, in some detail or other, with subjective expectations; when this happens, users should ask what information or assumptions guide the expectations but are not available to the algorithms.

GRAPHCOR (Hood 1986) automates many aspects of graphic correlation. It requires data in measured stratigraphic sections (Edwards 1982) and solves the sequencing and the scaling tasks simultaneously. This may be the best program with which to learn numerical correlation. It shares with manual graphic correlation the advantage of transparency; i.e., it is the easiest method to monitor at every step and is therefore the least likely to let a mistake pass unnoticed.

RASC (Agterberg & Nel 1982, Agterberg 1990, Agterberg & Gradstein 1999) is a fast implementation of permutable all-event matrices to find average ranges by probabilistic ranking. Speed and reliability are achieved by limiting the kinds of data that can be used; events must be observed in sufficient sections to generate multiple local estimates of pairwise ordering. Solutions include statistical estimates of their quality. The scaling task is solved after the sequencing task.

BIOGRAPH (Savary & Guex 1991, 1999; Guex 1991) is well suited to data that are dominated by isolated faunas. It uses graph theory to solve the permutation

of coexistence matrices. Information from measured sections provides polarity to the sequence of faunal associations determined from coexistence data.

CONJUNCT (Alroy 1994) is suited to the same types of information as BIOGRAPH. It optimizes not only coexistence matrices but also FAD-LAD matrices.

CONOP9 (Sadler 2003, Sadler et al. 2003) implements sequence permutations with the simulated annealing heuristic. Designed originally as a multidimensional implementation of graphic correlation, it now offers great flexibility to incorporate different data types, including isolated faunas. The measure of fitness may be changed to mimic all the methods based upon permutable matrices. The price of flexibility is slower execution than the more narrowly focused programs. Separate algorithms in the same program solve the scaling task.

PCSLOT (Clark 1985, 1995) determines optimal paths across immutable matrices and permits constraints. Although developed to correlate sequences of pollen assemblages, it has much wider applicability.

PAST (Hammer et al. 2003) is a versatile paleontologic utility program that implements versions of BIOGRAPH, RASC, and CONOP with a common data manager.

THE REMAINING CHALLENGES

Paleobiologists can extract considerable information about the phylogenetic sequence of taxa by analyzing the morphology of fossils, without recourse to stratigraphic information. But these insights do not yet aid the correlation task as much as they might. To date, more effort has been committed to questions concerning the place of stratigraphic information in cladistic analyses of morphology than to the possibility that the resulting cladograms provide independent evidence of sequence that can improve biostratigraphy. It is already possible to make better use of the biological relationships of fossil species as determined by cladistic analyses. Cladograms can reveal the order of FADs that best fits the morphologic information. Pairwise ordering of FADs derived from cladistic analyses can easily be added to the stratigraphic entries in a permutable event matrix. Ordered strings of FADs extracted from a cladogram branch can provide additional constraints for permutable sequences.

Two looming challenges posed by modern stratigraphic databases may not easily be reconciled. First, it is desirable to integrate more data types into a single method. Every method, regardless of the data to which it is suited, must seek a sequence of events. Consequently, the best way to suit all the data is to invert the problem, working through a suite of permutable sequences and achieving iterative improvements as judged by the fit between the sequences and the data. But the second challenge is to manage considerably larger data sets without loss of speed. The flexibility of the inverse approach sacrifices speed. The fastest algorithms are those that are tailored to specific data types and work forward from the data to the best solution.

A third challenge concerns the problem of reworked fossils. Large integrated databases will combine taxa that are prone to reworking with those that are not. Decisions about the likelihood of reworking, or the most palatable assumptions concerning reworking, currently force a dichotomous choice between methods that seek maximal ranges and those that seek probable ranges. No method yet embodies a satisfactory theory of reworking that can obviate this unfortunate choice. As in the past, answers to all these challenges might be discovered by recognizing analogies with problems in other disciplines and adapting their numerical methods.

ACKNOWLEDGMENTS

The development of CONOP9 and experiments to incorporate cladograms into biostratigraphy have been supported by NSF awards EAR 9219731 and EAR 9980372.

**The Annual Review of Earth and Planetary Science is online at
<http://earth.annualreviews.org>**

LITERATURE CITED

- Agterberg FP. 1990. Automated stratigraphic correlation. *Dev. Palaeontol. Stratigr.* 13:1–424
- Agterberg FP, Gradstein FM. 1999. The RASC method for ranking scaling of biostratigraphic events. *Earth Sci. Rev.* 46:1–25
- Agterberg FP, Nel LD. 1982. Algorithms for the ranking of stratigraphic events. *Comput. Geosci.* 8:69–90
- Alroy J. 1992. Conjunction among taxonomic distributions the Miocene mammalian biochronology of the Great Plains. *Paleobiology* 18:326–43
- Alroy J. 1994. Appearance event ordination: a new biochronological method. *Paleobiology* 20:191–207
- Alroy J. 2000. New methods for quantifying macroevolutionary patterns processes. *Paleobiology* 26:707–33
- Blackham M. 1998. The unitary association method of relative dating its application to archaeological data. *J. Archaeol. Method Theory* 5:165–207
- Clark RM. 1985. A fortran program for constrained sequence-slotting based on minimum combined path length. *Comput. Geosci.* 11:605–17
- Clark RM. 1995. Depth-matching using PC-SLOT version 1.6. *News. INQUA Work. Group Data-Handling Methods* 13. <http://kv.geo.uu.se/inqua/news12/n13-mc.htm>
- Cooper RA, Crampton JS, Raine JJ, Gradstein FM, Morgans HEG, et al. 2001. Quantitative biostratigraphy of the Taranaki Basin, New Zealand: a deterministic probabilistic approach. *Bull. Am. Assoc. Pet. Geol.* 85:1469–98
- Cooper RA, Sadler PM. 2004. The Ordovician period. See Gradstein et al. 2004, in press
- Cooper RA, Sadler PM, Melchin MJ. 2004. The Silurian period. See Gradstein et al. 2004, in press
- Dell RF, Kemple WG, Tovey CA. 1992. Heuristically solving the stratigraphic correlation problem. *Proc. Ist Industr. Engr. Res. Conf.* 1:293–97
- Edwards LE. 1982. Quantitative biostratigraphy: the methods should suit the data. In *Quantitative Stratigraphic Correlation*.

- ed. JM Cubitt, RA Reyment, pp. 45–60. Chichester, UK: Wiley
- Gradstein FM, Agterberg FP. 1998. Uncertainty in stratigraphic correlation. In *Sequence Stratigraphy: Concepts Applications*, ed. FM Gradstein, KO Sandvik, NJ Milton, *Norw. Pet. Soc. Spec. Publ.* 8:9–29. Amsterdam: Elsevier
- Gradstein FM, Ogg J, Smith A, eds. 2004. *A Geologic Time Scale*. Cambridge, UK: Cambridge Univ. Press. In press
- Gordon AD, Clark RM, Thompson R. 1988. The use of constraints in sequence slotting. In *Data Analysis and Informatics*, ed. E Diddy, pp. 353–64. Amsterdam: North Holland
- Gordon AD, Reyment RA. 1979. Slotting of borehole sequences. *Math. Geol.* 11:309–27
- Guex J. 1977. Une nouvelle méthode d'analyse biochronologique. *Bull. Lab. Géol. Univ. Lausanne* 224:309–22
- Guex J. 1991. *Biochronological Correlations*. Berlin: Springer. 252 pp.
- Hammer O, Harper DAT, Ryan PD. 2003. PAST—PAlaeontological STatistics, version 1.16. <http://folk.uio.no/ohammer/past/>
- Harries P, ed. 2003. *High Resolution Approaches in Stratigraphic Paleontology*. Dordrecht: Kluwer Acad. Publ. 474 pp.
- Hawkins DM, Merriam DF. 1974. Zonation of multivariate sequences of digitized geologic data. *Math. Geol.* 6:263–69
- Hood KC. 1986. *GRAPHCOR—Interactive graphic correlation software*. Version 2.2, copyright 1986–1995, KC Hood
- Kemple WG, Sadler PM, Strauss DJ. 1995. Extending graphic correlation to many dimensions: stratigraphic correlation as constrained optimization. See Mann & Lane 1995b, pp. 65–82
- Macleod N, Sadler PM. 1995. Estimating the line of correlation. See Mann & Lane 1995b, pp. 51–64
- Mann KO, Lane HR. 1995a. Graphic correlation: a powerful stratigraphic technique comes of age. See Mann & Lane 1995b, pp. 3–13
- Mann KO, Lane HR, eds. 1995b. Graphic correlation. *Spec. Publ. Soc. Econ. Paleontol. Mineral.* 53. 263 pp.
- Pels B, Keizer JJ, Young R. 1996. Automated biostratigraphic correlation of palynological records on the basis of shapes of pollen curves evaluation of next-best solutions. *Palaeogeogr. Palaeoclimatol. Palaeoecol.* 124:17–37
- Ryan PD, Ryan MDC, Harper DAT. 1999. A new approach to seriation. In *Numerical Palaeobiology: Computer-Based Modeling Analysis of Fossils and Their Distributions*, ed. DAT Harper, pp. 433–49. Chichester, UK: Wiley
- Sadler PM. 2001. Best-fit intervals: measures of the resolving power of paleontological first- last-appearance events. *Paleobios* 21(S1):111 (Abstr.)
- Sadler PM. 2003. *Constrained optimization approaches to the paleobiologic correlation seriation problems: a users' guide and reference manual to the Conop family of programs*. Version 7.0. Copyright 1998–2003, PM Sadler
- Sadler PM, Cooper RA. 2003. Best-fit intervals consensus sequences: comparison of the resolving power of traditional biostratigraphy and computer-assisted correlation. See Harries 2003, pp. 49–94
- Sadler PM, Kemple WG, Kooser MA. 2003. CONOP9 programs for solving the stratigraphic correlation and seriation problems as constrained optimization. See Harries 2003, pp. 461–65
- Savary J, Guex J. 1991. Biograph: un nouveau programme de construction des correlations biochronologique basées sur les associations unitaires. *Bull. Lab. Géol. Univ. Lausanne* 313:317–40
- Savary J, Guex J. 1999. Discrete biochronological scales unitary association: description of the BioGraph computer program. *Mém. Géol. Lausanne* 34:1–281
- Shaw AB. 1964. *Time in Stratigraphy*. New York: McGraw Hill. 365 pp.
- Thompson R, Clark RM. 1993. Quantitative marine sediment core matching using a modified sequence-slotting algorithm. In *High*

- Resolution Stratigraphy*, ed. EA Hailwood, RB Kidd, pp. 39–49. London: Geol. Soc. London
- Tipper JC. 1988. Techniques for quantitative stratigraphic correlations: a review annotated bibliography. *Geol. Mag.* 125:475–94
- Webster M, Sadler PM, Kooser MA, Fowler E. 2003. Combining stratigraphic sections museum collections to increase biostratigraphic resolution: application to Lower Cambrian trilobites from southern California. See Harries 2003, pp. 95–128
- Zhang T, Plotnick RE. 2001. Determining the line of correlation using genetic algorithms. *Geol. Soc. Am. Abstr. Prog.* 33(6):141 (Abstr.)