Definition of Homogeneous Segments on the Basis of Condition Data: General Approaches and Specific Application to Rail Maintenance

ROEMER M. ALFELOR AND SUE MCNEIL

Automated systems for acquiring sequential data on the condition of many transportation facilities generate a huge volume of data to be processed. Aggregation of automatically collected and processed data is required in order to make the information appropriate and useful for maintenance management. The problem of data aggregation and techniques that can be adopted to divide linear structures into homogeneous segments for modeling deterioration and assigning maintenance actions are described. The objective is to use the disaggregate sequential condition data from a data acquisition system to determine points of transition from one homogeneous segment to another. The theory and general methodologies are explained in the context of transportation facilities, and an application to maintenance-based aggregation of rail surface condition data is described.

The deteriorated state of the nation's infrastructure has received considerable attention in the past decade (*1,2*). New technologies have been developed to quickly and consistently provide measures of facility condition to enable better quantification of the extent and degree of deterioration and provide an accurate assessment of the actual condition so that funds may be allocated and transportation facilities repaired or rehabilitated effectively. Because transportation facilities such as highways, pipelines, and railroads are linear in structure and extend continuously over several miles, such state-of-the-art automated systems result in large volumes of condition data. Not only does the volume of data create problems for physical storage despite recent advances in storage media, but the data also become incomprehensible.

One way to reduce the quantity of data is to use samples. However, population data are desirable because deterioration does not occur uniformly over the entire length of facilities. Spatial variations in condition occur as a result of variations in traffic loadings, construction quality, subsurface conditions, curvature, and environmental factors to which the facilities are exposed.

An alternative solution is to aggregate the data into physically homogeneous segments. As condition data are used to develop deterioration models and identify defective sections or segments that require maintenance, replacement, or rehabilitation, the aggregation procedure can account for physical constraints and economies of scale in the maintenance process. Maintenance, replacement, and rehabilitation are intended to ensure safe operations and adequate performance of the facility.

R. M. Alfelor, The Urban Institute, Transportation Studies Program, 2100 M St., N.W., Washington, D.C. 20037. S. McNeil, Department of Civil Engineering, Carnegie Mellon University, Schenley Park, Pittsburgh, Pa. 15213.

In this paper, the concepts and parameters associated with data aggregation are discussed for the two main objectives of data collection: modeling deterioration and planning maintenance. A review of existing approaches to data aggregation is presented, and general steps for maintenance-based aggregation are proposed. A hypothetical example that demonstrates the proposed procedures is also presented. A specific railroad application for determining grinding strategies based on rail surface data is also discussed.

AUTOMATED CONDITION ASSESSMENT

Automation of data collection and condition assessment for transportation facilities has taken a major step during the last decade. For example, pavement condition data can now be obtained using a wide variety of technologies such as video cameras, laser range finders, acoustics, infrared, and many others (*3*). Application of these technologies to highway management has potential for increased productivity and reduced costs. The processed data are used to identify the appropriate maintenance strategies and develop models of pavement deterioration, which could be incorporated in pavement management systems.

Similarly, existing technologies for obtaining railroad defect data are designed to acquire rail flaws including geometry, wear, corrugation, internal cracks, and transverse profile irregularities. This information is used to make decisions on grinding, rail and plug replacement, lubrication, and rail relays. Recently, railroads have experienced unusually high instances of fatigue defects on the surface of the rail. These surface flaws, consisting of cracks, spalls, nicks, slivers, and batters, have been found to increase the dynamic impacts of the wheels on the rail, thereby reducing its service life and requiring premature replacement (*4*). Grinding usually eliminates these surface defects. The growing problem with surface defects becomes apparent as more and more railroads resort to aggressive grinding strategies to remove the flaws. Grinding now constitutes a sizeable amount of rail maintenance budgets (*5*).

In most cases, sensors are used to collect the data at high speed. These data are processed and interpreted (either in real-time or at a later date) using signal or image processing to provide condition measures such as amount of head wear on the rail or percentage of pavement covered by alligator cracking. These automatically collected and processed condition data are then used to determine maintenance strategies and understand patterns of deterioration. Because condition data are collected and processed sequentially and are disaggregative, an incomprehensible volume of data is available. For example, every digitized image of the rail surface corresponding to approximately 1 ft in length requires around 200 kilobytes of disk space. Therefore, it is necessary to combine or cluster these data into longer segments having the same level of deterioration or requiring the same level of maintenance in order to be useful. The process is called data aggregation.

DATA AGGREGATION: DEFINITION OF HOMOGENEOUS SEGMENTS

The nature and intended use of the data influence the complexity and approach to the problem of data aggregation. The following issues must be addressed:

- The expected use of the data will determine the method of aggregating the condition data. If the data are used for estimating deterioration models, it is desirable to have segments of similar lengths to prevent aggregation bias in the estimation. If the data are used for specifying maintenance decisions, minimum or maximum lengths may be specified. For example, machines that grind the rail must grind a minimum length. In contrast, if internal defect rates for rail are at such a high level that relay is more economic than plug replacement, a maximum length has been specified for plug replacement.
- The aggregation problem should be able to treat condition data as either deterministic or stochastic. Data collection and processing introduces error into the data (*6*), and the aggregation method should be able to take this uncertainty into consideration. However, analytical procedures that include the impacts of uncertainty in the aggregation process and overall life-cycle cost computations are not trivial.
- The complexity of the problem is also influenced by the dimension of the condition data that require aggregation. In some situations it may be necessary to use a vector of condition measures. For example, the decision to remove a portion of rail in one location and move it to another is based on the amount of wear or loss in cross-sectional area and the number of fatigue defects detected. In the case of highway pavements, resurfacing and overlay are often based on the amount of cracking and rutting on the surface of the road. The approaches for aggregating segments based on multivariate condition indicators are much more involved than those based on single condition measures.
- The complexity of the problem increases with the number of categories of maintenance activities into which the aggregate segments are classified.

These issues are discussed in the following sections.

Aggregated data may be used for modeling deterioration or planning maintenance. In some cases, however, aggregation for both purposes is done by predefining segments that have uniform attributes other than condition. Such attributes include alignment, construction, and traffic loading, which tend to have abrupt changes but remain constant for a significant length. The aggregate condition and maintenance needs of each segment are subsequently determined. This is a special case of aggregation, and although it provides a convenient approach to assigning maintenance or modeling deterioration, it is not efficient because the type and severity of defects within a segment are not homogeneous because of the effects of other variables not taken into consideration. The aggregation problem considered in this research is based on condition and not parameters such as alignment or traffic.

There is also the issue of sensor limitations. For example, an optical rail profile measuring system produces a digitized rail profile at 20 to 30 ft intervals (*7*). This is the minimum unit of measurement for aggregation, even if the rail profile is theoretically nonuniform within the 20 to 30 ft interval. Maintenance and deterioration modeling will then have to be made at least at this level of aggregation. The same is true for measuring rutting or cross profile, roughness, and longitudinal profile of pavements, which cannot be done continuously using existing sensors. However, video cameras can obtain continuous images of the rail or pavement surface, which usually results in disaggregate surface condition data.

Deterioration models are often estimated on a unit length basis (i.e., number of cracks per foot, number of internal defects per mile). This means that, for example, for a frame-by-frame video condition data, each frame represents an individual sample and the deterioration model should be estimated at this level of disaggregation. However, at this level, it may be difficult to relate condition to causal variables, such as traffic or construction. A more typical deterioration model will be based on continuous data, such as percentage defective or a more defect-specific measure. This may be derived from binary defect/nondefect frame-by-frame condition data, which can be aggregated into longer unit segments and the condition data for each unit length calculated as the percentage of frames defective or the sum of continuous data for individual frames. This aggregate condition data can be used to model deterioration as a function of construction, maintenance, and traffic parameters. It is clear that the unit length of segment required to model deterioration may be arbitrarily chosen.

However, if segments must be categorized as either good or bad for defining problematic segments, considering nonuniform traffic, curvature, construction, and maintenance, aggregate condition data based on fixed lengths is no better than aggregate condition data for predetermined segments. For this purpose, the aggregation procedure requires clustering the data samples into segments with or without minimum length requirements, and the segments need not be of uniform length. The problem becomes one of identifying specific locations on the facility where the condition changes from good to bad or from one level or severity of deterioration to another and then finding the explanatory variables to which the condition can be attributed. The aggregation procedure becomes similar to that required for assigning maintenance.

On the basis of maintenance decision-making, the data points or samples can be categorized according to the total percentage of defect (or whatever extent measure), without regard to the type of surface defect. The objective of categorization in this case is to put different segments (made up of adjacent samples) in different categories corresponding to specific types and levels of maintenance. Hence, one does not make a distinction between a spall and a shell that extend with the same depth beyond the rail surface.

The maintenance-based categorization is different from deterioration modeling because there are constraints on the minimum length of a segment. For example, rail maintenance strategies include grinding, plug replacement, rail relay, and rail replacement. Each type of maintenance is appropriate for a particular extent of defect and length of defective segment. In the case of rail grinding, a typical grinding train will grind at least 500 ft of rail surface defects as suggested by one railroad. This is a constraint on the minimum length of rail segments to grind, and therefore, the data should be aggregated to lengths of at least this much. Rail surface grinding is the focus of this research, but aggregation is also a relevant issue for other types of condition data, maintenance, and track renewal.

APPROACHES TO DATA AGGREGATION

Approaches to definition of homogeneous segments may be classified by (a) the way the raw data are analyzed, (b) the type of condition data, or (c) the number of maintenance categories associated with the aggregate data. Based on the way the data are analyzed, aggregation that does not use predefined segments of fixed lengths can be classified simply as either on-line or off-line. An on-line technique sequentially analyzes data from one end to the other end, and transition points between homogeneous segments are determined one by one using a window of fixed or varying size. Off-line search (also called retrospective search) looks at the entire data set and determines the transition points simultaneously. Depending on the length of data to be processed, on-line search has the advantage of not having to deal with a large set of data in a single step. However, it may not lead to the best solutions. The type of condition data aggregated may be binary or multivariate, categorical or continuous, and deterministic or stochastic. The type of data has significant bearing on the complexity of the aggregation procedures. Aggregation of multivariate categorical stochastic data is the most difficult to implement. Finally, the number of categories associated with the data refers to the number of maintenance actions into which the aggregated data are assigned. For example, one might be interested in rail grinding, relay, and plug replacement as maintenance activities.

Existing approaches to aggregation of condition data illustrate the diverse nature of the process. Three aggregation procedures, namely rule-based aggregation (on-line), control chart (on-line), and change-point analysis (on-line or off-line) are described next.

Rule-Based Aggregation Method

A rule-based aggregation method is an on-line procedure through which the transition point is determined by comparing running averages or other derived measures with a given threshold or set of thresholds. These thresholds may be derived analytically or arbitrarily. For example, application of a rule-based system to aggregation of rail condition data was implemented to determine homogeneous wear segments from transverse profiles of the rail from an optical rail wear measurement system (*8*). The algorithm uses a technique whereby a window of specified fixed length is moved along the track and the homogeneous segments are determined from the running averages and standard deviations of rail wear. The bases for combining segments are the differences in the averages and deviations of the wear measurements for adjacent segments. These thresholds are set arbitrarily. The final results consist of segments classified on the basis of wear as either requiring rail relay or not. Rule-based aggregation such as this is convenient for problems for which the thresholds can be determined by expert judgment or analysis or based on common practice. However, in many aggregation problems, the objective is not simply to satisfy the rules but to maximize or minimize a performance function, which is usually an economic measure (e.g., benefits or costs).

Control Charts

Control charts are commonly used in assembly lines to check product quality (*9*). It is an on-line process based on the concept that each product sample should conform to certain specifications in terms of measurable properties. It is assumed that somewhere along the manufacturing process, the quality of production deteriorates and products become rejectable. Using control charts, the point in time at which the process becomes out of control is determined by looking at the distribution of the sample properties with time.

To put the control-chart analysis in the context of the highway or railroad, each section of a specified length might represent one sample. The vector of parameters describing each section consists of the following:

- Defect measure (e.g., number of cracks),
- Curvature,
- Construction quality,
- Previous maintenance, and
- Traffic loading.

Assuming several miles of road that are homogeneous in terms of the above parameters except the defect measure, control charts can be used to determine which sections have experienced unusually high frequencies of surface defects and consequently greater-than-expected deterioration. The idea behind a control chart is that the transition from acceptable to unacceptable samples is gradual. In other words, the transition point does not correspond to an abrupt jump in quality. This may not be the case for many transportation facilities because defects tend to occur randomly and may be found in isolated locations.

Change-Point Analysis

In statistics, problems of inferring transition points in a sequence of data are referred to as change-point problems (*10*). This procedure differs from control charts because the changes are abrupt as opposed to gradual. The simplest illustration of the change-point problem is that of determining one change-point from a series of observations. Change-point analysis can be off-line or on-line, depending on whether the entire segment is analyzed for all change-points (off-line) or a running window is used in the search (on-line).

To illustrate, let $y_1, y_2, y_3 \ldots y_n$ be the series of n observations. Assuming only one change-point at location r, $1 < r < n$, the best location of this change-point is determined by calculating the propensities of the change-point being at all possible locations. The location for which the propensity is maximum represents the most likely location of the change-point. This procedure is equivalent to the maximum likelihood technique or Bayes' ratios for analyzing pairwise change-point models (*11*).

The propensity referred to in the change-point problem could be the likelihood of condition or any other performance function like cost impacts of the aggregate data. The problem of clustering frames into longer segments for rail surface defect data involves many change-points because the sample may consist of miles and miles of track. Literature on change-point analysis states that it is virtually impossible to solve problems involving more than one change-point (*12*). The change-point analysis can also be formulated as a mathematical optimization problem using an objective function, which is explained later.

Change-point analysis and its mathematical programming formulation is a complex optimization procedure, and the solution becomes prohibitive for large problems with many possible change-points. However, it provides the most optimal aggregation strategy if solution procedures allow.

GENERAL STEPS IN AGGREGATING CONDITION DATA

The following steps are proposed for maintenance-based aggregation of condition data.

1. Determine costs and life-cycle impacts of maintenance. These include costs of not performing maintenance or allowing the facility to deteriorate and fail prematurely.
2. Determine the threshold condition value that warrants maintenance. This value may be obtained from maintenance experts on the basis of current standards. Because administrators have different approaches to doing maintenance, the threshold values will likewise vary. An alternative approach is to determine these thresholds analytically using economic analysis. It is apparent that these maintenance standards will dictate the type and measure of condition data to collect.
3. Identify the constraints on maintenance. These constraints may be equipment-related (as the case of minimum rail grinding length of 500 ft because of the length of the grinding train) or system-related (uniformity in the alignment of the road).
4. Collect and process sequential defect data and, if necessary, apply smoothing algorithm to remove noise.
5. Use the threshold values (maintenance standards) described in Step 2 to identify the type of maintenance to use for each data point corresponding to a unit length of facility. This step does not take into consideration the constraints identified in Step 3.
6. Apply decision analysis to the entire facility to satisfy the constraints while minimizing the cost of incorrect maintenance decisions made or maximizing the benefits derived from maintenance. This is an optimization problem that can be solved analytically using a variety of techniques including heuristics and integer programming formulations. The complexity of the solution procedures depends on the type of constraints and the number of maintenance strategies, which are related to the measures of condition. Obviously, the longer the facility being analyzed, the longer it takes to arrive at a solution.

The following generic example will help illustrate the above procedures. It is a more general problem than the one that will be solved later. Consider the hypothetical conditions stated next, which may have resulted from application of Steps 1–3.

- The thresholds for percent defective for doing maintenance to each data sample are as follows:

M_1 = percent defect < 30 percent,
M_2 = 30 percent ≤ percent defect ≤ 50 percent,
M_3 = 50 percent ≤ percent defect ≤ 65 percent, and
M_4 = 65 percent < percent defect.

- M_1 corresponds to a particular maintenance action associated with a range of percentage area defective. The level of maintenance is a function of the severity of the defect.
- M_2 uses an equipment whose configuration will correct at least 20 adjacent samples. Therefore, if one sample is maintained by M_2, 20 adjacent samples will be affected and subjected to the same maintenance.
- M_4 requires replacing the sample, but the minimum length for replacement is 50 ft (governed by the standard lengths of replacements).
- M_1 is the do-nothing alternative, and M_3 is a spot-maintenance activity that can be performed on individual samples.

Hypothetical condition data are shown in Figure 1 for every unit of collected data (Step 4). Each data point is subject to error, assuming the data collection technology and the data processing techniques are both imperfect. Noise-removal algorithms in the form of data smoothing eliminate some of these imperfections (Step 4). The idea behind a smoothing algorithm is that sequential data points are spatially dependent. For instance, autocorrelation of pavement condition data is discussed elsewhere (*13*). In the case of rail defects, it has been reported that defects tend to cluster at some locations (*14*). Smoothing not only accounts for minor errors in technology or measurement but also makes the clustering of data into homogeneous segments easier. There are many ways of smoothing data (*15*), but one must be careful in using them because they can smooth out abrupt changes in condition, which may be important in understanding the behavior of the facility.

Figure 2 shows the same data points after applying the 3R running median smoothing procedure, which determines the running median of data points taken three at a time (*15*). The continuous curve in Figure 2 is then used in the succeeding steps. Given the ranges of percentage defective (percent of defect) for which the defined maintenance activities are applied, the entire facility is divided into different segments

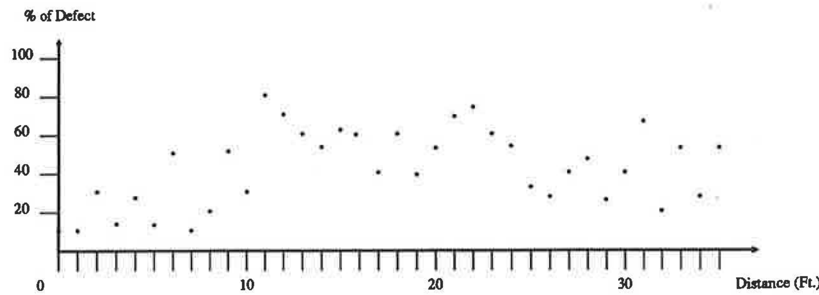

FIGURE 1 Hypothetical raw defect data.

corresponding to different maintenance levels as shown in Figure 3 (Step 5). However, the constraints imposed by M_2 and M_4 may not be satisfied by the segmentation shown in Figure 3. Hence, segments will have to be combined. There are many possible ways of doing this; Figure 4 shows one possible combination (Step 6). The objective is to come up with the best segmentation given a criterion and the set of constraints. The criterion may be to minimize the total cost of performing and not performing maintenance on the different segments of the facility. This is an optimization/decision analysis problem that is equivalent to determining the changepoints in a series of observations.

This example can be used to represent the problem of aggregating rail condition data for different maintenance actions; where M_2, M_3, and M_4 may correspond to light grinding, corrective grinding, and plug replacement, respectively, and the constraints on length are changed to their more realistic values.

The general steps for maintenance-based aggregation illustrated in the generic example were adopted in formulating the simpler and more specific problem of using rail surface condition data for one maintenance action: corrective grinding. A description of the rail-grinding problem and the procedures adopted for aggregating rail surface condition data for the purpose of doing maintenance are presented in the following example.

AGGREGATING RAIL SURFACE CONDITION DATA

A prototype optical system using video camera and image processing subroutines was developed at Carnegie Mellon University to acquire continuous images of the rail surface (*16*). The system can identify the presence or absence of rail surface defects on a frame-by-frame basis. An automated defect recognition system was also developed that can process and classify continuous images of the rail surface (*17*). Initial rail surface data were obtained on a 10-mi track and processed in the laboratory using the automated recognition system.

The sequential data obtained from the inspection system were then aggregated to identify segments that needed to be ground. Each frame sample of video data corresponds to approximately 1 ft of actual railhead. The standard identified for performing maintenance is that only defective frames or samples should be ground. However, there is a minimum grinding length that is dictated by the length of the grinding train. This means that a minimum length will be ground. This information is used, along with the costs of incorrect grinding decisions, in formulating the problem as a set-packing optimization problem (*18*).

A set-packing problem is an integer programming problem in which an optimal combination of feasible subsets of a choice set is sought such that an objective is maximized and that each element of the choice set belongs to at most one of the subsets chosen. In the case of the rail grinding problem, the choice set consists of a sequence of video (frame) data, and a subset represents a group or pack of adjacent frames that will either be ground or not. Within each pack, individual frames can be either defective or not, and the higher the proportion of defective frames in a pack chosen for grinding, the better the grinding decision is in terms on maximizing the correct decisions.

The rail grinding problem is a set-packing problem because of the constraint on the minimum grinding length. If such constraint does not exist, then the optimal solution is just to

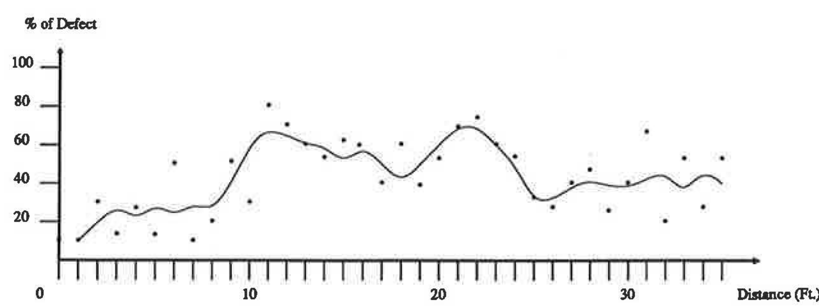

FIGURE 2 Approximate 3R running median smooth.

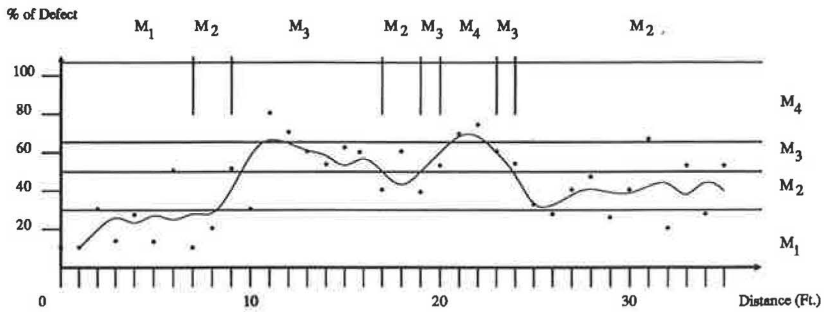

FIGURE 3 Categorized segments based on maintenance level.

grind all the defective frames. However, the constraint becomes a parameter for the minimum size of a pack. Hence, each pack becomes a feasible grinding subset, and the optimal combination of feasible grinding subsets, which maximizes the total benefits due to grinding, is sought. Two aspects of this problem make it different from the conventional set-packing problems. First, it is much, much larger in dimension because of the number of packing combinations and the number of frames involved (again, about 1 ft of rail per frame). Second, the problem has a structure in which each pack is a segment of sequential data.

Solutions exist for solving the set-packing integer programming problem, but then again the dimensions of those problems that have been encountered in operations research literature are much less than the dimension involved in determining rail-grinding strategies.

To solve the problems realistically, two rule-based heuristic algorithms were developed that utilize the structure of the problem. The heuristic solutions can be described as on-line rule-based techniques for solving the set-packing problem. A window that slides from one end to the other end of the track is used in analyzing feasible packs that represent segments of sequential condition data. As the window moves along the rail, segments to grind are determined locally using the criterion for grinding, which is the minimum grinding length, and the minimum threshold for the proportion of defective frames in the pack, which is determined analytically.

The heuristic procedures differ from the general set-packing optimization formulation in that the latter is an off-line procedure whereas the former use a threshold value for the average percentage defective of segments to grind. This threshold value is calculated on the basis of comparisons of cost of grinding and not grinding defective and nondefective rail segments. Therefore, instead of minimizing the overall life-cycle costs of maintenance for the entire track, the algorithms seek on-line local solutions as the window moves along the track. Like most heuristic solution methods, this procedure does not guarantee optimal results. However, although the results obtained by using the heuristics are not necessarily optimal, they are more efficient than the optimal approaches resulting from solving the set-packing problem with respect to execution time and perform reasonably well compared with greedy solutions. Also, the existing approaches for solving set-packing problems are not guaranteed to converge to a solution.

For a lengthy discussion on the aggregation of rail surface condition data, the formulation of the grinding problem as a set-packing problem, and the description as well as application of the heuristic algorithms, the reader is referred to other work by Alfelor (*17*). The effects of uncertainty were also considered in the analysis. The heuristic procedures for rail surface data aggregation were modified to account for uncertainty. Indeed, the solutions to the aggregation problem are shown to be sensitive to uncertainty and imperfect information about the condition of the rail.

CONCLUSION

Maintenance of linear structures requires definition of pieces or segments that represent uniform condition and hence require a particular form of maintenance. Automated inspection of these structures provide the data needed in defining homogeneous segments.

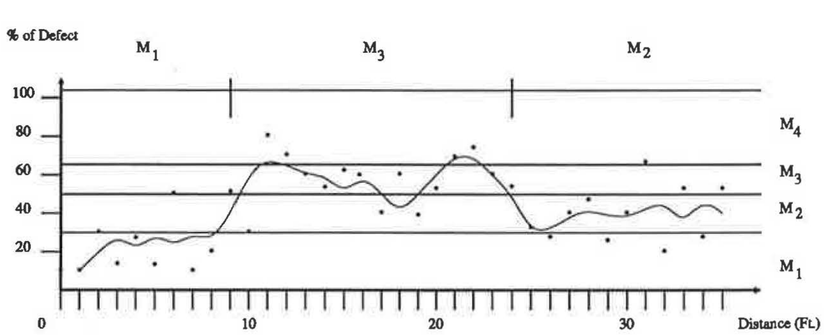

FIGURE 4 Small segments combined.

This research provided a detailed analysis of the problem of defining homogeneous segments for modeling deterioration and assigning maintenance. The general solution procedures are described for any type of linear facilities. An example is illustrated in the case of rail maintenance planning, specifically rail grinding.

Further research and illustrations are necessary for solving the same problem when the condition data are continuous and when the number of categories (maintenance actions) to which the segments are assigned is increased, which is true for most transportation facilities. The same can be said about dealing with the uncertainty in data collection. Moreover, the impacts of making incorrect decisions in terms of overall life-cycle cost and performance of the facility need to be explored.

ACKNOWLEDGMENT

This work was supported by an award from the National Science Foundation.

REFERENCES

1. National Council on Public Works Improvement. *Fragile Foundation: A Report on America's Public Works.* U.S. Government Documents, Washington D.C., 1988.
2. Science, Information and Natural Resources Division, Office of Technology Assessment. *Delivering the Goods: Public Works Technologies, Management and Finance.* Special Report, U.S. Congress, April 1991.
3. D. H. Mendelsohn. Automated Pavement Crack Detection: An Assessment of Leading Technologies. In *Proc., Second North American Conference on Managing Pavements*, Ministry of Transportation, Ontario; FHWA, U.S. Department of Transportation, Nov. 1987.
4. A. M. Zarembski. The Impact of Rail Surface Defects. In *Railway Track and Structures,* Nov. 1984.
5. W. Glavin. *Rail Grinding—The BN Experience.* American Railway Engineering Association, Washington, D.C., 1989.
6. F. Humplick. Theory and Methods of Analyzing Infrastructure Inspection Output: Application to Highway Pavement Surface Condition Evaluation. Ph.D. thesis. Massachusetts Institute of Technology, Cambridge, 1989.
7. C. Martland, J. Sussman, and S. McNeil. *Analytic Techniques for More Effective Rail Scheduling.* Summary Report. Burlington Northern Railroad Company, 1989.
8. D. Acharya. Using Expert System and Network Optimization Technique for Rail Relay Scheduling. Ph.D. thesis. Massachusetts Institute of Technology, Cambridge, 1990.
9. R. V. Hogg and J. Ledolter. *Engineering Statistics.* Macmillan Publishing Company, New York, N.Y., 1987.
10. A. F. Smith. A Bayesian Approach to Inference about a Change-Point in a Sequence of Random Variables. *Biometrika,* Vol. 62, No. 2, 1975.
11. A. F. Smith. *Change-Point Problems: Approaches and Applications.* University of Nottingham, England, 1975.
12. S. D. Silvey. The Lindisfarne Scribes' Problem. In *Journal of the Royal Statistical Society B, 20,* 1957.
13. P. Ullidtz. *Pavement Analysis.* Elsevier Publishing Company, New York, N.Y., 1987.
14. D. D. Davis and T. A. Gudiness. *Rail Planning: A Progress Report of the Track Maintenance Research Committee.* Technical Report WP-128. Association of American Railroads, Washington, D.C., 1986.
15. J. W. Tukey. *Exploratory Data Analysis.* Addison-Wesley, Reading, Mass., 1977.
16. S. McNeil, B. Motazed, R. Alfelor, T. Short, and J. Eshelby. Automated Collection and Detection of Rail Surface Defects. In *1991 International Heavy-Haul Workshop: Maintenance of Way and Train Operations—Balancing the Conflicts*, International Heavy-Haul Association, 1991.
17. R. M. Alfelor. Analysis of Automatically Collected Rail Surface Defect Data for Rail Maintenance Management. Ph.D. thesis. Carnegie Mellon University, Pittsburgh, Pa., 1991.
18. M. W. Padberg. Covering, Packing and Knapsack Problems. *Annals of Discrete Mathematics, 4,* 1979.

Publication of this paper sponsored by Committee on Railway Maintenance.