

CLASSIFICATION ANALYSIS FOR SIMULATION OF MACHINE BREAKDOWNS

Lanting Lu

School of Mathematics
 University of Southampton
 Southampton, SO17 1BJ, U.K.

Russell C.H. Cheng

School of Mathematics
 University of Southampton
 Southampton, SO17 1BJ, U.K.

Christine S.M. Currie

School of Mathematics
 University of Southampton
 Southampton, SO17 1BJ, U.K.

John Ladbrook

Dunton Technical Center
 Ford Motor Company
 Laindon, Essex, SS15 6EE, U.K.

ABSTRACT

Machine failure is often an important factor in throughput of manufacturing systems. To simplify the inputs to the simulation model for complex machining and assembly lines, we have derived the Arrows classification method to group similar machines, where one model can be used to describe the breakdown times for all of the machines in the group and breakdown times of machines can be represented by finite mixture model distributions. The Two-Sample Cramér-von Mises statistic is used to measure the similarity of two sets of data. We evaluate the classification procedure by comparing the throughput of a simulation model when run with mixture models fitted to individual machine breakdown times; mixture models fitted to group breakdown times; and raw data. Details of the methods and results of the grouping processes will be presented, and will be demonstrated using an example.

1 INTRODUCTION

Ford uses simulation modeling to evaluate new designs for assembly and machining lines and to improve the efficiency of existing lines. Our work has focused on an existing assembly line within a Ford manufacturing plant, where nearly 100 different machines are involved in the assembly process. The breakdown times of many machines follow similar distributions and so, to reduce the number of input distributions and increase the data available for fitting each of these distributions, we propose a method of grouping machines based on the breakdown data available. The

grouping is such that two machines with statistically significantly different breakdown times data cannot be placed in the same group. We then fit finite mixture distributions to the grouped breakdown data.

Finite mixture models are particularly flexible and useful, and [McLachlan and Peel \(2000\)](#) describe their structure and applications. A finite mixture model has probability density function,

$$h(x) = \sum_{i=1}^k w_i f_i(x|\theta_i), \quad (1)$$

where

$$0 < w_i \leq 1 \text{ for } i = 1, \dots, k \quad (2)$$

and

$$\sum_{i=1}^k w_i = 1 \quad (3)$$

are the weights of the components whose individual densities are $f_i(x|\theta_i)$ for $i = 1, \dots, k$. The parameter k is the number of components in the finite mixture model. We assume that all of the component densities take the same form, i.e.

$$f_i(x|\theta_i) = f(x|\theta_i), i = 1, \dots, k. \quad (4)$$

It is assumed that none of the θ_i nor the number of components k are known in the model. Fitting such models is therefore a non-standard statistical problem. Some of the issues this raises are given in [Cheng \(1998\)](#) and the

procedure we use for fitting finite mixture models to data is described in [Cheng and Currie \(2003\)](#).

Other methods exist for fitting data that are not distributed according to standard stochastic models. These are generally based on using flexible families of distributions, such as the Bézier distribution ([Wagner and Wilson 1996](#), [Wagner and Wilson 1993](#), [Law 2006](#)) or the Johnson family (see Chapter 12 of [Johnson and Kotz 1970](#), or page 297 of [Law 2006](#)). The Bézier distribution exploits the properties of Bézier curves and allows the modeler to fit the cumulative distribution function $F(x)$ to a wide range of distributions of data, its flexibility being due in part to the fact that the number of parameters to be used is not fixed. Johnson distributions are based on transformations of normal variables and, although they offer a wide range of shapes of distributions, do not cope as well with multimodality. The advantage of the use of finite mixture models is that they provide a good description of multimodal data, using parameters that have an intuitive meaning. They are also easy to implement in most standard simulation packages using a two-stage approach, where the component is sampled in the first step and then a random number is sampled from the component density.

In Section 2 we describe the details of the classification of the machines by their breakdown data, including a description of the Arrows classification method. This has similarities with cluster analysis and a comparison with the cluster analysis method is also given. We evaluate the classification using a simulation model of an assembly line with individual machine breakdown input distributions and group breakdown input distributions, comparing the output with that generated using raw data as the input to the model. This is described in Section 3. Finally, we conclude and discuss potential future work in Section 4.

2 CLASSIFICATION OF THE MACHINES

In order to simplify the inputs to the simulation model, we aim to classify the machines involved in the assembly line into a smaller number of groups, based on their breakdown data, such that one mixture model distribution can be used to describe the distribution of breakdown times for each group of machines.

We measure the similarity of the breakdown data of any two machines using the Two-Sample Cramér-von Mises goodness-of-fit statistic ([Anderson 1962](#)). Bootstrapping is used to determine the p-values of the statistics, which are stored in the similarity matrix. The similarity matrix is then used to classify the machines into groups. We have derived a method for classifying the data from the similarity matrix, termed the Arrows classification method, but Cluster Analysis ([Anderberg 1973](#)) could also be used. In this section, we describe the estimation of the similarity matrix and how it is used to classify the machines.

2.1 Estimating the Similarity Matrix

The similarity matrix we produce is made up of the p-values of every pair of the machines, where the p-value gives the probability that the breakdown data for the two machines are drawn from the same distribution. The p-values are estimated using bootstrapping as described below, following the calculation of the Two-Sample Cramér-von Mises goodness-of-fit statistic ([Anderson 1962](#)).

We wish to measure the similarity of two samples of breakdown data (x_1, x_2, \dots, x_n) , and (y_1, y_2, \dots, y_m) for machines X and Y respectively, and so obtain a measure of the goodness-of-fit. We consider using the Anderson-Darling statistic ([Stephens 1974](#)) but this requires some information about the hypothesized distribution in order to calculate the goodness-of-fit ([Stephens 1974](#)). Thus, we use the Cramér-von Mises test. The Cramér-von Mises T criterion for testing that the two samples come from the same unspecified continuous distribution is

$$T = (nm/(n+m)) \int_{-\infty}^{\infty} (F_n(x) - G_m(x))^2 dH_{n+m}(x), \quad (5)$$

where $F_n(x)$ is the *empirical distribution function* (EDF) of the first sample; that is, $F_n(x) = (\text{no. of } x_i \leq x)/n$; $G_m(x)$ is the EDF of the second sample and $H_{n+m}(x)$ is the EDF of the two samples together; that is, $(n+m)H_{n+m}(x) = nF_n(x) + mG_m(x)$.

Let r_i and s_j be the ranks in the pooled sample of the ordered observations of the two samples X and Y, respectively, where $i = 1, 2, \dots, n$ and $j = 1, 2, \dots, m$. Then

$$F_n(x) - G_m(x) = i/n - (r_i - i)/m \quad (6)$$

at the i th x -observation and

$$F_n(x) - G_m(x) = (s_j - j)/n - j/m \quad (7)$$

at the j th y -observation. Thus we can write the criterion T as

$$T = \frac{U}{nm(n+m)} - \frac{4nm-1}{6(n+m)}, \quad (8)$$

where

$$U = n \sum_{i=1}^n (r_i - i)^2 + m \sum_{j=1}^m (s_j - j)^2. \quad (9)$$

To test the null hypothesis that the two samples are drawn from the same distribution, all of the observations are ordered, the ranks $r_1 < r_2 < \dots < r_n$ of the n observations from the first sample and the ranks $s_1 < s_2 < \dots < s_m$ of the m observations from the second sample are then determined and U is computed. If T is too large, we reject

the null hypothesis, that the samples are drawn from the same distribution.

Tabulated criterion values are not very extensive and do not cover the samples that we are dealing with and so we use bootstrapping to determine the distribution of T , $\Phi(T)$, and hence the significance level. To carry out the bootstrapping, let $Z = (z_1, z_2, \dots, z_{n+m})$ be the pooled sample of breakdown data from machines X and Y . The EDF of Z is denoted by $H_{n+m}(z)$. We generate two samples of size n and m from the original pooled set of observations, Z , with replacement and call this our *bootstrap sample*, written as $X^* = (x_1^*, x_2^*, \dots, x_n^*)$ and $Y^* = (y_1^*, y_2^*, \dots, y_m^*)$. By comparing X^* and Y^* , we can calculate the Cramér-von Mises statistic T^* for the bootstrap sample. In order to estimate $\Phi(T)$, we generate a number, B , pairs of bootstrap samples from Z : $(X^{*1}, Y^{*1}), (X^{*2}, Y^{*2}), \dots, (X^{*B}, Y^{*B})$ and calculate the statistic T^{*j} for each pair of samples. The EDF of the sample $T^* = (T^{*1}, T^{*2}, \dots, T^{*B})$ is then written as

$$\Phi_B(T) = \frac{(\text{no. of } T^{*j} \leq T)}{B} \quad (10)$$

The bootstrap distribution $\Phi_B(T)$ will converge to $\Phi(T)$ with probability one as B tends to infinity (Cheng 2001) and we use $\Phi_B(T)$ as our estimate of $\Phi(T)$. (We consider $B = 5000$ as a large enough number for our bootstrap analysis.)

The Bootstrapping Process is

For $j = 1$ to B

For $i = 1$ to n

Draw x_i^{*j} from Z (with replacement)

Next i

For $i = 1$ to m

Draw y_i^{*j} from Z (with replacement)

Next i

Calculate T^{*j} by comparing X^{*j} with Y^{*j}

Next j

Form the EDF of T^* , $\Phi_B(T)$.

The p-value describing the fit of data from machine M_i to data from machine M_j is then obtained from $\Phi_B(T)$. This procedure is carried out for all pairs of machines to form the similarity matrix. The whole process is illustrated in Figure 1.

For classification purposes we shall use a specific threshold, p_0 say. If a p-value is less than p_0 we consider the data from the two machines are significantly different. If the p-value is greater than p_0 we may accept the performance of the two machines as being the same. Thus, for example, if we set $p_0 = 0.10$, then, Figure 2 shows that the p-value corresponding to T is smaller than 0.10, which means that the data from the two machines being compared are significantly different and have not been drawn from the same distribution. In contrast, Figure 3 shows that the p-value of T is over 0.90, which means that the data from the two machines being compared can be assumed to have

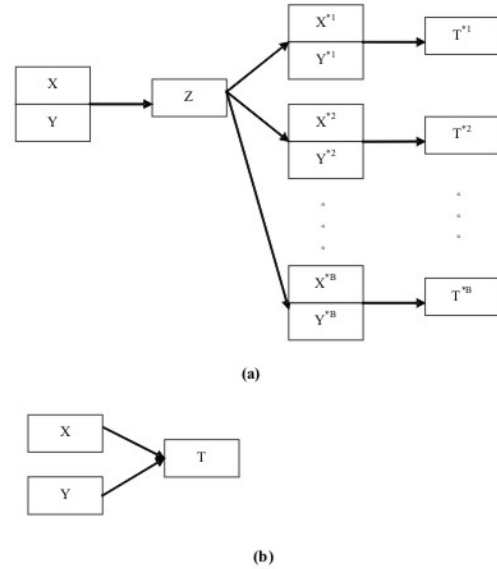


Figure 1: (a) The bootstrapping process used to determine the null distribution of T , $\Phi(T)$, and (b) the evaluation of the Cramér-von-Mises statistic for the original samples, which is compared with $\Phi(T)$ to determine the p-value for the similarity of the two machines.

been drawn from the same distribution, with a probability of more than 0.90.





Figure 2: M_1 vs. M_2 , $p_{12} < 0.10$.

2.2 Classification

We consider the p-value introduced in Section 2.1 as a measure of the similarity of the breakdown data of any two machines. Therefore, the matrix of p-values is denoted the similarity matrix and forms the basis of the classification process.

The Arrows classification method that we describe in Section 2.2.1 has similarities with complete-linkage hier-

Figure 3: M_1 vs. M_2 , $p_{12} > 0.90$.

archical cluster analysis (Anderberg 1973). The cluster analysis algorithm proceeds iteratively, combining the two most similar machines or groups of machines at each iteration, where the distance between any two groups is defined to be the greatest distance (in this case, the smallest p-value) from any member of one group to any member of the other group. By selecting a threshold distance it is therefore possible to ensure that all of the machines in a group have significant similarities. The use of a threshold distance or similarity is also a characteristic of the Arrows method. Where the two methods differ is that the clustering method searches the whole matrix to find the most similar groups to merge while the Arrows method aims to keep together machines, which have what we term a *one-arrow connection*. Two machines have a one-arrow connection if the data from one machine is most similar to the data from the second machine, i.e. machine i and machine j have a one-arrow connection if p_{ij} , the p-value comparing machine M_i with machine M_j , is the biggest in only *one* of row i or row j , but not both. The Arrows classification method allows us to more easily control the similarity level in the resultant groups than cluster analysis and is described further in the following section.

2.2.1 Arrows Classification Method

We currently use $p_0 = 0.10$ as a threshold p-value for assuming that two sets of data are similar enough to be grouped together. If the p-value for the fit between the breakdown data of a pair of machines is greater than 0.10 then they can be put in the same group; otherwise, the data are assumed to be significantly different. Increasing the p_0 to, e.g. 0.20, may increase the average similarities within groups but may also increase the number of groups.

We follow the process below to determine the groups.

1. Search the similarity matrix,

- (a) If p_{ij} , the p-value comparing machine M_i with machine M_j , is the biggest in both row i and row j , we say M_i and M_j have a *double-arrow connection* and put machine i and machine j into one group.
- (b) If p_{ij} , is the biggest in only one of row i or row j , we say M_i and M_j have a *one-arrow connection*, and put machine i and machine j into one group.

Search the similarity matrix, until no more new machine groupings can be made.

2. Process the p-values of the remaining pairs of machines in each group, e.g. for group C_a ; if the values are all bigger than p_0 , the threshold, keep C_a ; for pairs with p-values less than 0.10, use the following decision process to determine which machine in the pair to keep.

- (a) If p_{ij} is the biggest in column j or column i or both, keep M_i and M_j in C_a
- (b) Take out the machine(s) with the weakest connection with the others until there is no pair of machine with p-value under 0.10 in C_a , where the strength of a connection of a machine is measured by the average of all of the p-values of this machine compared with the rest of the machines in C_a .
- (c) Thus, we have one group C_a .

3. Try to combine some of the primary groups. If all pairs of machines in primary group C_a and primary group C_b have p-values bigger than 0.10, combine these two groups. If we can combine C_a and C_b or combine C_a and C_d , compare the average of the sum of the p_{ik} , where M_i is in C_a and M_k is in C_b , with the average of the sum of the p_{il} , where M_i is in C_a and M_l is in C_d . Combine the two groups with stronger connections. Search until all of the primary groups are processed.

The above classification procedure has been implemented in Visual Basic for Applications. Although this method has been devised to classify machines, it could be used to classify other data. It has already proved to be very useful in grouping hospital length of stay data by specialty.

2.3 Example

We illustrate the classification method using an example of 20 machines involved in one of Ford's assembly lines. The similarity matrix of the breakdown data from the 20 machines is given in Table 1.

For these 20 machines the classification process proceeds as follows (machines are denoted by their ID):

Table 1: Actual Similarity Matrix of the breakdown data from the 20 machines.

	M01	M02	M03	M04	M05	M06	M07	M08	M09	M10	M11	M12	M13	M14	M15	M16	M17	M18	M19	M20
M01	—	0.02	0.00	0.00	0.09	0.00	0.02	0.00	0.09	0.12	0.55	0.52	0.08	0.00	0.02	0.00	0.06	0.01	0.16	0.12
M02	0.02	—	0.62	0.00	0.02	0.00	0.00	0.00	0.11	0.07	0.02	0.03	0.01	0.00	0.17	0.00	0.00	0.01	0.02	0.08
M03	0.00	0.62	—	0.00	0.02	0.00	0.00	0.00	0.02	0.05	0.00	0.00	0.00	0.00	0.15	0.00	0.00	0.00	0.04	0.06
M04	0.00	0.00	0.00	—	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
M05	0.09	0.02	0.02	0.00	—	0.00	0.04	0.01	0.01	0.23	0.15	0.20	0.58	0.00	0.18	0.00	0.16	0.01	0.82	0.90
M06	0.00	0.00	0.00	0.00	0.00	—	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.00
M07	0.02	0.00	0.00	0.00	0.04	0.00	—	0.17	0.00	0.32	0.00	0.00	0.32	0.01	0.00	0.00	0.80	0.00	0.19	0.03
M08	0.00	0.00	0.00	0.00	0.01	0.00	0.17	—	0.00	0.32	0.00	0.00	0.22	0.00	0.00	0.00	0.59	0.00	0.33	0.01
M09	0.09	0.11	0.02	0.00	0.01	0.00	0.00	0.00	—	0.03	0.06	0.06	0.01	0.00	0.01	0.00	0.00	0.00	0.02	0.00
M10	0.12	0.07	0.05	0.00	0.23	0.01	0.32	0.32	0.03	—	0.28	0.25	0.85	0.05	0.15	0.03	0.25	0.08	0.55	0.39
M11	0.55	0.02	0.00	0.00	0.15	0.00	0.00	0.00	0.06	0.28	—	0.36	0.25	0.00	0.03	0.00	0.04	0.00	0.38	0.23
M12	0.52	0.03	0.00	0.00	0.20	0.00	0.00	0.00	0.06	0.25	0.36	—	0.19	0.00	0.09	0.00	0.01	0.00	0.33	0.32
M13	0.08	0.01	0.00	0.00	0.58	0.00	0.32	0.22	0.01	0.85	0.25	0.19	—	0.02	0.04	0.00	0.57	0.02	0.55	0.46
M14	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.05	0.00	0.00	0.02	—	0.00	0.00	0.37	0.00	0.05	0.00
M15	0.02	0.17	0.15	0.00	0.18	0.00	0.00	0.00	0.01	0.15	0.03	0.09	0.04	0.00	—	0.00	0.00	0.08	0.37	0.41
M16	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.03	0.00	0.00	0.00	0.00	0.00	—	0.09	0.00	0.01	0.00
M17	0.06	0.00	0.00	0.00	0.16	0.00	0.80	0.59	0.00	0.25	0.04	0.01	0.57	0.37	0.00	0.09	—	0.00	0.35	0.09
M18	0.01	0.01	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.08	0.00	0.00	0.02	0.00	0.08	0.00	0.00	—	0.45	0.10
M19	0.16	0.02	0.04	0.00	0.82	0.02	0.19	0.33	0.02	0.55	0.38	0.33	0.55	0.05	0.37	0.01	0.35	0.45	—	0.73
M20	0.12	0.08	0.06	0.00	0.90	0.00	0.03	0.01	0.00	0.39	0.23	0.32	0.46	0.00	0.41	0.00	0.09	0.10	0.73	—

1. Step 1 (see Figure 4)
Form 8 groups based on identifying the single and double arrow connections.
2. Step 2 (see Figure 4)
Identify 6 pairs of machines initially grouped together that have significantly different breakdown data ($M09, M03$; $M18, M05$; $M18, M20$; $M18, M15$; $M07, M14$; $M08, M14$).
Therefore, machines $M09, M03, M18, M05, M20, M15, M07, M14, M08$ may be taken out of the existing groups. In order to decide which machine or machines should be removed to make sure there are no red curves, the following steps are considered.
 - (a) As the priority is to keep pairs of machines that have double-arrow connections (see Section 2.2.1), $M03, M05, M20$ and $M07$ should be kept in their groups. As a result, we can decide to remove $M09$ to make sure there is no red curve in the second group as shown in Figure 4.
 - (b) $M18$ and $M14$ have the weakest connection with the others in the third and fourth group respectively as shown in Figure 4, thus $M18$ and $M14$ should be taken out of their groups.
 - (c) The resultant groups have no red curves, i.e. no pairs of machines with p-values less than p_0 .
3. Step 3 (see Figure 5)
Combine the fourth and fifth group as no pair of members are significantly different.

Finally 10 groups are obtained. With the complete clustering method, the same 10 groups are found if merging is not allowed for p-values < 0.10 .

3 EVALUATION OF THE CLASSIFICATION METHOD

3.1 Breakdown Input for Simulation Model

We carry out a simulation to evaluate the classification of the machines by analyzing the simulation output of an assembly line made up of the 20 machines described in the example above, using three different methods for generating breakdown times. We use WITNESS (Lanner Group) to simulate the assembly line, and use their “BUSY TIME” mode to describe the occurrence of breakdowns. This only allows machines to breakdown when they are working. A Negative Exponential distribution is used to simulate the Time Between Failure (Ladbroke 1998), where we calculate the Mean Time Between Failure (MTBF) for a machine to be

$$MTBF = \frac{TT - TTR}{\text{No. of Failures}}, \quad (11)$$

where TT is the time period over which our raw breakdown times data are collected, and TTR is the total time a machine is broken down for during the data collection period. To calculate TTR , we split the data into n bins, with thresholds b_1, b_2, \dots, b_n , and so

$$TTR = \sum_{i=1}^{n-1} \frac{(b_{i+1} - b_i)}{2} F_i, \quad (12)$$

where $F_i, i = 1, \dots, n$ is the number of observations in bin i .

The three methods we consider for sampling breakdown times are: (1) sampling from raw historical data; (2) sampling from a different mixture model distribution for each machine; (3) sampling from the fitted mixture models for the groups of machines found in Section 2.3, where the group fitted mixture models is obtained by fitting finite mixture distributions to the breakdown data of all of the machines in the group.

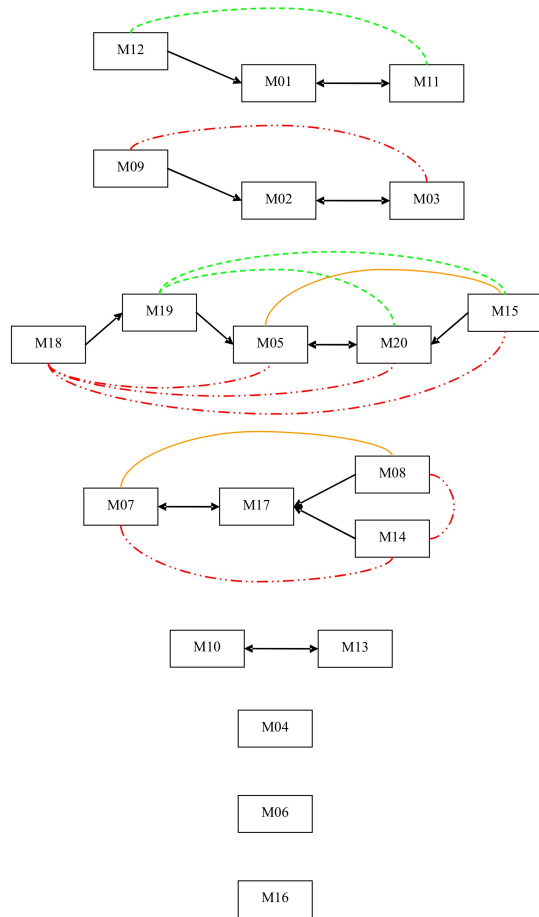


Figure 4: Steps 1 and 2 from the Example, showing groups with double-arrow and one-arrow connections and the strength of the connections within each group. Red curve (---): p-value of the two connected machines is not significant; Yellow curve (—): p-value of the two connected machines is on the borderline; Green curve (- - -): p-value of the two connected machines is significant.

3.2 Output Evaluation

We set the warm up period to be 1440 minutes using the Time-Series Inspection method (Robinson 2004) and run 100 replications for each of the three different models, with a results collection period of 55 days. In order to evaluate the whole line's behavior, we take the simulation output to be the number of parts being shipped in those 55 days. Figure 6 shows the number of parts shipped for each of the three methods used for sampling breakdown times.

As we can see from the plot, the inter-quartile ranges for the three methods overlap, showing a degree of similarity between the methods. However, at the 95% significance level we find that the outputs are different. Figure 7 shows the histogram of the raw breakdown time data and the

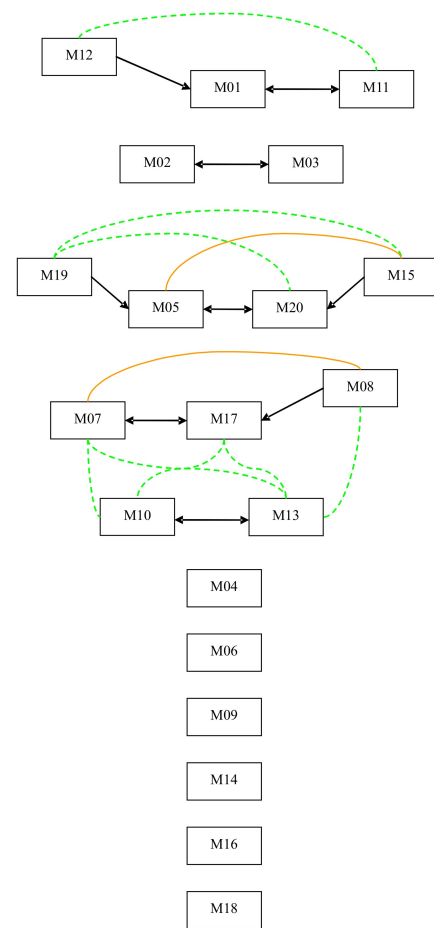


Figure 5: Step 3 of the Example in which we try to combine the primary groups without Red curve connections.




Figure 6: Boxplot of the number of parts shipped per 55 days using the three methods for sampling breakdown times. The central line shows the medium number and the box describes the inter-quartile range.

fitted mixture model for machine *M01*. The histogram has some high value outliers for which the breakdown time is extremely long, resulting in the whole assembly line being

down for a long period. The fitted mixture model is much smoother. Moreover, the two models are using the same distribution to simulate time between failures, so the fact that the model using the raw data as its input has a lower efficiency is quite reasonable. This could also explain the differences between the models using individual machine input and using group input. Ongoing work is being carried out to improve the modeling of the times between failures. A data set coming from a group of machines will also have more data points than those from individual machines, and so the model fitted to the whole group's breakdown data may provide a better description of the relatively large breakdown times than the individual fitted models. As shown in the boxplot, the model using raw data is more similar to the model using group inputs than to the model using individual machine inputs.

4 CONCLUSIONS AND FUTURE WORK

The classification process described in this paper could be applied to grouping data from a wide range of applications, in addition to manufacturing. For example, we have already used it to classify hospital procedures by their length-of-stay data. The proposed Arrows classification method, which forms part of this process, produces similar results to cluster analysis, while making it much easier to control the similarity level in each resultant group. We currently use 0.10 as the threshold p-value for assuming two sets of data are similar enough to be grouped together. Increasing the threshold may increase the number of groups but also would increase the average similarities within resultant groups. The choice of threshold will be investigated further for this example by running the simulation model with groups generated by the different thresholds.

Grouping like procedures serves to increase the data available for fitting input distributions as well as simplifying the inputs required for the simulation model. We have tested the grouping method using a simulation model with three methods of sampling breakdown times: raw historical data; individual fitted finite mixture models and group fitted finite mixture models. The results show that the model output is similar for each of these three methods, justifying the use of the classification method. The current method of modeling the time between failures may be influencing the model output and improving this will be an area of future work.

We do not include the time required for setting up machines in the model nor any modeling of labor. These are likely to produce some small changes to the output of the model but should not affect the results of the comparisons we have made between the methods used to model breakdowns. Testing a more detailed assembly line model should confirm this.

ACKNOWLEDGMENTS

We thank Ford for their support, especially the Operational Research Department and Witness for their advice, especially Tony Waller.

REFERENCES

- Anderberg, M. R. 1973. *Cluster Analysis for Applications*. New York: Academic Press.
- Anderson, T. 1962. The distribution of the two-sample Cramér-von Mises criterion. *The Annals of Mathematical Statistics* 33:1148–1159.
- Cheng, R. C. H. 1998. Bayesian model selection when the number of components is unknown. In *Proceedings of the 1998 Winter Simulation Conference*, ed. D. J. Medeiros, E. F. Watson, J. S. Carson, and M. S. Manivannan, 653–659. IEEE.
- Cheng, R. C. H. 2001. Analysis of simulation output by re-sampling. *International Journal of Simulation Systems, Science & Technology* 1:51–58.
- Cheng, R. C. H., and C. Currie. 2003. Prior and candidate models in the Bayesian analysis of finite mixtures. In *Proceedings of the 2003 Winter Simulation Conference*, ed. S. Chick, P. J. Sánchez, D. Ferrin, and D. J. Morrice, 392–398. IEEE.
- Johnson, N., and S. Kotz. 1970. *Continuous Univariate Distributions - 1*, Volume 2. New York: John Wiley.
- Ladbrook, J. 1998, September. Breakdowns modelling - an inquest. Master's thesis, University of Birmingham.
- Law, A. M. 2006. *Simulation Modeling and Analysis*. 4th ed. McGraw-Hill series in industrial engineering and management science. New York: McGraw-Hill.
- McLachlan, G. J., and D. Peel. 2000. *Finite Mixture Models*. Wiley series in probability and statistics. Applied probability and statistics section. New York: Wiley.
- Robinson, S. 2004. *Simulation: the Practice of Model Development and Use*. Chichester: John Wiley.
- Stephens, M. A. 1974. EDF statistics for goodness of fit and some comparisons. *Journal of the American Statistical Association* 69:730–737.
- Wagner, M. A. F., and J. R. Wilson. 1993. Using univariate Bézier distributions to model simulation input processes. In *Proceedings of the 1993 Winter Simulation Conference*, ed. G. Evans, M. Mollaghasemi, E. Russell, and W. Biles, 365–373. IEEE.
- Wagner, M. A. F., and J. R. Wilson. 1996. Recent developments in input modeling with Bézier distributions. In *Proceedings of the 1996 Winter Simulation Conference*, ed. J. Charnes, D. Morrice, D. Brunner, and J. Swain, 1448–1456. IEEE.




Figure 7: Histogram of raw breakdown data and the fitted finite model's Probability Distribution Function for the first group of machines in the example.

AUTHOR BIOGRAPHIES

LANTING LU is a second year Ph.D student of Operational Research in the School of Mathematics in the University of Southampton. In addition, she has an MSc in Operational Research from the University of Southampton. Her research interests include mathematical modeling of machine breakdown and simulation models. Her email address is [<L.Lu@maths.soton.ac.uk>](mailto:L.Lu@maths.soton.ac.uk) and her web page is [<www.maths.soton.ac.uk/~ll/>](http://www.maths.soton.ac.uk/~ll/).

CHRISTINE S.M. CURRIE is a lecturer of Operational Research in the School of Mathematics in the University of Southampton, where she also obtained her Ph.D. In addition, she has an MPhys from Oxford University and an MSc in Operational Research from the University of Southampton. She is currently the Book Review editor for the Journal of Simulation and co-chair of the Simulation Special Interest Group in the UK Operational Research Society. Her research interests include mathematical modeling of epidemics, Bayesian statistics, revenue management, variance reduction methods and optimization of simulation models. Her email address is [<christine.currie@soton.ac.uk>](mailto:christine.currie@soton.ac.uk) and her web page is [<www.soton.ac.uk/~ccurrie>](http://www.soton.ac.uk/~ccurrie).

RUSSELL C.H. CHENG is Professor of Operational Research in the School of Mathematics in the University of Southampton. He has an M.A. and the Diploma in Mathematical Statistics from Cambridge University, England. He obtained his Ph.D from Bath University. He is a former Chairman of the U.K. Simulation Society, a Fellow of the Royal Statistical Society, Fellow of the Institute of Math-

ematics and Its Applications, Member of the Operational Research Society. His current research interests include: design and analysis of simulation experiments, variance reduction methods and parametric estimation methods and he has over seventy publications in these and related areas. He was a former Associate Editor of the ACM Transactions on Modeling and Computer Simulation and for Management Science and Editorial board member of the American Journal of Mathematical and Management Sciences. He is currently a technical board member of Simulation News Europe and is was a Joint Editor of the IMA Journal of Management Mathematics. His email address is [<rchc@maths.soton.ac.uk>](mailto:rchc@maths.soton.ac.uk) and his web page is [<www.maths.soton.ac.uk/staff/Cheng>](http://www.maths.soton.ac.uk/staff/Cheng).

JOHN LADBROOK has worked for Ford Motor Company since 1968 where his current position is Simulation Technical Specialist. In 1998 after 4 years research into modeling breakdowns he gained an M.Phil (Eng.) with the University of Birmingham. In his time at Ford, he has served his apprenticeship, worked in Thames Foundry Quality Control before training to be an Industrial Engineer. Since 1982 he has used and promoted the use of Discrete Event Simulation. In this role he has been responsible for sponsoring many projects with various universities including this project. He is also the Chairman of the WITNESS Automotive Special Interest Group. His email address is [<jladbroo@ford.com>](mailto:jladbroke@ford.com).