Designed experiments for semi-parametric models and functional data with a case-study in Tribology

David C. Woods[1], Christopher J. Marley and Susan M. Lewis

Southampton Statistical Sciences Research Institute,
University of Southampton, Southampton, SO17 1BJ, UK
[1]Corresponding author: David Woods, email: `D.Woods@southampton.ac.uk`

Abstract

Experiments with functional data are becoming ubiquitous in science and engineering with the increasing use of online monitoring and measurement. Each run of the experiment results in the observation of data points that are realised from a smooth curve. Although large quantities of data may be collected from each run, it may still only be possible to perform small experiments with a limited number of runs.

We describe statistical methodology for an example from Tribology, concerning the wear-testing of automotive lubricants. Here, we investigated how lubricant properties and process variables affected the shape of a functional response measuring wear. Novel techniques were developed for the initial design of a screening study where the levels of some of the factors could not be set directly. A two-stage semi-parametric modelling approach was applied, using a varying coefficient model and principal components. New methods for the design of follow-up experiments for such models were also developed and applied. In addition to the new methodology, we present conclusions from the case study about which factors had substantial effects, and how they influence the shape of the wear curves.

Keywords: *D*-optimality, hierarchical modelling, principal components, screening, varying coefficient modelling.

1. Introduction

In many industrial and scientific experiments, each run can now produce substantial amounts of data collected using automatic monitoring and measurement systems. Often, it can be assumed that these data are generated by a smooth underlying function (Ramsay and Silverman, 2005) and that the measurement processes are sufficiently effective for function to be accurately reconstructed. Then, the data can assumed to be functional, with the output from each run of the experiment being a smooth function, typically not having a simple parametric form. Further, these functions may vary between runs of the experiment, potentially due to aleatoric (i.e. random) variability and systematic variability which results from the application of different combinations of values of the controllable variables.

In this short paper, we use such an example from Tribology to motivate and demonstrate methodology for the design of experiments and two-stage modelling of functional data. In Section 2, the Tribology study is introduced and the experiment described. Section 3 describes our two-stage modelling strategy, and presents some

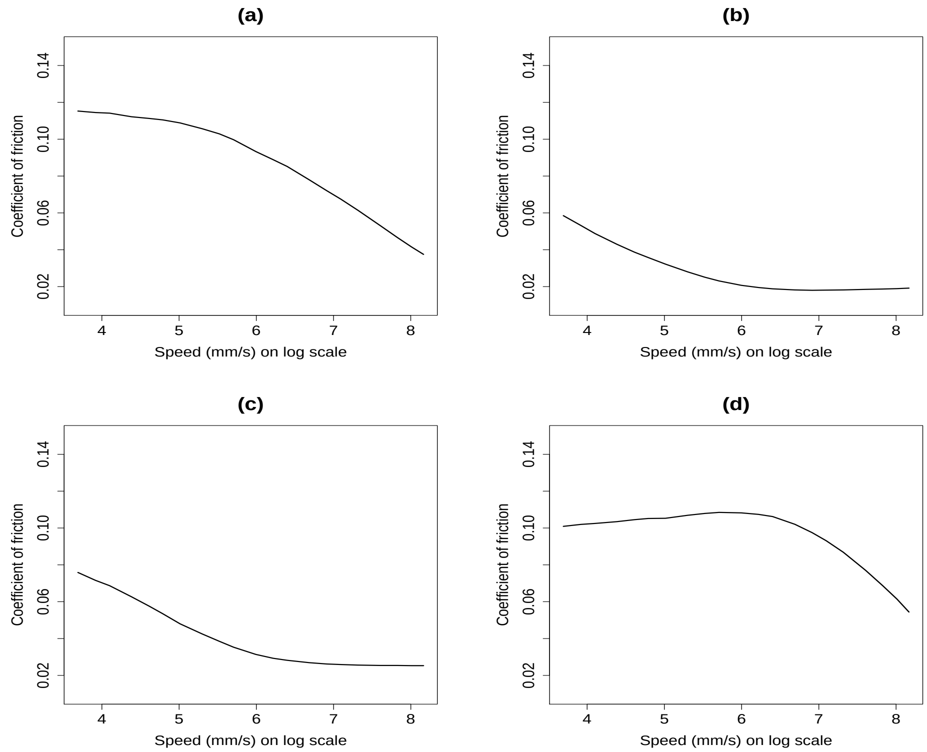

Figure 1: Four example Stribeck curves obtained using different lubricants and different process conditions

analysis of the Tribology data. Finally, we briefly discuss some issues for the design of experiments in Section 4.

2. Motivating example

Tribology is the study of interacting surfaces in relative motion. Our example experiment investigates friction between two surfaces (a ball and a disc) for a variety of different lubricant types (described via their chemical properties) and process conditions. In this study, friction is measured by the coefficient of friction as represented by a Stribeck curve (Stribeck, 1901; Dowson, 1998, pp. 344-345). For a given combination of lubricant and process conditions, this curve describes the relationship between the coefficient of friction and a dimensionless function of speed, load and viscosity. Figure 1 shows example Stribeck curves from four different runs of our experiment, where coefficient of friction is plotted against log speed. These curves are examples of functional data.

In addition to providing the coefficient of friction, Stribeck curves also help to determine when the system under study enters different lubrication regimes, and hence when different lubricant properties may be required. For example, the curves in Figure 1 (a) and (d) are in the boundary/mixed region, whilst (b) and (c) are in the mixed/elasto-hydrodynamic region. A predictive model for Stribeck curves, in terms of lubricant properties and process variables, could therefore have considerable value in the lubricant development process.

We briefly describe the modelling of data from a 20 run screening experiment performed to investigate the impact of 13 factors (11 lubricant properties and two process variables) on the shape of the Stribeck curve. The experiments were run on a Mini Traction Machine (MTM). This involved a ball rotating against a spinning disc for approximately 6 hours. Lubricating oil was placed between the ball and the disc before the start of each run and a different disc and ball surface was used for each run. The Stribeck curve was obtained by slowly reducing the speed of the disc at the end of the experiment. Measurements of the coefficient of friction were obtained automatically by the MTM for 21 speeds per run. The Stribeck curves can then be displayed as a plot of coefficient of friction against log(speed).

3. Two-stage modelling

A two-stage semi-parametric procedure was used to model the data with, at the first stage, a model for each curve being constructed in terms of principal components (PCs; see, for example, Jolliffe, 2002). At the second stage, the PC loadings were related to the lubricant properties and process variables. A similar two-stage approach has been applied to a variety of other applications, including engine mapping (Grove et al., 2004) and linguistics (Aston et al., 2010).

3.1. Stage 1: principal component models for each Stribeck curve

At the first stage, a common model is fitted to each curve. Let $\mathbf{G}$ be a 21×20 matrix containing the observed Stribeck curve data points. We used singular value decomposition to express $\mathbf{G}$ as

$$\mathbf{G} = \mathbf{U}\mathbf{Z}\mathbf{V}^T\,, \tag{1}$$

where $(1/\sqrt{20})\mathbf{U}\mathbf{Z}$ gives the principal components and $\sqrt{20}\mathbf{V}$ contains the loadings on each principal component for each run. The principal components are ordered so that they sequentially explain the highest proportion of variation in the observed Stribeck curves.

In this case, the first four components explain 98.6% of the overall variation and hence we model the Stribeck curves using linear combinations of the first four principal components. The weights used in the linear combinations are the loadings from (1).

Let $\mathbf{D}_j$ be the jth principal component (dimension 21×1) and l_{ij} be the loading for the ith run on the jth principal component. Then the model for the observed data for the ith run, $\mathbf{Y}_i$, can be written as

$$\mathbf{Y}_i = \sum_{j=1}^{4} \mathbf{D}_j l_{ij} + \mathbf{E}_i\,,$$

where $\mathbf{E}_i$ is a 21×1 vector of independent errors from $N(0, \sigma^2)$.

3.2. Stage 2: linear modelling of the PC loadings

In the second stage of the modelling procedure, we model the loadings for the principal components in terms of the variables in the experiment. We fit the following models for $i = 1, \ldots, 20$ and $j = 1, \ldots, 4$.

$$l_{ij} = \mathbf{x}_i \boldsymbol{\beta}_j + f_{ij}\,, \tag{2}$$

where $\mathbf{x}_i$ denotes the ith row of the model matrix, $\boldsymbol{\beta}_j$ is a vector of unknown coefficients to be estimated from the data and f_{ij} is the second stage error on the loading for the jth principal component for the ith run. The vector of errors for the ith run, $(f_{i1}, f_{i2}, f_{i3}, f_{i4})$, is distributed as $N(0, \Delta)$. Note that this distribution does not depend on i.

All the second stage models were fitted simultaneously using maximum likelihood estimation, allowing for the possibility of correlation between the l_{ij} and also additional within-run variation. Initially, a model containing only main effects was fitted. However, this model was deemed inadequate after studying residuals and measures of predictive quality. Hence, it was required to consider two-variable interactions using model selection techniques. We employed the Dantzig selector (Candes and Tao, 2007), which has been shown to be particularly effective for the analysis of supersaturated designs with $n < p$ (Phoa et al., 2009; Marley and Woods, 2010; Draguljić et al., 2013). Separate model selection exercises were conducted for each of the four PC loadings, and main effects were added, as necessary, to the identified models to ensure marginality and ease of interpretation. Four fitted curves from the second stage model are given in Figure 2 with approximate 95% Wald confidence intervals (Woods, 2003, ch. 4).

All four of the example fits closely match the shape of the observed Stribeck curves. The confidence intervals were judged by subject experts to be narrow enough for the model to be useful. In the vast majority of the 20 runs, the fitted model captures very accurately the shape of the curve. Discrepancies between the fitted curve and the observed curve tend to come in the form of an upward or downward shift. This is due to the nature of the two-stage fitting procedure. For example, a small error in the prediction of the loading for the first principal component will result in an upward or downward bias.

4. Design issues

Design of the initial 20 run screening experiment was complicated by 11 of the controllable variables being properties of the formulated lubricants that cannot be set independently (see also Scinto et al., 2011). The values taken by these properties are determined by the values taken by 11 different ingredients. In addition, although 20 runs of the experiment were possible, only 12 unique lubricants could be used. A

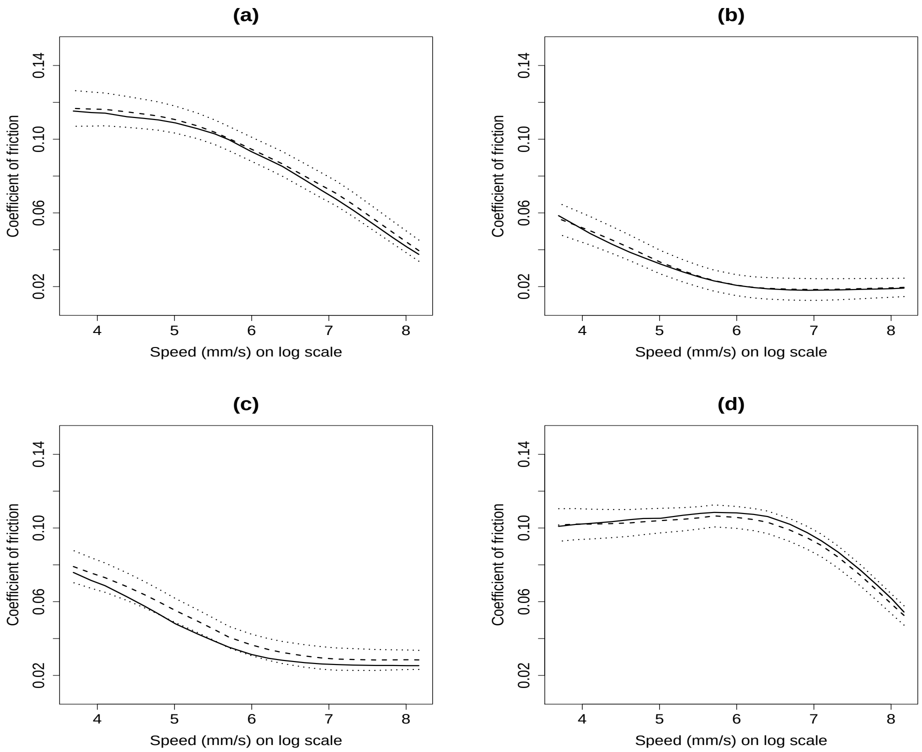

Figure 2: Four example Stribeck curves (solid line) and their second-stage fits (dashed line) with approximate 95% confidence intervals (dotted lines).

novel coordinate exchange algorithm was used to find a D-optimal design, that is, the amount of each ingredient for each run along with settings of the two process variables (load and disc roughness).

After the initial screening experiment, 10 further runs were made in order to improve the model. Marley (2010) gives further details of both designed experiments.

Acknowledgments

The first author was partly supported by a Fellowship from the UK Engineering and Physical Sciences Research Council (EPSRC). The second author was supported by a PhD studentship provided by EPSRC and The Lubrizol Corporation. The authors would like to thank Andrew Rose, Michael Gahagan, John Durham, Oliver Smith and Robert Wilkinson from Lubrizol for running the experiments and engaging in invaluable discussions.

References

Aston, J. A. D., Chiou, J.-M., and Evans, J. P. (2010), "Linguistic pitch analysis using functional principal component mixed effect models," *Journal of the Royal Statistical Society Series C*, 59, 297–317.

Candes, E. and Tao, T. (2007), "The Dantzig selector: statistical estimation when p is much larger than n," *The Annals of Statistics*, 35, 2313–2351.

Dowson, D. (1998), *History of Tribology*, London and Bury St Edmunds: Professional Engineering Publishing, 2nd ed.

Draguljić, D., Woods, D. C., Dean, A. M., Lewis, S. M., and Vine, A. E. (2013), "Screening strategies in the presence of interactions (with discussion)," *Technometrics*, in press.

Grove, D. M., Woods, D. C., and Lewis, S. M. (2004), "Multifactor B-spline mixed models in designed experiments for the engine mapping problem," *Journal of Quality Technology*, 36, 380–391.

Jolliffe, I. T. (2002), *Principal Component Analysis*, New York: Springer, 2nd ed.

Marley, C. J. (2010), "Screening experiments using supersaturated designs with application to industry," Ph.D. thesis, University of Southampton.

Marley, C. J. and Woods, D. C. (2010), "A comparison of design and model selection methods for supersaturated experiments," *Computational Statistics and Data Analysis*, 54, 3158–3167.

Phoa, F. K. H., Pan, Y. H., and Xu, H. (2009), "Analysis of supersaturated designs via the Dantzig selector," *Journal of Statistical Planning and Inference*, 139, 2362–2372.

Ramsay, J. O. and Silverman, B. W. (2005), *Functional data analysis*, New York: Springer, 2nd ed.

Scinto, P. R., Wilkinson, R. G., and Lin, D. K. J. (2011), "Screening for fuel economy: a case study of supersaturated designs in practice," *Quality Engineering*, 23, 15–25.

Stribeck, R. (1901), "Kugellager für beliebige belastungen," *Zeitschrift des Vereines Deutscher Ingenieure*, 45.

Woods, D. C. (2003), "Designing experiments for polynomial spline models," Ph.D. thesis, University of Southampton.