Modelling Simul. Mater. Sci. Eng. **8** (2000) 687–706. Printed in the UK PII: S0965-0393(00)15624-4

Accelerated learning using Gaussian process models to predict static recrystallization in an Al–Mg alloy

T J Sabin†, C A L Bailer-Jones‡ and P J Withers§

Department of Materials Science and Metallurgy, University of Cambridge, Pembroke Street, Cambridge CB2 3QZ, UK

E-mail: tanya.sabin@risoe.dk, calj@mpia-hd.mpg.de and mblsspw2@fs2.mt.umist.ac.uk

Received 13 April 2000, accepted for publication 12 July 2000

Abstract. This paper describes an investigation into the suitability of Gaussian process models for predicting the microstructure evolution arising from static recrystallization. These methods have the advantage of not requiring a prior understanding of the micromechanical processes. They are wholly empirical and use a Bayesian framework to infer the probability distribution of data, given a 'training set' comprising observed outputs for known inputs. Given the evidence from the training set, they can make a prediction and assess its certainty, taking into account the noise in the data. In addition, non-uniform deformation geometries were chosen to provide the training data, both to approximate typical manufacturing processes with complex strain paths and to investigate whether learning could be accelerated by using only a small number of test samples containing a distribution of deformation histories. The model was trained and tested on data from samples of a cold-deformed and annealed aluminium–magnesium alloy.

1. Introduction

For many years there has been great interest in trying to predict the microstructure evolution in thermo-mechanically processed metals. A number of approaches have been taken, including Avrami-based models [1, 2] and state-variable-based models [3, 4]. Both types of model have the advantage that they incorporate some physical understanding into the form of the model, but they also have a number of disadvantages. They are especially difficult to apply in situations that are either inherently very complex, such as rolling in a four stand tandem mill, or which are far from ideal, for example recrystallization in coarse-grained samples. Furthermore, they often require input parameters that are difficult to measure, such as dislocation density and subgrain misorientation.

Recently, there have been developments in empirical modelling in other fields, using neural networks and similar devices to map nonlinear patterns in multi-parametric 'training' data. Neural networks, in particular, have been employed with some success in materials science applications, such as the prediction of weld toughness in steels [5] and damage in forged composites [6]. These models adjust internal parameters to minimize differences between their predictions and the measured output values for corresponding inputs in the training data set. In this paper we evaluate an alternative method, called a Gaussian process model, which

† Present address: Materials Department, Risø National Laboratory, PO Box 49, DK-4000 Roskilde, Denmark.
‡ Present address: Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany.
§ Present address: Manchester Materials Science Centre, Grosvenor Street, Manchester, M1 7HS, UK.

0965-0393/00/050687+20$30.00 © 2000 IOP Publishing Ltd

infers a probability distribution over all of the training data and then interpolates to make predictions [7–11]. These models have the advantage that one can train them simply on the available input parameters and then examine the extent to which they are able to make reliable predictions.

Materials models generally include some parameters that must be measured by experiment. Consequently, almost all models require some degree of training. For micromechanical models, this is often to calculate the value of certain constants, but for approaches that contain no information on the form of the relationship before training, this stage is more extensive. As a result, it usually requires a much larger data set than is needed to 'tune' methods based on models with a predetermined form. Traditionally, the training data set comprises a series of samples, each of which has experienced different strains, constant strain rates and deformation temperatures. In this paper, training is based on samples that have experienced non-uniform straining conditions. This means that a single sample contains many locations that have each experienced a different deformation trajectory. By calculating the details of the strain trajectory point by point, each sample can thus provide a number of strain trajectory/microstructure data points. In principle, this increases the efficiency of data collection with respect to the amount of material, time and cost. Also, because real forming processes never apply constant and uniform strain rates throughout, these trajectories can be more representative of the strain histories actually encountered in industrial workpieces. This might lead to greater reliability, as well as open the opportunity of training microstructural models on the basis of previous forming experience, without the need for separate classical training trials.

The aim of this paper is to evaluate the Gaussian process modelling, incorporating training strategies based on sparse and noisy data of the type typically found in many applications, including metal processing. This has been carried out by studying the microstructure caused by the cold deformation and subsequent annealing of Al alloy samples. To test the efficacy of the approach, the training and testing of the microstructural model was undertaken using two different deformation geometries.

2. Experimental procedure

The material used for this study was a high-purity aluminium alloy containing 1 wt% magnesium. The alloy was cast, hot-rolled and homogenized by *Alcan Aluminium UK Ltd* so that the starting grain size for the deformation tests was approximately 1 mm. This is typically a difficult grain size for modelling purposes because large grain sizes can lead to inhomogeneous nucleation of recrystallization. This makes a statistical approach particularly apposite.

2.1. Training phase

The deformation testing that was required in order to provide the training data for the model was undertaken using cylindrical samples deformed in approximately plane strain. A large range of strains was produced by compressing long cylindrical rods of diameter 10 mm with the compression direction perpendicular to the axis of the cylinder (geometry I in figure 1(a)). This deformation geometry is essentially in plane strain since there is negligible strain along the cylinder axis. Two deformation tests of this type were performed to give approximate height reductions of 20 and 40%.

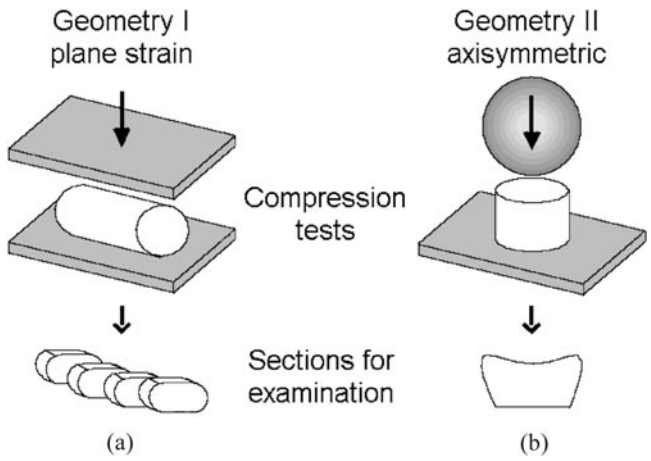

Figure 1. Geometries for the two non-uniform deformation tests.

2.2. Testing phase

It is possible to test the reliability of the trained Gaussian process on unseen deformation and annealing trajectories taken from the plane strain compression tests. A more general test of the model, however, is provided by applying it to another deformation geometry. With this aim in mind, short cylinders were deformed in axisymmetric compression with a ball-bearing (geometry II in figure 1(b)). The indenter was 20 mm in diameter and the cylinders were 15 mm in diameter with initial heights of either 7.5 or 10 mm; the indentation was to a depth of approximately 2 mm.

2.3. Heat treatment and characterization

After the deformation, samples were sectioned as indicated in figure 1. These were annealed for different lengths of time (1–60 min) and at five different temperatures (between 325 and 375 °C) to give rise to a selection of recrystallized and partially recrystallized samples. As a result of the non-uniform distribution of straining, the extent of the recrystallization and the recrystallized grain size varied as a function of the position through each sample. The samples were then polished and anodized for metallographic examination. The grain size was evaluated at several locations through each sample. These locations corresponded to areas of approximately uniform strain, as predicted by finite-element (FE) models (figure 2). For the plane strain samples, these locations are marked on figure 1(a), and symmetrically related areas were used to increase the number of measurements. For geometry II, measurements were made along bands parallel to the strain contours. In some less annealed samples, there was no recrystallization in regions of low strain.

2.4. FE modelling

Each heat-treated section contained a range of different strain histories. Provided a satisfactory FE model can be constructed, it is possible to link the observed point by point variation in the microstructure to the point by point variation in the deformation history. The FE models of the training and testing deformation geometries were generated using DEFORM^{TM} [12]. Local strain distributions predicted for each compression geometry are shown in figure 2.

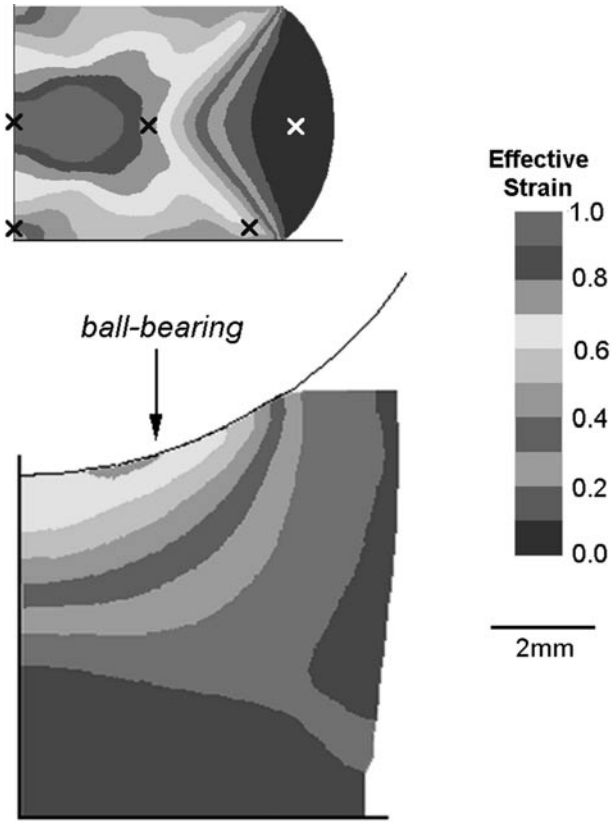

Figure 2. FE maps showing the predictions of local strain for the two deformation geometries used. For geometry I the approximate measurement locations are marked with crosses; measurements made in symmetrically equivalent positions were added to the same data point.

3. Results

As would be expected, given the non-uniform straining, a considerable variation in the extent of recrystallization and recrystallized grain size was observed in the metallographic sections following annealing. Figure 3 illustrates the progression of the recrystallization in the plane strain compressed testing and training samples annealed at 350 °C.

By combining the FE modelling with the metallographic analysis, it was possible to construct two deformation/microstructure data sets, one for each of the two deformation geometries. The recorded inputs were the local effective strain ($\varepsilon(x, y)$), annealing temperature (T) and annealing time (t); the output was the local mean *recrystallized* grain size ($d_r(x, y)$).

The plane strain data set (geometry I) comprised 64 input–output vectors ($\varepsilon(x, y)$, T, t, $d_r(x, y)$) with strains between 0.0 and 1.0 and annealing conditions as described in section 2.3. The data were obtained from 36 sections taken from the 20 and 40% deformed cylinders. A subset of these (set A), comprising 45 points from samples annealed at 325, 350 and 375 °C, was used as the basic training data set. The remaining 19 points (set B) were annealed at 335 and 360 °C. Eleven indentation tests (geometry II) produced 26 points with strains between 0.0 and 0.7 and annealing conditions within the same range as for the plane strain tests. These are referred to as set C.

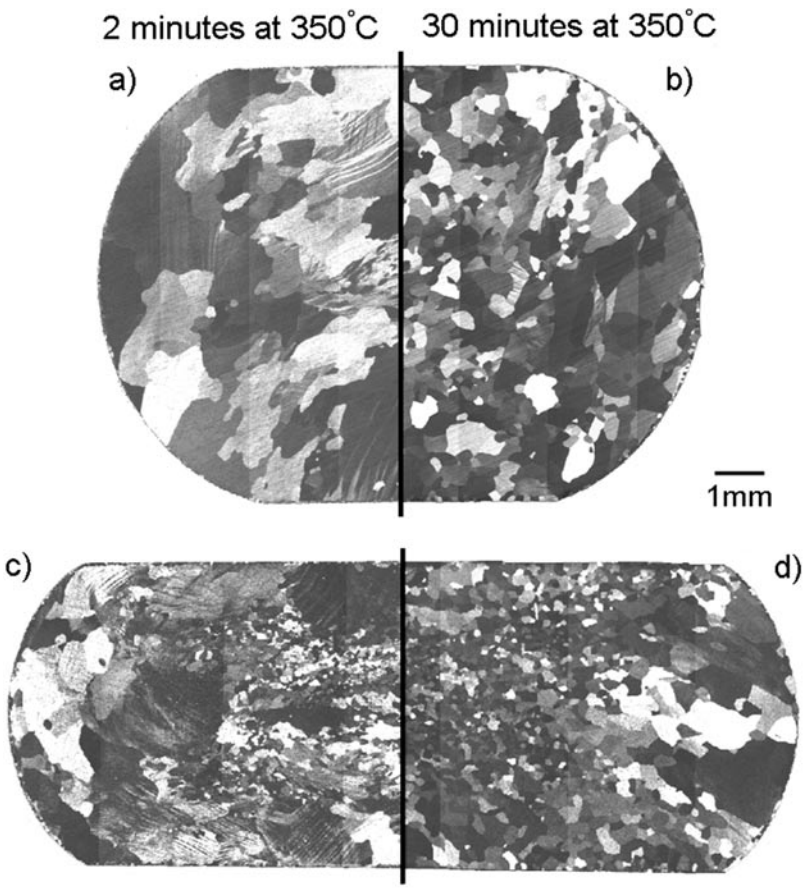

Figure 3. Evolution in the microstructure during the annealing of the cylinders compressed in plane strain (geometry I) and annealed at 350 °C. (a), (b) Reduced by 20% and (c), (d) reduced by 40%; (a), (c) annealed for 2 min and (b), (d) for 30 min. For each deformation the most highly-strained areas, near the centres of the samples, recrystallized first. It can be seen that (a) shows very little recrystallization while (b), with approximately 2.5 times greater maximum strain in its centre, has considerably more new grains. After 30 min, recrystallization was almost complete in both samples except in the regions of lowest strain.

Note that compared with traditional forgings the samples were very small. This meant that there was an inherent limitation on the statistics of the grain size. Too small a sampling area would lead to too few grains in the count, giving rise to large scatter. Too large an area would mean that a single strain trajectory could not be sensibly related to all the grains within the sampling area. As a result, the method was inherently noisy, and presented an important test of the approach. In the regions of small recrystallized grain size, up to 180 grains were measured and gave rise to uncertainties of a few per cent of the mean. Conversely, in regions for which the measured grain size was large, it was possible to measure only a small number of grains (<40), yielding larger errors. Furthermore, in some low-strain regions, especially in the plane strain samples, it was very difficult to distinguish between large recrystallized grains and original grains that were of a similar size, but were largely undeformed and did not show features characteristic of deformation such as elongated shapes or shear bands.

4. Gaussian process models

Similarly to neural networks, Gaussian process models are capable of identifying a relationship between inputs and outputs in a set of training data. This is then used as a basis for making predictions of outputs for new input sets. Whereas a neural network directly infers a relation predicting a single output value for given inputs, a Gaussian process model infers a joint probability distribution over all possible outputs for all inputs. A certain level of random Gaussian noise is assumed. The probability distribution is a multi-dimensional Gaussian, hence the name. This form may not be representative of all systems, but it enables Bayes' theorem (see the appendix) to be implemented in a simple way. This offers a number of benefits to the modeller, as explained by MacKay [13, 14]. The advantages over conventional modelling techniques include: implicit use of Occam's razor to limit model complexity; an objective method of comparing models using a likelihood calculated from the training data; and the estimation of the certainty of any predictions. A brief explanation of Bayesian theory and its application in Gaussian process models is given in the appendix; more rigorous, mathematical descriptions of Gaussian processes have been written by various workers, for example [7–11, 15, 16].

The inference of a joint probability distribution over the data involves inferring a number of quantities called hyperparameters [8, 9]. These indicate the precision and relevance of the various input parameters for predicting the outputs. This aspect of training is explained in the appendix.

Figure 4 illustrates how the training and prediction are performed for a simple two-parameter problem. If errors in the data can be estimated, they can be entered into the model to guide the inference of the hyperparameters. This is not necessary, but may help to reject models that are unrealistically optimistic (fitting data too closely) or pessimistic (over-simple with high uncertainty). It should be noted that where there is a high density of training data supporting the predictions, the confidence of the model is high and thus the Gaussian distribution is sharp. Conversely, for inputs outside its range, or in sparsely populated regions of input space, predictions will simply be made with large uncertainty so that the user is aware of their unreliability. As a consequence, for some practical applications it may be best to have an even distribution of training data across the input space.

A number of parameters are available for a Gaussian process model which indicate the quality of its fit to the given data set. In training and testing data sets, the measured output for a set of inputs is often referred to as the target, τ. A standard measure of how well a model such as a neural network has been trained is to compare the predictions of the outputs, p_i, with the targets for the ith elements in an N-dimensional data set, using a formula such as

$$E = \sum(\tau_i - p_i)^2. \tag{1}$$

In this work, the training and testing errors quoted are the root mean square (rms) errors given by $(E/N)^{1/2}$ where N is the number of data points. If the model has generalized well, the errors should be similar for the training and testing data. If the training error is significantly smaller, it suggests that the model was over-fitted to noisy points in the training data, so that it cannot make reasonable predictions for unseen data.

The Gaussian process model predicts not only the most likely output (the mean of the Gaussian distribution), but also the probability distribution (the standard deviation of the Gaussian) and thus another quantitative indicator of fit is possible. This is called the *likelihood* and it is a measure of the strength with which data are predicted. Instead of taking merely the difference between the target and prediction, it evaluates the probabilities of the predictions

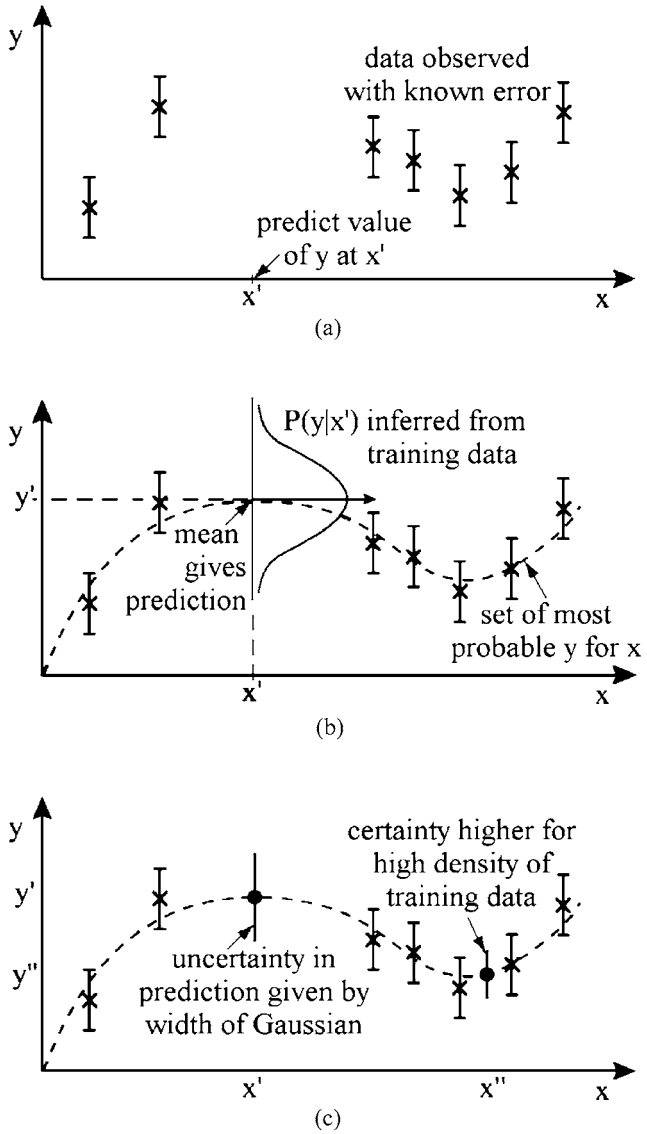

Figure 4. Training and prediction using a Gaussian process model. Plot (a) shows experimental observations of an output y that varies with some input x. It is known that there is some error in each measurement, indicated by the error bars. To make a prediction at x', the Gaussian process model first infers a Gaussian probability distribution across the training data space. The Gaussian drawn in (b) is the projection of this distribution at x' onto the page. The mean of the Gaussian gives the most probable value of y for any x, so that a prediction of y' can be made. It also allows an estimate of the uncertainty to be made; one standard deviation of the distribution has approximately a 67% certainty of containing the target value and thus an uncertainty bar may be drawn equal to the width of the Gaussian. It should be noted that in regions of high data density, the certainty of the model is increased, the Gaussian sharper and so the uncertainty contours are closer together (c).

(see the appendix). The logarithm of the likelihood, L, is given by

$$L = -\sum_k \left[\frac{(\tau_k - p_k)^2}{2e_k^2} + \ln e_k \right] \tag{2}$$

for all k points in the data set, where e_k is the standard deviation of the Gaussian at the kth point.

It should be noted that the greatest likelihood does not necessarily correspond to the smallest training error. This is because the e_k term makes it strongly dependent on the confidence of the model in various regions of the data space. For example, the training error may be low if dominated by a few points where the prediction almost equals the target. However, if the confidence in those regions of the data space is very low, a higher likelihood may be obtained by a model that makes poor predictions of those points, but slightly better ones of other points in high-confidence space. The likelihood is less subject to the influence of noise than is the error, and gives a measure of the 'plausibility' of the model. The likelihood should be considered together with error when evaluating the model as a high likelihood will result from a poor model as long as a large uncertainty, e, is also predicted.

5. Application of the Gaussian process model

5.1. Formulation of the training data sets

In order that each input parameter is initially treated with equal importance, it is necessary to normalize the input data. This is especially important in data sets such as that used here because all the strains lie between zero and one, while annealing temperatures lie between 325 and 375 °C, yet there is no reason to believe that temperature is more important than strain. All the inputs were scaled linearly to lie between +0.5 and −0.5. As it is well established that recrystallized structures often show a log–normal grain size distribution [17], the mean of the logarithms of the grain sizes was used as the output target variable. This also has the advantage of limiting the spread of output values to within an order of magnitude.

5.2. Training and testing of the Gaussian process model

The first test of any model is to assess its predictive capability. This is best achieved by training it on one set of data and then testing it on previously unseen data. First, a model was trained on just the 45 points belonging to set A, with a view to testing it on the 19 points (set B) which had experienced annealing temperatures not seen during training. It was also tested for its predictions on the training data. The correlation between the predicted and target (measured) outputs for the test data is plotted in figure 5 and the rms errors in the predictions of the Gaussian process model are given in table 1. The normalized log likelihood for each set of test data is also given, calculated by dividing the value of L (equation (2)) by the number of points in the test set. It is clear from the figure that the model has not significantly over-fitted the training data set, in that the predicted against measured graph for the training data (figure 5(a)) is not noticeably better than that for the unseen test data (figure 5(b)). This is because during training the model sacrifices an exact fit to the measured outputs in favour of simpler models with better prospects of generalization. It is also reflected in the fact that the testing error in table 1 is only about 10% worse than the training error. This is an important aspect of the current model; normally one does not want to hold back large portions of the training data set for testing and evaluation of uncertainty since this impoverishes the training. Instead, it is better to have modelling algorithms that inhibit over-fitting, as here. The described results suggest that the model has learned the essence of the coupling between the inputs and the recrystallized grain size. The discrepancy between the gradients of the best-fit line and the predicted equals measured line shown on figure 5(b) will be discussed later.

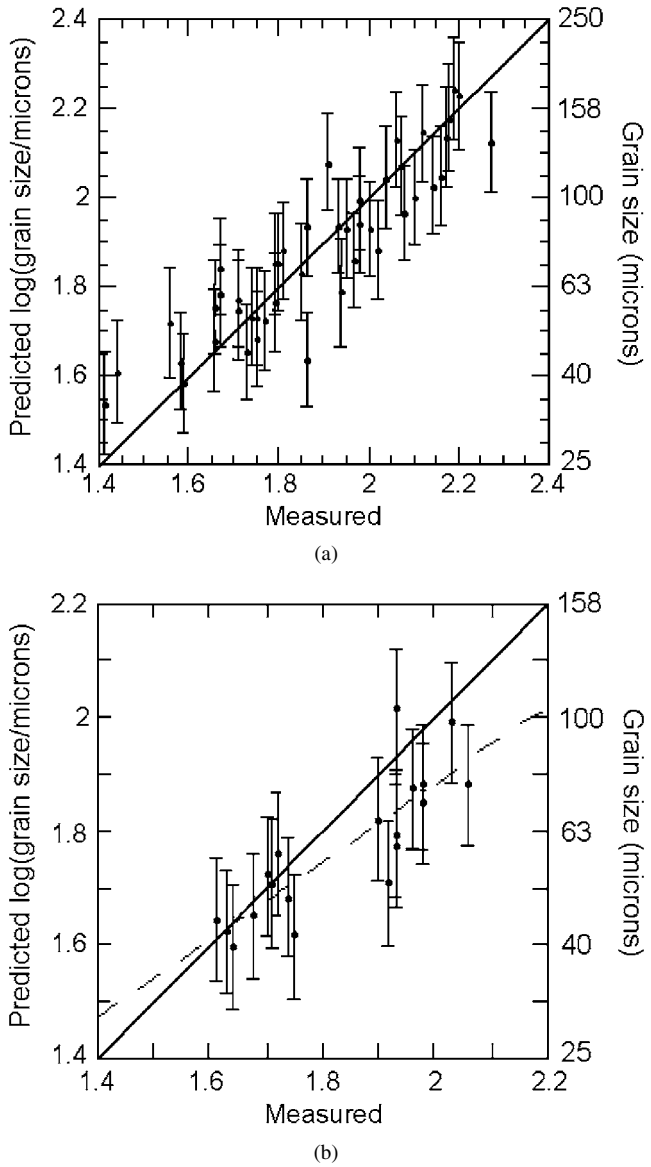

Figure 5. Plots of the predicted against the measured grain sizes for a model trained on set A. (a) The 'predictions' for the 45 plane strain training points comprising set A and (b) testing on the remaining 19 plane strain points annealed at temperatures outside the experience of the training set (set B). The solid line is the 1:1 gradient, the broken line the best-fit line for the points. The uncertainty bars are for one standard distribution, and thus approximately 67% of measurements would be expected statistically to fall within them.

A more general test of the trained Gaussian process model is to apply it to a data set (set C) for a second, previously unseen geometry (geometry II). The results of this test are shown in table 1 and figure 6(a).

It appears from the plots and from the test errors that the predictions for geometry II are better than those (set B) made for the same geometry as the training data-set, although all of the models are consistent with the uncertainty predictions of the Gaussian process. Indeed, the test

Table 1. A summary of the Gaussian process model for various training and testing scenarios.

Training set	Test set	Training error (log μm)	Test error (log μm)	No of test data points	L	Normalized log likelihood
A	A	0.092	0.092	45	99	2.20
A	B	0.092	0.102	19	42	2.21
A	C	0.092	0.071	26	58	2.23
A + B	C	0.081	0.087	26	61	2.35

error is about 30% smaller than that achieved for set B. This may be because the indentation samples had smaller strain gradients and, therefore, correlating grain size with local strain was more accurate, leading to points closer to the smooth relationship predicted by the Gaussian process. Alternatively, it could be simply good fortune that the data are fitted particularly well here. After training on set A, the likelihood per point is less than 1% greater for testing on set C than set B. This is insignificant compared with the errors in the data.

For the tests described above, the model was trained on just part of the available data obtained for geometry I, i.e. set A. Next, the Gaussian process model was retrained using the complete data set (64 points) collected on geometry I before applying it to the new geometry. The results for this strategy are shown in figure 6(b). It can be seen from table 1 that the training error for this Gaussian process model is less than for the model that was trained on set A alone, perhaps because the additional points increase the self-consistency of the data. At the same time, the likelihood of the test data is increased. As explained in section 4, this indicates an increase in the accuracy of the predictions, modified by their strength. Hence this decreased error suggests that the new model either predicts the test points more closely or has increased confidence due to the higher density of training data. The uncertainty bars in figure 5(b) are in fact smaller than in figure 5(a), indicating a higher confidence.

In figures 5 and 6, the broken line is the best-fit line for the predictions. In table 2, gradients are given, with their errors, for all of the training and testing combinations used rather than just those for which the best-fit line is plotted. In each case the gradient is significantly less than one. This means that, in general, the Gaussian process models do not predict a sufficiently wide variation in grain size. For the conditions that lead to small grains, that is low temperature, short time and high strain, they predict sizes that are larger than those measured. For low strains and more extensive annealing the predicted grain sizes are too small. If this were due to errors in the FE model used to calculate the strain histories, then the difference would be larger for the unseen geometry, since the strain distributions for sets A and B were produced using identical FE models. This is not true for the model trained on set A alone. The low gradient might be due in part to noise in the data causing the inferred maximum likelihood probability distribution to be skewed compared with expectations based on test error. It could also be because, at the extremes of the training data, the density of information is low and so the uncertainty is high. The Gaussian process is set up so that its predictions tend to the mean outside the training set and so predictions at large and small grain sizes can be pulled towards intermediate values [18]. This is likely to be the most significant cause of the discrepancy. This tendency towards the mean away from the data is due to the specific form of the covariance function we used (see the appendix), and is not a feature of Gaussian process models in general. In principle, the covariance function could be altered to provide for linear extrapolations of the predicted function away from the data, if that is what prior physical information on the specific problem in hand assumes.

The importance of this deviation is, in any case, arguable, because it is clear from the plots that lines with gradients steeper than one could also be drawn within the uncertainty bounds.

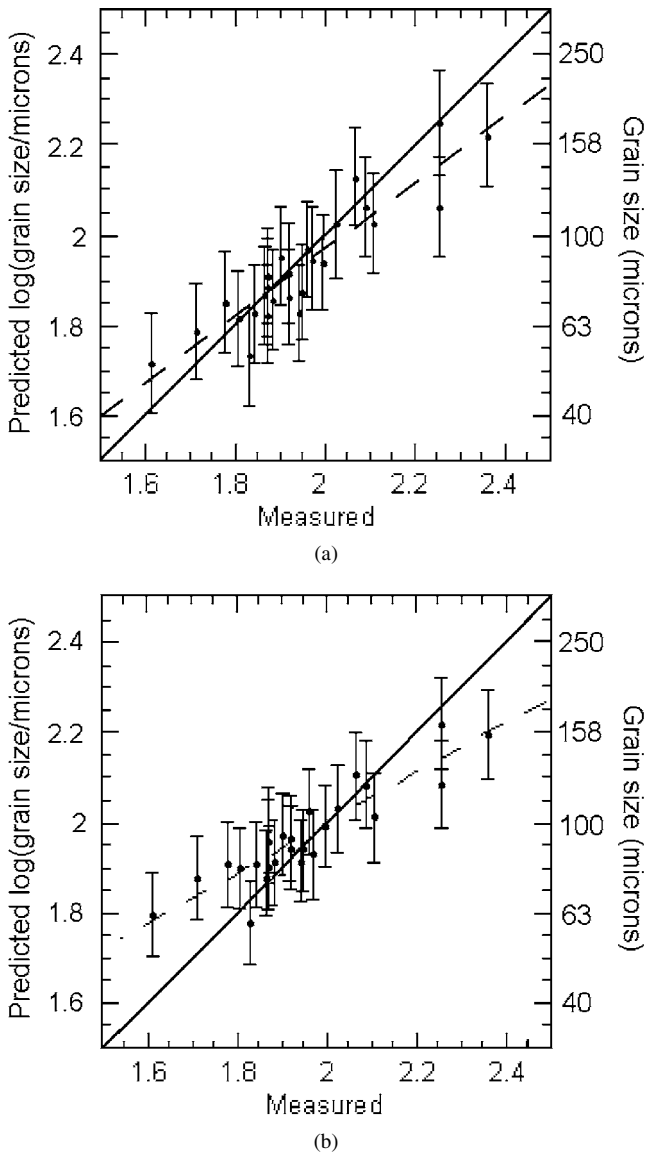

Figure 6. Plots of predicted against measured grain sizes for models tested on the indentation data (set C, geometry II) after training on (a) the 45 plane strain points comprising set A (geometry I) and (b) all 64 plane strain points comprising set A plus B.

The model is clearly good enough for process design: it is capable of indicating the conditions required to obtain larger or smaller recrystallized grains.

5.3. Mapping microstructural variation

The training against testing comparisons highlight the predictive capability of the model. It is possible to encode the trained Gaussian process model into a metal forming programme so that microstructure maps can be produced automatically for direct comparison with the

Table 2. The gradients of the best-fit lines for comparison of predicted with measured values for each combination of training and test sets used.

Training set	Test set	Figure	Best-fit gradient
A	A	5(a)	0.80 ± 0.006
A	B	5(b)	0.69 ± 0.03
A	C	6(a)	0.74 ± 0.02
A + B	A + B	8(a)	0.78 ± 0.004
A + B	C	6(b)	0.56 ± 0.01
A + B + C	A + B + C	8(b)	0.78 ± 0.003

Table 3. Analysis of the Gaussian process models trained and tested on the same data sets. The first two sets are for plane strain data and the third includes the axisymmetric data. The likelihood increases with the amount of training data, as there are more points to support the inferred probability distribution. The noise is each Gaussian process model's estimate from the smoothness of the training data.

Data set	Training error = test error ($\log \mu$m)	Noise	No of data points	L	Normalized log likelihood
A	0.092	0.101	45	99	2.20
A + B	0.081	0.089	64	150	2.34
A + B + C	0.086	0.091	90	212	2.36

observed metallographic sections (figure 7). Such maps produced for various die geometries and strain and temperature histories allow the optimization of the microstructure across forged components without the need for costly forging trials. The software used also enables mapping of other parameters, including the uncertainty of the predictions. This could be particularly useful for designing tests to obtain more training data for the Gaussian process model: forgings could be chosen to provide data across ranges with high uncertainty, where they will be of the greatest use for improving the model.

In some senses, we would ideally like to run this process backwards, i.e. predict process parameters that would give a desired microstructure. However, this is clearly more complex as there may be many combinations of inputs that would give the same outputs, and there is then the possibility of predicting strains and temperatures that would in fact be impracticable.

5.4. A comparison of different training approaches

The training and testing comparisons have shown that the Gaussian process model is a reliable indicator of recrystallized grain size. An important question then arises as to what training should yield the most reliable model for subsequent application. For example, should the results from both geometries be combined? Table 3 indicates how the models compare. In this case, the training and test error are the same since the training and test data are the same. The noise given is that estimated by the Gaussian process model, from the smoothness of the variation in measured grain size with the input parameters in the training data set. It can be seen that the noise is comparable with the errors; that is, it is directly related to the closeness of the predictions to the measurements.

The normalized log likelihood increases with the amount of training data, since the inferred relationship between the inputs and outputs is supported more strongly by a higher density of points, i.e. the Gaussian probability distribution is sharper, giving greater confidence in the model predictions of targets. Increasing the training set from A to A plus B improves this

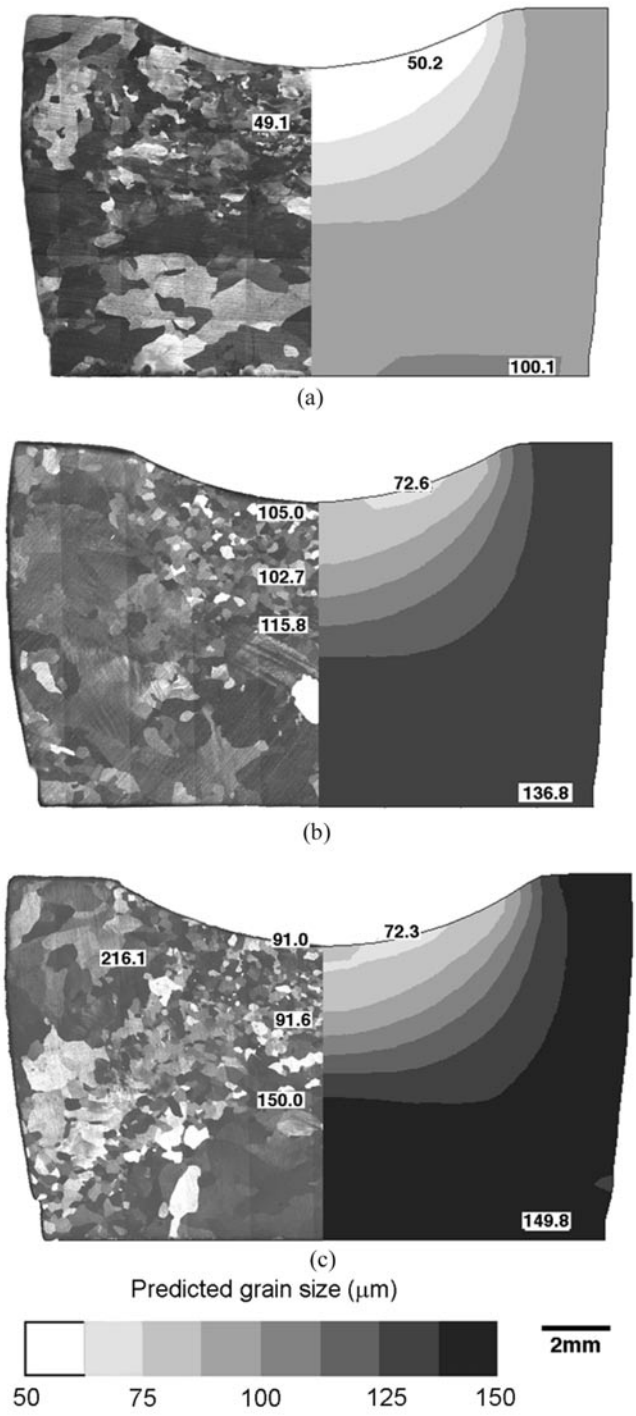

Figure 7. Comparison of predictions and measurements of the grain size in the representative indentation samples for a Gaussian process model trained on all plane strain data. The samples were annealed for 10 min at (a) 325 °C, (b) 360 °C and (c) 375 °C.

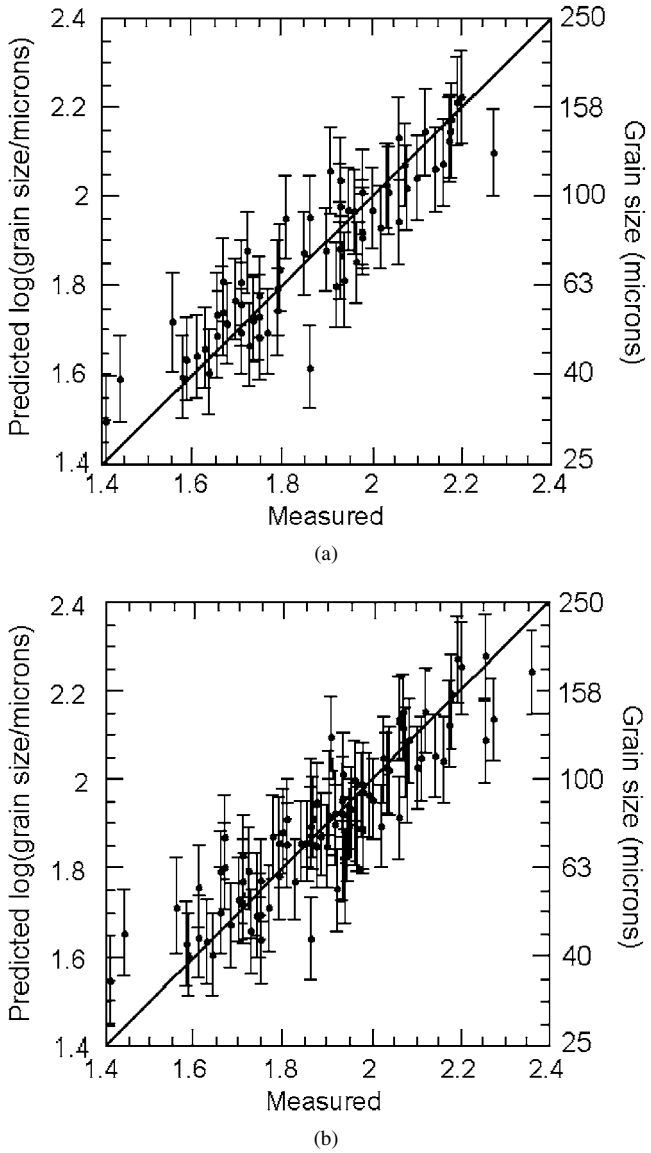

Figure 8. Plot showing predicted against measured grain sizes after the Gaussian process model has been trained and tested on (a) the complete plane strain data set A plus B (geometry I) and (b) on set A plus B plus C (all plane strain and all indentation points).

factor by approximately 7%, although the further improvement for the addition of set C is insignificant. This suggests that provided one believes (as in this case) that the data set is governed by a single underlying relation, it is best to combine all the data across the sets to train the Gaussian process model.

When the model is trained on plane strain data alone, the error is decreased by the addition of a second set of points. This can be explained by an improvement in the fit of the Gaussian to the data, so that more points are predicted more accurately. When the axisymmetric set is also included, the test error is greater since more points are poorly predicted. This is probably due to some discrepancy between the best-fit relationships for the plane strain and axisymmetrical

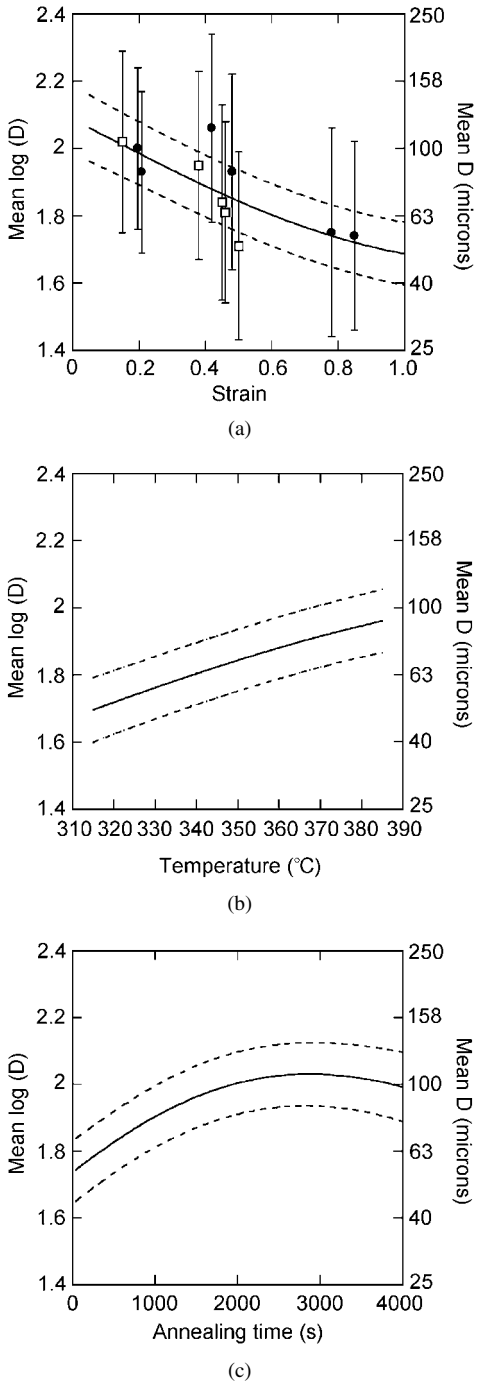

Figure 9. Recrystallized grain size predictions of the Gaussian process model after training on all data from non-uniform deformation tests holding two of the three input variables constant and varying the third. The constant conditions are a strain of 0.5, a temperature of 350 °C and an annealing time of 10 min. In (a) the symbols represent the experimental data: the full circles from geometry I tests and the squares from geometry II. It can be seen that the model has generalized well and has predicted a smooth curve rather than attempting to fit training data exactly.

geometries or differences in the way the FE software treats the different conditions. It could also be affected by underlying physical processes that are not accounted for in the FE method, for example, texture development, leading to different rates of hardening and anisotropy in the samples.

The results of the testing on the training sets are plotted in figures 5(a) and 8. The bars marked on the plots represent the standard deviation of the probability distribution and can be regarded as contours of certainty about the prediction. Probability theory states that approximately two-thirds of data should fall within these bars. For all three tests, a greater fraction of points have error bars crossing the predicted equals measured line. This suggests that the Gaussian process model might have overestimated its own uncertainty and that in fact the predictions are better than expected.

A further observation that can be made from the plots is that the size of the uncertainty bars decreases as the amount of training data increases. As the model becomes more certain, due to increased evidence, the Gaussian probability distribution becomes sharper and the standard deviation, which gives the bar sizes, decreases. This is in agreement with the expectation from the likelihoods calculated and confirms that optimal training is obtained by using the maximum data set.

5.5. The predicted variation in grain size with input variables

One of the main advantages of the Gaussian process models is that they perform multivariate nonlinear analysis. Nevertheless, once trained it is helpful to plot out slices through the multi-parametric domain in order to compare the modelled response with basic metallurgical understanding.

The model trained on the complete set of data (sets A plus B plus C) was interrogated to investigate the predicted relationships between the grain size and the individual inputs. This was accomplished simply by holding two of the inputs constant and varying the third. In effect, this is taking two-dimensional sections through the four-dimensional space of the model. Figure 9 shows typical predictions. They can be seen to be qualitatively similar to the observations of others on static recrystallization [19]. It is also clear from figure 9(a)) that the model has inferred a smooth relationship between the input and outputs despite the noisy training data; that is, it has generalized well.

It is notable that the grain size is predicted to decrease at long annealing times (figure 9(c)). This is not realistic, but the uncertainty is large enough that it would be possible to draw the expected increasing trend well within the bounds marked. It is a result of the sparsity of data at long annealing times and the choice of covariance function, which pulls predictions towards the mean (see the appendix).

6. Discussion

The training and testing of a Gaussian process model to predict statically recrystallized grain sizes has been demonstrated. After training, it was able not only to simulate the development of recrystallized grains in samples similar to those used in training, but also to make good predictions for a different deformation geometry. The success of the model can be measured not only by the closeness of its predictions to the measured values, but, in this example where the data are limited and very noisy, by the fact that the predicted levels of accuracy are commensurate with the experimental data. The model does not over-fit; instead it generalizes well to unseen data, despite the noise in the training set, and predicts patterns of behaviour that agree with accepted trends.

There are other benefits in using this type of model. First, given an appropriate training set, it is applicable to any alloys, regardless of any differences in the micromechanical mechanisms operating. Conventional models are generally only useful for a particular alloy system or even a single alloy. The Gaussian process model can also be easily retrained on a wider data set as further data become available.

There are a number of advantages of using non-uniform deformation geometries for training. First, unstructured approaches such as neural nets and Gaussian process models always require more data than approaches that start with a model which then needs to be calibrated. Non-uniform deformation geometries reduce the matrix of specially prepared samples needed for training, by each one providing a spread of strain trajectories. It may be possible to use the results of forging trials that are similar to the target geometry, reducing the influence of errors. Another approach is to use FE modelling to calculate optimized training forging geometries that will give a large variation in strain trajectories, while being similar to the industrial product. This ensures that the training set is well distributed throughout the input space and prevents the model from being naive. Finally, non-uniform strain trajectories enable the model to take into account strain path effects that monotonic constant strain rate testing cannot.

In the present study, it appears that the model could be enhanced by improving the quality of the training data. There is clearly a large amount of noise in both the microstructure measurements and the strains with which they are correlated. Defining the friction factors at the die–sample interfaces is particularly difficult in non-uniform geometries, such as those used in the experimental work. In future work it may be profitable to include training data from more easily controlled and modelled processes such as wedge rolling. There may also be other deficiencies in the FE model such as inaccurate material properties. However, the Gaussian process generalizes well over the data space examined and it is arguable that the increase in effort needed to significantly improve the model certainty would not be cost or time effective for industrial purposes, where there are inherently large errors in measurements. Although the likelihood is increased by increasing the amount of training data, the decrease in the error when an additional 19 training points (over 40% extra) were used was relatively small. Furthermore, it must be remembered that the model is not a physical causal model but a statistical model capable of establishing trends in the data, which can be exploited to make predictions. It is thus, probably, best aimed at providing general guidelines on how to increase the microstructural control within certain statistical bounds of certainty rather than the exact deformation conditions required to produce a specific average grain size precisely.

Ideally, metallurgists would like to understand the underlying physical mechanisms that determine microstructure and so develop understanding-based models to perfectly predict grain sizes and orientations. We recognize that Gaussian processes, neural networks and similar techniques are far from this ideal as they neglect physical processes. To deal with the effects of such phenomena as texture development and inhomogeneous slip, which have been ignored in this work, a Gaussian process model would have to be given more input parameters and considerably more training data to capture relevant patterns.

Physically-based state variable models have been developed to give more accurate predictions of the recrystallized grain size for the alloy used here, although without the prediction of uncertainty [4]. However, the approach taken by Furu *et al* [4] was tested only within the bounds of the data used to tune the model and for relatively simple deformation histories, and may prove difficult to apply to industrial workpieces. It required large amounts of accurate substructural data and so was time consuming to develop.

For complex problems such as this, the Gaussian process model has the potential to provide guidance in the development of physical models. It can be trained on a large number of inputs,

some of which may not be relevant, and can give any measurable quantity as an output, for example subgrain size or misorientation. Given a suitable training set, the significance and effect of varying individual inputs or combinations can be investigated, as in section 5.5 of this work. This could lead to more efficient testing to provide additional data for physical models and indicate the nature of the relationships between the parameters.

This application clearly demonstrates the advantages of a probabilistic approach to the modelling of real data. First, the estimates of uncertainty allow users of the model to determine how far they may rely on the predictions for any new processing conditions. This is likely to find applications in industry where the uncertainty of many factors, such as heat transfer and friction coefficients, make it difficult to produce accurate input data. In addition, the model does not require sub-structural information, which might be difficult or expensive to measure in an industrial plant.

7. Conclusions

The main conclusions of this work can be summarized as follows.

(i) The new type of statistical model investigated, which requires no prior physical understanding, is capable of being trained to make reliable predictions of metal microstructures after static recrystallization.

(ii) The approach is capable of finding patterns in limited and noisy data sets such as might be obtained from industrial workpieces. The models generalize well and are not prone to over-fitting; they predict reasonable patterns of behaviour.

(iii) The approach is able to calculate a level of uncertainty in making a given prediction. High uncertainty is predicted for deformation conditions outside the range of training, regions of data space for which the training set was sparsely populated or those for which the data was particularly noisy. Confidence is consistent with the test data. This feature allows the user to take appropriate cautionary measures when applying the model to new conditions. It is likely to be particularly valuable in applications such as this, where noise is inherent in making measurements.

(iv) It is believed that this type of model could be a valuable aid to the development of understanding and of physically-based models for complex metallurgical systems.

(v) It has been shown that used in conjunction with a FE model, non-uniform geometries can lead to reliable training using complex training data. Such training schedules can be closer to the end application of the model and can accelerate the training process.

Appendix

Bayes' theorem states that the posterior probability of a condition is given by the product of the prior probability and the likelihood in the light of the evidence. For a simple true/false problem, this can be written as follows:

$$\overset{\text{posterior}}{\boldsymbol{P}(\mathrm{B}|\mathrm{A},\mathrm{H})} = \frac{\overset{\text{likelihood}}{\boldsymbol{P}(\mathrm{A}|\mathrm{B},\mathrm{H})}\ \overset{\text{prior}}{\boldsymbol{P}(\mathrm{B}|\mathrm{H})}}{\underset{\text{evidence}}{\boldsymbol{P}(\mathrm{A}|\mathrm{H})}} \tag{A1}$$

$\boldsymbol{P}(\mathrm{B}|\mathrm{A},\mathrm{H})$ is the posterior probability that statement B is true, given that condition A is observed and that hypothesis H is correct. $\boldsymbol{P}(\mathrm{A}|\mathrm{B},\mathrm{H})$ is the probability of observing A if B is true and H is correct, which is called the likelihood. $\boldsymbol{P}(\mathrm{B}|\mathrm{H})$ is the prior probability of B

being true, without having made any observations. $\boldsymbol{P}(\mathrm{A}|\mathrm{H})$ is the evidence: the probability of observing A if hypothesis H is correct.

A more useful illustration of Bayes' theorem in the context of this paper is its application to find the hyperparameters, Θ, in a Gaussian process model, and thus enable it to make predictions of outputs from inputs. In a Gaussian process model, the aim is to choose model parameters for which the probability of the training data is maximized.

Let us consider a set of training data containing N points, comprising a set of targets (e.g. a list of grain sizes), τ_N, with their corresponding vectors of inputs (e.g. strains, annealing temperatures and times), $\{x_N\}$. A Gaussian process model is defined by an N-dimensional covariance matrix, C_N, that describes the closeness to each other of outputs for different inputs, taking into account the noise, the significance and the scaling of each parameter and so on. This allows predictions of the output to be made, based on differences between new inputs and those seen in the training data. Each element of the matrix is given by C_f, which is a function of the inputs and the hyperparameters, of which noise is one. For the element ij in the covariance matrix, $C_{ij} = C_f(x_i, x_j, \Theta)$. C_f is called the covariance function and it may take a range of forms, chosen by the user, including those that simply indicate that similar inputs give rise to similar outputs and those that embody periodicity [9]. The determination of hyperparameters is similar to finding the optimum values for parameters in a conventional equation in order to fit data. In this work, the covariance function used was [8]

$$C_f = \theta_1 \exp\left[-\frac{1}{2} \sum_{l=1}^{l=L} \frac{(x_i^{(l)} - x_j^{(l)})^2}{r_l^2} \right] + \theta_2 + \sigma_n^2 \delta_{ij} \tag{A2}$$

where the set of hyperparameters is $\Theta = \{\theta_1, \theta_2, r_l, \sigma_n\}$ and δ_{ij} is a delta function with a value of zero for all $i \neq j$. These are explained fully in [8], but described briefly here. The first term allows the closeness of two outputs, τ_i and τ_j, to be related to the closeness of their corresponding inputs; generally if the inputs are close, it is assumed that the outputs will be also. The length scale r_l for the lth input parameter (e.g. temperature) indicates how much the output will vary relative to any changes in an input; if it is large, then large changes in input will be needed to affect the output. θ_2 allows the function to be offset from zero in the case that the mean of the data and the mean of the Gaussian process are different. The final term is a noise model, with σ_n the standard deviation of noise, which is assumed to be Gaussian.

To find the optimal hyperparameters, equation (A1) becomes

$$\boldsymbol{P}(\Theta|\tau_N, \{x_N\}, C_f) = \frac{\boldsymbol{P}(\tau_N|\{x_N\}, C_f, \Theta)\boldsymbol{P}(\Theta)}{\boldsymbol{P}(\tau_N|\{x_N\}, C_f)}. \tag{A3}$$

Since the evidence (the denominator in equation (A3)) is independent of the hyperparameters, it is effectively constant for a given data set. The prior may be non-informative or use prior knowledge about the process. To find the optimal hyperparameters we must maximize the posterior probability (which, in the case of a uniform prior, is equivalent to maximizing the likelihood). Therefore for a particular training data set and covariance function, the Gaussian process will select the hyperparameters that give the best predictions of the training data.

In this work it was seen that the predicted outputs tended towards the mean outside the bulk of the data, causing significant errors in the predictions for large and small grain sizes. This was due to our choice of covariance function, which was perhaps not the most appropriate given our prior knowledge. It should be possible to improve the results by encoding a C_f which would, for example, assume a smaller rate of change in outputs away from the training data or prescribe an increase in an output with a parameter, such as the annealing time. There is no mathematical reason for this, it is an assumption based on prior physical knowledge. Prior knowledge should be built into the covariance function where possible, but care should be

taken as there may be limits; for example grain size can never become larger than the sample. Bayesian theory offers the possibility of choosing the most appropriate function from several alternatives by taking that for which the evidence is greatest.

Acknowledgments

The authors would like to thank Mark Gibbs for use of his Gaussian process model software and David MacKay for helpful discussion. Thanks are also due to the EPSRC (CBJ) and INCO Alloys Limited (TJS) for financial support, and to Alcan Aluminium UK Limited for providing the material.

References

[1] Puchi E S, Beynon J and Sellars C M 1988 *Proc. Int. Conf. on Physical Metallurgy of Thermomechanical Processing of Steels and other Metals: THERMEC '88* ed I Tamura (Tokyo: The Iron and Steel Institute of Japan)
[2] Hodgson P D and Gibbs R K 1992 *ISIJ Int.* **32** 1329–38
[3] Sandström R and Lagneborg R 1975 *Acta Metall.* **23** 481–8
[4] Furu T, Shercliff H R, Baxter G J and Sellars C M 1999 *Acta Mater.* **47** 2377–89
[5] Bhadeshia H K D H, MacKay D J C and Svensson L E 1995 *Meas. Sci. Technol.* **11** 1046–51
[6] Roberts S M, Kusiak J, Liu Y L, Forcellese A and Withers P J 1998 *J. Mater. Proc. Technol.* **80–81** 507–12
[7] Williams C K I and Rasmussen C E 1996 Gaussian processes for regression *Advances in Neural Information Processing Systems 8* ed D S Touretzky, M C Mozer and M E Hasselmo (Boston, MA: MIT Press). Obtainable from http://bayes.imm.dtu.dk/gp.html
[8] Bailer-Jones C A L, Bhadeshia H K D H and MacKay D J C 1999 *Meas. Sci. Technol.* **15** 287–94
[9] Gibbs M N and MacKay D J C 1997, unpublished, webpage http://wol.ra.phy.cam.ac.uk/ mng10/GP/GP.html
[10] Gibbs M N 1998 Bayesian Gaussian processes for regression and classification *PhD Thesis* University of Cambridge. Obtainable from http://wol.ra.phy.cam.ac.uk/ mng10/GP/GP.html
[11] Bailer-Jones C A L, MacKay D J C, Sabin T J and Withers P J 1998 *Aust. J. Intell. Inform. Process. Syst.* **5** 10–17 Obtainable from http://www.mpia-hd.mpg.de/homes/calj/publications.html
[12] DEFORM 1992 *DEFORM System Users' Manual Version 3.0* Scientific Forming Technologies Corporation, Columbus, OH
[13] MacKay D J C 1992 *Neural Comput.* **4** 415–47
[14] MacKay D J C 1995 *Network: Computat. Neural Syst.* **6** 469–505
[15] Neal R M 1997 Monte Carlo implementation of Gaussian process models for Bayesian regression and classification *Report* 9702 Department of Statistics, University of Toronto. Obtainable from http://www.cs.toronto.edu/ radford/papers-online.html
[16] Bailer-Jones C A L, Sabin T J, Roberts S M, MacKay D J C and Withers P J 1997 *Proc. Australasia Pacific Forum on Intelligent Processing and Manufacturing of Materials: IPMM '97 (Gold Coast, Australia)* ed T Chandra *et al* . Obtainable from http://www.mpia-hd.mpg.de/homes/calj/publications.html
[17] Marthinsen K, Lohne O and Nes E 1989 *Acta Metall.* **37** 135–45
[18] Shercliff H R 1999 Private communication
[19] Humphreys F J and Hatherly M 1995 *Recrystallization and Related Annealing Phenomena* 1st edn (Oxford: Pergamon, Elsevier)