LOAD SYNTHESIS OF HELICOPTER DYNAMIC COMPONENTS: IMPROVEMENTS TO LINEAR REGRESSION IN THE AZIMUTH DOMAIN

John Vine*, Xiaobo Yu*
***Defence Science and Technology Organisation**
John.Vine@dsto.defence.gov.au;Xiaobo.Yu@dsto.defence.gov.au

Keywords: *linear regression, load synthesis, load monitoring, dynamic components*

Abstract

This paper reports improvements to the accuracy and practicality of an approach using linear regression in the azimuth domain to synthesize loads on helicopter dynamic components. Major improvements result from the addition of flight parameters as predictor variables, and the ability to perform concurrent regression at multiple azimuth locations. Their effectiveness is demonstrated by improvements in the predictive accuracy and model parsimony of load synthesis models developed for 160 representative level flight runs selected from a Black Hawk helicopter strain survey.

Nomenclature

$\hat{\beta}_i$	regression coefficient of *i*th predictor
$\hat{\boldsymbol{\beta}}$	vector containing $\hat{\beta}_i$
ε_j	residual error of model at *j*th event
$\boldsymbol{\varepsilon}$	vector containing ε_j
ϕ	azimuth angle of main rotor in degrees
ϕ_o	azimuth angle of main rotor at $t_R = 0$
N_R	ratio between measured and nominal main rotor rotational speeds in percentage
$\mathbf{R_p}$	correlation coefficient matrix
R^2	multiple correlation coefficient
RSS	residual sum of squares
V_h	maximum horizontal speed of helicopter
$\mathbf{X}$	matrix containing x_{ij}
a	element of inverse correlation
f_{MR}	nominal (N_R = 100%) main rotor frequency in Hz
m	number of events
n	number of candidate predictors
p	number of selected predictors for a predictive model
t_R	elapsed time at event recording, in seconds
x_{ij}	normalized value of *i*th candidate predictor at *j*th event
y_j	normalized response variable at *j*th event
$\mathbf{y}$	vector containing y_j

Subscripts:

Te	taken from test data set
Tr	taken from training data set
i	index of candidate predictor variable
j	index of event within a data set

1 Introduction

Applying individual loads monitoring to fatigue-critical dynamic components on helicopters can lead to reductions in maintenance costs and improvements in fleet management. These benefits arise from the ability to use actual loads instead of assumed loads in the calculation of component fatigue lives, thereby improving the accuracy of the calculated life for a specific component.

Fatigue damage in helicopter dynamic components is generally managed using the safe-life methodology and lives are calculated based on an assumed load spectrum [1]. This spectrum is a combination of an assumed usage spectrum and loads generated by either flight tests or detailed analysis. It can be significantly different from the actual loading that a component will experience [2], which means that the actual fatigue life of the component may vary significantly from the calculated life [1].

By applying individual loads monitoring to a helicopter, measured instead of assumed load spectrums can be used to calculate fatigue lives. These lives would have greater accuracy and allow for the fatigue life of components to be managed on an individual aircraft basis. These refined lives would result in improved safety for those cases where the monitored fatigue life is less than the originally-calculated safe-life; and potential reductions in maintenance costs where the monitored life exceeds the original.

Although beneficial, directly monitoring loads on helicopter dynamic components is not an easy task. During flight testing these components can be monitored with strain gauges whose signals are passed via sliprings from the rotating components to the recording system in the aircraft. However, these systems are ill-suited for monitoring larger numbers of aircraft in an operational environment due to their excessive maintenance requirements. If a fleet-wide direct load monitoring program is to be implemented, alternative monitoring techniques must be used.

Wireless technology is one potential alternative. Using wireless transmission, measurements made on dynamic components can be transferred to nearby data acquisition systems without requiring costly intermediate assemblies. Although potentially an attractive solution, high power requirements associated with current wireless technologies prevent their immediate use in direct load monitoring applications. Novel approaches to the power problem are being developed [3]; however these solutions cannot currently generate enough power to allow for continuous monitoring at high-sampling rates, without resorting to techniques such as duty sampling [4].

Another potential monitoring approach, and the focus of this paper, is load synthesis. This area of research, also known as load prediction, involves developing transfer functions which calculate loads in hard-to-measure locations, from easily measured inputs. The benefits of this approach depend on the inputs to the transfer functions. If input data is not monitored by existing instrumentation, installation of the additional instrumentation may lessen the benefit of the approach. Hence, load synthesis approaches generally place emphasis on eliminating or minimising the number of inputs that require additional instrumentation.

Many investigations have used load synthesis to predict helicopter loads with varying degrees of success. One investigation compared the ability of three independent approaches to predict main rotor pitch-link loads using airframe-mounted strain gauges [5]. Despite the use of relatively simple linear models, the accuracy of each approach proved promising, highlighting the presence of linear relationships between dynamic component and static airframe loads. Some more recent efforts make use of non-linear models to predict dynamic component loads using flight state and control parameter data as input [6]. As this data is recorded by existing instrumentation, its use minimises the requirement for additional instrumentation. This approach [6] demonstrated reasonable predictive accuracy, although adjustments had to be made to counter some underestimation of loads.

The work detailed within this paper is an extension of an approach which used linear regression in the azimuth domain to predict main rotor pitch-link loads [5]. This approach was selected because it demonstrated reasonable predictive capability, and meanwhile had two areas where obvious improvements can be made. One of the areas relates to the limiting of candidate predictors to the strain gauge and accelerometer measurements on stationary components. As noted in [5], the reaction-less effects prevent some load features being transmitted through the swash-plate to stationary components, therefore the information from the stationary components alone are insufficient to build a deterministic model. The other area relates to model parsimony. Ref [5] did not investigate a technique to guide all azimuth locations to select a common set of these inputs. Since different inputs were selected at different azimuth locations, total input requirements could be excessive.

To address the above deficiencies, two changes are proposed in the present study: one is to expand the candidate predictor variables to include flight state and control parameters, and the other is to develop a concurrent regression technique that allows a common set of predictors to be selected over multiple azimuth locations. The effects of these changes on predictive accuracy and model parsimony are to be investigated using 160 level flight runs of a Blackhawk flight strain survey.

2 Data Preparation

2.1 Essential Role of Data

The ultimate purpose of the present investigation on load synthesis is to develop a generic approach that can be used on various helicopter platforms. The development of such an approach is heavily based on, and driven by, available experimental data. This is especially the case when the predictive model is to be revealed through a data mining process like linear regression.

The data used in the present study was selected from a Black Hawk Flight Strain Survey detailed below.

2.2 Black Hawk Flight Strain Survey

This survey was jointly conducted in 2000 by the United States Air Force (USAF) and Australian Defence Force (ADF). The major objective was to utilize a large number of strain gauges and accelerometers to measure the strains in areas known to have cracks or other forms of structural distress [7]. Not including duplicate gauges installed for redundancy, a total of 421 strain gauges were installed on the test aircraft. The outputs of these gauges were combined to provide 217 channels of airframe strain gauge data, and 20 channels of dynamic component data. In addition to the strain gauge channels, 18 channels of accelerometer data were measured at various locations on the airframe, and 28 standard flight state and control system parameters were recorded.

In total, 65 hours of useful flight test data was recorded during the program, comprising 3759 runs of 98 unique manoeuvres. The flight tests were conducted for varying aircraft configurations, gross-weight values, altitude, and centre of gravity locations.

2.2 Subset for This Study

This study was based on a subset of the Flight Strain Survey data, consisting of 160 unique runs of level flight, evenly drawn from 8 manoeuvre groups corresponding to level flight at speeds from 0.3Vh to 1.0Vh. These 160 runs were representative of all configuration, gross-weight, centre of gravity, and altitude points tested during the flight test program.

A number of choices were made to select this subset; these included: (i) restricting runs to level flight manoeuvres, which limited variability in load relationships and simplified model development; (ii) including runs conducted at various flight speeds, which allowed for the development of generalized models - the even distribution of these runs (20 runs from each speed regime) prevented bias; and (iii) a range of gross-weight, configuration and centre of gravity test points were included to develop a generalized model.

The main rotor pitch-link load was selected as the target for prediction for two main reasons. Firstly, pitch-link loads form a basis for the fatigue substantiation of numerous dynamic components [8], and therefore refinements in their prediction have a wide impact. Secondly, pitch-link loads are complicated load signals, which ensured the prediction problem was not unrealistically simple.

2.3 Re-sampling to Azimuth Domain

The strain data consists of measurements of various outputs (e.g. strain, loads, and accelerations) in the time-domain. Due to variations of main rotor speed and inconsistency of starting azimuth locations, the time variable alone is not a valid indicator of the azimuth location across all selected runs. The survey data therefore need to be re-sampled into the azimuth domain before a predictive model can be developed in that domain. Unfortunately, for this dataset, the main rotor azimuth was not recorded, and therefore needs to be derived from other recorded data based on Eq. (1):

$$\phi(t_R) = \phi_o + 360 \times \int_0^{t_R} N_R f_{MR} dt_R \qquad (1)$$

Where, t_R is recorded time, starting from zero for each run, $N_R f_{MR}$ is actual main rotor frequency, ϕ_o is azimuth location at $t_R = 0$, and $\phi(t_R)$ is the derived azimuth location in degrees. ϕ_o was calculated for each run from the first harmonic of the main rotor shaft bending moment, based on the assumption that the direction of thrust of the main rotor was similar between runs.

Once all time samples within the subset were assigned an azimuth, they could be restructured and categorised according to that value, as illustrated in Fig. 1 using a 2-second pitch-link measurement as an example. For the work detailed within this paper, 360 azimuth "bins" were used, each 1° wide. In the subsequent regression analysis (see Section 3.1), each of the data points within a bin corresponds to an individual "event".

Along with the re-sampling process, minor data correction and elimination was performed to account for missing values within the data. This involved either removing entire data channels if they were missing excessive amounts of data, or removing isolated events corresponding to missing values.

2.4 Partition of Runs for Cross-Validation

A 10-fold cross validation was adopted in this study for a more effective use of the strain survey data. The process involved partitioning the data into 10 folds. Of these 10 folds, one was retained for evaluating the model, and the remaining nine were used as training data. This process was repeated 10 times with each fold used once as the testing data.

To realistically assess the predictive ability of a model produced by a regression process, the testing data must not be seen by the model during training and must approximate the unseen events that the model may experience. As helicopter flight data is periodic in nature, events from within the same run may be nearly identical and must not be included in both training and test data sets. This required proper partitioning of the available data, to ensure that a run was contained wholly within one of the ten folds and not split amongst them, and that each fold was representative of the fold being withheld as testing data.

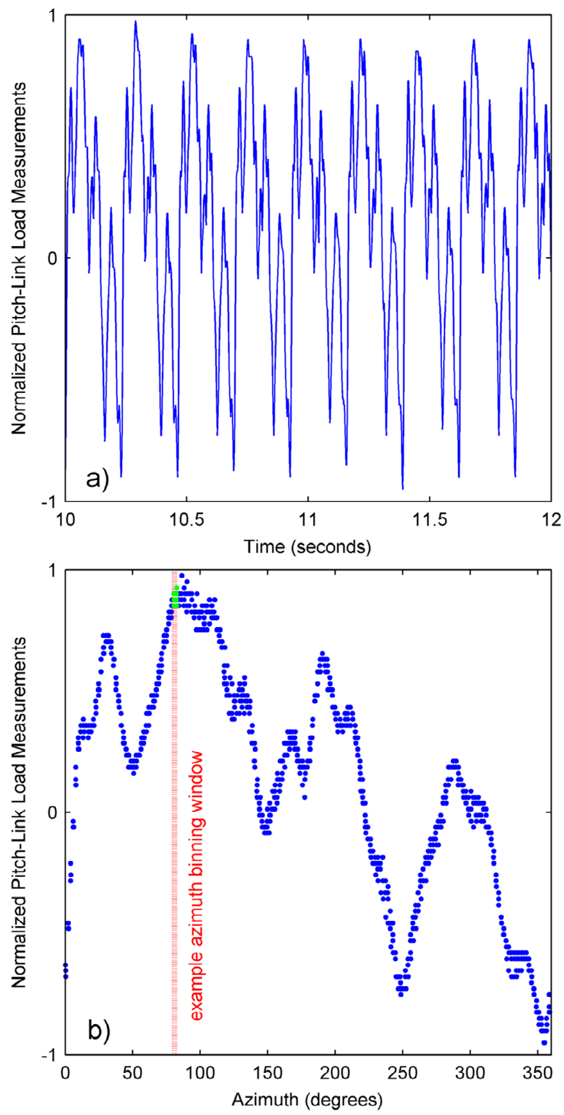

Fig. 1. Illustration of pitch-link measurements: (a) in the time domain, and (b) re-sampled to the azimuth domain.

The construction of the 10-folds data sets is illustrated in Fig. 2, where the tabulated numbers are indexes of level flight runs. Each of the 10 folds contained 16 level flight runs, 2 for each speed regime, with the 2 runs randomly selected from a set of 20 runs at each speed.

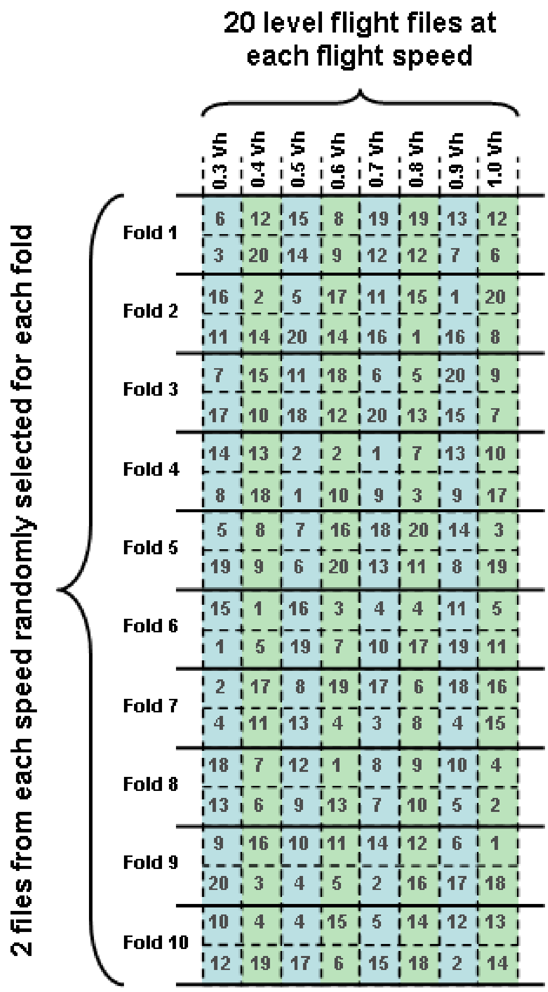

Fig. 2. Illustration of the 10-fold datasets.

3 Regression Approaches

3.1 Overview

The fundamental assumption of this study is a linear relationship between the dependent variable y (here the main-rotor pitch link force), and a collection of independent predictor variables x_i (consisting of strain gauge and accelerometer measurements on the stationary components and, unless otherwise stated, flight state and control parameters[1]). At each azimuth location, Eq. (2) gives a mathematical description of this assumption:

$$y_j = \sum_{i=1}^{n} \hat{\beta}_i x_{ij} + \varepsilon_j, \qquad j = 1, .., m \qquad (2)$$

where, y_j and x_{ij} are normalized values of y and x_i, respectively, at the j-th "event", n is the total number of candidate predictors, m is the total number of "events", $\hat{\beta}_i$ (i = 1, .., n) are to-be-determined coefficients, and ε_j is a residual term accounting for noise and modelling errors. Normalization was performed individually on y and x_i, such that the values were centred and scaled to have a unit length of 1.

When $\mathbf{y}$, $\hat{\boldsymbol{\beta}}$, and $\boldsymbol{\varepsilon}$ are vectors containing y_j, $\hat{\beta}_i$, and ε_j, respectively; and $\mathbf{X}$ is a matrix containing x_{ij}, Eq. (2) can be re-written into a matrix form as in Eq. (3),

$$\mathbf{y} = \mathbf{X}\hat{\boldsymbol{\beta}} + \boldsymbol{\varepsilon} \qquad (3)$$

Eq. (3) can be established, and solved, which leads to a load synthesis model at each of the 360 azimuth locations. To obtain the capability of predicting "unseen" events, the models should only rely on a small subset of the predictor x_i variables, which means that most $\hat{\beta}_i$ are equal to zero. Also, to retain model parsimony, a common set of variables should be selected across all azimuth locations.

The above requirements, especially the second one, are not trivial and cannot be readily fulfilled by existing advanced regression approaches, such as Ridge [9], LASSO (least absolute shrinkage and selection operator) [10] or Elastic-Net [11]. These regression approaches only apply to "independent regressions", where variable selection and model development at one azimuth location are independent from those at other azimuth locations.

A new approach called "concurrent regression" was developed in this study to synchronise variable selection across all azimuth locations, while $\hat{\beta}_i$ values relating to selected variables were independently determined at each azimuth location. This development was made based on two forms of stepwise linear regression, as they were more

[1] The inclusion of flight state and control parameters relates to the first change introduced by the present study.

adaptable than other regression approaches to incorporate the concurrent regression algorithm.

Stepwise regression is a technique which uses feature selection to set most $\hat{\beta}_i$ in Eq. (2) equal to zero. The technique operates by ranking each candidate predictor according to specific criteria and allowing only the highest ranked predictors to have non-zero $\hat{\beta}_i$. This process selects and deselects predictors iteratively with the final models, often chosen using forms of cross-validation, representing the selection of predictors that best predicts the response variable.

3.2 Independent Regression

Two stepwise regression techniques were used within this paper, with their sole difference being the selection criterion used. These techniques were initially used for independent regression and later adapted for concurrent regression.

3.2.1 Stepwise Linear Regression – Adjusted R^2

The selection criterion used within the first technique was the adjusted multiple correlation coefficient, R^2. This value was calculated for each potential predictor subset, and subsequently used to rank predictors. As proposed in [12], R^2 was calculated from the inverse of the correlation matrix, $\mathbf{R_p}$, as elaborated in Eqs. (5) and (6). $\mathbf{R_p}$ was calculated from a matrix defined as $[\mathbf{y}_{\mathbf{TR}}, \text{potential subset of } \mathbf{X}_{\mathbf{TR}}]$, with p being the number of predictors within the subset. To discount potential improvements in R^2 due to chance, R^2 was adjusted according to Eq. (7).

$$(\mathbf{R_p})^{-1} = \begin{bmatrix} a_{11} & .. & a_{1p} \\ .. & .. & .. \\ a_{p1} & .. & a_{pp} \end{bmatrix} \tag{5}$$

$$R_k^2 = 1 - \frac{1}{a_{kk}} \tag{6}$$

$$adjusted\ R_k{}^2 = 1 - \left(1 - R_k{}^2\right)\frac{m-1}{m-p-1} \tag{7}$$

After selection, $\hat{\beta}_i$ values for predictors selected by the technique were calculated using an ordinary least-square approach, while unselected predictors had $\hat{\beta}_i$ equal to zero. The predictive error of the resultant model, $\boldsymbol{\varepsilon}$ within Eq. (3), was then calculated using the fold of the flight data withheld from training.

The above process was repeated ten times, with a different set of data withheld each time, and the mean error was then calculated to assess predictive capability.

3.2.3 Stepwise Linear Regression – RSS

For this technique, the selection criterion *RSS* (residual sum of squares) as defined in Eq. (8) was used to rank and select predictive variables:

$$RSS = \left|\mathbf{y}_{\mathbf{Tr}} - \mathbf{X}_{\mathbf{Tr}}\hat{\boldsymbol{\beta}}\right|^2 \tag{8}$$

As $\hat{\boldsymbol{\beta}}$ is calculated using the ordinary least square approach from the training datasets, Eq (8) can be transformed into:

$$RSS = \mathbf{y}_{\mathbf{Tr}}^{\mathbf{T}}\mathbf{y}_{\mathbf{Tr}} - \mathbf{y}_{\mathbf{Tr}}^{\mathbf{T}}\mathbf{X}_{\mathbf{Tr}}\left(\mathbf{X}_{\mathbf{Tr}}^{\mathbf{T}}\mathbf{X}_{\mathbf{Tr}}\right)^{-1}\mathbf{X}_{\mathbf{Tr}}^{\mathbf{T}}\mathbf{y}_{\mathbf{Tr}} \tag{9}$$

3.3 Concurrent Regression

The purpose of concurrent regression was to select a common set of predictive variables for all 360 azimuth locations. To accommodate this purpose, a concurrent variable selection technique was proposed in this study.

The essence of this technique was to combine ranking from all azimuth locations and select predictor variables that contribute the most across all azimuths. This technique can be integrated with the stepwise linear regression approaches by substituting the variable selection criterion with an average of the same criterion across all azimuth locations.

In a 10-fold cross-validation, the above technique allows for the selection of a common set of predictive variables for each set of training data. However, different sets of predictors might be chosen when different sets of the cross-validation data were used. To account for this, an extra stage of subset reduction was added to the concurrent regression technique.

This new stage of subset reduction operated by combining the predictor subsets selected by each cross-validated model, and discarding the least used predictor. Once a predictor had been eliminated, regression models at each azimuth were rebuilt using this reduced set of predictors. The predictive accuracies of these models were evaluated, recorded, and the process repeated. This stage of subset reduction continued until only one predictor was left.

Once this final stage of subset reduction was complete, the predictive accuracies for the models at all iterations were analysed. The models and subset with the best performance were selected as the final output of the concurrent regression process.

4 Results and Discussion

4.1 Effect of Including Standard Flight State and Control Parameters

The effect of including standard flight state and control parameters was assessed using the azimuth-by-azimuth independent regressions. For each regression approach described in Section 3.2., two classes of regression models were developed: a base-line model that only used strain gauge and accelerometer data as candidate predictors, and an expanded model which also included flight state and control parameters.

Fig. 3 shows the residual error of each expanded model as a percentage of the residual error of the corresponding base-line model, these residual errors were used as a measure of predictive accuracy for each model. It can be seen that the average residual error across all azimuth locations decreased when the additional parameters were included. This improvement in accuracy was seen in each of the regression techniques, with the stepwise regression using adjusted R^2 showing the most improvement. This improvement in accuracy demonstrates that even when using simple linear relationships, the flight state and control parameters provide predictive information not present in the other input data.

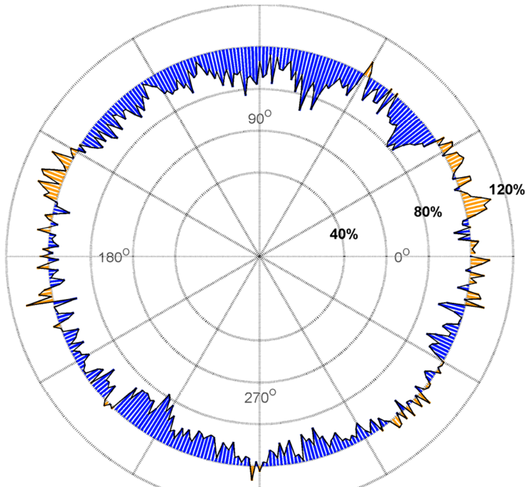

a)

Radial axis: relative error;
Angular axis: azimuth location

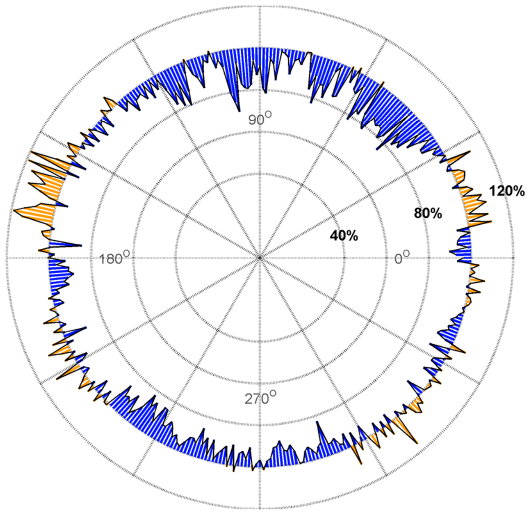

b)

Fig. 3. Polar plots showing expanded model error as percentage of base-line model error. a) Stepwise – adjusted R^2, b) Stepwise – RSS.

Although the overall accuracy improved, it can be seen in Fig. 3 that some azimuth locations experienced a reduction in predictive accuracy for the expanded model. As no information was removed from the expanded model, it may be asked how the models at these azimuth locations can perform worse than the base-line model. This question is not easy to answer because of the complexity of the regression and error estimation processes. A possible explanation may relate to the method used to calculate model accuracy. The residual errors displayed in Fig. 3 are calculated based on how well each model predicts unseen flight test data. As such, any over-fitting of the model to training data would result in poorer performance on unseen data, and therefore an increase in predictive error. It is possible that, for those azimuth locations with poorer performance, the additional predictors tuned the model to aspects of the training data not seen in the validating data, reducing the predictive accuracy at these locations.

In addition to the overall improvements in accuracy, the expanded models rely on fewer strain gauge measurements than the base-line models. This is the case for each of the regression approaches as tabulated in Table 1. As discussed earlier, this reduction in strain gauge usage increases the practicality of the model by reducing the requirement for additional instrumentation.

Table 1. Effects of including flight state and control parameters within regression models.

	adjusted R^2	RSS
Average improvement in predictive accuracy	4.87%	3.88%
Average reduction in strain gauges required	2	4

4.2 Concurrent Regression

The effect of performing concurrent regression was investigated for both regression techniques, with both showing similar levels of improvements. Nevertheless, the technique using the *RSS* variable selection criterion performed slightly better in all aspects and therefore its results are detailed as follows.

In order to compare changes to accuracy and predictor selection as a result of concurrent regression, an independent regression model was developed using the same *RSS* stepwise regression technique. Both models included all flight state and control parameters as candidate predictors. The independent regression model differed from those discussed in section (4.1) by limiting the maximum number of predictor variables used at each azimuth location. This limit was imposed to mimic the same limit which exists in the concurrent regression techniques for computational reasons.

Fig. 4 compares predictor variable selections between the concurrent and independent regressions. A variable is deemed to be selected if it is used by at least one of the 360 azimuth models. For the independent regression, a total of 173 variables were selected. Although nearly all the candidate predictors were selected, most of these were only used in a small number of azimuth models. In contrast, for the concurrent regression, a total of 31 predictive variables were selected across all azimuth locations and each selected variable was used in all 360 azimuth models.

Fig. 5 compares predictive errors between the concurrent and independent regressions. It shows that the concurrent regression led to improved predictive accuracy at most of the azimuth locations with reduced accuracy at a small number of azimuth locations. Overall, the average predictive accuracy was improved despite the reduction of 142 predictor variables. This accuracy improvement may seem counter intuitive at first; however, the reduction of 142 predictor variables must be put into perspective. Although the independent regression model used a total of 173 predictors, Fig. 4 shows that these were spread across all azimuth models. In fact, no one azimuth model used more than 30 predictors. This means that while the concurrent model selected less predictors overall, the efficiency with which they were used was greatly increased. Additionally, although the predictors selected by the independent model were optimized for each location, the highly correlated nature of the flight load data means that, for any predictor, there are a number of alternatives that would perform nearly as well. If one imagines a group of highly correlated predictors, the independent regression may pick a different predictor from this group for each azimuth location. This is in contrast to the concurrent regression which would pick one predictor from the group that best informs all azimuth locations. While this predictor may not perform as well as those chosen specifically at

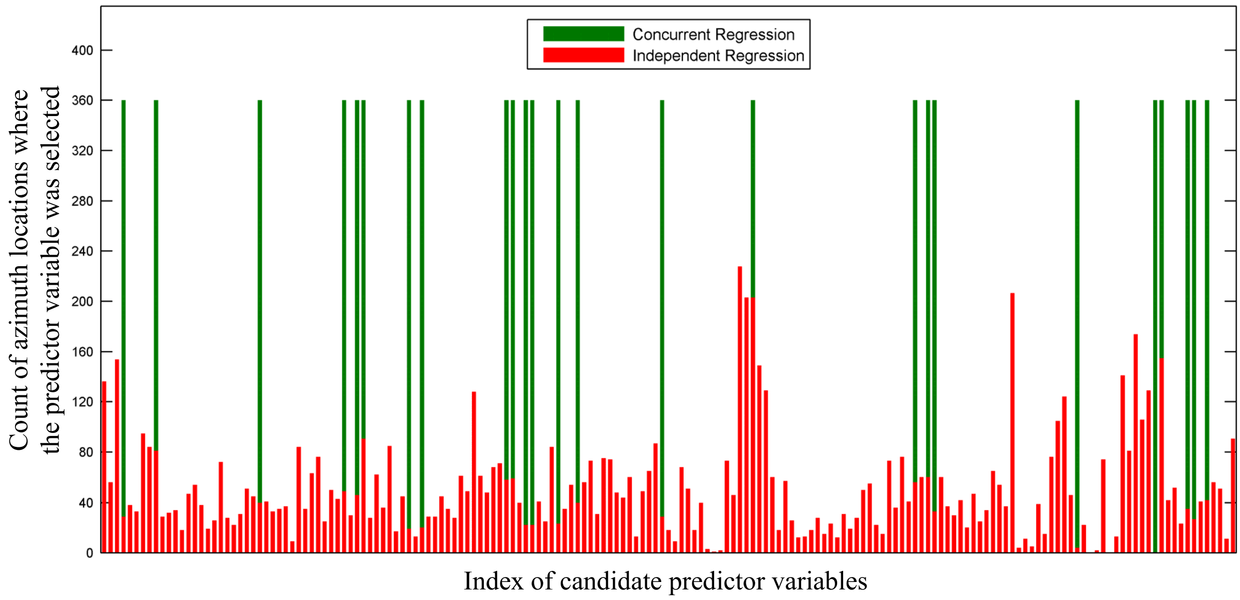

Fig. 4. Selected predictor variables, and frequency of selection across azimuth locations for each regression technique.

each location, for those highly correlated it would perform similarly.

While the above reasons account for a lack of degradation in predictive accuracy, they do not explain the improvement, as seen in Fig. 5. The improvement to predictive accuracy can be explained by a reduction of over-fitting in the model. As shown in Fig. 6 there is a gradual reduction in model error, and therefore an improvement in accuracy, as predictors are eliminated. This gradual reduction in error reaches a minimum at the selected model (N=31) and increases with further elimination of predictors. This gradual reduction in error can be explained as a reverse of the over-fitting effect described in section 4.1. Predictors are eliminated based on their frequency of selection amongst cross-validation folds, with the predictors used least being the first eliminated. The fact that these predictors are chosen by few of the cross-validation tests implies that they are selected because of features specific to their training data. As they bias the model towards the training data, they simultaneously degrade the predictive accuracy of the model towards unseen data. Their elimination from the common subset would therefore improve the predictive accuracy of the overall model.

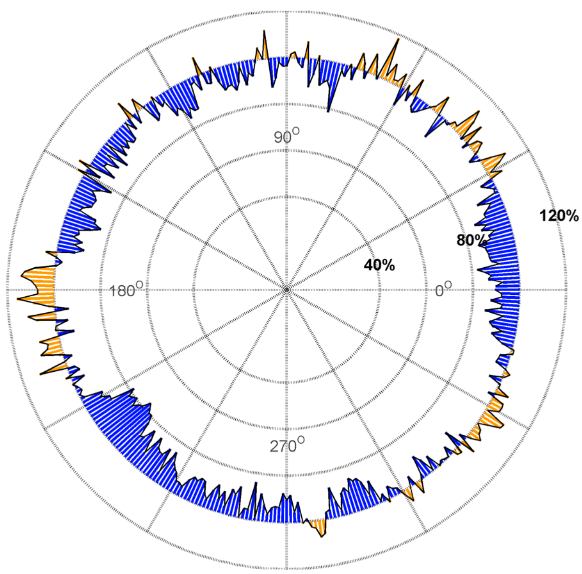

Fig. 5. Polar plot showing predictive errors of concurrent regression model (N=31) as a percentage of predictive errors of the independent regression model (N=173).

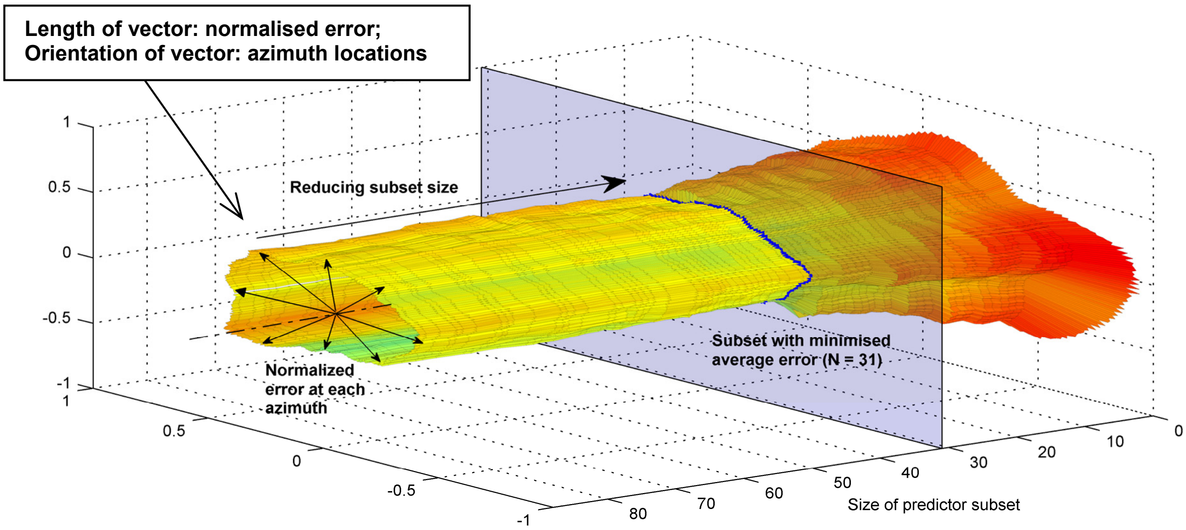

Fig. 6. Surface plot displaying normalized error for all azimuths and predictor subsets.

4 Conclusion

The synthesis of loads on helicopter dynamic components using linear regression in the azimuth domain was improved by two changes to the approach: one was to include the standard flight state and control parameters as candidate predictive variables, and the other was to perform concurrent regression at multiple azimuth locations. The effects of these changes were evaluated in building load synthesis models using two stepwise linear regression techniques to predict main rotor pitch link loads of 160 representative level flight runs selected from a Blackhawk flight strain survey.

It was shown that the inclusion of flight state and control parameters within the model training data improved the predictive accuracy of the model. Also, by adopting the concurrent regression technique, models with improved predictive accuracy could be developed using a significantly reduced set of predictor variables.

5 Acknowledgements

The authors would like to thank their colleagues Luther Krake, Domenico Lombardo, and Christopher Dore for advice and support provided throughout this project.

6 References

1. Lombardo, D. C. (1993) Helicopter Structures—A Review of Loads, Fatigue Design Techniques and Usage Monitoring AR-00-137,
2. Sikorsky Aircraft (1993) Results of the Commonwealth of Australia S-70A-9 Usage Survey and Dynamic Component Impact Assessment. SER-701962,
3. Arms, S. W., et al. (2006) Energy Harvesting Wireless Sensors for Helicopter Damage Tracking. In: Annual Forum Proceedings - American Helicopter Society,
4. Arms, S. W., et al. (2011) Flight Testing of Wireless Sensing Networks for Rotorcraft Structural Health and Usage Management Systems. In: Australian International Aerospace Congress,
5. Yu, X., Wright, C. and Heller, M. (2008) Remote Synthesis of Loads on Helicopter Rotating Components use Linear Regression, Load Path and Statistical Analyses. In: European Rotorcraft Forum, Liverpool, UK
6. Valdès, J. J., Cheung, C. and Wang, W. (2011) Evolutionary Computation Methods for Helicopter Loads Estimation. In: Congress on Evolutionary Computation, New Orleans
7. Georgia Tech Research Institute (2001) Joint USAF-ADF S-70A-9 Flight Test Program, Summary Report. A6186,
8. Sikorsky Aircraft (1986) Dynamic Components UH60A Fatigue Substantiation Summary Report. SER 70131,
9. Hoerl, A. E. and Kennard, W. R. (1970) Ridge Regression: Biased Estimation for Nonorthogonal Problems. Technometrics 42 (1) 02/2000 80-86
10. Tibshirani, R. (1996) Regression shrinkage and selection via the lasso Journal of the Royal Statistical Society Series B 58 (1) 267-288
11. Zou, H. and Hastie, T. (2005) Regularization and variable selection via the elastic net. Journal of the Royal Statistical Society Series B 67 (2) 301-320
12. Raveh, A. (1985) On the Use of the Inverse of the Correlation Matrix in Multivariate Data Analysis. The American Statistician 39 (1) 02/1985 39-42

Copyright Statement