An Approach to Fault Detection Using a Unified Linear Gaussian Framework

R. FUENTES, A. HALFPENNY, E. CROSS, R. BARTHORPE and K. WORDEN

ABSTRACT

Presented in this paper is an approach to fault diagnosis based on a unifying review of linear Gaussian models. The unifying review draws together different algorithms such as PCA, factor analysis, hidden Markov models, Gaussian mixture models and linear dynamical systems as simple variations of a single linear Gaussian model. The focus in this work is on PCA, factor analysis and linear dynamical systems in order to compare static and dynamic data models. One of the advantages of using a unified framework for these already well known models is the quantification of uncertainty, and the ability to evaluate the probability of a dataset given a model. To highlight the applicability of such an approach to a fault diagnosis application and to establish in some way its robustness to environmental variability, a set of data collected from the suspension system of a ground vehicle is used.

INTRODUCTION

It can be argued that one of the basic axioms of SHM is that signal processing and statistical pattern recognition and classification are necessary to extract any damage information from measured signals [1]. The objective of this paper is to present an application of a unifying theory of linear Gaussian models [2] to fault detection. To demonstrate some of the features of using this approach it has been applied to fault detection on a ground vehicle suspension system during a set of trials over different road surfaces, speeds, handling and loading conditions, which establishes the robustness of the methods presented here to these types of environmental variability.

The unifying theory allows application of several widely used *static data* algorithms in SHM, namely Principal Component Analysis (PCA), factor analysis and Gaussian Mixture Models (GMMs) as well as some *dynamic data* models such as Linear Dynamical Systems (LDS) and Hidden Markov Models (HMM) as simple variations of a basic model. In this paper the application of two static data models is presented, namely PCA and factor analysis and a linear dynamic system model. In doing so, it is demonstrated how these methods can be applied to fault detection, and how a unified approach to the system identification of the parameters can be useful in the hope of establishing some advantages of using a unified model for these already well known algorithms within SHM.

UNIFIED MODEL – THEORETICAL BACKGROUND

Basic Model

The aim of this paper is to introduce, in the context of SHM, the application of a unified approach to linear Gaussian models through the use of a single generative linear model. The basic model used is the following discrete time linear Gaussian model:

$$x_{t+1} = Ax_t + w_t , \qquad w \sim \mathcal{N}(0, Q) \tag{1a}$$
$$y_t = Cx_t + v_t , \qquad v \sim \mathcal{N}(0, R) \tag{1b}$$

The main idea behind this model is to represent the *p*-dimensional vector y_t at time t, as a linear combination of another *k*-dimensional vector x_t (called the hidden *state* vector) and y_t (called the *measurement* vector). The measurement and the state vectors are related by the $p * k$ matrix C as shown in equation (1b). Equation (1a) models the dynamics of the state vector with first order Markov dynamics through matrix . In other words, the state vector at time t depends linearly on the state vector at time $t-1$ and the linear map between them is provided by matrix A, called the *state transition matrix*. Both the state and the measurement vectors are modeled as corrupted by zero-mean Gaussian distributed noise, with covariance matrices Q and R respectively.

There are two main questions to be answered in using this model, both readily applicable to its application to SHM. The first question is, given measured data and known model parameters, what are the hidden states and what is the confidence that the measured data behaves according to those parameters? The second question is, given further measured data, what are the model parameters? These two different problems are termed within the machine learning community as *inference*, and *system identification/learning.* In this paper, the Expectation Maximization (EM) algorithm [3] will be used for system identification. As presented in [2], EM provides a unified approach to system identification across the different variations of the basic model, and in doing so provides us with a quantification of the uncertainty in the model parameters in the form of the noise structure encoded in matrices Q and R. The EM algorithm will not be discussed here, but an excellent reference can be found in [4]. The next section will briefly discuss the background theory behind the problem of inference, and how it applies within a fault detection application.

Modeling Dynamic Data – Linear Dynamical System

It is worth starting the discussion with the case of modeling the data as belonging to a linear dynamical system, which makes use of equation (1) as it is. This will later better illustrate the fact that PCA and factor analysis are in a way a simplification of the dynamical system. It will be assumed in this section that the parameters of the model are already known. They can be estimated by making use of the EM algorithm (which in fact makes use of the inference discussed here) and so the focus here will not be on system identification but on the inference problem.

INFERENCE

The inference problem is that of computing the probability distributions of the hidden state sequence $X = [x_1 \dots x_T]$ conditioned on the measurement sequence $Y = [y_1 \dots y_T]$. The measurement vector y_t is simply a vector containing the sensor readings of all *p*-degrees of freedom at time *t*. Through the product rule of probability:

$$P(X|Y) = \frac{P(Y,X)}{P(Y)} \tag{2}$$

where the marginal can be evaluated using the sum rule

$$P(Y) = \sum_{X} P(Y,X) \tag{3}$$

the probability of the states conditioned on the measurements can be estimated. First, an expression for the joint probability must be obtained. The output of a linear system excited by a Gaussian distributed excitation will also be Gaussian distributed [4]. If it is assumed that the initial distribution of the state vector x_1 is Gaussian, then it can be inferred that all the other state vectors will be Gaussian distributed too. The following conditional expectations, inferred from (1) can be written:

$$P(x_{t+1}|x_t) = \mathcal{N}(x_{t+1}|Ax_t, Q) \tag{4a}$$
$$P(y_t|x_t) = \mathcal{N}(y_t|Cx_t, R) \tag{4b}$$

In other words, the expectation of x_{t+1} is conditioned on x_t and the state transition matrix A, and corrupted by noise with covariance *Q*. Similarly, the expectation of y_t is conditioned on x_t and *C*, the linear map between them. Note that (4b) already has a very practical meaning; given known model parameters for the system and the hidden states, the likelihood of a new measurement can be readily assessed. This will be used later on when discussing the suspension system of the ground vehicle. First, the hidden states need to be inferred. Based on (4), the joint probability of the sequence of measurements and states is:

$$P(X,Y) = P(x_1) \prod_{t=1}^{T-1} P(x_{t+1}|x_t) \prod_{t=1}^{T} P(y_t|x_t) \tag{5}$$

Since $P(x_{t+1}|x_t)$ and $P(y_t|x_t)$ are Gaussian distributed, as per (4), the log of the right hand side of (5) can be evaluated in order to express it explicitly as:

$$\log P([x_1 \dots x_T], [y_1 \dots y_T]) \tag{6}$$
$$= -\frac{1}{2}\sum_{t=1}^{T}[\,(y_t - Cx_t)'\,R^{-1}\,(y_t - Cx_t)] + \log|R|]$$
$$-\frac{1}{2}\sum_{t=1}^{T-1}[\,(x_{t+1} - Ax_t)'\,Q^{-1}\,(x_{t+1} - Ax_t) + \log|Q|\,]$$
$$-\frac{1}{2}(x_1 - \mu_1)'Q_1^{-1}\,(x_1 - \mu_1) - \log|Q_1| - \frac{T(p+k)}{2}\log 2\pi$$

where $Q_1$and μ_1 are the covariance and expectations of the initial state vector. The above expression can be used in order to sum over the joint probability required in equation (3). Inference of the hidden state now requires the evaluation of the expectation of the state:

$$\widehat{x_t} = E[x_t|Y] \tag{7}$$

This is done using the above probabilities implemented in an extension of the Kalman filter, the Rauch-Tung-Striebel smoother [5]. Note also the caret on $\widehat{x_t}$, this denotes an *estimate*. A further note on distinguishing between filtering and smoothing; A filter evaluates the above expectation from the distribution $P(x_t\,|[y_1 \dots y_t])$ whereas the smoother evaluates the expectation under $P(x_t \mid Y)$ which includes all of the measurement points before and after time step *t*.

It is also worth noting that the eigenvalues of the state transition matrix, *A*, contain the natural frequencies of the underlying dynamical system, and its eigenvectors can be used together with *C* to extract the mode shapes of the structure [6].

Modeling Static Data – PCA and Factor Analysis

When modelling the data as static, it is assumed that it has no dynamics; there is no temporal dependence between the states at different times. The state transition matrix for this case is 0, which reduces (1) to:

$$x = w \qquad w \sim \mathcal{N}(0, Q) \tag{8a}$$
$$y = Cx + v\,, \qquad v \sim \mathcal{N}(0, R) \tag{8b}$$

Note above that we have dropped the time index since no temporal dependence is being assumed between the states. The above equations would still hold if the measurement sequence $[y_1 \dots y_T]$ was to be randomly reshuffled. Through the use of (6) to evaluate the integral in (3) it is possible to arrive at an expression for the likelihood of a measurement under the static data assumption:

$$P(y) = \mathcal{N}(y|\,0, CQC' + R) \tag{9}$$

where there is degeneracy in the matrix product *CQC'*, so *Q* is restricted to be the identity matrix so that

$$P(y) = \mathcal{N}(y|\,0, CC' + R) \tag{10}$$

Note now that this has practical values again; if the noise structure of the model and the linear transformation *C* (from the measurements to the hidden states) are both known, it is possible to evaluate the likelihood of any given measurement. In order to infer the hidden states given the measurements, Bayes' rule is used and the result for the marginal in equation (10), [2] [7], is:

$$P(x|y) = \frac{P(y|x)P(x)}{P(y)} \tag{11a}$$

$$= \frac{\mathcal{N}(y|\,Cx, R)\,\mathcal{N}(x|\,0, I)}{\mathcal{N}(y|\,0, CC' + R)} \tag{11b}$$

$$= \mathcal{N}(x|\,\beta y, I - \beta C\ ) \tag{11c}$$

where

$$\beta = C'(CC' + R)^{-1} \tag{11d}$$

The above evaluation provides an expression for the expectation of the hidden states$E[x|y] = \beta y$, dependent on the measurement vector as well as a quantification of the uncertainty in such estimate in the form of the covariance $cov[x|y] = I - \beta C$.

PROBABILISTIC PRINCIPAL COMPONENT ANALYSIS

Principal component analysis is used extensively as a dimensionality reduction technique in SHM. The main idea of PCA is to express a set of multivariate data in terms of a set of reduced (principal) components which can be obtained through an axes rotation (a linear transformation), which is specified by the eigenvectors of the covariance of the original dataset. This definition of PCA does not capture any uncertainties in the estimation of the principal axes. Roweis [7] and Bishop [8] introduced an alternative approach through the use of a linear Gaussian model, which is probabilistic in nature. It will be referred to this here as Probabilistic Principal Component Analysis.

The basis for the linear model for PCA is equation (8). Here, the columns of the measurement matrix *C*, are the principal axes and the hidden states/variables *x*, are the principal components. The uncertainty is captured by the measurement noise covariance *R*. Both *C* and *R* can be learned using the EM algorithm. The covariance matrix *R* needs to be constrained somehow in order for the learning not to choose $C = 0$, and to explain all of the variability of the data as noise. It turns out that choosing $R = \alpha I$, that is, restricting it to be a diagonal matrix with equal valued entries yields a matrix *C* that is representative of the principal axes. Note that this means that the model assigns a "global" uncertainty value; it is a single number for the whole dataset. Once *C* and *R* are known inference can be performed using (11) and the likelihood of new data can be evaluated using (10)

FACTOR ANALYSIS

Factor analysis has long been recognised as closely related to PCA. Under this unifying view, factor analysis can also be represented by equation (8). For factor analysis, R is constrained in a slightly different way to model the noise structure differently. In fact, R it is still constrained to be diagonal, but the entries are free to take on different values. In practice this means that it assumes that the noise from different sensors is uncorrelated, but quantifies the uncertainty in the estimation of the parameters for each degree of freedom being measured. The parameters C and R can be estimated using the EM algorithm, and once these are known, the hidden states, or 'factors' can be estimated using equation the likelihood of new data can be evaluated using equation (10).

APPLICATION TO GROUND VEHICLE SUSPENSION SYSTEM

These linear Gaussian models are very useful for SHM applications [9]. The methods described above will be demonstrated here on a set of data collected from the suspension system from a military ground vehicle during proving ground trials. The main idea is to demonstrate the use of PCA, factor analysis and linear systems from a probabilistic perspective. This is achieved by system identification using the EM algorithm on one of the test runs and subsequent use of inference and the use of equations (4) and (10) to evaluate the likelihood of the data at each time step. This is in contrast to previous usage of the same methods [some quotes here] which mostly try and detect faults from the hidden states (the factors, principal components, or Kalman states) or by performing some type of parameter estimation or system identification for each new run and then comparing these parameters.

Acceleration data was collected from the vertical direction from transducers placed at the damper tops, wheel centres and rocker midpoints. The data was collected during proving ground trials over 10 different road surfaces and a range of different speeds. This range of speeds and road surfaces accounts for a wide range of environmental variability. In order to present the results concisely only 14 test runs were chosen, which includes 4 different road surfaces and a range from 15 to 25 miles per hour runs. Details of these are shown in Table I. Figure 1 shows the results for all of these runs as boxplots of the likelihoods evaluated using different methods at each time step. The likelihoods assess what is the probability is of the data given the model. Note that in Figure 1 these are log-likelihoods.

The first column in Figure 1 shows these likelihoods evaluated for the original data in order to assess the robustness of these methods to environmental variability. There are no actual faults present in the original data and so a pseudo-fault and a sensor fault have been simulated. The simulated sensor fault is simply a change in gain in one of the accelerometers following the case studies by Isermann and colleagues [10], [11], [12] . Adding noise or an offset to one of the sensors produces similar results but this is not presented here for conciseness. The pseudo-fault is a single-degree-of-freedom FIR filter applied to one of the damper tops in order to simulate a failing shock absorber [13].

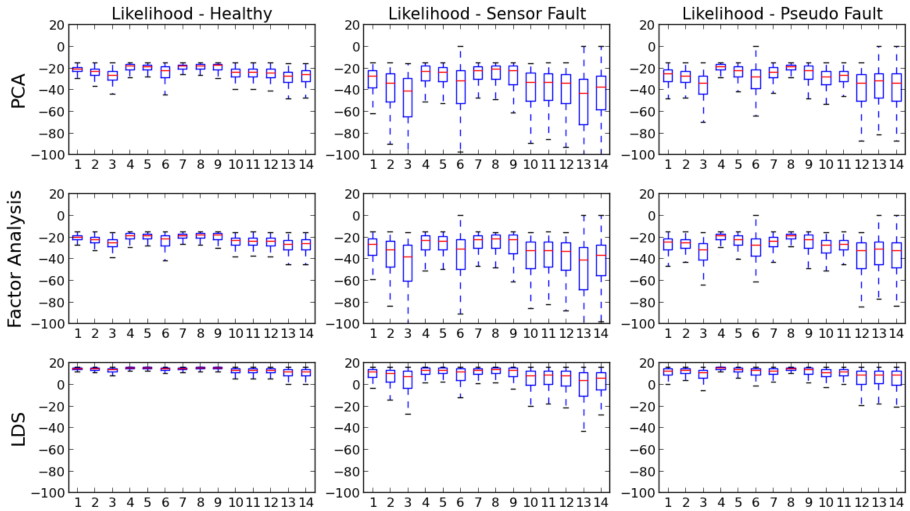

Figure 1. Box plots from likelihoods evaluated using factor analysis, PCA and linear dynamics on the ground vehicle suspension. The horizontal axis is the run number which corresponds to that of Table 1. The likelihoods are evaluated for the “healthy” condition, a simulated sensor fault and a simulated shock absorber fault.

TABLE I – ROAD PROFILES AND SPEEDS

Run id	Road Profile	Vehicle Speed (mph)
1	Belgian Block	15
2	Belgian Block	20
3	Belgian Block	25
4	Bumps	10
5	Bumps	15
6	Bumps	20
7	City Road	10
8	City Road	15
9	City Road	20
10	City Road	30
11	City Road	35
12	Gravel	25
13	Gravel	30
14	Gravel	35

DISCUSSION OF RESULTS AND CONCLUDING REMARKS

So far, a unifying review of linear models, introduced in [2] has been described and its applicability to SHM has been highlighted. It has also been demonstrated how the likelihood of data modeled under this framework can be evaluated using static and dynamic data models. It has been shown how the static data models correspond to the well-known methods of PCA and factor analysis, and the dynamic data model corresponds to a linear dynamical system. An example was drawn by applying these methods to test data from a ground vehicle suspension system. A key result observed from this, in Figure 1, is that evaluating data likelihoods under these "learned" models is somewhat robust to variations in environmental conditions. As expected, the likelihoods evaluated by the linear dynamics system model show smaller variance in the healthy state than both static data cases. It is also clear that PCA and factor analysis yield similar results, which is expected since the model for the data is very similar, with the only difference being that (probabilistic) PCA models the noise structure globally whereas factor analysis does it for each individual degree-of-freedom.

The static data models are less robust to environmental variability than the dynamic data model. This is expected as the data being modeled is dynamic in nature, and so a signal model that conditions current values based on previous values will do better than a model that does not. Although the models are robust to changes in environmental conditions, they are also flexible enough to capture *change* in the data. The simulated faults are not necessarily representative of a real fault, but serve the purpose of demonstrating how change in the data affects its likelihood (with respect to the models) and how using this framework, even in a crude way as was done with PCA and factor analysis can capture change in the data.

REFERENCES

[1] K. Worden, C. Farrar, G. Manson and G. Park, "The Fundamental Axioms of Structural Health Monitoring," *Proceedings of the Royal Society A,* 2007.

[2] S. Roweis and Z. Ghahramani, "A Unifying Review of Linear Gaussian Models," *Neural Computation,* vol. 11, no. 2, 1999.

[3] A. Dempster, N. Laird and D. Roubin, "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society,* vol. 39, no. 1, pp. 1-38, 1977.

[4] C. M. Bishop, Pattern Recognition and Machine Learning, Springer, 2006.

[5] H. Rauch, F. Tung and C. T. Striebel, "Maximum likelohood estimates of linear dynamic systems," *AIAA J,* vol. 3, no. 8, pp. 1445-1450, 1965.

[6] F. J. Cara, J. Carpio, J. Juan and E. Alarcon, "An approach to operational modal analysis using the expectation maximization algorithm," *Mechanical Systems and Signal Processing,* vol. 31, pp. 109-129, 2012.

[7] S. Roweis, "EM Algorithms for PCA and SPCA," in *in Advances in Neural Information Processing Systems,* MIT Press, 1998, pp. 626 - 632.

[8] E. M. Tipping and C. M. Bishop, "Probabilistic Principal Component Analysis," *Journal of the Royal Statistical Society, Series B,* vol. 61, no. 3, pp. 611-622, 1999.

[9] K. Worden and C. Farrar, Structural Health Monitoring: A Machine Learning Perspective, 2013: John Wiley & Sons.

[10] M. Börner, H. Straky, T. Weispfenning and R. Isermann, "Model based fault detection of vehicle suspension and hydraulic brake systems," *Mechatronics,* vol. 12, no. 8, pp. 999-1010, 2002.

[11] R. Isermann, Fault-Diagnosis Systems; An Introduction from Fault Detection to Fault Tolerance, Springer, 2006.

[12] T. Weispfenning, "Fault Detection and Diagnosis of Components of the Vehicle Vertical Dynamics," *Meccanica,* vol. 32, pp. 459-472, 1997.

[13] Bosch Automotive Handbook, 8th edition, Bentley Publishers, 2011.