A nonlinear state-space model for the forecasting of field failures

Un modello state-space non lineare per la previsione di guasti post vendita

Antonio Pievatolo

Abstract We consider time series of field failure data (warranty claims) of domestic appliances, manufactured by different plants, with the aim of forecasting failures within the warranty period. The monthly failure profiles over two-year periods display variation across monitoring epochs and also batch-to-batch variation. A nonlinear state space model is developed to jointly represent the variation of the underlying failure rate parameters and the observed occurrence of failures, obtaining a dynamic Poisson-Lognormal model with a meaningful covariance structure of failure rates between monitoring epochs. An adaptation of the auxiliary particle filter is used for parameter learning and forecasting. A series of examples with data from two different production plants show that it is possible to obtain a small forecasting error for claims having very different patterns.

Abstract *Analizziamo serie storiche di riparazioni in garanzia di elettrodomestici, prodotti da impianti differenti, per prevedere la frequenza di guasti entro la fine del periodo di garanzia. I profili di guasto mensili in periodi di due anni mostrano variazioni sia tra periodi di monitoraggio sia tra lotti di produzione. Sviluppiamo un modello state space per rappresentare allo stesso tempo la variazione del parametro del tasso di guasto sottostante e l'apparizione dei guasti, ottenendo un modello dinamico Poisson-Lognormale con una appropriata struttura di covarianza tra periodi di monitoraggio. Utilizziamo un adattamento dell'auxiliary particle filter per la stima dei parametri statici e per la previsione. Alcuni esempi con dati da due impianti di produzione mostrano che è possibile ottenere un piccolo errore di previsione per serie di guasti post-vendita con profili molto diversi.*

Key words: Warranty claims, State-space model, Poisson-Dirichlet model, Poisson-Lognormal model, Particle learning, Covariance of failure rates

---

Antonio Pievatolo
CNR IMATI, Via A. Corti, 12, Milano, e-mail: antonio.pievatolo@cnr.it

1 Introduction

Warranty data have long been recognized as a source of useful information for several purposes, some of which are prediction of future claims, comparison of groups of products and identification of faulty design or material defects or undetected production line problems. [5] provided one of the first review articles on this subject.

Claims are often aggregated by epoch index (e.g. by month) for every production batch, so that available data may be represented by n_t, for the number of items produced at epoch t, and y_{tj}, for the number of claims at epoch $t+j-1$ for batch t, $j = 1,\dots,d$, where d is the duration of the warranty period.

[5] proposed to model the claim arrival process of a given production batch as a sequence of independent Poisson random variables: $Y_{tj} \sim Poisson(n_t\lambda_j)$, as $j = 1,\dots,d$, where λ_j is the expected number of failures per produced item in batch t at epoch $t+j-1$. [9] used the same model for the early detection of reliability problems. Other works are focussed on the modelling of failure times instead ([3]; [4]).

The model based on independent Poisson counts has the main drawback that λ_j does not depend on t so it cannot adequately describe batch-to-batch variation, which could be attributed to: material defects due for example to a change of supplier; changes in the production line that may affect reliability; other unmeasured batch to batch heterogeneity.

This deficiency could be addressed following [6], who introduced a rate function for car warranty data, depending on a unit-specific random effect Z_t and a usage function modelled as $U_{tj} = jZ_t$. This approach requires to select a parametric form for the rate function and does not include possible dependence among unit failure rates at close t values. Furthermore it was developed for situations in which times of occurrence of failures of each unit are available, which is not our present situation.

In this work we have proposed a state space modelling framework in which the observation equation for (conditionally independent) describes a Poisson distribution, whereas the batch-to-batch heterogeneity and the possible dependence among failures of units from batches close in time is represented by a state transition equation on failure rates. In this way the failure rates are regarded as generated by a stochastic process and no assumption on their form is needed. On the other hand, we will focus on the specification of a meaningful dependence structure among failure rates and on a particle learning algorthm, which will provide forecasts on future failure rates before the end of the warranty period and also learn about unknown model parameters.

Our methodology has been applied to claims regarding home appliences manufactured by a multinational company with markets and production plants in several countries.

2 A state space model

Let $y_t = (y_{t1}, \ldots, y_{td})^T$ and $\lambda_t = (\lambda_{t1}, \ldots, \lambda_{td})^T$, where, unkile the introductory section, λ_t now denotes a vector of failure rates. We consider a state space model for claims

$$\begin{aligned} y_{tj} &\sim Poisson(n_t \lambda_{tj}), \quad j = 1, \ldots, d \\ \lambda_t | \lambda_{t-1} &\sim f(\lambda_t | \lambda_{t-1}; \theta) \end{aligned} \tag{1}$$

where θ is an unknown parameter. This model is flexible enough to be able to describe batch-to-batch variation, dependence between batches, dependence between claim numbers in different epochs for the same batch, variation of claim reporting patterns. The task is now to select the form of the state transition equation, which we will do by examining observations from two production plants, where epochs are months, in Figure 2. The left-hand panels show a substantial batch-to-batch variation of observed failure rates for entire batches (number of failures in two years divided by the batch size). The right-hand panels highlight the within-batch variation, that is, how the overall number of failures in two years is distributed over the epochs (months) for all observed batches, with variablity in the shape of the curves.

A model which separates these two types of variation is the following Poisson-Dirichlet model

$$\begin{aligned} y_{tj} &\sim Poisson(n_t \mu_t p_{tj}), \quad j = 1, \ldots, d \\ \log(\mu_t) &= \log(\mu_{t-1}) + \sigma w_t \\ p_t &= \gamma p_{t-1} + (1-\gamma) q_t \end{aligned} \tag{2}$$

where μ_t is a scalar, $q_t \sim Dir(M\eta)$, a d-dimensional Dirichlet distribution with mass parameter M and shape parameter $\eta = (\eta_1, \ldots, \eta_d)$ such that $\sum_j \eta_j = 1$, and $\gamma \in (0,1)$, $w_t \sim N(0,1)$. The unknown parameter θ includes σ, η and M. The second equation describes the dynamics of the overall batch failure rate, whereas the third one describes the within-batch variation, also defining a covariance structure among epochs in the same batch via the Dirichlet error q_t.

A more tractable version of this model, with a view to sequential Bayesian parameter update, can be obtained by collapsing the two state equations into one. Letting $\lambda_{tj} = \mu_t p_{tj}$, we derive the following:

$$\log \lambda_{tj} = \log \lambda_{t-1,j} + \varepsilon_{tj} \tag{3}$$

where $\varepsilon_{tj} = \log \lambda_{t-1,j} + \log(\gamma + (1-\gamma) q_{tj}/p_{t-1,j}) + \sigma w_t$ and the vector ε_t has a non-diagonal covariance matrix, which we can approximate via error-propagation formulas from the covariance structure of q_t. By doing so, we find

$$Var(\varepsilon_{tj}) \simeq \frac{(1-\gamma)^2 \eta_j (1-\eta_j)}{M+1} \frac{1}{(\gamma p_{t-1,j} + (1-\gamma)\eta_j)^2} + \sigma^2$$

which is further approximated by

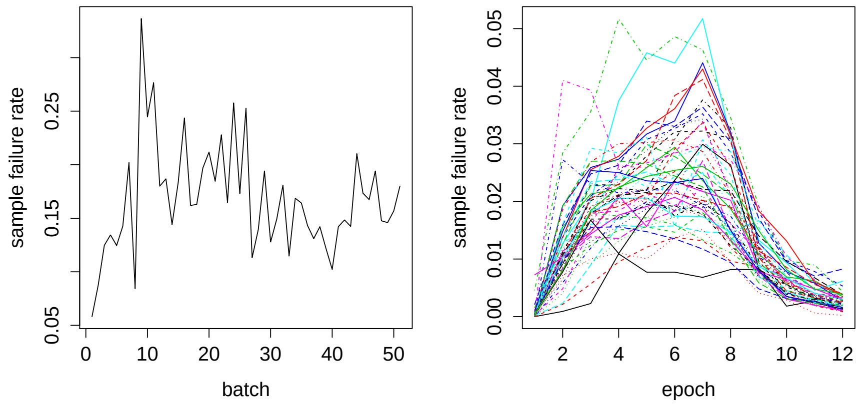

First plant

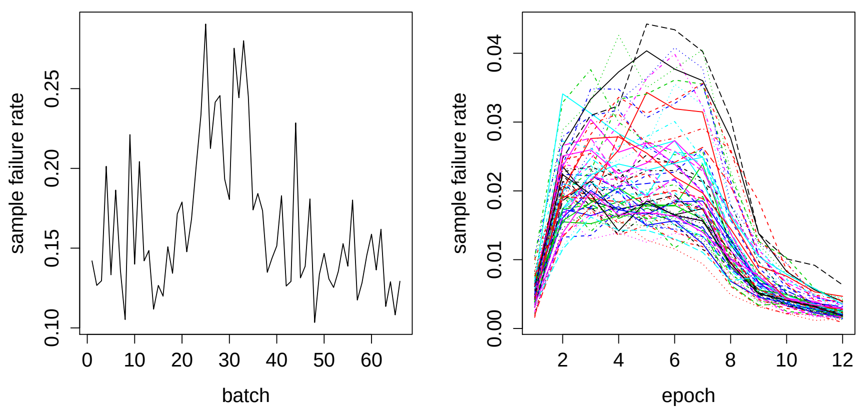

Second plant

Fig. 2: First column: plot of pairs $(t, \sum_j y_{tj}/n_t)$; second column: plots of pairs $(j, y_{tj}/n_t)$, for all available values of t

$$\frac{(1-\gamma)^2}{M+1}\frac{(1-\eta_j)}{\eta_j} + \sigma^2 = \tau^2 \frac{(1-\eta_j)}{\eta_j} + \sigma^2$$

using η_j for $p_{t-1,j}$. Continuing with the covariances, by the error propagation formulas and using η_j for $p_{t-1,j}$, $Cov(\varepsilon_{tj}, \varepsilon_{tr}) \simeq -\tau^2/2 + \sigma^2$ and finally $E(\varepsilon_{tj}) \simeq 0$.

Then, the new Poisson-Lognormal state-space model is defined as

$$
\begin{aligned}
y_{tj} &\sim Poisson(n_t \mathrm{e}^{\alpha_{tj}}), \quad j = 1, \ldots, d \\
\alpha_t &= \alpha_{t-1} + \varepsilon_t
\end{aligned}
\tag{4}
$$

where $\alpha_t = \log \lambda_t$, $\varepsilon_t \sim N_d(0, \Sigma)$ and

$$
\begin{aligned}
\Sigma_{jj} &= \tau^2 \frac{(1-\eta_j)}{\eta_j} + \sigma^2, j = 1, \ldots, d \\
\Sigma_{jr} &= -\frac{\tau^2}{2} + \sigma^2, j, r = 1, \ldots, d,\ j \neq r\ .
\end{aligned}
\tag{5}
$$

Model (4)-(5) is not equivalent to the original model (2) and the normality of the error term has been assumed ex post, however this approximation procedure has provided a justification for applying a certain covariance structure to the logarithms of the failure rates. Furthermore, conditional on α_{t-1}, the model for (y_t, α_t) is the multivariate Poisson-lognormal model of [1], for which $Var(Y_{tj}) > E(Y_{tj})$, allowing for overdispersion relative to the Poisson distribution.

3 Particle filtering for claim forecasting

Let $\alpha^t = (\alpha_1, \ldots, \alpha_t)$ and $y^t = (y_1, \ldots, y_t)$. The filtering distribution at time t is

$$f(\alpha_t | y^t)$$

and it encodes the state of information on the claim rate vector given the observed claims. However, the needed warranty claim rate information for batch t is already provided by $\sum_j y_{tj}/n_t$, and the filter becomes useful for an early assessment of the overall claim when only a part of y_t is observed. So we seek to obtain

$$f(\alpha_t | y^{t-1}, y_{t1}, \ldots, y_{tr_t}), \quad r_t < d,$$

where r_t is the latest observed epoch for batch t, taking advantage of the covariance structure of our state-space model.

By combining the methodology of the auxiliary particle filter (APL) of [8] and of the particle learning method of [2], we have obtained a new modified APL which also includes parameter learning on Σ and converges to the correct filtering and posterior distributions. The filter has not been applied to model (4)-(5), but to a relaxed version to allow for parameter updating using conjugacy. In particular, an inverse-Wishart initial distribution has been assigned to Σ_0, that is, $\Sigma \sim iW(\Sigma_0, \nu_0)$. Then, given that Σ is independent from y^t given α^t, its posterior distribution, conditional on the past history of the observations and of the states, is

$$\Sigma | s_t \sim iW(\Sigma_0 + \sum_{r=1}^{t} (\alpha_r - \alpha_{r-1})(\alpha_r - \alpha_{r-1})^T, \nu_0 + t) \tag{6}$$

where $s_t = \sum_{r=1}^{t} (\alpha_r - \alpha_{r-1})(\alpha_r - \alpha_{r-1})^T$. At time $t+1$ this distribution can then be updated by knowing only the sufficient statistics s_t and $(\alpha_{t+1} - \alpha_t)$, which, together, give s_{t+1}. The covariance structure (5) is not discarded, but is used to provide the Σ_0 parameter of the initial distribution of Σ, which, as experiments showed, must been assigned carefully.

The derivation of our modified APL is not shown here, for reasons of space. Instead we describe now the filtering algorithm, which provides, at each time t, a discrete approximation of the filtering density as a set of support points with associated weights. Additional yet undefined notation used for the algorithm, is L, the likelihood function determined by the observation equation in model (4) and the function f, which denotes a conditional density.

Given a step-t sample $(\alpha_t, s_t, \Sigma)^{(i)}$ and weights $w_t^{(i)}$, the modified APL goes through the following steps:

1. *Resampling step*: sample N index values k_i, $i = 1, \ldots, N$, using weights proportional to $L(\mu_{t+1}^{(k)}; y_{t+1}) w_t^{(k)}$, putting $\mu_{t+1}^{(k)} = \alpha_t^{(k)}$
2. *Propagation step*: sample N new particles $\alpha_{t+1}^{(i)}$ from $f(\alpha_{t+1} | \alpha_{t,ADJ}^{(k_i)}, \Sigma^{(k_i)})$ where $\alpha_{tj,ADJ}^{(k_i)} = \log(y_{t+1,j}/n_{t+1})$ if $j \le r_{t+1}$, whereas $\alpha_{tj,ADJ}^{(k_i)} = \alpha_{tj}^{(k_i)}$ if $j > r_{t+1}$
3. Compute weights

$$w_{t+1}^{(i)} \propto \frac{L(\alpha_{t+1}^{(i)} | y_{t+1}) f(\alpha_{t+1}^{(i)} | \alpha_t^{k_i}, \theta^{(k_i)})}{L(\mu_{t+1}^{(i)} | y_{t+1}) f(\alpha_{t+1}^{(i)} | \alpha_{t,ADJ}^{k_i}, \theta^{(k_i)})}$$

update sufficient statistics $s_{t+1}^{(i)} = s_t^{(k_i)} + (\alpha_{t+1}^{(i)} - \alpha_t^{(k_i)})(\alpha_{t+1}^{(i)} - \alpha_t^{(k_i)})^T$ and sample $\Sigma^{(i)} \sim iW(\Sigma_0 + s_{t+1}^{(i)}, \nu_0 + t + 1)$

4 Data examples

Because of the abundance of training data, the prior can be pretty informative. Therefore the initial distribution for Σ is $iW((R+1)\Sigma_0, R+d+2)$, with large R. With this parameterization, $E(\Sigma) = \Sigma_0$ a priori. The initial Σ_0 is computed using $\eta_j = \sum_t y_{tj} / \sum_t n_t$ from a training sample (such as another plant), then the filtered claim rate for batches with r_t observed claim epochs, $r_t < d$, is estimated as

$$\frac{1}{n_t} \left(\sum_{j=1}^{r_t} y_{tj} + \sum_{j=r_t+1}^{d} \sum_{i=1}^{N} w_t^{(i)} n_t \lambda_{tj}^{(i)} \right) .$$

For predictive interval forecasts we resample particles and draw

$$\sum_{j=r_{t+1}}^{d} y_{tj}^{(i)} \sim Poisson(n_t \sum_{j=r_{t+1}}^{d} \lambda_{tj}^{(i)})$$

for every resampled particle. Then we take sample quantiles. The result of the forecasting procedure, including prediction intervals, is displayed in Figure 4, showing a good performance even with very little information on the current production batch.

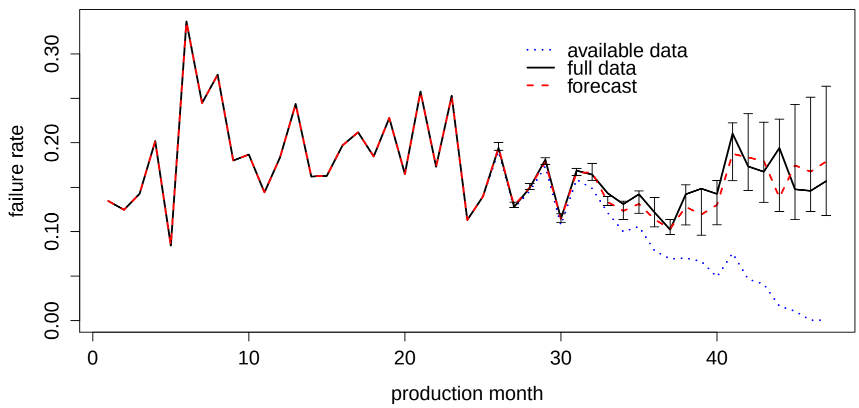

First plant

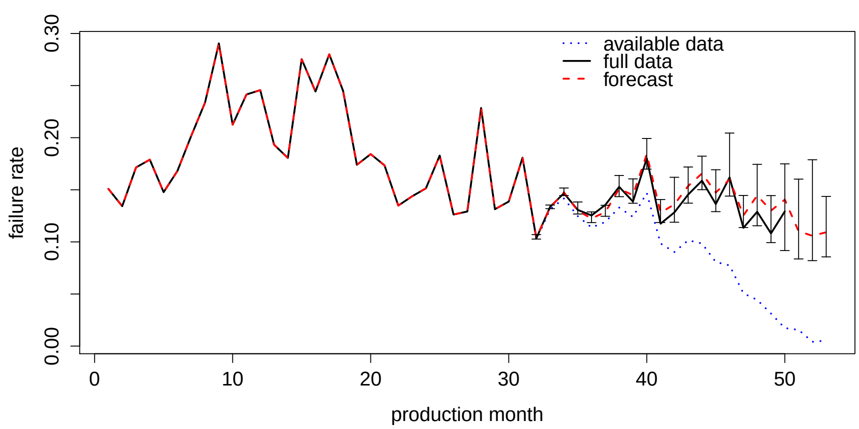

Second plant

Fig. 4: Forecasts of claims for the two example plants versus production batch. Solid line: complete data; dashed line: forecast; dotted line: ratio between available number of claims and batch size at the time of the forecast, representing the available information (the lower the less)

5 Conclusions

The state-space Poisson-Lognormal model developed in this work has shown a good potential for making early prediction of future claims from customers during the warranty period of a domestic appliance, thanks to the design of an appropriate covariance structure of within-batch failure rates and to parameter learning. This approach is ideal for the sequential monitoring and prediction of claims when they occur as counts in predefined monitoring intervals, without the need of any detailed modelling of known disturbances such as claim reporting delays and of the usage pattern of appliances. Experiments indicate that a good elicitation of the initial value of the covariance matrix is requested, which is possible in the present situation because of the abundance of data, therefore future work can be directed to the exploration of results using a less concentrated initial distribution, as well as to other parameter learning strategies that improve the convergence of the particle filter.

Acknowledgements This work is a follow-up to [7]

References

1. Aitchison, J., Ho, C.H.: The multivariate Poisson-log normal distribution. Biometrika, **76**, 643-653 (1989)
2. Carvalho, C.M., Johannes, M.S., Lopes, H.F., Polson, N.G.: Particle learning and smoothing. Statistical Science, **25**, 88-106 (2010)
3. Hong, Y., Meeker, W.Q.: Field-failure and warranty prediction based on auxiliary use-rate information. Technometrics, **52**, 148-159 (2010)
4. King, C., Hong, Y., Meeker, W.Q.: Product Component Genealogy Modeling and Fieldfailure Prediction. Qual. Reliab. Eng. Int., **33**, 135-148 (2017)
5. Lawless, J.F.: Statistical analysis of product warranty data. Int. Stat. Rev., **66**, 41-60 (1998)
6. Lawless, J.F., Crowder, M.J., Lee, K.A.: Analysis of reliability and warranty claims in products with age and usage scales. Technometrics, **51**, 14-24 (2009)
7. Pensa, R.: Particle filters: un'applicazione per la previsione di guasti di elettrodomestici. Master thesis. Politecnico di Milano (2015)
8. Pitt, M.K., Shephard, N.: Filtering via simulation: Auxiliary particle filters. Journal of the American statistical association, **94**, 590-599 (1999)
9. Wu, H., Meeker, W.Q.: Early detection of reliability problems using information from warranty databases. Technometrics, **44**, 120-133 (2002)