Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira, Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria

doi: 10.3850/ESREL2026061419_esrel26-p26420-cd

CBM Policy Optimization Based on Lévy Copula Modeling of Latent Degradation

Connor Stewart-Green

School of Mathematics, Statistics, Chemistry and Physics, Murdoch University, Australia.
E-mail: cstewartgreen@gmail.com

Soudabeh Shemehsavar

School of Mathematics, Statistics, Chemistry and Physics, Murdoch University, Australia.
E-mail: Soudabeh.Shemehsavar@murdoch.edu.au

Graeme Hocking

School of Mathematics, Statistics, Chemistry and Physics, Murdoch University, Australia.
E-mail: G.Hocking@murdoch.edu.au

Within practical systems, the continuous buildup of wear and tear is an unavoidable reality. Industries subject to such systems must account for degradation to maximize runtime while minimizing costs. Among common methods, Condition-Based Maintenance (CBM) uses observable system conditions to make decisions. This is considered both cost effective and reliable with regard to planning system maintenance. Latent degradation refers to any degradation that is hidden within a system and can cause breakdowns without warning. CBM can use correlated performance characteristics, called markers, to make predictions about latent degradation for decision making. Recent evidence suggests that a copula-based approach may be better at capturing the dependence structure between latent degradation and markers compared to previous methods. This study investigates the use of the Lévy copula function to model the dependence between observed markers and latent degradation, enabling prediction of latent degradation from system markers.

Keywords: Condition-Based Maintenance, Latent Degradation, Marker Process, Lévy Copula

1. Introduction

Within industrial contexts there are few matters as pervasive as maintenance. Whether you are considering large-scale assembly lines, power generating facilities, or a simple light bulb, maintenance is needed everywhere. To turn a profit, maintenance must be optimized to maximize system availability while minimizing costs. How this balance is determined varies by system and can depend on many external factors such as labor cost, material cost, and site remoteness. The many factors of reliability can be simplified through prediction of degradation alongside industry collaboration.

Many techniques for identifying the optimal maintenance policy have been developed, with Condition-Based Maintenance (CBM) (Yang et al. (2022)), and Time-Based Maintenance (Liu et al. (2024)) being common examples. CBM policies make predictions based on the physical condition of a system, allowing for real-time predictions that can account for jumps in deterioration. Due to this, CBM is considered more reliable and robust, compared to alternative methods (Omshi et al. (2023)).

CBM methods involve the use of observable conditions to predict degradation pathways. This can be done either directly or indirectly. One example of indirect degradation is called *latent degradation*, this refers to any degradation which cannot be directly observed and remains dormant (or latent) within a system until a breakdown occurs. Latent degradation can come in a variety of forms, applicable to systems that are impractical to inspect, such as satellites, systems where maintenance is expensive, such as wind turbines, or complex systems with many sub-systems, such as cardiovascular health. CBM can use correlated performance characteristics, called markers, to infer information about latent degradation (Omshi et al. (2023), Shemehsavar (2012)).

To predict latent degradation through its correlated markers the dependence structure must be accounted for. One promising method to achieve this is through the copula function. A copula is a type of function used to describe the dependence structure between correlated random variables Sklar (1959), Strauch and Balaz (2023). The type of copula chosen is dependent on the type of data that the components provide. We choose a copula-based method over other methods of describing dependence (such as linear approaches) as it is comparatively more robust. Copulas are able to capture a wider range dependence structures compared to other methods, producing better estimates of the underlying failure mechanism compared to less robust alternatives. For this paper, we consider an Archimedean Lévy copula.

Our objective is to explore the viability of the Lévy copula function as a method of predicting a systems' accumulated latent degradation based on its markers, particularly within the context of maintenance optimization.

2. General Framework

The methods we use in this paper are based on the conditional copula approach outlined in Li et al. (2016). We adapt this method to a gamma-Gaussian process within the specific context of latent degradation.

Consider a system subject to latent degradation that can be described by a bivariate stochastic process $\{(X_t, Y_t), t \geq 0\}$, where X_t represents the latent degradation process, and Y_t represents the marker process, both at time t. The degradation process X_t is modeled as a gamma process with shape $\alpha > 0$ and scale $\beta > 0$ parameters. This is considered an appropriate representation of latent degradation due to its monotonicity. The lack of research on non-identical latent and marker processes led us to focus on a system with a Gaussian marker process Y_t with mean $\mu \in \mathbb{R}$ and variance $\sigma^2 > 0$. In practical contexts, the marker process might be analogous to observable vibration data that is related to the underlying degradation process. For example, lubrication drying up results in larger observed vibrations but the lubrication level itself can be difficult/expensive to measure. Hence, vibration data can be used to assess lubrication without inspecting it directly. One promising dataset could be Saxena et al.

We describe the relationship between X_t and Y_t via a Lévy copula. We use a Lévy copula to maintain the properties expected of a Lévy process, such as a gamma process (Mercier and Veridier (2022)). The general form of an Archimedean Lévy copula is:

$$\begin{aligned} U_{(X_t,Y_t)}(x,y) &= \mathbb{C}(U_{X_t}(x), U_{Y_t}(y)) \\ &= \phi^{-1}(\phi[U_{X_t}(x)] + \phi[U_{Y_t}(y)]), \end{aligned} \tag{1}$$

where $U(\cdot)$ is the integral of the marginal Lévy measure such that $U(n) = \int_n^\infty \nu(du)$, $\nu(.)$ is the corresponding Lévy measure, $\phi(\cdot)$ is the copula dependent generator function and $\phi^{-1}(\cdot)$ is its inverse. We are primarily interested in a Lévy Clayton copula that considers $\phi(n) = n^{-\theta}$ and $\phi^{-1}(n) = n^{-1/\theta}$, with θ being a correlation coefficient with $\theta \to 0$ tending towards independence and $\theta \to \infty$ tending towards complete dependence.

Define the tail integral of X_t as:

$$U_{X_t}(x) = \alpha t E_1\left(\frac{x}{\beta}\right), \tag{2}$$

where $E_1(n) = \int_n^\infty \frac{e^{-t}}{t} dt$.

There is no corresponding tail measure for Y_t as the Lévy measure of the Gaussian process is 0. Instead we use the marker process $Z_t = Y_t^2$. This marker process is a non-central Chi-squared process (central if $\mu = 0$) with one degree of freedom. The tail integral for Z_t is

$$U_{Z_t}(z) = t\left(0.5E_1\left(\frac{z}{2\sigma^2}\right) + \frac{\mu^2}{2\sigma^2} e^{z/2\sigma^2}\right). \tag{3}$$

Referring to (1), we consider the specific form of a Clayton Lévy copula defined as:

$$\mathbb{C}(u,v) = (u^{-\theta} + v^{-\theta})^{-1/\theta}, \tag{4}$$

where $u = U_{Z_t}(z)$ and $v = U_{X_t}(x)$.

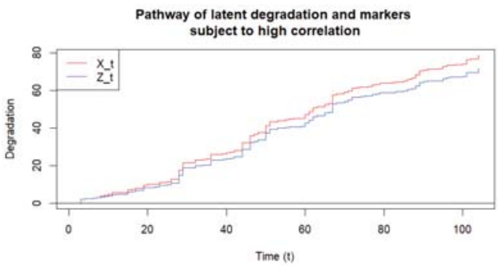

Fig. 1. Example trajectories of Z_t and X_t with high correlation over time, using $\alpha = 0.5$, $\beta = 2$, $\mu = 0$, $\sigma = 1$, and $\theta = 5$.

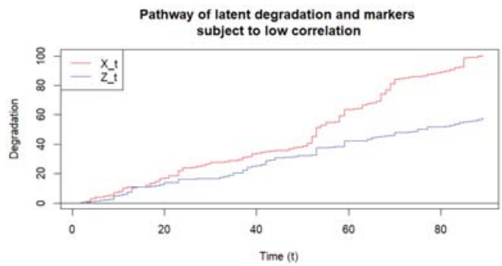

Fig. 2. Example trajectories of Z_t and X_t with high correlation over time, where $\alpha = 0.5$, $\beta = 2$, $\mu = 0$, $\sigma = 1$, and $\theta = 0.5$.

Since only the marker process is observable, we use (4) to derive the conditional copula:

$$\mathbb{C}(v|u) = \frac{d}{du}\mathbb{C}(u,v) = (1+(\frac{u}{v})^{\theta})^{-(\theta+1)/\theta} \quad (5)$$

As $\mathbb{C}(v|u)$ is a purely positive two-dimensional copula, the correlated jump v of $(X_t)_{t\geq 0}$ can be sampled by substituting a uniformly distributed random variable as:

$$v = \mathbb{C}^{-1}(m|u) = u(m^{-\theta/(1+\theta)} - 1)^{-1/\theta}, \quad (6)$$

where m is uniformly distributed on $[0, 1]$.

Since $v = U_{X_t}(x)$, $x = U_{X_t}^{-1}(v)$, the series representation for a bivariate gamma Gaussian process with Clayton Lévy copula is derived as:

$$\begin{aligned} Z_t &= \sum_{n=1}^{k} U_{Z_t}^{-1}(P_n)\mathbb{1}_{[0,t]}(a_n), \\ X_t &= \sum_{n=1}^{k} U_{X_t}^{-1}(P_n(m_n^{-\theta/(1+\theta)} - 1)^{-1/\theta})\mathbb{1}_{[0,t]}(a_n), \end{aligned} \quad (7)$$

where $t \in [0, \mathbb{T}]$, $\mathbb{T}$ is the system's lifespan, $\{P_n\}_{n\in N}$ is a sequence of standard Poisson arrival times, $\{m_n\}_{n\in N} \sim N(0,1)$, and k is a truncation index such that $\sum_{n=1}^{k} P_n < \tau$, where $N(\tau) = \inf\{n, P_n > \tau\}$.

Since both formulae in (7) involve an inverse exponential integral, which has no closed form, we approximate our outputs numerically.

Fig. 1 and Fig. 2 present examples of the evolution of the latent and marker processes using a fixed parameter set $\alpha = 0.5$, $\beta = 2$, $\mu = 0$, and $\sigma = 1$, while varying the level of correlation between them.

3. Maintenance Policy

For systems subject to latent degradation, inspections reveal information only about the marker process on which maintenance decisions are made. Inspections occur at time t with a fixed inspection interval of T. Smaller values of T correspond to more frequent inspections, with less frequent inspections having a larger T; as $T \to 0$, inspections approach continuous monitoring.

Based on the information obtained from an inspection, decisions are made on what action to perform on the system. To differentiate between actions, a preventive threshold M is assigned, representing the level of degradation that can be tolerated before a preventive maintenance action is triggered. Additionally, a corrective threshold L is assigned, corresponding to the maximum degradation level the system can handle before a breakdown occurs. Since L is fixed and $M < L$, the maintenance policy can be defined as follows:

- If $Z_t < M$ at time of inspection, no action is performed on the system.
- If $M \leq Z_t$ and $X_t < L$ at time of inspection, a preventive replacement is performed on the system.
- If $X_t \geq L$ at any point, a breakdown occurs and a corrective replacement is performed.

Since the breakdown point cannot be altered, our decision variables are the inspection interval T and the preventive threshold M.

To optimize the policy with respect to these decision variables, we use long-run expected cost:

$$EC = \lim_{t \to \infty} \frac{C(t)}{t} = \frac{E(C(\mathbb{T}))}{E(\mathbb{T})} \tag{8}$$

where $\mathbb{T}$ is the cycle length of the system and $C(t)$ its cost at time t, defined as:

$$C(t) = c_i N_i(t) + c_p N_p(t) + c_c N_c(t) \tag{9}$$

where $N_x(t)$ denotes the number of times a particular action occurs and c_x is the corresponding cost. Likewise, the subscripts i, p, and c refer to inspection, preventive replacement, and corrective replacement respectively.

To determine the optimal values of the decision variables that minimize the expected cost, we employ a genetic algorithm.

4. Maintenance Optimization

In this section, we investigate the effect of each model parameter on the system behavior. As a reference point, we first define the following set of control parameters: $\alpha = 0.5$, $\beta = 2$, $\mu = 0$, and $\sigma = 1$. The control costs for each action are taken as: $c_i = 5$, $c_p = 140$, and $c_c = 240$. For optimization ranges, the preventive threshold is investigated for values $20 \leq M \leq 90$ and the inspection interval for $1 \leq T \leq 100$. The corrective threshold is fixed at $L = 100$, and the copula dependence parameter is set to $\theta = 5$. We optimize through a genetic algorithm, treating M and T as decision variables, with all other parameters being fixed. Table 1 shows how the optimized expected cost behaves as θ changes. As higher values of θ correspond to stronger correlation, it can be understood that information provided by the marker process Z_t more accurately reflects latent degradation X_t. The results in Table 1 support this interpretation: the expected cost decreases as the marker becomes more informative. Similarly, higher preventive thresholds and longer inspection intervals are preferred when the marker is more reliable. This indicates that, as the marker process becomes more accurate, the ideal policy is more lenient. Table 2 depicts how changes in the parameters of X_t affect the expected cost. Note that for a gamma distribution, the mean is given by $\alpha\beta$; therefore, increases in either α or β result in quicker degradation. Faster degradation of X_t increases the likelihood of failure, which in turn raises the expected cost. The results in Table 2 confirm this trend: higher values of α or β lead to higher expected costs. Interestingly, increases in α demonstrate a sharper increase in expected cost relative to β. This may be attributed to the larger contribution of β to the variance of the gamma distribution $\alpha\beta^2$.

Table 1. Effect of correlation parameter.

θ	EC	M	T
0.1	0.4383	55.07	64.64
1	0.3460	57.51	72.26
10	0.3064	65.88	81.58

Table 2. Effect of latent degradation parameters.

α	EC	M	T
0.5	0.3014	55.41	74.66
1	0.8127	63.05	66.82
5	2.138	60.98	55.64
β	EC	M	T
0.5	0.2878	60.00	66.03
1	0.3001	44.79	74.97
5	0.4687	21.92	85.33

Regarding the optimal decision variables M and T, some notable patterns emerge. For M, increases in α do not exhibit a clear trend, whereas β displays a clear inverse relationship with M. This behavior is intuitive: faster degradation favors a wider warning range, leading to a lower preventive threshold. The absence of a clear trend between α and M may suggest that, for $\alpha \geq 1$, degradation becomes sufficiently rapid that failures can occur almost immediately, making the optimal threshold less sensitive to changes in α. For the inspection interval T, the relationship is reversed: α exhibits an inverse relationship with T, while β shows a direct relationship. Table 3 summarizes the in-

Table 3. Effect of Marker process parameters.

μ and $\sigma = 1$	EC	M	T
0.5	0.2871	71.13	66.18
1	0.3691	60.59	67.74
5	1.0600	46.47	70.77
σ and $\mu = 0$	EC	M	T
0.5	0.3610	67.94	61.10
1	0.4073	60.84	61.39
5	0.5070	60.96	65.09
σ and $\mu = 1$	EC	M	T
0.5	0.4934	56.46	66.89
1	0.3271	70.59	76.42
5	0.6976	79.06	74.06

fluence of the marker process parameters on the optimal decision variables. In particular, it reports the effects of μ and of σ under both $\mu = 0$ and $\mu = 1$. The reason for splitting the σ examination into two cases is the additive portion of Eq. (3) that is only present when $\mu \neq 0$. Within the Gaussian framework, μ represents the process mean and σ denotes the standard deviation. Consequently, an increase in μ is expected to accelerate the onset of warning signals, whereas an increase in σ is anticipated to introduce greater variability in the marker observations, potentially yielding both earlier and later warnings.

Table 3 demonstrates the anticipated trends; however, a number of notable observations can be made. As μ increases, both the expected cost and the optimal inspection interval rise, whereas the optimal preventive threshold decreases. This suggests that higher values of μ increase the probability of early warning signals, which in turn raises the likelihood of false positives and premature preventive actions. Such behavior is consistent with the observed increase in expected cost and the corresponding reduction in T, as implied by the long-run cost formulation in (8). However, the preference for smaller M suggests that there is still a risk of breakdowns occurring in between inspections.

In the case $\mu = 0$, σ is expected to exhibit trends analogous to those observed for β in Table 2, since when $\mu = 0$, Z_t is functionally a specific gamma process (see Eqs. (2) and (3)). A comparison of Table 3 with Table 2 confirms that, when $\mu = 0$, σ broadly follows the same qualitative patterns as β. Notably, σ exhibits a substantially smaller impact on the expected cost, M and T when compared to β. While differences in computational effort may contribute, a more plausible explanation is that modifications to the warning system, particularly when preventive actions are relatively inexpensive, have a smaller impact on the optimization outcome than comparable changes applied directly to the degradation rate.

The $\mu \neq 0$ case exhibits markedly different behavior in σ compared with $\mu = 0$. This difference is expected, as when $\mu \neq 0$ there is a significant additive component present in the tail integral in Eq. (3). In particular, from Eq. (3), as $\sigma \to 0$ the additive component of the tail integral $U_{z_t}(z)$ diverges, whereas as $\sigma \to \infty$ the tail integral $U_{z_t}(z)$ approaches that of a gamma process. This behavior leads to a noticeable increase in the mean observed marker as $\sigma \to 0$, resulting in a higher likelihood of false positives occurring. Table 3 supports this interpretation, showing a sharp increase in the expected cost for $\sigma = 0.5$, despite the trend observed in the $\mu = 0$ case. With respect to M and T, the preventive threshold decreases as σ increases while T exhibits no clear trend.

Finally, we examine the effect of varying individual cost parameters while keeping the other parameters fixed at their control values. Using the prior definitions of the control values:

$$\alpha = 0.5,\ \beta = 2,\ \mu = 0,\ \sigma = 1,\ \theta = 5,$$
$$c_i = 5,\ c_p = 140,\ c_c = 240,$$

This allows us to evaluate how each cost component influences the expected overall cost and the optimal decision variables. In particular, we consider the following scenarios:

(1) The inspection cost is high, $c_i = 20$.
(2) Preventive actions are inexpensive, $c_p = 14$.
(3) Preventive actions are expensive, $c_p = 230$.
(4) Breakdown events are costly, $c_c = 2400$.

For a system subject to high inspection costs, we

Table 4. Impact of cost on EC.

Cost environment	EC	M	T
Control	0.2671	61.33	66.36
$c_i = 20$	0.5466	66.16	87.81
$c_p = 14$	0.1053	37.04	74.04
$c_p = 230$	0.4393	70.10	84.52
$c_c = 2400$	0.4099	44.43	50.15

would expect the optimal policy to favor fewer but more effective inspections, implying a relatively low M and a large T. In our model, when inspections are expensive, larger values of T are indeed preferred; however, the optimal M remains similar to the control case. When preventive actions are inexpensive, the optimal policy would be expected to maximize the likelihood of preventive replacement, which corresponds to a lower preventive threshold. Table 4 confirms this behavior in our model. If the cost of preventive action is comparable to that of a breakdown, a run-to-failure policy may be optimal. In our model, the increase in both M and T suggests that a run-to-failure strategy is indeed preferred under these conditions. Finally, when breakdown costs dominate all other sources of cost, the optimal policy should prioritize preventive actions. In this case, both M and T are expected to be small in order to reduce the probability of failure, which is also observed in our results.

5. Concluding Remarks

Based on the results presented in this paper, we conclude that the conditional Lévy copula-based model for predicting latent degradation using a Gaussian marker process is viable. Each parameter behaves within expected ranges and yields optimized maintenance policies that are consistent with system logic. In future work, we aim to develop a new policy that incorporates direct prediction of latent degradation, rather than relying solely on the marker process, which may represent a conservative and costly approach. The results of the new policy will be compared with those of alternative policies, and their relative performance will be evaluated. Another direction for future work would involve the development of a benchmark model against which the copula-based model can be compared. This benchmark model would differ only in its dependence structure, using a linear approach instead of copula-based one. This would allow us to better understand how changes in dependence structure affect results. It would also be important to consider how the model outlined in this paper sits within the context of real data. From the NASA and PHM society data repositories, the Carbon Fiber-Reinforced Polymer (CMFP) Composites dataset Saxena et al. appears to be a promising dataset to work with.

References

Li, H., E. Deloux, and L. Dieulle (2016). A condition-based maintenance policy for multi-component systems with Lévy copulas dependence. *Reliability Engineering & System Safety 149*, 44-55.

Liu, B., J. Pang, H. Yang, and Y. Zhao (2024). Optimal condition-based maintenance policy for leased equipment considering hybrid preventive maintenance and periodic inspection. *Reliability Engineering & System Safety 242*, 109724.

Mercier, S. and G. Veridier (2022). On the modeling of dependence between univariate L´evy wear processes and impact on the reliability function. *Applied Stochastic Models in Business and Industry 40*, 549-573.

Omshi E. M., and S. Shemehsavar, and A. Grall (2023). An intelligent maintenance policy for a latent degradation system. *Reliability Engineering System Safety 242*, 109739.

Saxena, A., K. Goebel, C.C. Larrosa, and F.K. Chang (). NASA Prognostics Data Repository, NASA Ames Research Center, Moffett Field, CA. CFRP Composites Data Set . Link: https://data.phmsociety.org/nasa/

Shemehsavar, S. (2012). A Bivariate Gamma Model for a Latent Degradation Process. *Communications in Statistics - Theory and Methods 42*, 1924-1938.

Sklar, A. (1959). Fonctions de Repartition à n Dimensions et Leurs Marges. *Publications de l'Institut de Statistique de l'Universite de Paris*

8, 229–231.

Stewart-Green, C., S. Shemehsavar, and G. Hocking (2025). *Maintenance Optimization for Latent Degradation Systems.* Murdoch University.

Strauch, O. and V. Balaz (2023). Copulas. *Uniform Distribution Theory 18*, 147-200.

Yang, J., X. Zhao, and M. Han (2022). Joint optimization of imperfect condition-based maintenance and lot sizing via an availability-cost hybrid factor. *Engineering Reports 4*.