Bayesian Online Changepoint Detection Algorithm for Degradation Monitoring and Prognostics of Electrical Devices

Roman Mukin, Kai Hencken, Marco Cilibrasi

Corporate Research Center, ABB Switzerland Ltd, Segelhofstrasse 1K, 5405 Baden
roman.mukin@ch.abb.com
kai.hencken@ch.abb.com
cilibrasimarco@gmail.com

ABSTRACT

Electrical devices that undergo frequent switching operations can suffer mechanical wear over numerous cycles. Identifying early signs of degradation through routine measurements is important for condition monitoring and predictive maintenance. This work applies Bayesian Online Changepoint Detection (BOCD) to multivariate signal data from an electrical device to detect the onset of degradation and estimate remaining useful life (RUL) under uncertainty.

Classical approaches such as rolling-window trend fitting react slowly to regime changes and require careful tuning. BOCD instead maintains an online posterior over the run length and recursively updates a predictive model as new observations arrive. We extend the classical formulation to a multivariate regression setting with operation-dependent trends, conjugate matrix-normal inverse-Wishart priors with the regression-coefficient precision calibrated from prescribed alarm bounds to confine initial predictions to the operational envelope, and a Weibull hazard function that encodes wear-out behaviour. The resulting algorithm performs online inference without storing the full measurement history while providing calibrated predictive uncertainty.

The framework is first evaluated on synthetic degradation scenarios covering abrupt, gradual, stepwise, and accelerating fault modes, and then applied to real mechanical endurance-test data from a run-to-failure experiment. On the synthetic scenarios, abrupt changepoints are detected without delay and gradual degradation is flagged once the cumulative drift exceeds the noise level, with the lower, median, and upper RUL bounds converging consistently at end of life. On real endurance-test data, the algorithm maintains a stable RUL estimate throughout the healthy device life despite noise-induced changepoints, and provides a timely end-of-life warning once sufficient post-changepoint evidence confirms the degradation trend.

1. INTRODUCTION

Electrical switching devices are subjected to repeated mechanical and electrodynamic stresses during operation. Over many operating cycles, these stresses can progressively alter the motion of internal components and eventually lead to degraded switching performance or mechanical failure. Detecting such degradation early is therefore important for condition monitoring and predictive maintenance (Vachtsevanos, Lewis, Roemer, Hess, & Wu, 2006; Hoffmann et al., 2020), especially when measurements are noisy and the onset of damage is not directly observable.

In this work, we consider degradation monitoring of an electrical device through multivariate signal measurements extracted from routine switching recordings. These signals provide indirect signatures of the device's mechanical state (Runde, Ottesen, Skyberg, & Ohlen, 1992; Landry, Léonard, Minsou, & Ouellette, 2008). The central challenge is to identify the transition from healthy behavior to an evolving faulty regime as early and as reliably as possible, while accounting for measurement variability and uncertainty.

Earlier analyses of noisy endurance-test data often relied on local polynomial fitting and rolling-window trend estimation as straightforward baselines for extracting degradation trends from measured signals (Mosallam, Medjaher, & Zerhouni, 2014; Hu, Baraldi, Di Maio, & Zio, 2015; Niknam, Kobza, & Hines, 2017). These approaches are attractive because they are simple to implement and can suppress part of the measurement noise. Their performance, however, depends strongly on the chosen window size and smoothing settings. Small windows tend to react to noise and produce unstable trends, whereas large windows smooth the data excessively and can delay the detection of meaningful changes in device behavior. For degradation monitoring, where changepoints may correspond to the onset of a fault mechanism, this sensi-

tivity is a significant limitation.

Bayesian Online Changepoint Detection (BOCD) (Adams & MacKay, 2007) avoids this limitation by maintaining a posterior distribution over the run length (the number of observations since the most recent changepoint) and updating it recursively as new data arrive. This formulation lends itself to prognostics (Si, Wang, Hu, & Zhou, 2011): it supports online inference, incorporates uncertainty naturally, and does not require storing the full measurement history.

Recent work has demonstrated that integrating changepoint detection into RUL estimation pipelines can improve prognostic accuracy (Arunan, Sarda, Telesca, & Khalid, 2024). Combined with a regression model for the post-changepoint dynamics, BOCD provides a probabilistic route from online change detection to online prognostics. Figure 1 illustrates this conceptually: the upper panel shows a signal that transitions from a healthy to a degraded regime at a changepoint, with the regression model extrapolating the degraded trajectory toward a failure threshold; the lower panel shows the corresponding run-length posterior, which resets at the changepoint and thereby signals the regime change. Once the posterior favors a degraded regime, the regression component estimates the probability that prescribed operational bounds will be violated at future times. This forecast-based view yields both failure-probability assessment and Remaining Useful Life (RUL) estimation under uncertainty.

The contribution of this paper is the application of BOCD to multivariate signal monitoring for electrical-device degradation. We extend the BOCD framework with a multivariate Bayesian regression model using conjugate matrix-normal inverse-Wishart priors and a Weibull hazard function that encodes wear-out behavior. The resulting algorithm estimates changepoints online and provides probabilistic RUL bounds by extrapolating the degraded trend toward threshold crossing with calibrated uncertainty.

2. METHOD DESCRIPTION

2.1. Problem definition for remaining useful life monitoring

We consider a real-time monitoring setting in which a sequence of multivariate observations is acquired through repeated system operations. At each new operation, the algorithm receives the current observation vector along with its operation index and updates its belief about whether the system remains in the same operating regime or has entered a degraded one. The objective is not only to detect the onset of degradation as early as possible, but also to translate the evolving signal trend into prognostic information.

The framework proceeds in two coupled stages. In the first stage (bottom panel of Figure 1), BOCD maintains a posterior distribution over the run length r_k, which encodes how many operations have elapsed since the most recent changepoint. This posterior is updated recursively at each new operation using the predictive probability of the current observation for each candidate segment length (Section 2.2). A sharp concentration of posterior mass near $r_k = 0$ signals a regime change; during stable operation, the mass migrates toward larger run lengths.

In the second stage (top panel of Figure 1), the same per-segment regression posteriors that drive changepoint inference are reused for forecasting. Each run-length hypothesis carries its own estimate of the signal trend; these are combined into a mixture predictive distribution weighted by the run-length posterior (Eq. (22)). Monte Carlo samples drawn from this mixture at future horizons are compared against predefined operational bounds to obtain a failure-probability curve, from which RUL quantiles are extracted (Eq. (24)). Because both stages share the same probabilistic model, changepoint detection and prognostics are updated jointly rather than in a separate pipeline.

2.2. Bayesian Online Changepoint Detection (BOCD)

2.2.1. Problem setup

From the outset, we formulate BOCD for multivariate observations rather than for a single scalar signal. Let $\mathbf{z}_k \in \mathbb{R}^d$ denote the observation vector measured at operation k, where each of its d components represents a distinct signal channel. Let

$$\boldsymbol{\phi}_k = [1, k, k^2, \ldots, k^{p-1}]^\top \tag{1}$$

be a polynomial basis defined on the operation index. We define the *run length* variable r_k as the number of operations since the most recent changepoint before operation k. The key inference task is to update the posterior distribution $P(r_k \mid \mathbf{z}_{1:k})$ online as each new observation is recorded.

Within each regime, the observation vector is modeled by a multivariate regression law with polynomial mean,

$$\mathbf{z}_k \mid M, \Sigma \sim \mathcal{N}(M^\top \boldsymbol{\phi}_k, \Sigma), \tag{2}$$

where $M \in \mathbb{R}^{p\times d}$ contains the regression coefficients and $\Sigma \in \mathbb{R}^{d\times d}$ is the covariance matrix. A changepoint corresponds to a reset of the segment parameters, so BOCD directly tracks changes in the multivariate signal's evolution. Changepoint detection and multivariate regression are therefore part of the same probabilistic model rather than separate stages.

2.2.2. Recursive update of the run length distribution

Following (Adams & MacKay, 2007), the joint distribution of run length and observations at operation k is derived by first marginalizing over the previous run length and then applying the chain rule twice:

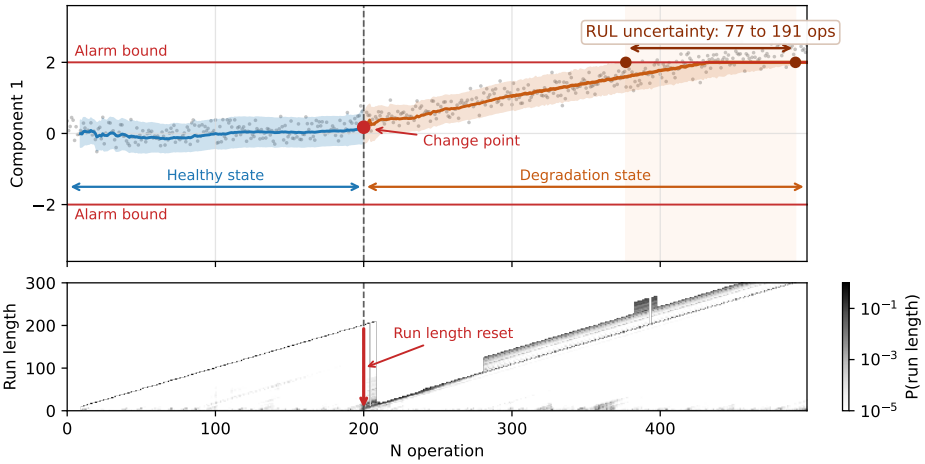

Figure 1. Conceptual illustration of online RUL prediction with a changepoint model and regression. Top: the signal follows the piecewise trend used in the synthetic BOCD example, with a changepoint at $N = 200$ operations separating a healthy and a degradation regime. From a future point, the regression model extrapolates the degraded trajectory toward a failure threshold and yields an uncertainty interval for the predicted failure operation. Bottom: the run-length posterior distribution, shown as a heatmap, resets at the changepoint, marking the transition to a new regime.

$$
\begin{aligned}
P(r_k, \mathbf{z}_{1:k}) &\overset{\text{marginalization}}{=} \sum_{r_{k-1}} P(r_k, r_{k-1}, \mathbf{z}_{1:k}) \\
&\overset{\text{chain rule}}{=} \sum_{r_{k-1}} P(r_k \mid r_{k-1}, \mathbf{z}_{1:k}) \\
&\quad \cdot P(r_{k-1}, \mathbf{z}_{1:k}) \\
&\overset{\text{chain rule}}{=} \sum_{r_{k-1}} P(r_k \mid r_{k-1}, \mathbf{z}_{1:k}) \\
&\quad \cdot P(\mathbf{z}_k \mid r_{k-1}, \mathbf{z}_{1:k-1}) \\
&\quad \cdot P(r_{k-1}, \mathbf{z}_{1:k-1}).
\end{aligned} \tag{3}
$$

Two conditional independence assumptions of the BOCD model simplify this expression. First, the changepoint process is independent of the observations given the run length, so that $P(r_k \mid r_{k-1}, \mathbf{z}_{1:k}) = P(r_k \mid r_{k-1})$. Second, given the run length r_{k-1}, only the observations within the current segment are relevant, hence $P(\mathbf{z}_k \mid r_{k-1}, \mathbf{z}_{1:k-1}) = P(\mathbf{z}_k \mid \mathbf{z}_{k-r_{k-1}:k-1})$. Substituting these into Eq. (3) yields the BOCD recursion:

$$
\begin{aligned}
P(r_k, \mathbf{z}_{1:k}) = \sum_{r_{k-1}} & P(r_k \mid r_{k-1}) \cdot \\
& \cdot P(\mathbf{z}_k \mid \mathbf{z}_{k-r_{k-1}:k-1}) \cdot \\
& \cdot P(r_{k-1}, \mathbf{z}_{1:k-1}).
\end{aligned} \tag{4}
$$

The three factors in Eq. (4) have a direct interpretation:

- $P(r_k \mid r_{k-1})$ is the hazard model, which defines the probability that a changepoint occurs at operation k given the previous run length;
- $P(\mathbf{z}_k \mid \mathbf{z}_{k-r_{k-1}:k-1})$ is the predictive probability of the newly observed signal vector given the data in the current segment;
- $P(r_{k-1}, \mathbf{z}_{1:k-1})$ is the joint distribution from the previous operation, which carries forward the posterior information.

2.2.3. Predictive probability and conjugate priors

The predictive probability appearing in Eq. (4) is computed by marginalizing over the segment parameters:

$$
\begin{aligned}
P(\mathbf{z}_k \mid \mathbf{z}_{k-r_{k-1}:k-1}) = \int & p(\mathbf{z}_k \mid M, \Sigma) \cdot \\
& \cdot p(M, \Sigma \mid \mathbf{z}_{k-r_{k-1}:k-1})\, dM\, d\Sigma.
\end{aligned} \tag{5}
$$

This integral can be evaluated in closed form when a conjugate prior is placed on the segment parameters. We briefly summarize the progression from the scalar to the multivariate regression case, since the latter is used throughout this work. The derivations follow standard Bayesian linear modeling as presented in (Murphy, 2012).

Consider scalar observations $x_i \sim \mathcal{N}(\mu, \sigma^2)$ with both μ and σ^2 unknown. The conjugate prior is the Normal-Inverse-Gamma:

$$
\sigma^2 \sim \text{Inv-Gamma}(\alpha_0, \beta_0), \tag{6}
$$

$$
\mu \mid \sigma^2 \sim \mathcal{N}(\mu_0, \sigma^2/\kappa_0). \tag{7}
$$

After n observations with sample mean $\bar{x}$, the posterior hy-

perparameters are

$$\kappa_n = \kappa_0 + n, \qquad \mu_n = \frac{\kappa_0 \mu_0 + n\bar{x}}{\kappa_n}, \tag{8}$$

$$\alpha_n = \alpha_0 + \tfrac{n}{2}, \tag{9}$$

$$\beta_n = \beta_0 + \tfrac{1}{2}\textstyle\sum_{i=1}^{n}(x_i - \bar{x})^2 + \frac{\kappa_0 n}{2\kappa_n}(\bar{x} - \mu_0)^2, \tag{10}$$

and the predictive distribution is a univariate Student-t:

$$x_{n+1} \sim t_{2\alpha_n}\left(\mu_n,\ \frac{\beta_n(\kappa_n + 1)}{\alpha_n\, \kappa_n}\right). \tag{11}$$

This is the model originally used in (Adams & MacKay, 2007).

The scalar model generalizes to d-dimensional observations with a polynomial mean. Given the regression model $\mathbf{z}_k \mid M, \Sigma \sim \mathcal{N}(M^\top \boldsymbol{\phi}_k, \Sigma)$ introduced in the problem setup, the conjugate prior on (M, Σ) is the matrix-normal inverse-Wishart (MNIW):

$$\Sigma \sim \mathcal{W}^{-1}(\Psi_0, \nu_0), \tag{12}$$

$$\mathrm{vec}(M) \mid \Sigma \sim \mathcal{N}(\mathrm{vec}(M_0),\ \Sigma \otimes V_0^{-1}). \tag{13}$$

Let $\Phi \in \mathbb{R}^{n \times p}$ denote the design matrix whose rows are the basis vectors $\boldsymbol{\phi}_i^\top$, and let $Z \in \mathbb{R}^{n \times d}$ collect the corresponding observations. After n data points in the current segment, the posterior hyperparameters become

$$V_n = V_0 + \Phi^\top \Phi, \tag{14}$$

$$M_n = V_n^{-1}(V_0 M_0 + \Phi^\top Z), \tag{15}$$

$$\nu_n = \nu_0 + n, \tag{16}$$

$$\Psi_n = \Psi_0 + Z^\top Z + M_0^\top V_0 M_0 - M_n^\top V_n M_n. \tag{17}$$

The predictive distribution for the next operation is then multivariate Student-t:

$$\mathbf{z}_{k+1} \sim t_{\nu_n}\left(M_n^\top \boldsymbol{\phi}_{k+1},\ \frac{1 + \boldsymbol{\phi}_{k+1}^\top V_n^{-1} \boldsymbol{\phi}_{k+1}}{\nu_n}\, \Psi_n\right). \tag{18}$$

This conjugate structure keeps the predictive step in Eq. (5) analytically tractable while allowing each run-length hypothesis to carry a multivariate regression model with uncertain covariance (see (Murphy, 2012) or (Murphy, 2007) for a full derivation of the MNIW conjugacy). The scalar Student-t result above is recovered as a special case when $d = 1$ and $p = 1$.

In both cases the prior hyperparameters M_0, ν_0, and Ψ_0 are set to diffuse values, while V_0 is calibrated from the prescribed alarm bounds (see Sections 3.1 and 4), so that the posterior is dominated by the observed data after only a few operations.

2.2.4. Hazard function

The hazard function governs the prior probability that a changepoint occurs at operation k, given the current run length:

$$\begin{aligned} P(r_k = 0 \mid r_{k-1}) &= H(r_{k-1}), \\ P(r_k = r_{k-1} + 1 \mid r_{k-1}) &= 1 - H(r_{k-1}). \end{aligned} \tag{19}$$

Because H enters the BOCD recursion in a multiplicative manner, the choice of hazard model encodes a prior belief about how segment durations are distributed. In practice, a system that has operated for many cycles may become progressively more likely to enter a degraded regime, reflecting wear-out behavior. To capture this, a Weibull hazard function is used, providing a run-length-dependent changepoint probability. Using the incremental cumulative hazard of a Weibull distribution with scale $\alpha > 0$ and shape $\beta > 0$, the hazard at run length r is defined as

$$H(r) = 1 - \frac{S(r+1)}{S(r)}, \quad S(r) = \exp\left[-(r/\alpha)^\beta\right], \tag{20}$$

where $S(r)$ is the Weibull survival function, the probability that no changepoint has occurred by operation r.

For $\beta > 1$ the hazard increases with run length, biasing BOCD toward shorter segments as the elapsed time grows. Conversely, $\beta < 1$ produces a decreasing hazard that favors longer segments. The special case $\beta = 1$ recovers an exponential inter-changepoint distribution. By adjusting α and β, the practitioner can encode domain-specific knowledge about the expected onset of degradation without altering the rest of the inference machinery.

2.2.5. Changepoint detection and online inference

At each operation, the joint distribution is normalized to obtain the posterior over run lengths:

$$P(r_k \mid \mathbf{z}_{1:k}) = \frac{P(r_k, \mathbf{z}_{1:k})}{\sum_{r_k} P(r_k, \mathbf{z}_{1:k})}. \tag{21}$$

The quantity $P(r_k = 0 \mid \mathbf{z}_{1:k})$ can be interpreted as the probability that the current operation marks a changepoint. More generally, the full posterior indicates how strongly the model supports different segment lengths and therefore different explanations of the current operating regime. Because each run-length hypothesis is associated with its own multivariate regression posterior, changepoint inference and signal-trend estimation are updated simultaneously.

2.2.6. Future predictions

Once BOCD has been updated up to operation k, the future evolution of the signal is represented as a mixture over the

current run-length posterior,

$$P(\mathbf{z}_{k+h} \mid \mathbf{z}_{1:k}) = \sum_{r_k} P(\mathbf{z}_{k+h} \mid r_k,\, \mathbf{z}_{k-r_k:k})\, P(r_k \mid \mathbf{z}_{1:k}), \tag{22}$$

where h denotes the number of future operations ahead. The conditioning on $\mathbf{z}_{k-r_k:k}$ makes explicit that each predictive component uses only the observations within the segment defined by r_k; by the conditional independence assumption of BOCD, data before the most recent changepoint do not influence the prediction. Each component of this mixture corresponds to a different run-length hypothesis and carries its own regression posterior, so the forecast reflects uncertainty about both the current regime and its trend.

2.2.7. Failure probability and remaining useful life estimation

At a future operation $k+h$, the probability that one or more signal components fall outside their prescribed bounds is estimated via Monte Carlo sampling. Drawing N_{MC} samples from $\mathcal{N}(\hat{\boldsymbol{\mu}}_{k+h},\, \hat{\Sigma}_{k+h})$, where $\hat{\boldsymbol{\mu}}$ denotes the saturated mixture mean and $\hat{\Sigma}$ the mixture covariance, the failure probability is

$$P_{\mathrm{fail}}(k+h) = 1 - \frac{1}{N_{\mathrm{MC}}} \sum_{i=1}^{N_{\mathrm{MC}}} \mathbf{1}\big[\mathbf{z}^{(i)} \in [\mathbf{l},\, \mathbf{u}]\big]. \tag{23}$$

This quantity summarizes, at each future horizon, the fraction of predicted trajectories that violate at least one operational bound. RUL is then estimated by scanning a grid of future operations from the current time k up to a configurable horizon. RUL estimates are defined as the first future horizons at which the failure probability exceeds prescribed quantile thresholds:

$$\mathrm{RUL}_q = \min\big\{h > 0 : P_{\mathrm{fail}}(k+h) > q\big\}, \tag{24}$$

with $q \in \{0.05, 0.50, 0.95\}$ yielding upper, median, and lower RUL bounds, respectively. Because each future time step is evaluated independently from the full mixture distribution, the resulting failure probability curves are smooth and do not require temporal coherence between successive forecast steps. This forward-prediction step extends BOCD from change detection to prognostics: the same online posterior that identifies degradation is used to quantify failure risk and estimate remaining useful life.

2.3. Interpretation of run-length uncertainty in probabilistic forecasting

The run-length posterior does not identify a single changepoint location; instead, it encodes a distribution over plausible segment lengths. A large posterior mass near small values of r_k indicates that a recent changepoint is likely, whereas mass concentrated at larger values of r_k suggests that the current regime has persisted for many operations. The posterior vector $P(r_k \mid \mathbf{z}_{1:k})$ should therefore be interpreted as a probabilistic summary of changepoint uncertainty rather than as a hard reset index.

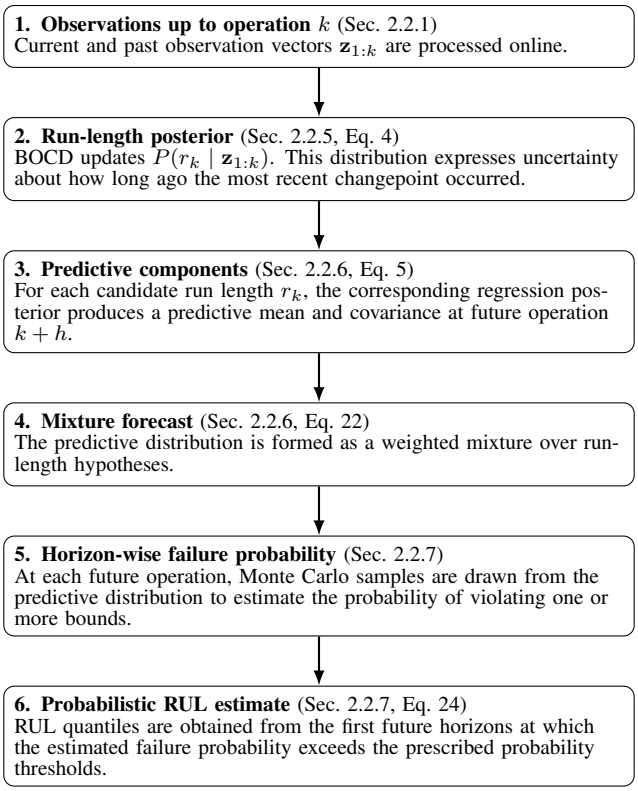

Figure 2. Step-by-step interpretation of how the run-length posterior encodes changepoint uncertainty and how it is transformed into horizon-wise failure probabilities and probabilistic RUL estimates.

This uncertainty directly shapes the forecast. Because each run-length hypothesis maintains a separate regression posterior, different hypotheses yield different predictive means and covariances. At a future operation $k+h$, the forecast is represented as the mixture in Eq. (22), now with each component explicitly conditioned on the segment data. Consequently, the slope of the predicted trend is governed by the data within the segment lengths that receive significant posterior support. This influence is soft and probabilistic: the algorithm does not select a single reset location, but combines multiple candidate post-changepoint segments according to their posterior probabilities. After saturation is applied to the predictive mean, Monte Carlo samples are drawn independently at each forecast horizon to estimate the marginal failure probability. Run-length uncertainty therefore affects RUL through the mixture weights that shape these horizon-wise failure probabilities.

Figure 2 summarizes how this posterior is transformed into Monte Carlo samples in the current implementation.

The present implementation uses the posterior inferred from observations up to operation k to build a predictive mixture at each future horizon. For each value of h, Monte Carlo samples are drawn from the corresponding predictive distribution and used only to estimate the marginal failure probability $P_{\text{fail}}(k+h)$. The sampled points are not fed back into BOCD, so the run-length posterior remains the one inferred at the forecast origin throughout the RUL calculation.

This is consistent with the probabilistic RUL definition in Eq. (24). The method does not attempt to generate recursively updated future histories; instead, it evaluates how the horizon-wise failure probability evolves under the current posterior uncertainty. Run-length uncertainty enters directly through the weights $P(r_k \mid \mathbf{z}_{1:k})$, which determine how strongly different segment hypotheses contribute to the predictive distribution and therefore to the resulting RUL bounds.

2.4. Computational considerations.

The BOCD update at each operation maintains a posterior vector of run-length hypotheses, bounded by a user-specified window W. When the number of active hypotheses would exceed W, the oldest two entries are merged by summing their probability mass before the new hypothesis is appended, keeping memory and per-step compute at $\mathcal{O}(Wp^3)$, where p is the polynomial degree. For the linear model ($p = 2$) and $W = 500$ used in the synthetic experiments, each BOCD update requires a small number of matrix inversions and is suitable for step-by-step online processing: no future observations are needed, and the posterior is updated immediately upon arrival of each new observation. The dominant runtime cost in the present implementation is the Monte Carlo failure-probability scan across future horizons (Section 2.2.7), which is separable from the detection stage and can be reduced independently by adjusting the number of samples or the forecast horizon. The synthetic experiments use 500 operations and a bivariate signal as a controlled benchmark, enabling ground-truth evaluation across all four degradation scenarios. The window mechanism ensures that the per-step cost does not grow with dataset length, so the same algorithm applies without modification to longer operation sequences.

3. EVALUATION ON SYNTHETIC DATA

The BOCD framework described in Section 2 is evaluated on a suite of synthetic degradation scenarios. Using synthetic signals allows both the changepoint detection and the RUL estimation to be assessed against known ground truth. The scenarios are designed to represent qualitatively different degradation mechanisms that can arise in practice.

3.1. Synthetic scenarios setup

Each scenario generates a bivariate signal $\mathbf{z}_k = [z_{k,1},\, z_{k,2}]^\top$ over $N = 500$ operations. The clean signal is constructed from two component functions $c_1(t)$ and $c_2(t)$ that encode the degradation pattern, and additive Gaussian noise with standard deviation $\sigma = 0.2$ is applied to both components. Prescribed bounds are $[-2, 2]$ for both components. All scenarios share the same algorithmic configuration: a Weibull hazard with scale $\alpha = 500$ and shape $\beta = 2.0$, polynomial degree 1, Monte Carlo sample size $N_{\text{MC}} = 1000$, and a prediction horizon of 500 operations.

The prior mean $M_0 = 0$, $\nu_0 = 0.1$, and $\Psi_0 = 0.01\, I$ are set to diffuse values. The regression coefficient precision matrix V_0 is calibrated automatically from the prescribed alarm bounds rather than fixed to a generic diffuse value: each diagonal entry is set so that the 2σ-interval of the prior predicted regression mean at the maximum horizon $N = 500$ operations stays within the alarm bounds $[-2,\, 2]$, after whitening the observations by the noise standard deviation $\sigma = 0.2$. For a linear model, this yields $V_0 = \text{diag}(25,\ 10^{-4})$, giving a weakly informative prior on the intercept and a tighter prior on the slope to prevent diverging extrapolations.

The following degradation scenarios are considered, each targeting a different degradation mechanism:

1. **Late sudden failure.** Component 1 remains healthy; Component 2 is quiescent until operation 400, then ramps sharply to the lower bound by operation 500.
2. **Slow creep.** Component 1 drifts linearly from zero to 2.5 starting at operation 200; Component 2 stays constant. The drift rate is approximately 0.008 per operation.
3. **Step drift.** Component 1 undergoes step jumps at operations 150 and 300; Component 2 steps down at operation 200 and then drifts further.
4. **Exponential runaway.** Component 1 grows exponentially from operation 300 with time constant $\tau = 50$; Component 2 drifts linearly throughout the entire life.

3.2. Synthetic results

Figure 3 presents the BOCD output for each of the four synthetic scenarios. Each subfigure shows, from top to bottom, the observed signal with detected changepoints, the posterior run-length distribution, the evolving failure probability, and the predicted RUL with uncertainty bands.

A feature shared by all four scenarios is a wide RUL confidence band at early operations that narrows progressively as observations accumulate. Because V_0 is calibrated from the alarm bounds (Section 3.1), the prior predicted mean is confined to the operational envelope from the very first operation, so the median RUL starts near the prediction horizon of 500 operations rather than collapsing to low values. However, the prior does not eliminate coefficient uncertainty — it only anchors the mean. At early time steps, each run-length hypothesis is conditioned on only a handful of observations, so the posterior variance of the regression coefficients is large,

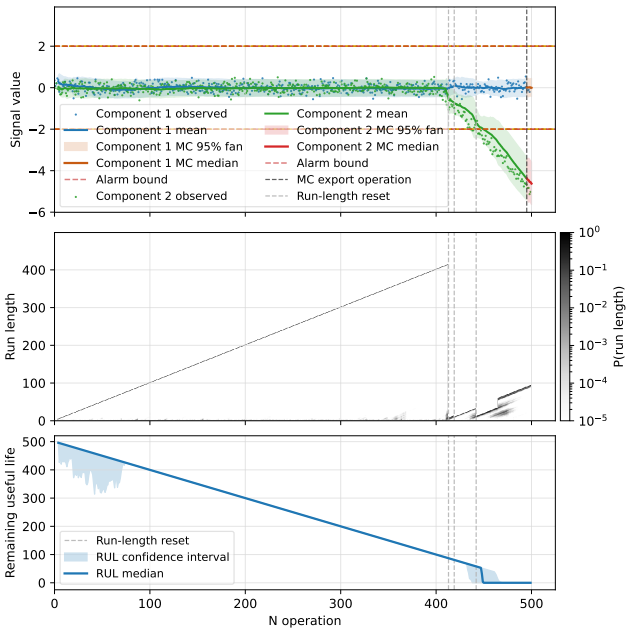

(a) Late sudden failure.

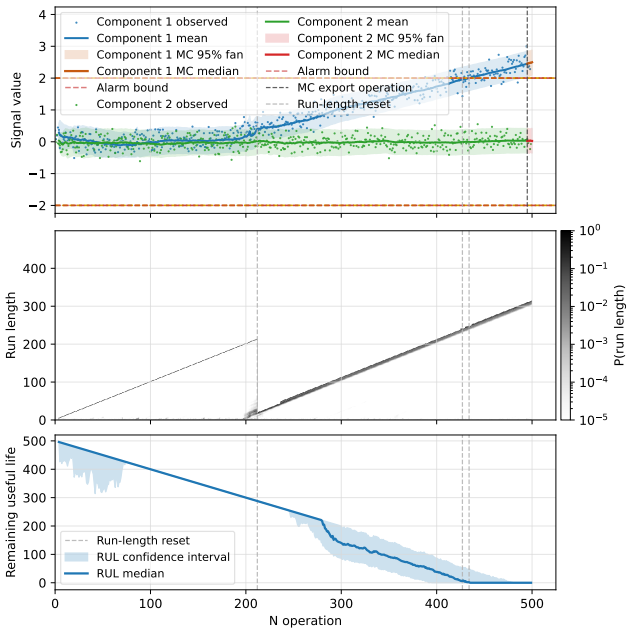

(b) Slow creep.

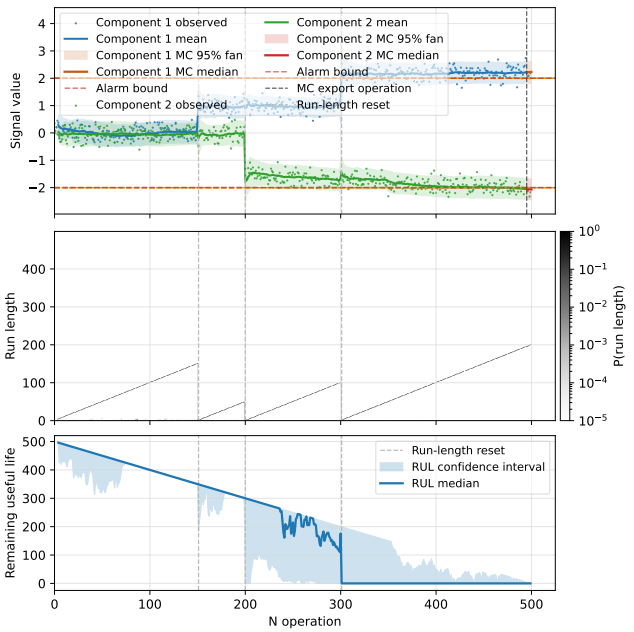

(c) Step drift.

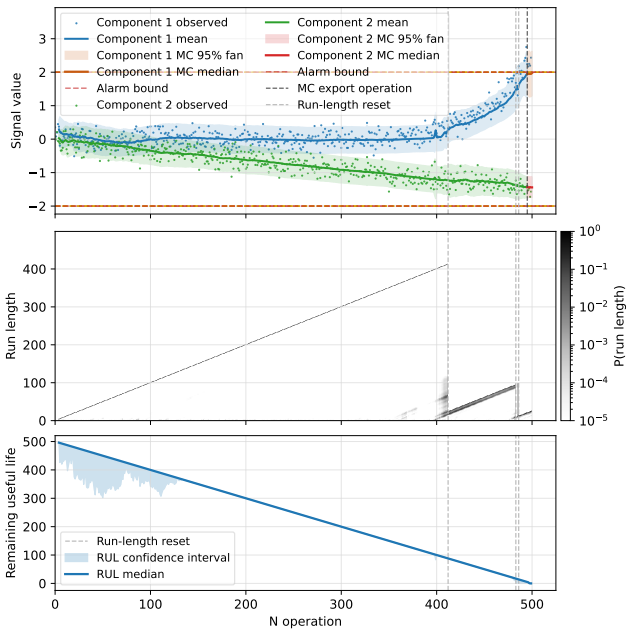

(d) Exponential runaway.

Figure 3. Synthetic scenario results. (a) The abrupt degradation onset near operation 400. (b) The gradual drift from operation 200. (c) Multiple step changes are identified as distinct changepoints with the failure probability adjusting after each transition. (d) The accelerating exponential degradation from operation 300.

see Eq. (22). When this uncertain model is extrapolated, the resulting predictive distribution is wide: while most samples remain within the alarm bounds, the confidence band is substantially spread. As additional observations accumulate within the current segment, the posterior variance contracts, the extrapolated predictive distribution narrows, and the RUL confidence band tightens around the median. This initialization transient is an inherent property of online Bayesian estimation and does not indicate a false alarm; the bound-aware V_0 ensures that it manifests only as elevated uncertainty rather than as a spuriously low RUL estimate.

Late sudden failure (Figure 3a). During the first 400 operations, both signal components remain within the nominal range, and the run-length posterior grows steadily, consistent with a single uninterrupted regime. The failure probability remains near zero, and the median RUL tracks the remaining distance to the prediction horizon. Once the abrupt ramp begins at operation 400, the run-length resets sharply, failure probability rises steeply, and the RUL collapses to near zero within approximately 50 operations. During this transient, the previously narrow confidence band widens as the algorithm adapts to the new, rapidly evolving regime.

Slow creep (Figure 3b). A run-length reset near operation 200 signals the onset of the gradual upward drift in Component 1. After this changepoint, the regression model captures the linear trend and failure probability begins to rise steadily. Median RUL decreases at a moderate rate throughout the drift phase, and the uncertainty band narrows as the trend estimate is refined by accumulating observations. A second run-length adjustment occurs later as the signal approaches the upper bound, further tightening the RUL estimate.

Step drift (Figure 3c). Each step change triggers a distinct run-length reset; the BOCD framework correctly resolves multiple successive changepoints. After the first step at operation 150, the failure probability increases, and the RUL decreases accordingly. At operation 300 the second step produces a further reset and an additional drop in RUL. Between steps the run-length grows linearly and failure probability plateaus, as expected for locally stationary segments. Confidence bands widen briefly at each transition before reconverging as the new segment accumulates data.

Exponential runaway (Figure 3d). During the first 300 operations, the signal is nearly stationary, and the RUL remains high with a gentle downward slope. Once the exponential growth of Component 1 becomes significant, the run-length resets and the failure probability accelerates. Median RUL drops from approximately 300 operations to near zero over the interval 300–400. In the runaway region, the previously tight confidence band expands because the rapidly increasing trend introduces greater extrapolation uncertainty.

4. APPLICATION TO REAL ENDURANCE-TEST DATA

The BOCD framework is applied to mechanical endurance-test data from an electrical switching device, tracking four signal components (Signal 1 through Signal 4) over a certain operation trajectory. The algorithmic parameters are adapted to the longer life and higher noise level of the real dataset compared to the synthetic scenarios: a Weibull hazard with scale $\alpha = 7500$ and shape $\beta = 2.0$, a sliding window of up to 5000 observations, an initialization buffer of $N_{\text{init}} = 50$ operations, and every fourth measurement processed to reduce computational cost. Alarm bounds are derived from the first 50 reference operations. Future predictions are extrapolated up to 1000 operations ahead using $N_{\text{MC}} = 2000$ Monte Carlo samples and polynomial degree 1. The prior hyperparameters are $M_0 = 0$, $\nu_0 = 0.1$, and $\Psi_0 = 0.01\,I$. The regression-coefficient precision matrix V_0 is calibrated from the prescribed alarm bounds for each signal following the procedure in Section 3.1, keeping the initial predicted mean within the operational envelope.

As a practical addition outside the core BOCD formulation, RUL estimation is suppressed during an initialization buffer of $N_{\text{init}} = 100$ operations, over which the regression posterior is too diffuse for meaningful extrapolation. During this phase the RUL rises to the prediction horizon as the posterior contracts; once the buffer is filled, the RUL saturates at the maximum prediction range, consistent with the warm-up behavior observed in the synthetic scenarios (Section 3.2).

The run-length panel reveals numerous changepoints throughout the healthy phase, triggered by the natural variability of the measurements, i.e. small shifts in the signal components that, while statistically significant under the BOCD model, do not indicate degradation. After each reset, the run length grows again, producing the characteristic sawtooth pattern. None of these changepoints propagate into a sustained RUL decrease, confirming that the prognostic layer distinguishes measurement variability from true degradation.

Near the end of life, Signal 2 and 4 drift beyond the alarm bound and trigger a run-length reset. The RUL does not decrease immediately: the new segment initially has too few observations to confirm the trend, and most Monte Carlo trajectories remain within bounds. Only after several hundred more operations does the segment after the changepoint accumulate sufficient data for the failure probability to increase and the Remaining Useful Life (RUL) to decrease. This process uses the same mechanism that suppresses false alarms during the healthy phase. In the final phase, the RUL collapses rapidly to near zero as the signal continues to deteriorate and the confidence interval narrows around the extrapolated trend, reaching zero shortly before the device is taken out of service.

5. CONCLUSION

The results confirm that an online changepoint posterior can serve directly as a prognostic engine for electrical switching devices, with degradation detection and RUL estimation unified within a single recursive update. The framework extends the classical BOCD formulation to a multivariate setting with polynomial regression, conjugate matrix-normal inverse-Wishart priors, and a Weibull hazard function that encodes wear-out behavior.

On synthetic data, the algorithm responds to abrupt regime changes within a few operations, identifies gradual drift once it becomes statistically distinguishable from noise, and pro-

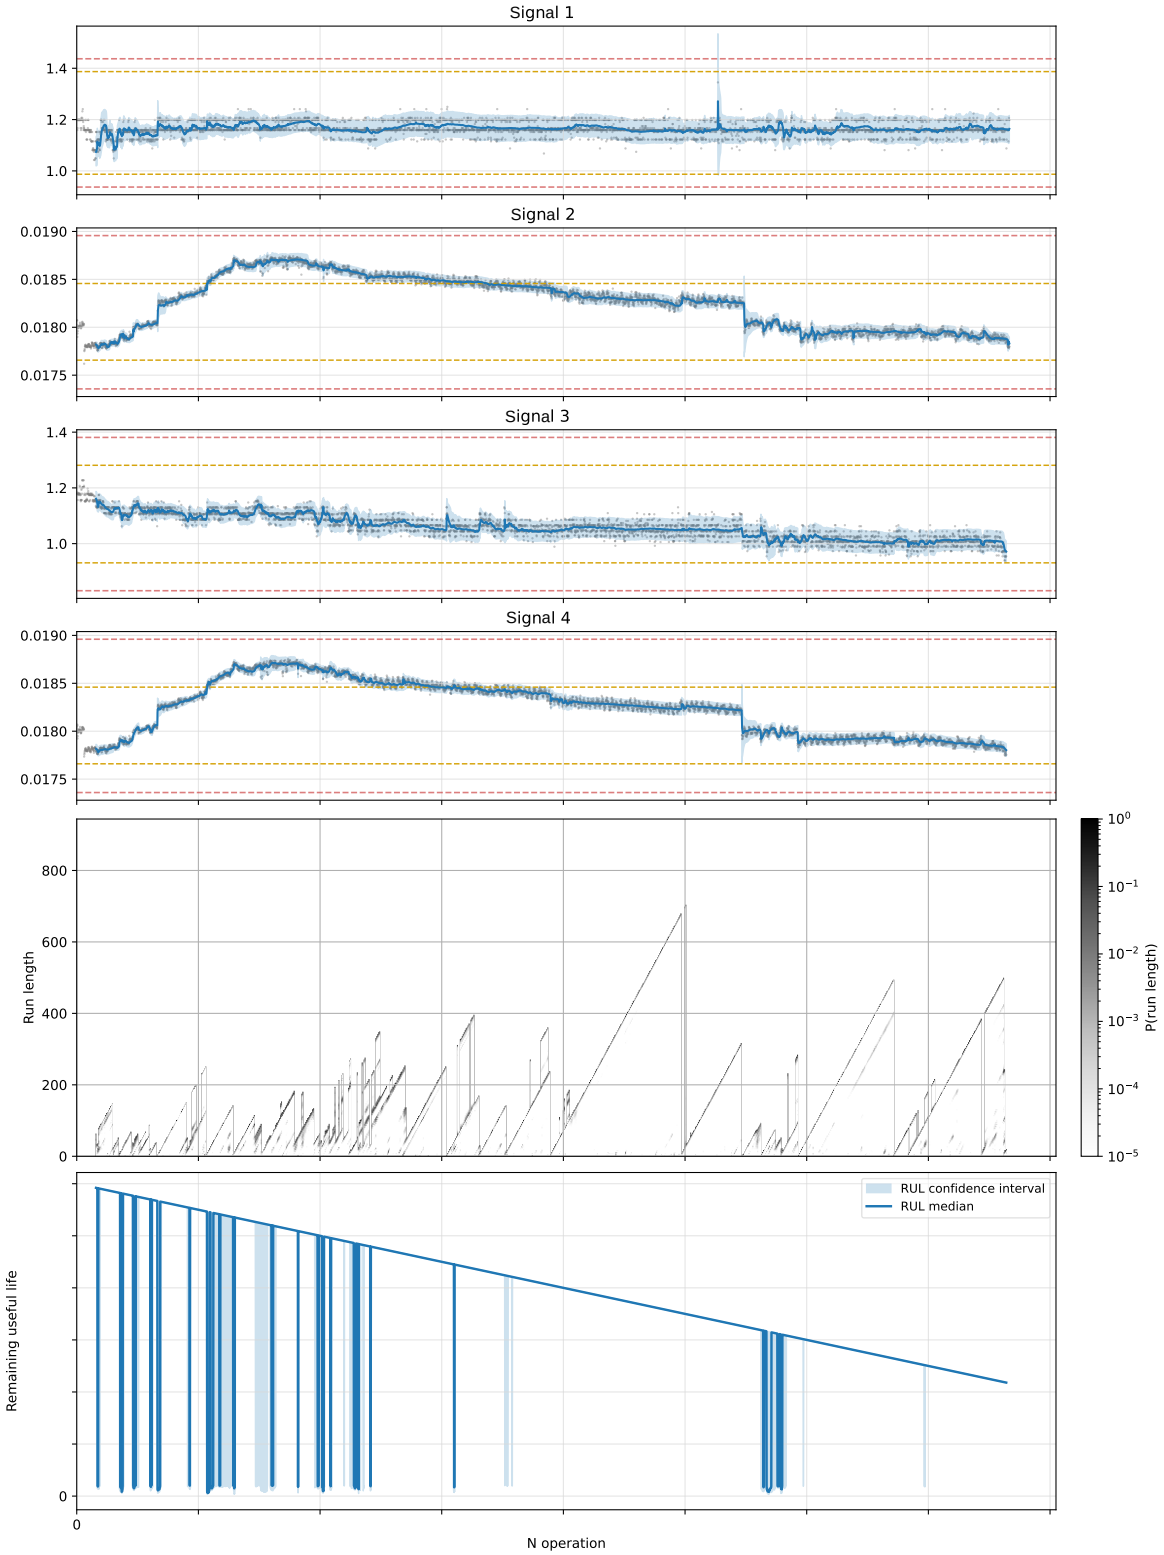

Figure 4. BOCD results for a run-to-failure endurance test. From top to bottom: Signal 1 through Signal 4, each shown with a 95% confidence band, warning bounds (orange dashed), and alarm bounds (red dashed); run-length posterior heatmap; and probabilistic RUL with confidence interval.

duces probabilistic RUL bounds that reflect the full posterior uncertainty. At the end of life, the lower, median, and upper RUL bounds converge consistently. An initialization transient is present in all scenarios: the RUL confidence band is initially wide because the posterior has not yet accumulated enough observations to constrain the segment trend. Because V_0 is calibrated from the prescribed alarm bounds for each monitored signal, the prior predicted mean is confined to the operational envelope from the outset, so the median RUL starts near the prediction horizon rather than collapsing to low values. This effect is an inherent property of online Bayesian estimation and should not be interpreted as a fault indicator; it vanishes once sufficient data constrain the posterior.

On real endurance-test data, the algorithm converges quickly after the initialization phase and maintains a stable RUL throughout the healthy device life despite noise-induced changepoints. When degradation begins, the onset is detected by a run-length reset, but the RUL decreases only once the post-changepoint segment has accumulated enough observations to confirm the trend. This caution prevents premature alarms while still issuing a prognostic warning before failure.

REFERENCES

Adams, R. P., & MacKay, D. J. (2007). Bayesian online changepoint detection. *arXiv preprint arXiv:0710.3742*.

Arunan, A., Sarda, K., Telesca, D., & Khalid, S. (2024). A change point detection integrated remaining useful life estimation model under variable operating conditions. *Control Engineering Practice*, *144*, 105840.

Hoffmann, M. W., Wildermuth, S., Gitzel, R., Boyaci, A., Gebhardt, J., Kaul, H., ... Tornede, T. (2020). Integration of novel sensors and machine learning for predictive maintenance in medium voltage switchgear to enable the energy and mobility revolutions. *Sensors*, *20*(7), 2099. doi: 10.3390/s20072099

Hu, Y., Baraldi, P., Di Maio, F., & Zio, E. (2015). A particle filtering and kernel smoothing-based approach for new design component prognostics. *Reliability Engineering & System Safety*, *134*, 19–31.

Landry, M., Léonard, F., Minsou, C., & Ouellette, R. (2008). An improved vibration analysis algorithm as a diagnostic tool for detecting mechanical anomalies on power circuit breakers. *IEEE Transactions on Power Delivery*, *23*(4), 1986–1994.

Mosallam, A., Medjaher, K., & Zerhouni, N. (2014). Time series trending for condition assessment and prognostics. *Journal of Manufacturing Technology Management*, *25*(4), 550–567.

Murphy, K. P. (2007). *Conjugate Bayesian analysis of the Gaussian distribution* (Tech. Rep.). unpublished: University of British Columbia.

Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. The MIT Press.

Niknam, S. A., Kobza, J. E., & Hines, J. W. (2017). Techniques of trend analysis in degradation-based prognostics. *The International Journal of Advanced Manufacturing Technology*, *88*, 2429–2441.

Runde, M., Ottesen, G. E., Skyberg, B., & Ohlen, M. (1992). Vibration analysis for diagnostic testing of circuit breakers. *IEEE Transactions on Power Delivery*, *7*(4), 1806–1813.

Si, X.-S., Wang, W., Hu, C.-H., & Zhou, D.-H. (2011). Remaining useful life estimation – a review on the statistical data driven approaches. *European Journal of Operational Research*, *213*(1), 1–14.

Vachtsevanos, G. J., Lewis, F., Roemer, M., Hess, A., & Wu, B. (2006). *Intelligent fault diagnosis and prognosis for engineering systems* (Vol. 456). Wiley Online Library.

BIOGRAPHIES

Roman Mukin has been a research scientist in the theoretical and computational methods group at ABB Corporate Research since 2021. He obtained a PhD in Theoretical Physics from Lomonosov Moscow State University in 2007. He was a postdoc from 2011 to 2013 at Vienna University of Technology, Austria. From 2013 to 2021, he worked as a researcher at the Paul Scherrer Institute in Switzerland. His research interests include combining physical modeling with statistical methods to solve problems related to industrial devices, developing reliability engineering and prognostics approaches, and applying computational fluid dynamics and numerical methods to arc physics modeling.

Marco Cilibrasi is a Machine Learning Engineer with a Master's degree in Physics from Scuola Normale Superiore in Pisa (Italy) and a Ph.D. in Computational Science from the University of Zurich (Switzerland). His doctoral research focused on Computational Fluid Dynamics (CFD) and Machine Learning for scientific data analysis. He subsequently worked at ABB Corporate Research Center in Baden (Switzerland), where he specialized in Predictive Maintenance and Prognostics & Health Management. Marco currently works as a Machine Learning Engineer at Artificialy, an AI consulting company based in Lugano (Switzerland).

Kai Hencken is a corporate research fellow at the ABB corporate research center. He obtained a PhD in Theoretical Physics from the University of Basel in 1994. He was a postdoc at the University of Washington from 1995 to 1997 and at the University of Basel from 1997 to 2005, where he received his Habilitation in 2000 and has been a lecturer since. In 2005, he joined the theoretical Physics group at ABB corporate research. His research interests include the combination of physical modeling with statistical methods to solve problems related to industrial devices, as well as developing diagnostics and prognostics approaches.