A Prognostic Framework for Railway Track Geometry: Tamping Detection, Settling-Aware Estimation, and Spatially Resolved RUL Prediction

Abdelhamid Ghoul[1], Wolfgang Lachnit[2], Mohammed Amin Adoul[3], and Wolfgang Birk[1,4]

[1] *Automatic Control, Department of Computer Science, Electrical and Space Engineering, Luleå University of Technology, 97187 Luleå, Sweden*
abdelhamid.ghoul@associated.ltu.se

[2] *Schweizerische Südostbahn AG, Geschäftsbereich Infrastruktur, Bahnhofplatz 1a, 9001 St. Gallen, Switzerland*
wolfgang.lachnit@sob.ch

[3] *Operation and Maintenance, Department of Civil, Environmental and Natural Resources Engineering, Luleå University of Technology, 97187 Luleå, Sweden*
mohammed.amin.adoul@associated.ltu.se

[4] *Predge AB, Vastra Varvsgatan 11, 97236 Luleå, Sweden*
wolfgang.birk@predge.se

ABSTRACT

Accurate estimation of railway track geometry degradation and reliable prediction of remaining useful life (RUL) are essential for cost-effective infrastructure management. This paper presents a self-contained Kalman filtering framework that integrates unsupervised tamping detection and Settling-aware post-tamping state management, jointly addressing state estimation, short-term forecasting, and RUL prediction for longitudinal level and twist parameters using only historical onboard monitoring (OBM) and measurement train (MT) data, without requiring traffic load, environmental, or maintenance metadata. The framework comprises four integrated stages: (i) source-specific outlier detection exploiting the distinct noise characteristics of OBM and MT instruments; (ii) unsupervised tamping detection through adaptive jump analysis on monthly representative values; (iii) a damped Kalman filter with adaptive noise modelling and spatial fusion of neighboring positions, augmented by a Settling-aware state management mechanism that detects the rapid post-tamping consolidation phase and injects physically informed velocity priors to prevent filter lag; and (iv) a RUL prediction module that propagates the final Kalman state forward under damped dynamics until the predicted geometry violates the Alert Limit (AL) or Immediate Action Limit (IAL) defined by EN 13848-5. The complete pipeline is evaluated on 50 consecutive track positions spanning a 50-metre segment of Südostbahn line 870 (Switzerland), using 6 MT and 33 OBM observations per position collected over five years (2016–2021). Results demonstrate accurate estimation through both degradation and maintenance phases, with six-month forecast confidence bands. The multi-position RUL analysis classifies 98% of positions as safe beyond a 2-year horizon and identifies the remaining positions with finite RUL values, enabling spatially targeted maintenance prioritization.

1. INTRODUCTION

The geometric integrity of railway tracks is a primary determinant of operational safety, ride quality, and infrastructure longevity. Among the various track geometry parameters monitored by European railways, longitudinal level and twist irregularities drive the majority of maintenance interventions for speed classes up to 280 km/h (Birk, Westerberg, Larsson-Kråik, & Lachnit, 2021). According to European Standard EN 13848-5, when these parameters exceed prescribed Alert Limits (AL) or Immediate Action Limits (IAL), corrective actions ranging from speed restrictions to track closure become mandatory. The ability to accurately estimate the current geometry state and forecast the time remaining until such thresholds

are violated, commonly referred to as the remaining useful life (RUL), is therefore central to condition-based maintenance strategies.

Modern track geometry monitoring relies on two complementary data sources. Dedicated measurement trains (MT) provide high-precision inspections but at sparse intervals, typically once or twice per year. Onboard monitoring (OBM) systems mounted on in-service vehicles offer temporally dense coverage without disrupting traffic, but at the cost of higher noise, irregular sampling, and susceptibility to localisation errors (Ghiasi, Khan, Sorrentino, Diaine, & Malekjafarian, 2024; Traquinho et al., 2023; Züger, Schlatter, Wolter, Nerlich, & Hunn, 2020). Recent evidence from operational networks suggests that the dense observation rate of OBM can improve the timeliness of condition tracking compared with MT alone (Yan, Hoelzl, Corman, Dertimanis, & Chatzi, 2025); however, exploiting both sources simultaneously requires careful data fusion, noise modelling, and anomaly treatment.

A range of modelling approaches have been explored for track geometry degradation. Classical Kalman filters perform well for recursive state estimation under Gaussian noise but are typically deployed with fixed parameters that limit their adaptability to time-varying degradation dynamics (Traquinho et al., 2023). Machine learning methods, including deep learning architectures such as LSTM and GRU networks, can capture complex spatio-temporal patterns but depend on large labelled training sets that are seldom available in operational settings (Chen, Zhang, & Yang, 2021; Liu, Du, He, & Zhang, 2024). Probabilistic approaches, notably Gaussian processes and hierarchical Bayesian models, provide principled uncertainty quantification but scale poorly over large networks with heterogeneous conditions (Andrade & Teixeira, 2015; Truong-Ba, Rebello, Cholette, Reddy, & Borghesani, 2025). A practical gap therefore persists: the need for a lightweight, metadata-free framework that provides both accurate state estimation and actionable maintenance forecasts directly from the available monitoring data.

A further challenge specific to ballasted track concerns the post-tamping consolidation phenomenon. After a tamping intervention restores the track geometry, the ballast undergoes rapid initial Settling before transitioning to slower long-term degradation, a two-phase behaviour well documented in the geotechnical literature (Selig & Waters, 1994; Sato, 2018; Dahlberg, 2004). Conventional estimation frameworks that reset the filter state to zero velocity after tamping detection incur a transient lag during this consolidation phase, precisely when accurate state tracking is most operationally relevant.

This paper addresses these challenges with a self-contained framework that requires only historical OBM and MT measurements to deliver the complete estimation–forecasting–RUL pipeline. The key contributions are:

1. A fully adaptive Kalman filter with source-specific outlier detection, unsupervised tamping identification, data-driven process and measurement noise estimation, and spatial fusion of neighbouring track positions.
2. A Settling-aware state management mechanism that automatically detects the fast post-tamping consolidation phase and injects physically informed velocity priors, eliminating the estimation lag that affects standard filter resets.
3. A RUL prediction module that propagates the damped Kalman dynamics forward up to 2 years to determine position-specific threshold-crossing times for EN 13848-5 AL and IAL, producing spatially resolved maintenance priority maps.

The framework is validated across 50 consecutive positions on Südostbahn line 870 (Switzerland), demonstrating scalable, position-level health assessment that directly supports spatially targeted maintenance planning.

2. USED FRAMEWORK

The framework processes multi-source track geometry measurements through a pipeline of four stages: preprocessing and data fusion, unsupervised tamping detection, Settling-aware Kalman estimation, and RUL prediction. Each stage is designed to operate without external metadata, relying solely on the statistical properties of the measurement time series.

2.1. Data Fusion and Preprocessing

OBM and MT measurements at the target track position p and its spatial neighbours $p \pm 1$ are merged into a single chronologically sorted time series. Each observation z_k carries a timestamp t_k and a source label. This multi-position extraction enables the spatial fusion described in Section 2.3.

2.1.1. OBM Outlier Detection

OBM data are cleaned using a monthly Z-score strategy. For each calendar month m, the mean μ_m and standard deviation σ_m of all OBM readings are computed. A point z_k is flagged as an outlier and replaced by μ_m if $|z_k - \mu_m| > 1.5\,\sigma_m$. This leverages the quasi-stationarity of degradation within a single month to isolate sensor-noise-driven anomalies while preserving the temporal density of OBM records (Grubbs, 1969).

2.1.2. MT Outlier Detection

The sparse nature of MT data demands a more conservative approach. An adaptive jump-detection criterion flags z_k as an outlier only if both the forward and backward deviations exceed five times the local variation *and* the deviations exhibit a trend reversal (i.e., the point is a spike rather than a level shift). Flagged MT points are linearly interpolated from their temporal neighbours, retaining the measurement timestamp to

preserve the chronological structure.

2.2. Unsupervised Tamping Detection

Tamping events produce abrupt geometry corrections that, if undetected, corrupt noise statistics and cause filter divergence. Since ground-truth maintenance logs are typically unavailable, we detect tamping from the measurement signal itself.

OBM readings are aggregated into monthly representative values (means), while MT points retain their original precision. The resulting mixed time series is scanned for jumps. A tamping event is flagged when the absolute change between consecutive representative values exceeds the threshold

$$\tau = \max\big(3 \cdot \mathrm{MAD}(\Delta z),\ \Delta z_{\text{range}}/10\big) \tag{1}$$

where $\mathrm{MAD}(\Delta z)$ is the median absolute deviation of consecutive differences (scaled by 1.4826 for robustness) and Δz_{range} is the overall data range. Two guards enhance robustness: (i) the jump direction must oppose the pre-event degradation trend estimated by linear regression on the preceding representative values, ensuring that only corrective interventions (not accelerated degradation) are flagged; and (ii) a 30-day guard is applied to opposite-direction jumps occurring shortly after a detected tamping event, since such jumps may reflect rapid post-tamping settling rather than a new corrective intervention. This guard is therefore not intended as a blanket exclusion of all successive maintenance detections within 30 days. After the first tamping is identified, its correction direction is locked so that subsequent events must share the same restoration polarity, a physically motivated constraint since tamping on a given rail consistently restores geometry in one direction. Corrective jumps with the same restoration direction remain eligible as candidate maintenance events, provided that they satisfy the jump-magnitude and trend-consistency criteria. In exceptional emergency scenarios, operator knowledge or maintenance logs would still be required to distinguish a genuine rapid intervention from settling-driven behaviour. Detected tamping times are then mapped to the nearest raw measurement timestamp to preserve chronological consistency.

Detected tamping timestamps $\mathcal{T}_{\text{tamp}}$ are used to exclude maintenance periods from noise estimation and to trigger the settling-aware state reset described in Section 2.4.

2.3. Adaptive Kalman Estimation

2.3.1. State-Space Model

The degradation state at time step k is represented by $\mathbf{x}_k = [\theta_k,\ v_k,\ a_k]^\top$, where θ_k is the track geometry value (longitudinal level or twist), v_k the degradation velocity, and a_k the acceleration. The state transition incorporates a velocity damping factor $\alpha \in [0, 1]$:

$$\mathbf{x}_{k+1} = A_k\,\mathbf{x}_k + \boldsymbol{\epsilon}_k, \quad A_k = \begin{bmatrix} 1 & \Delta t_k & \frac{1}{2}\Delta t_k^2 \\ 0 & \alpha & \Delta t_k \\ 0 & 0 & 1 \end{bmatrix} \tag{2}$$

where $\Delta t_k = t_{k+1} - t_k$ in days and $\boldsymbol{\epsilon}_k \sim \mathcal{N}(\mathbf{0}, Q_k)$. The damping factor α reflects the well-documented tendency of ballast Settling rates to attenuate over time (Selig & Waters, 1994; Sato, 2018), following the fading-memory filtering paradigm (Bar-Shalom, Li, & Kirubarajan, 2001). When $\alpha = 1$ the model reduces to constant-acceleration kinematics; when $\alpha < 1$ older velocity information is progressively discounted, preventing filter divergence under regime changes such as post-tamping recovery.

The process noise covariance is $Q_k = G_k G_k^\top \sigma_w^2$ with the noise gain vector $G_k = [\frac{1}{6}\Delta t_k^3,\ \frac{1}{2}\Delta t_k^2,\ \Delta t_k]^\top$. The process noise intensity σ_w^2 is estimated from second-order finite differences of the cleaned data, scaled by a factor $\xi = 0.1$ to compensate for noise amplification inherent in numerical differentiation (Simon, 2006).

2.3.2. Data-Driven Damping Estimation

The damping factor is estimated separately for the pre-tamping and post-tamping regimes via least-squares regression on consecutive finite-difference velocities $v_k = (z_{k+1} - z_k)/\Delta t_k$ computed from cleaned measurements (excluding tamping timestamps):

$$\hat{\alpha} = \frac{\sum_{k=1}^{N-1} v_{k+1}\, v_k}{\sum_{k=1}^{N-1} v_k^2} \tag{3}$$

where k denotes the time-step index within the considered estimation window, and N is the number of available measurements in the considered estimation window. This provides a regime-specific characterisation of velocity persistence, enabling the filter to adapt its dynamics to the degradation phase.

2.3.3. Adaptive Measurement Model

The measurement noise covariance adapts to the data source. For MT, a fixed $R_{\text{MT}} = 5 \times 10^{-3}$ reflects their high precision. For OBM, the monthly variance of first-order differences $R_{\text{OBM},m}$ captures time-varying noise characteristics. A fused observation model exploits spatial redundancy by jointly updating the state with measurements from positions $p - 1$, p, and $p + 1$:

$$H = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \tag{4}$$

$$R_{\text{fused},m} = \mathrm{diag}(R_{m,p-1},\ \lambda R_{m,p},\ R_{m,p+1}) \tag{5}$$

where $\lambda = 0.5$ reduces the covariance assigned to the target-position measurement, thereby giving the central position a higher weight than its neighbours. This spatial fusion dampens idiosyncratic noise fluctuations while preserving position

specificity, yielding more stable Kalman updates under the high-noise conditions typical of OBM data. The spatial window was intentionally limited to the immediate neighbours $p-1$ and $p+1$. Since the objective of the proposed framework is to estimate degradation and RUL at the individual position level, a broader spatial window could over-smooth local defects or maintenance-induced discontinuities and reduce the ability to identify position-specific hotspots. The immediate-neighbour fusion therefore provides a conservative compromise: it exploits local spatial redundancy to reduce isolated measurement noise while preserving the spatial resolution required for position-specific RUL mapping. The covariance weighting further ensures that the target position remains dominant in the Kalman update, while neighbouring measurements provide only local contextual support.

2.4. Settling-Aware State Management

After tamping, ballasted track undergoes rapid initial consolidation (Settling) before transitioning to slower long-term degradation (Selig & Waters, 1994). If the filter simply resets to zero velocity after tamping, it must "learn" the elevated Settling rate from subsequent observations, introducing a transient estimation lag during the operationally critical post-maintenance period.

2.4.1. Pre-Tamping Trend Estimation

For each detected tamping event c, a second-order polynomial is fitted to the representative values in a two-year window preceding the event. The resulting velocity $\hat{v}_{\text{pre}}^{(c)}$ and acceleration $\hat{a}_{\text{pre}}^{(c)}$ at the tamping time characterise the steady-state degradation baseline. Their signs are constrained to match the Settling direction (opposite to the tamping correction), enforcing the physical prior that degradation resumes in the direction that originally necessitated the intervention.

2.4.2. Settling Phase Detection

The Settling phase is identified by comparing post-tamping velocities against a rate threshold:

$$r_{\text{thr}} = \max\big(3 \cdot \bar{v}_{\text{steady}},\ 0.001\big) \tag{6}$$

where $\bar{v}_{\text{steady}}$ is the mean absolute pre-tamping rate. Starting from the first representative point after tamping, intervals where the velocity exceeds r_{thr} in the Settling direction are flagged as fast Settling. The phase terminates when the rate drops below r_{thr} or after 150 days, whichever occurs first.

2.4.3. Modified Kalman Updates During Settling

During the detected Settling phase, the filter state is modified at three key moments:

- **First Settling point:** The position state is set to the observed measurement; velocity and acceleration are initialised at $3\times$ the pre-tamping estimates, capturing the elevated consolidation rate.
- **Intermediate points:** Standard Kalman updates proceed, allowing the filter to learn the evolving Settling dynamics from the data.
- **Last Settling point:** After the standard update, velocity and acceleration are overwritten with the pre-tamping baseline values, ensuring a smooth transition to the steady-state degradation regime.

This three-phase injection strategy faithfully tracks the rapid post-tamping transient while ensuring convergence to physically plausible long-term dynamics.

The main advantage of the proposed Settling-aware mechanism is that it separates the post-tamping response into two physically different regimes: a short rapid consolidation phase and a subsequent steady degradation phase. In a conventional Kalman reset, the post-tamping velocity is usually initialized to zero, forcing the filter to infer the high initial Settling rate only after several new observations. This creates a delayed estimate and may underestimate the true short-term degradation immediately after maintenance. In contrast, the proposed mechanism injects velocity and acceleration priors derived from the pre-tamping trend and aligned with the physically expected Settling direction. Therefore, the filter starts the post-maintenance phase with a realistic dynamic state, tracks the fast consolidation without waiting for multiple updates, and then returns to the long-term degradation regime once the Settling phase ends. This improves post-intervention verification and provides a more reliable initial condition for short-term forecasting and RUL estimation.

2.5. Forecasting and Remaining Useful Life Prediction

2.5.1. Forecasting Module

The six-month forecast is initialised from the final Kalman-estimated geometry $\hat{\theta}_N$, with velocity v_m and acceleration a_m obtained by exponentially weighted moving averaging (EWMA) of the last three normal (non-Settling) representative estimates:

$$v_m = \frac{v_N + 0.5\,v_{N-1} + 0.25\,v_{N-2}}{1.75}, a_m = \frac{a_N + 0.5\,a_{N-1}}{1.5} \tag{7}$$

The EWMA formulation assigns decaying weights to older estimates exponentially, striking a balance between responsiveness to current trends and robustness against short-term fluctuations. The forecast state is propagated at a fixed 7-day step using the same transition matrix (Eq. 2) with the post-tamping damping factor α. The forecast covariance $P^{(f)}$ is initialised with position variance from the final Kalman update and sample variances of the velocity and acceleration estimates.

2.5.2. RUL Computation

The RUL with respect to a regulatory threshold L (AL or IAL per EN 13848-5) is defined as the time until the forecast geometry first violates the threshold, i.e., $|\theta^{(f)}_{N+j}| \geq L$. The RUL computation uses the same forecast trajectory generated in Section 2.5.1. Specifically, the predicted geometry component $\theta^{(f)}_{N+j}$ of the propagated Kalman state is evaluated against the regulatory threshold L at each forecast step. The forecast propagation is extended up to 24 months (2 years). At each step j, if $|\theta^{(f)}_{N+j}|$ first exceeds L, the nominal crossing time is refined by linear interpolation between steps $j-1$ and j:

$$\mathrm{RUL}_L = t_{N+j-1} + \frac{L - |\theta^{(f)}_{N+j-1}|}{|\theta^{(f)}_{N+j}| - |\theta^{(f)}_{N+j-1}|} \cdot \Delta\tau - t_N \quad (8)$$

where $\Delta\tau = 7$ days. If neither threshold is reached within 2 years, the position is classified as SAFE.

Equation (8) provides the nominal RUL computed from the deterministic forecast mean $\theta^{(f)}_{N+j}$. Forecast uncertainty is propagated through the forecast covariance $P^{(f)}_{N+j}$, whose first diagonal element gives the predicted geometry variance, $\sigma^2_{\theta,N+j} = P^{(f)}_{N+j}(1,1)$. When uncertainty-aware RUL bounds are required, confidence envelopes can be formed as

$$\theta^{+}_{N+j} = \theta^{(f)}_{N+j} + q\sigma_{\theta,N+j}, \qquad \theta^{-}_{N+j} = \theta^{(f)}_{N+j} - q\sigma_{\theta,N+j},$$

where $q = 1.96$ corresponds to a 95% confidence level. The same threshold-crossing rule in Eq. (8) is then applied to the envelopes: the earliest crossing of either $|\theta^{+}_{N+j}|$ or $|\theta^{-}_{N+j}|$ provides a conservative lower RUL bound, while the deterministic forecast provides the nominal RUL. If no crossing occurs within the 24-month horizon, the corresponding value is reported as SAFE.

The velocity damping factor $\alpha < 1$ naturally bounds the forecast trajectory, preventing unrealistic divergence and reflecting the physical tendency of Settling rates to attenuate over time.

2.6. From Estimation to Maintenance Decision Support

The four stages described above form an integrated pipeline that translates raw monitoring data into maintenance decision support without external metadata. Source-specific outlier detection and multi-source fusion of OBM and MT data, combined with spatial fusion of neighbouring positions, ensure that the estimation operates on robust, noise-reduced input throughout the pipeline. The tamping detection module (Section 2.2) provides automated maintenance history reconstruction, eliminating the dependency on incomplete or unavailable maintenance logs. The adaptive Kalman filter with velocity damping factor α (Section 2.3) captures regime-specific degradation dynamics, while the Settling-aware mechanism (Section 2.4) enables post-intervention verification by accurately tracking the rapid consolidation phase. The RUL module (Section 2.5) translates the estimated state into position-specific threshold-crossing times relative to EN 13848-5 limits. The damping factor $\alpha < 1$ ensures that these forecasts remain physically bounded, preventing unrealistic divergence and yielding planning-grade predictions over horizons of up to two years. By computing RUL independently per position and per geometry parameter, the framework produces spatially resolved priority maps that identify which positions require intervention and which parameter governs the maintenance timeline, supporting targeted planning rather than uniform inspection schedules.

3. CASE STUDY

3.1. Dataset and Experimental Setup

The framework is evaluated on a segment of Südostbahn (SOB) line 870 between Mogelsberg and Brunnadern (Switzerland). We analyse 50 consecutive track positions at 1 m spatial resolution (positions 42 775 to 42 824), covering three geometry parameters: left longitudinal level, right longitudinal level, and 3 m twist. Each position–parameter pair comprises 6 MT observations and approximately 33 OBM-derived indicator readings over the period July 2016 to September 2021. The OBM indicators are pre-processed and quality-assured longitudinal level values in the D1 wavelength range (3–25 m per EN 13848-1) provided by the operator's monitoring pipeline (Züger et al., 2020). Applicable EN 13848-5 thresholds are AL = 8 mm and IAL = 12 mm for longitudinal level, and AL = 3.5 mm and IAL = 4 mm for 3 m twist.

3.2. Single-Position Results (Position 42 800)

3.2.1. Right Longitudinal Level

Figure 1 shows the estimation and forecasting results for the right rail at position 42 800 m. The Kalman estimate (green circles) accurately tracks the gradual degradation from $+0.2$ mm to $+0.8$ mm over 2016–2020. The unsupervised tamping detector correctly identifies the mid-2020 maintenance event (yellow band), where the geometry drops by approximately 2 mm. Following the filter reset, the Settling-aware mechanism captures the elevated post-tamping consolidation (orange band) before transitioning to steady-state tracking. The six-month forecast (magenta) predicts gradual recovery toward $+0.2$ mm with a confidence band of ± 0.15 mm.

Figure 2 presents the same results relative to the EN 13848-5 safety thresholds. All estimated and forecasted values remain well within the regulatory limits, with a margin exceeding 7 mm to the Alert Limit.

3.2.2. Left Longitudinal Level

The left rail (Figure 3) exhibits similar behaviour. The mid-2020 tamping corrects the geometry by approximately 1.8 mm.

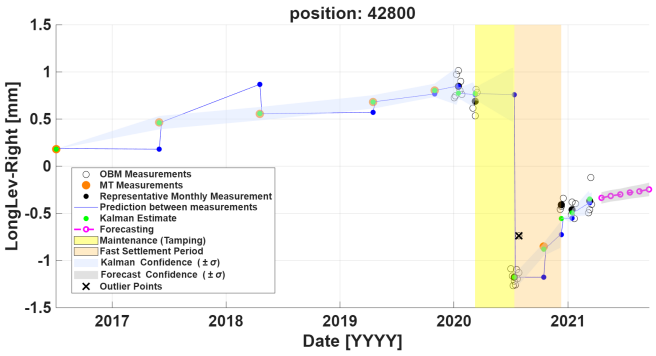

Figure 1. Kalman estimation and six-month forecasting for right longitudinal level at position 42 800.

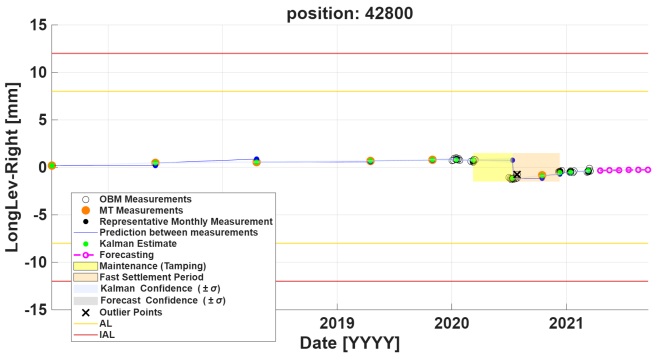

Figure 2. Right longitudinal level at position 42 800 with EN 13848-5 thresholds (AL = ±8 mm, IAL = ±12 mm). All values remain far from regulatory limits.

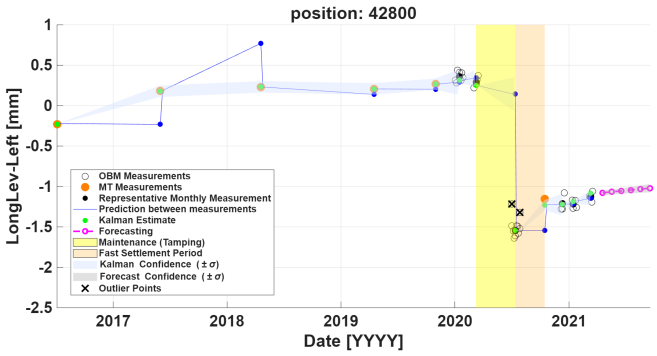

Figure 3. Kalman estimation and six-month forecasting for left longitudinal level at position 42 800.

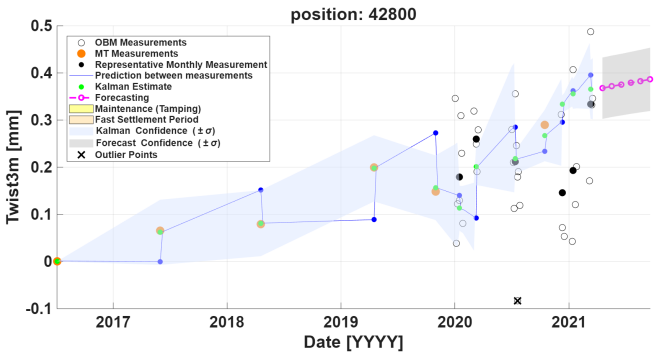

Figure 4. Kalman estimation and six-month forecasting for 3 m twist at position 42 800. No tamping event is detected for this parameter.

The Settling-aware filter correctly captures the rapid post-tamping consolidation, and the forecast indicates stable behaviour around -1.2 mm with a confidence band of ± 0.05 mm. The baseline method diverges by 0.5 mm within six months of the maintenance event.

3.2.3. Twist over 3 m Base

For 3 m twist (Figure 4), values remain near zero throughout the study period (0 to $+0.35$ ‰), with a degradation rate of approximately 0.07 ‰/year. No tamping event is detected for twist, which is physically consistent: since twist measures the differential cant variation between the two rails over a 3 m base, symmetric maintenance on both rails produces minimal twist change. The forecast predicts continued stability at $+0.35$ ‰with a confidence band of ± 0.15 ‰, far below the 3.5 mm Alert Limit.

3.3. Multi-Position Spatio-Temporal Analysis

To demonstrate the framework's scalability and its utility for network-level decision support, the full pipeline is applied independently to all 50 positions (42 775–42 824) and the results are visualised as spatio-temporal surfaces and position-level RUL summaries.

3.3.1. 3D Track Geometry Surface

Figure 5 presents the Kalman-estimated right longitudinal level as a three-dimensional surface, with time on one axis, position on the other, and geometry value on the vertical axis. The surface reveals the spatially localised tamping event around mid-2020 as a sharp downward correction affecting a cluster of positions near 42 800. The surrounding positions exhibit varying degradation rates and geometry levels, illustrating the spatial heterogeneity that motivates position-level analysis rather than segment-averaged approaches.

3.3.2. RUL Summary Across Positions

Figure 6 summarises the RUL estimates for right longitudinal level across all 50 positions. Green dots indicate positions classified as SAFE (threshold not reached within 2 years), while coloured bars represent finite RUL values in months. Under this 2-year horizon, nearly all positions are classified as safe for both AL and IAL, with only a single position exhibiting a finite RUL of 17 months to AL; no position reaches IAL within 2 years, and therefore no IAL bars are shown. The isolated at-risk position is consistent with the localised degradation hotspot observed in the 3D surface and suggests a position-specific substructure weakness or loading concentration that

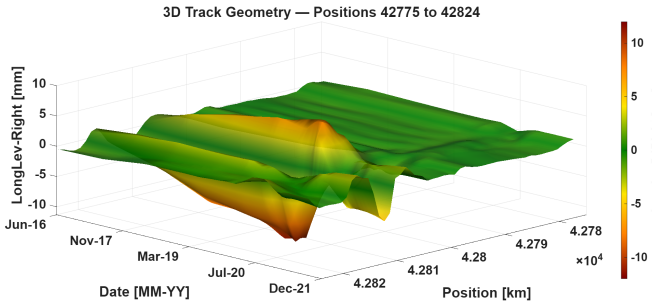

Figure 5. 3D spatio-temporal surface of estimated right longitudinal level across positions 42 775–42 824. The localised tamping event (mid-2020) is visible as a sharp correction near position 42 800.

Table 1. Summary of estimation and forecasting performance at position 42 800.

Parameter	Tamping detected	Forecast $\pm\sigma$	Margin to AL
Left long. level	Yes (mid-2020)	± 0.05 mm	> 6.8 mm
Right long. level	Yes (mid-2020)	± 0.15 mm	> 7.0 mm
Twist 3 m	No	± 0.15 ‰	> 3.1 mm

warrants root-cause investigation. For the left longitudinal level (Figure 7), all positions are classified as SAFE within the 2-year horizon, confirming stable behaviour across the analysed segment.

For twist (Figure 8), nearly all positions are classified as safe, consistent with the low degradation rates and the absence of maintenance-induced discontinuities for this parameter. Only one boundary position exhibits a finite RUL.

3.4. Discussion

The case study results validate the decision support pipeline described in Section 2.6.

The unsupervised tamping detector correctly identified the mid-2020 maintenance event on both rails at position 42 800 without external logs. The multi-source fusion of 33 OBM and 6 MT observations per position produced stable estimates throughout degradation and post-maintenance phases, with forecast confidence bands of ± 0.05–0.15 mm for longitudinal level.

The Settling-aware mechanism eliminated post-maintenance estimation lag, where observed post-tamping rates of 0.6–1.4 mm/year exceeded steady-state rates, confirming the two-phase consolidation behaviour (Selig & Waters, 1994; Sato, 2018). The filter immediately captured the fast Settling without requiring multiple updates to converge.

Under the 2-year horizon, 98% of positions were classified as SAFE. The single at-risk position (17 months to AL for right longitudinal level) is spatially consistent with the hotspot in Figure 5, confirming that the RUL maps identify position-specific weaknesses. All left longitudinal level positions remained safe, and the absence of tamping detection for twist is physically consistent with symmetric bilateral maintenance. Overall, no network-wide intervention is required within two years; targeted monitoring should focus on the identified hotspot locations. Table 1 summarises the key performance metrics at the reference position.

4. CONCLUSION

This paper presented a self-contained, Settling-aware Kalman filtering framework for railway track geometry estimation, forecasting, and remaining useful life prediction. The framework operates exclusively on historical OBM and MT data, requiring no traffic load, environmental, or maintenance metadata. Evaluated across 50 consecutive positions of the Südostbahn line 870 in Switzerland, the approach demonstrated three key capabilities.

First, the unsupervised tamping detection reliably identified maintenance events from the measurement signal alone, enabling automatic filter state management without external maintenance logs. Second, the Settling-aware mechanism eliminated the estimation lag observed during the rapid post-tamping consolidation phase by injecting physically informed velocity priors derived from pre-tamping trend analysis. Third, the RUL prediction module, bounded by the data-driven velocity damping factor, produced physically plausible threshold-crossing times that classified 98% of positions as safe (no threshold violation within 2 years).

From an operational perspective, the multi-position RUL map provides a compact, actionable summary for maintenance prioritisation. The risk positions enables infrastructure managers to investigate underlying causes and plan targeted interventions rather than relying on uniform inspection schedules.

Future work will incorporate additional geometry parameters (gauge, alignment, cant), investigate cross-parameter RUL coupling for joint maintenance optimisation, and validate the framework on independent railway networks with different operational and climatic profiles.

ACKNOWLEDGMENT

This work is supported by Kempestiftelserna, Project ID: 371940. The authors gratefully acknowledge the Kempe Foundation for supporting this research and Südostbahn AG for providing the measurement data.

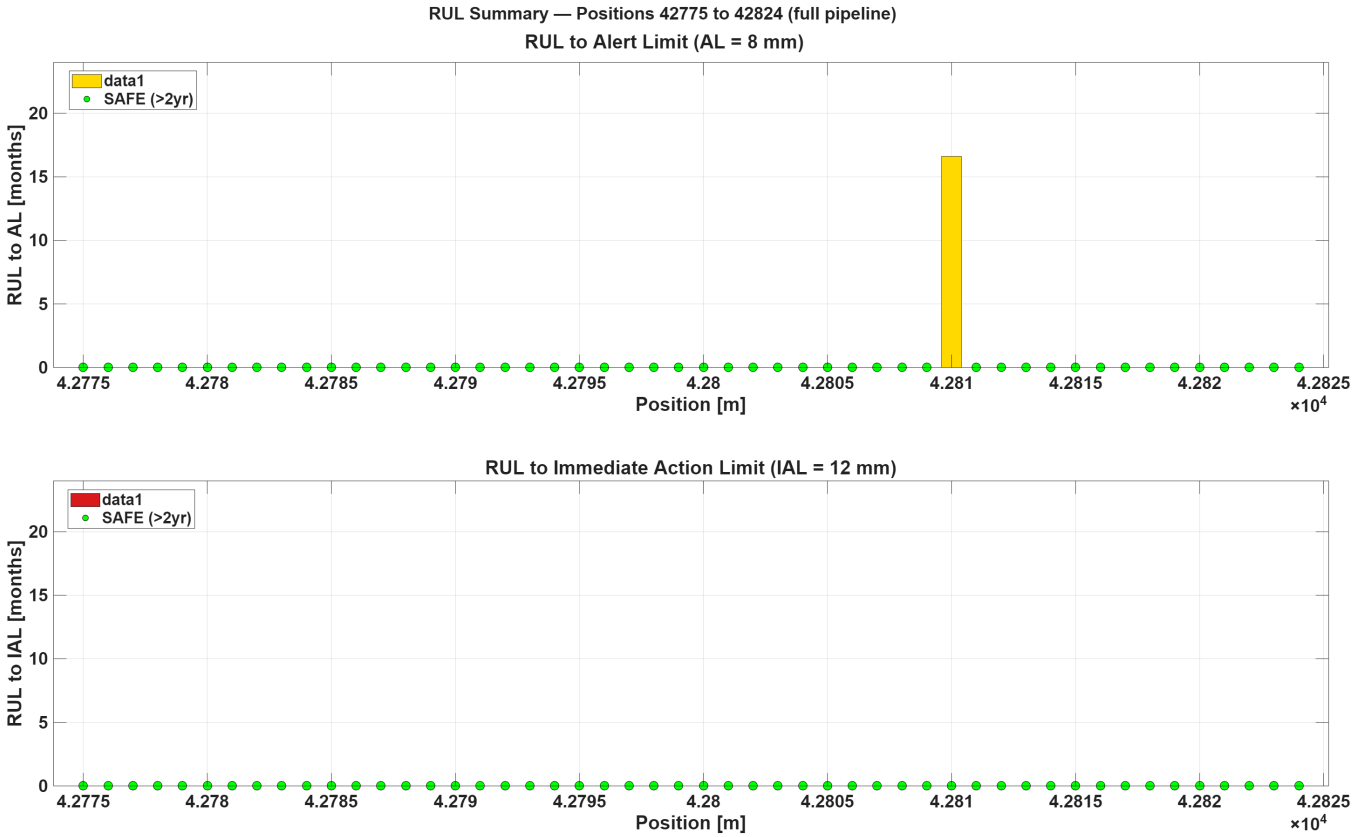

Figure 6. RUL summary for right longitudinal level across positions 42 775–42 824. Upper: time to Alert Limit (AL = 8 mm); lower: time to Immediate Action Limit (IAL = 12 mm). Green dots: safe (>2 years); bars: finite RUL in months.

NOMENCLATURE

θ_k	Track geometry state (longitudinal level or twist)
v_k	Degradation velocity
a_k	Degradation acceleration
α	Velocity damping (retention) factor
σ_w^2	Process noise intensity
R_k	Measurement noise covariance
λ	Spatial weighting coefficient
AL	Alert Limit (EN 13848-5)
IAL	Immediate Action Limit (EN 13848-5)
RUL	Remaining Useful Life
OBM	Onboard Monitoring
MT	Measurement Train

REFERENCES

Andrade, A. R., & Teixeira, P. F. (2015). Statistical modelling of railway track geometry degradation using hierarchical bayesian models. *Reliability Engineering & System Safety*, *142*, 169–183.

Bar-Shalom, Y., Li, X. R., & Kirubarajan, T. (2001). *Estimation with applications to tracking and navigation*. John Wiley & Sons.

Birk, W., Westerberg, J., Larsson-Kråik, P. O., & Lachnit, W. (2021). Track geometry estimation and prediction tool combining onboard monitoring and measurement vehicle data. In *2021 arema annual conference & expo.* (Virtual Conference)

Chen, Y., Zhang, Y., & Yang, F. (2021). Learn to predict vertical track irregularity with extremely imbalanced data. In *Asian conference on machine learning* (pp. 1493–1504). PMLR.

Dahlberg, T. (2004). *Railway track settlements — a literature review* (Tech. Rep. No. 463). Report for the EU project SUPERTRACK.

Ghiasi, R., Khan, M. A., Sorrentino, D., Diaine, C., & Malekjafarian, A. (2024). An unsupervised anomaly detection framework for onboard monitoring of railway track geometrical defects using one-class support vector machine. *Engineering Applications of Artificial Intelligence*, *133*, 108167.

Grubbs, F. E. (1969). Procedures for detecting outlying observations in samples. *Technometrics*, *11*(1), 1–21.

Liu, J., Du, D., He, J., & Zhang, C. (2024). Prediction of remaining useful life of railway tracks based on dmgdcc-gru hybrid model and transfer learning. *IEEE Transactions on Vehicular Technology*, *73*(6), 7561–7575.

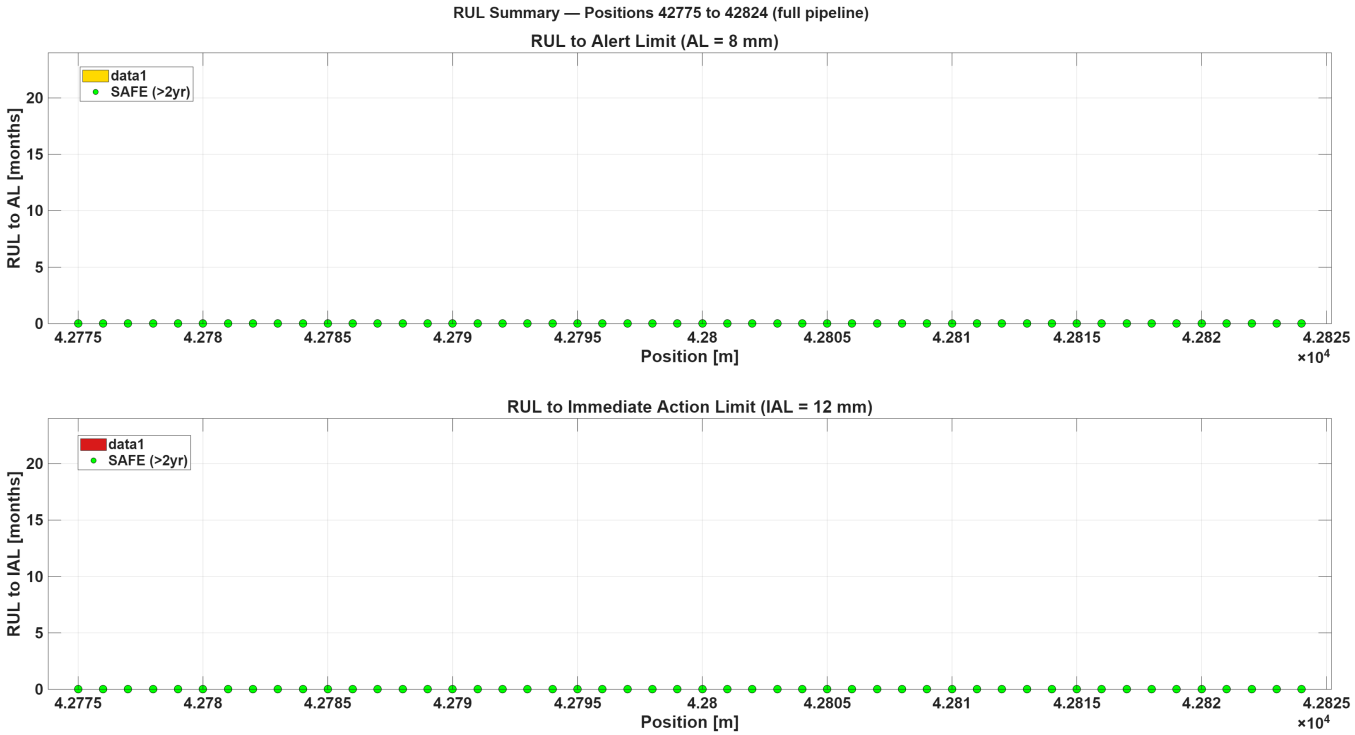

Figure 7. RUL summary for left longitudinal level across positions 42 775–42 824. Upper: time to Alert Limit (AL = 8 mm); lower: time to Immediate Action Limit (IAL = 12 mm). Green dots: safe (>2 years); bars: finite RUL in months.

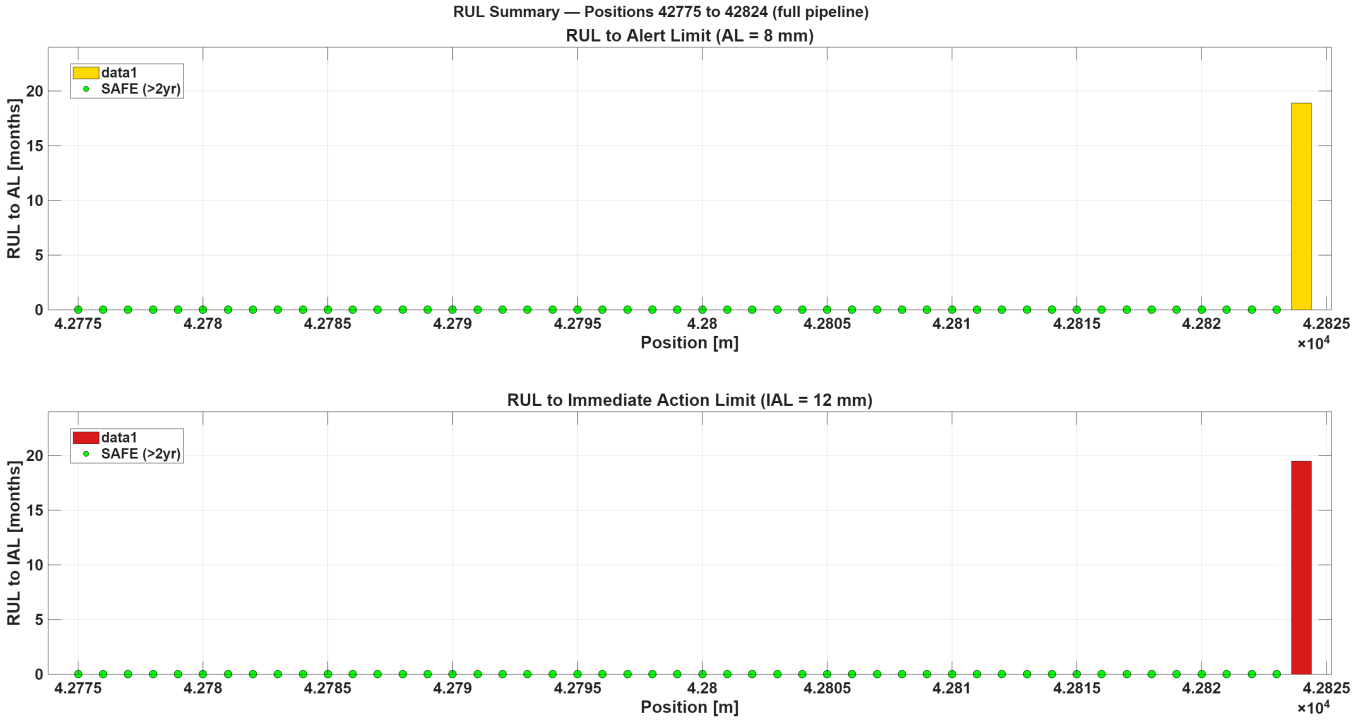

Figure 8. RUL summary for 3 m twist across positions 42 775–42 824. Nearly all positions are safe, confirming twist is not the governing maintenance parameter on this segment.

Sato, Y. (2018). Japanese studies on deterioration of ballasted track. In *Interaction of railway vehicles with the track and its substructure* (pp. 197–208). Routledge.

Selig, E. T., & Waters, J. M. (1994). *Track geotechnology and substructure management*. Thomas Telford.

Simon, D. (2006). *Optimal state estimation: Kalman, h infinity, and nonlinear approaches*. John Wiley & Sons.

Traquinho, N., Vale, C., Ribeiro, D., Meixedo, A., Montenegro, P., Mosleh, A., & Calçada, R. (2023). Damage identification for railway tracks using onboard monitoring systems in in-service vehicles and data science. *Machines*, *11*(10), 981.

Truong-Ba, H., Rebello, S., Cholette, M. E., Reddy, V., & Borghesani, P. (2025). Bayesian multivariate track geometry degradation modeling and its use in condition-based inspection. *Railway Engineering Science*, 1–25.

Yan, T. H., Hoelzl, C., Corman, F., Dertimanis, V., & Chatzi, E. (2025). Integration of on-board monitoring data into infrastructure management for effective decision-making in railway maintenance. *Railway Engineering Science*, *33*(2), 151–168.

Züger, S., Schlatter, C., Wolter, K. U., Nerlich, I., & Hunn, S. (2020). Onboard monitoring in der schweiz, ein gemeinschaftswerk dreier bahnen. *ZEVrail*, *144*(4).