Bridge Weigh-in-Motion as a regularized Bayesian inverse problem

Verification, validation against measured data, and model-error-aware uncertainty

Muhammet Emir Fil

B.Sc. Computational Engineering Science, RWTH Aachen • Technical report, June 21, 2026

Aim. Build the forward physics of a vehicle crossing a bridge, then run it backwards to recover the truck's axle weights (Bridge Weigh-in-Motion), and treat that recovery as an inverse problem: analytical verification of every forward block, regularization of the ill-posed part, an uncertainty estimate calibrated from model error rather than assumed, and validation against real measured bridge data. The core is numpy-only.

Abstract

A truck crossing a bridge leaves a measurable trace in the structure's response. Read forwards, that trace predicts the bridge's vibration; read backwards, it recovers the axle weights. The second reading is Bridge Weigh-in-Motion (B-WIM), an inverse problem that turns an instrumented bridge into a traffic-speed scale. This report builds the forward physics (an Euler–Bernoulli beam under a moving, coupled vehicle load), inverts it for the axle loads, and treats the result as an inverse problem: verification of every forward block against an analytical reference, regularization of the ill-conditioned tandem-axle split, an uncertainty estimate that is shown to be badly over-confident on a dynamic bridge and is then re-calibrated from physics, and validation against real measured data from the KW51 bridge. The core is numpy-only.

1 Motivation

Road and bridge damage grows roughly with the fourth power of axle load (AASHO Road Test): a truck at twice the legal axle weight does about 16× the structural damage. Static weigh stations are expensive and easy to evade. B-WIM instead weighs trucks at traffic speed by inverting the bending response of a bridge that is already instrumented. The same dynamic signal also carries information about the structure's health (drive-by, vibration-based monitoring). The engineering questions are therefore those of an inverse problem: how well-posed is the recovery, how much can a recovered number be trusted, and does the method survive contact with real data.

2 Forward model

Bridge. A simply-supported Euler–Bernoulli beam is discretised with 2-node Hermitian beam finite elements, two degrees of freedom per node (transverse deflection w, rotation θ). Element consistent mass $\mathbf{M}_e$ and bending stiffness $\mathbf{K}_e$ assemble into

$$\mathbf{M}\,\ddot{\mathbf{u}} + \mathbf{C}\,\dot{\mathbf{u}} + \mathbf{K}\,\mathbf{u} = \mathbf{F}(t), \tag{1}$$

with optional Rayleigh damping $\mathbf{C} = \alpha\mathbf{M}+\beta\mathbf{K}$. The reference structure is a 40 m highway overpass with $f_1 \approx 2.6\,\mathrm{Hz}$ and $\zeta = 2\,\%$.

Moving load. A wheel at position $x_p(t)$ in element e is mapped to nodal forces through the cubic Hermite shape functions $\mathbf{N}(\xi)$, $\mathbf{F}_e(t) = \mathbf{N}(\xi(x_p,t))^\top P(t)$, so $\mathbf{F}(t)$ is rebuilt every step as the load travels.

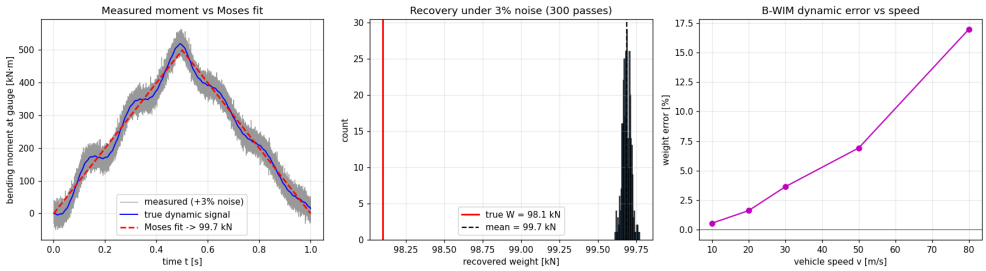

Figure 1: B-WIM round trip. A known truck is driven across the simulated bridge, the bending response is measured, and Eq. (2) recovers the axle loads. The recovered gross weight tracks the truth to sub-percent in the well-separated-axle case, which is the baseline before dynamics and tandem ill-conditioning enter.

Vehicle. The wheel force is not constant. The truck is a 2-DOF quarter-car (sprung mass on suspension k_s, c_s; unsprung mass on tyre k_t) coupled to the beam at the contact point, so the road profile and the bridge deflection feed back into $P(t)$.

Time integration. Two integrators are implemented and cross-checked: explicit RK4 on the state-space form, and implicit Newmark-β ($\beta = \frac{1}{4}, \gamma = \frac{1}{2}$, unconditionally stable). A modal model-order reduction keeping the first r modes provides a third, cheaper forward path.

Inverse model. With the section influence line $\mathbf{C}$ assembled from the beam's response to a unit load at each sampled position, the measured bending-moment history $\mathbf{m}$ is fit in least squares for the static axle loads $\mathbf{P}$ (Moses' method),

$$\hat{\mathbf{P}} = (\mathbf{C}^\top \mathbf{C})^{-1} \mathbf{C}^\top \mathbf{m}. \tag{2}$$

Equation (2) is extended below to a Tikhonov-regularized and a Bayesian estimator.

3 Verification

Each forward block is checked against an independent analytical reference (Table 1). These are agreements with closed-form solutions to many significant figures for the assembly, boundary conditions, and integration.

Table 1: Verification of the forward model against analytical references.

Check	Reference	Result
Static mid-span deflection	$PL^3/48EI$	rel. err $\approx 1 \times 10^{-13}$
Natural frequencies	$(n\pi/L)^2\sqrt{EI/\bar{m}}$	mode 1 rel. err $\approx 4 \times 10^{-7}$
Moving constant force	Frýba closed form	peak rel. err $\approx 6 \times 10^{-6}$
Quarter-car frequencies	analytic 2×2 eigenproblem	within FFT resolution
Damping ζ	free-vibration log-decrement	within 0.1 %
Newmark-β vs Frýba	closed-form moving force	match to $\approx 2 \times 10^{-4}$
MOR (3 modes) vs full order	full-order response	reproduces to 0.16 %

4 Validation and uncertainty

Verification is not validation. The four results below address whether the recovered answer can be trusted.

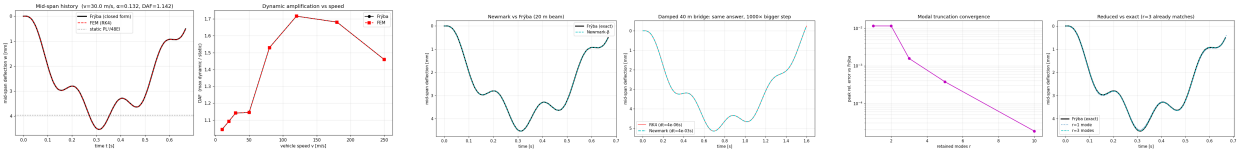

(a) Moving force vs Frýba **(b)** Newmark-β vs closed form **(c)** 3-mode MOR vs full order

Figure 2: Representative verification figures. The FEM mid-span deflection (markers) lies on the analytical curve (line) in each case, and the implicit and reduced-order integrators reproduce the explicit full-order response.

4.1 Model-error-aware uncertainty

Casting Eq. (2) as a linear-Gaussian inverse problem gives a posterior over the weight: the mean is the Moses point estimate, the covariance is $\sigma^2(\mathbf{C}^\top\mathbf{C})^{-1}$ with σ from the fit residual, and from these follow a 95 % credible interval and a probability of exceeding the legal limit.

A self-consistent coverage test, in which the synthetic data and the inverse both use the static influence line, returns 95.3 % coverage (Fig. 3, right). This is reassuring but only verifies the linear-Gaussian machinery; it is circular. A stronger test replaces the forward model with the full coupled dynamic FEM. The recovered weight is then biased by dynamic amplification (about +0.3 % quasi-static, rising to +5 % at 40 m/s), and because the residual σ measures goodness of fit (the signal still looks influence-line-shaped) rather than the dynamic scale error, the nominal 95 % interval's true coverage collapses to nearly zero. The RMS dynamic bias over the highway speed band, about 3.1 %, is the model-error term that restores calibration. Inflating the posterior by this term returns the coverage to nominal.

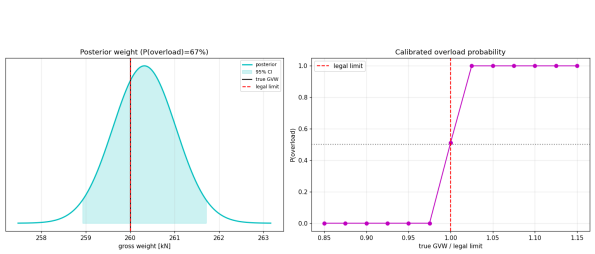

(a) Bayesian posterior with 95 % credible interval and P(overload).

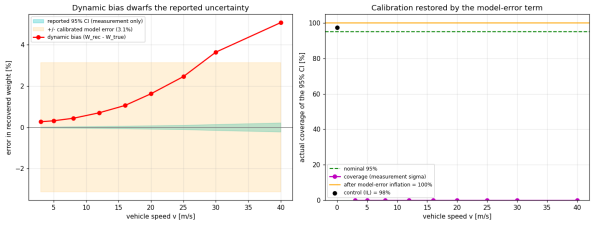

(b) Self-consistent 95.3 % coverage collapses under the dynamic forward model; a 3.1 % physics-derived term restores it.

Figure 3: Model-error-aware uncertainty. The result is the gap between the circular test and the honest one, and that the gap closes with a physics-derived model-error term.

4.2 Regularizing the ill-posed axle split

Moses' method recovers gross weight to better than 1 % but splits closely spaced (tandem) axles poorly. The cause is linear-algebraic: a tandem makes two columns of $\mathbf{C}$ nearly parallel, so $\mathrm{cond}(\mathbf{C}^\top\mathbf{C})$ rises from about 9 at 8 m spacing to about 2000 at 0.5 m, and the split amplifies noise. Tikhonov regularization,

$$\hat{\mathbf{P}}_\lambda = (\mathbf{C}^\top\mathbf{C} + \lambda^2\mathbf{I})^{-1}\mathbf{C}^\top\mathbf{m}, \tag{3}$$

with λ chosen at the L-curve corner ($\lambda^\star \approx 7.9$), cuts the per-axle scatter by about 6× (from 16 % to 2.6 %) for a small bias, while leaving the well-posed gross weight unchanged (Fig. 4). This is the bias–variance trade-off of an ill-posed inverse.

4.3 Validation against real data (KW51)

The 16-month tracked-modal-frequency record of the instrumented KW51 railway bridge (Leuven; public Zenodo dataset) is independent ground truth. It confirms the drive-by-SHM premise: a real,

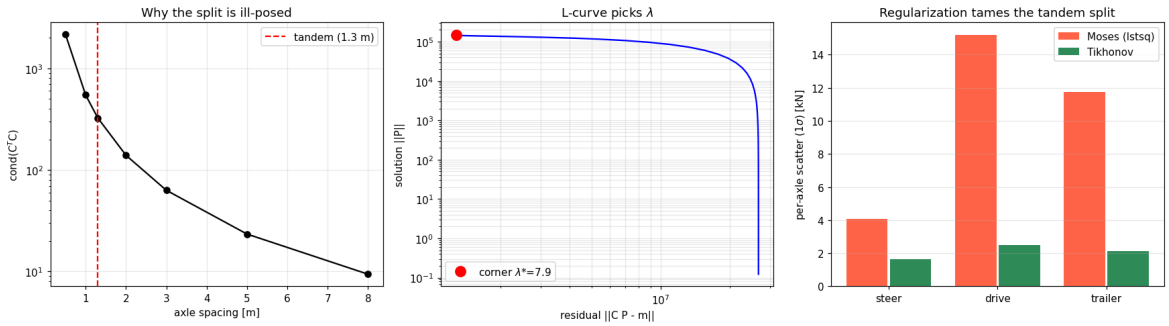

Figure 4: Ill-conditioning and its cure. Left: cond($\mathbf{C}^\top\mathbf{C}$) versus axle spacing. Middle: the L-curve corner selects $\lambda^\star$. Right: regularization cuts the tandem-split scatter by about 6× while the gross weight is untouched.

dated retrofit (stiffening) shifts the modal frequencies by up to +2.6 %. It also exposes the dominant confounder: temperature swings the same frequencies by a comparable amount (up to about 3.8 %), and for some modes the temperature swing exceeds the retrofit jump. Removing the temperature trend (fit on the healthy period) sharpens the change from a raw to a temperature-corrected residual, for example from $9.8\,\sigma$ to $17.3\,\sigma$ for one mode. This is the limit the synthetic damage study predicted, now confirmed in the field: a single global frequency is a blunt, environment-contaminated indicator that detects but cannot localize.

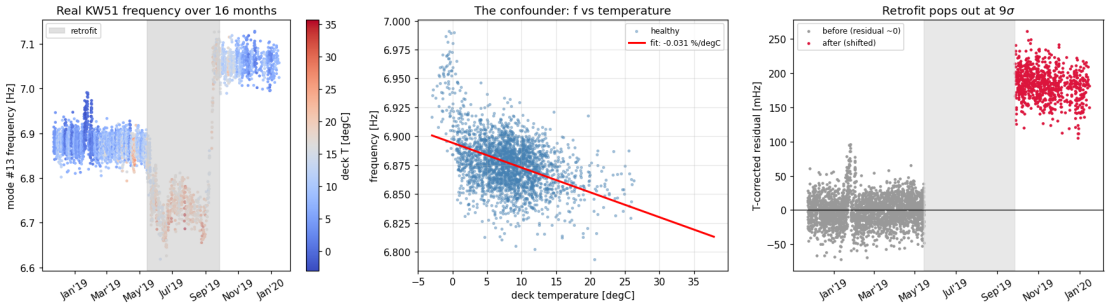

Figure 5: Validation on real KW51 data. The retrofit step is visible in the tracked modal frequencies, but temperature contaminates the signal at a comparable amplitude; a temperature correction is required before a stiffness change is detectable.

4.4 From weight to decision (fatigue economics)

A recovered weight is only useful if it drives a decision. Steel-detail fatigue follows an S–N law with slope m (Eurocode EN 1993-1-9, $m = 3$; AASHTO pavement, $m = 4$), so damage per pass scales as $(W/W_{\text{ref}})^m$: a 2× overload does 8 to 16× the damage. Over a simulated traffic stream the damage concentrates sharply, with the heaviest 10 % of trucks causing about 45 % of the fatigue damage and the overloaded minority about half, and B-WIM identifies exactly those trucks. An asset-depreciation estimate puts one heavy crossing at about €10 of bridge life against €0.80 for a legal pass, giving a per-crossing cost estimate for overload enforcement.

5 Limitations

- One simply-supported single span; linear, small-displacement Euler–Bernoulli theory.
- A single quarter-car (one wheel path); the multi-axle B-WIM uses constant axle forces.
- The drive-by study detects a global stiffness change but does not localize it, since a single fundamental frequency cannot by construction.

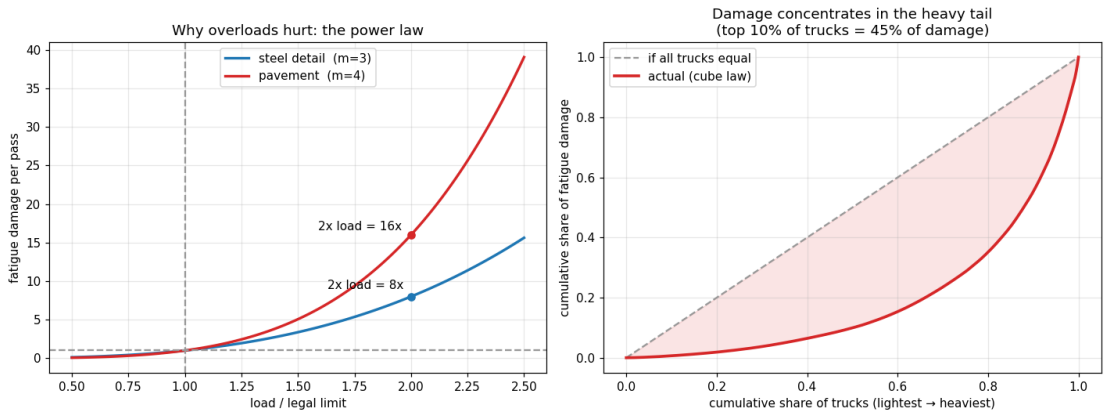

Figure 6: From a recovered weight to a decision. The fourth-power law concentrates fatigue damage in a few heavy trucks (Pareto), which B-WIM identifies.

- KW51 is a 115 m steel bowstring arch, not a simply-supported beam, so it validates the SHM behaviour and method rather than the FEM geometry.
- B-WIM here uses the analytical influence line; field systems calibrate it from a known truck and lose accuracy to road roughness, multiple vehicles, and temperature.

6 Reproducibility

The solver is numpy-only. Each result is produced by a `verify_*.py` script that prints the numbers and saves the figure shown here: `verify_v2.py` for Frýba, `verify_newmark.py`, `verify_mor.py`, `verify_coverage.py`, `verify_regularize.py`, `verify_kw51.py`, and `verify_fatigue.py`. An interactive web demo accompanies the report.

References

[1] AASHO Road Test (1962); fourth-power load-equivalency law (80 kN ESAL).

[2] L. Frýba, *Vibration of Solids and Structures under Moving Loads*, Noordhoff, 1972.

[3] F. Moses, "Weigh-in-motion system using instrumented bridges," *J. Transp. Eng.*, 105(3), 1979.

[4] E. OBrien et al., "A regularised solution to the bridge weigh-in-motion equations," *Int. J. Heavy Vehicle Systems*, 2009.

[5] P. C. Hansen, "The L-curve and its use in the numerical treatment of inverse problems," 1999.

[6] K. Maes & G. Lombaert, "Monitoring data for railway bridge KW51," Zenodo 3745914 (2020).

[7] COST 323, *European Specification on Weigh-in-Motion of Road Vehicles.*

[8] Eurocode EN 1993-1-9, *Fatigue strength of steel structures.*