Print copy. This is the print copy of an article whose version of record is a web page. That page, with the data and code behind every figure, is at
https://hertzforge.com/journal/2026-004/v1/
DOI of this version: 10.5281/zenodo.24254182

NEGATIVE RESULT

Multi-Shaker Shock Control: Ideas That Did Not Work in Simulation, and Why

Kevin Cross ORCID 0009-0002-6608-9426

Evidence	**Simulator.** Simulations only: eight simulated four-shaker structures the methods were developed on, nine held-out four-shaker structures, most with modeled bolted joints, and simulated frames with more controls than shakers; up to twenty random seeds and five levels of measurement noise. No hardware tests.
Version	Version 1, published 2026-11-15. First published 2026-11-15. Change note: First published.
Latest version	10.5281/zenodo.24254181
Cite as	Cross, K. (2026). Multi-Shaker Shock Control: Ideas That Did Not Work in Simulation, and Why. Hertzforge Journal 1, 2026-004, version 1. https://doi.org/10.5281/zenodo.24254182

Abstract

A shock test with several shakers is controlled from a model of the test article measured at a low level. Bolted joints soften as the level rises, so a drive computed from that model can over-test the article on the way to full level. This article reports, with numbers, the control ideas tried in simulation against that problem that did not work. A neural network, as the law's model or as the setter of how much softening to allow for, never beat a plain law built on a family of models, and where its prediction was low it put the article 5 to 8 dB above the full-level target. A modal model that changes inside each pulse passed seven of the eight structures it was developed on and failed held-out structures with modeled bolted joints that the plain law passes. A corrector that steps all controls together was not the main reason its law passed, and it fired full-level pulses 6 to 9 dB over target between the pulses that a judgment of the first, third and last full-level pulses looks at. A step rule did not replace the limit on softening, and two cures for failures under heavy measurement noise were worse, or only partly better, on the strict criteria. Twelve further ideas are listed with their evidence. The failures share causes: low-level pulses do not show the softening that matters, a law that fires a predicted drive is only as safe as its lowest prediction, and judging some pulses hides over-tests. Simulation only; nothing here has been run on shakers.

Contents

Summary 3
1. The Problem 3
2. The Test Model 4

2.1 The Structures 4
2.2 The Test Sequence 6
2.3 How Hard the Structures Are 7
3. How the Ideas Were Judged 8
4. The Laws Every Idea Was Compared With 9
5. A Neural Network in the Control Law 9
6. A Physics Model That Changes Inside the Pulse 12
7. A Coupled Corrector That Was Not the Main Reason Its Law Passed 12
8. A Step Rule in Place of the Softening Ceiling 14
9. Two Cures for Failures Under Heavy Noise 15
10. Other Ideas, With Weaker or Narrower Evidence 16
11. What the Failures Have in Common 18
12. Limits 18
How This Was Made 19
Acknowledgments 19
References 19
Version History 21

Summary

A shock test is judged on a shock response spectrum, the peak response of a row of simple oscillators to the acceleration measured at a point. With several shakers and several control accelerometers on a flexible test article, each control has its own target spectrum, and the controller works out the shaker drive voltages from a model of the article measured with a low level of random noise before the test. The test then climbs a ladder of levels to full level. Bolted joints soften as the level rises: the article's resonances move down and its damping goes up. A drive computed from the low-level model can then over-test the article on the way up.

This article reports the ideas for dealing with that which did not work in simulation, so that others need not repeat them. The laws that did work are described in a companion article in preparation; this one describes them only as far as needed to compare.

What did not work, and how far the evidence goes:

- **A neural network as the control law.** It never beat a plain law built on a family of models, and it over-tested where its prediction was low. As a predictor of the next level's over-test, given four low pulses, it was modestly better than the simple extrapolation it was compared with.
- **A model of the article that changes inside each pulse,** fitted from the test's own pulses. It passed seven of the eight structures it was built on and failed structures with modeled bolted joints that the plain law passes. It was also far too slow.
- **A corrector that adjusts every control at once.** The law it was part of passed mostly for other reasons, and the corrector fired full-level pulses 6 to 9 dB over target between the pulses that a judgment of the first, third and last full-level pulses looks at.
- **A rule that sets the size of the next level step from the last prediction error,** in place of an operator's limit on how much the article can soften. It was too bold or too timid.
- **Two cures for failures under heavy measurement noise.** Averaging the probe pulses made things worse. A slower climb cured the realistic structures and failed two stress-tier ones.
- **Twelve more ideas,** most of them tested more weakly, in one table.

What they have in common. Low-level pulses do not show the softening that matters. A law that fires a predicted drive is only as safe as its lowest prediction. Judging only some full-level pulses hides over-tests. A model built with the simulator's own joint law in it proves little about real joints. Results from one or five random seeds can mislead.

The evidence is simulation only, on four-shaker structures and a few frames. The strongest results come from structures held back from each method's design, run from each method's saved setting independently of its development. The simulated jointed structures the methods were developed on soften 3 to 15 %: the moderate one as much as well-tightened bolted joints have been measured to soften (about 4 % or less), the others as much as loose or friction joints under shaker drive or more, most with no rise in damping. Results are given separately for structures that soften as little as well-tightened joints do.

1. The Problem

A **shock response spectrum (SRS)** describes a shock by its effect rather than its waveform: it is the largest response, at each frequency, of a single-mode oscillator driven by the measured acceleration. Here the oscillators have a Q of 10, and the SRS is read at frequency points spaced in fractions of an octave. Two quite different waveforms can have the same SRS, and a shaker shock test is accepted when the measured SRS is within a tolerance, here 3 dB either way, of the target.

With one shaker, the controller measures the frequency response from drive to control, synthesizes a waveform that should meet the target, fires it at a reduced level, and corrects the drive by the SRS error as it steps up.

Published single-shaker practice corrects the magnitude of the estimated frequency response and keeps the correction factor below one because the article is nonlinear (Coleman, Smallwood and Skousen 2016).

With several shakers and several control accelerometers on the article itself (multiple-input multiple-output, MIMO), each control has its own SRS target and the timing between controls is left free. The controller uses a **model**: the matrix of frequency responses from every shaker to every control, measured before the test with low-level random noise on all shakers at once. The drive is computed from that model. The trouble is that a bolted structure does not keep that model. As the level rises, its joints slip a little more, its resonances move down and its damping rises. A drive computed for the resonances the model knows meets resonances that have moved. On one of the simulated structures used here, one that softens more than a well-tightened bolted joint (Section 2.3), a drive computed from the model and fired at −12 dB (a quarter of full level) put one control 15.4 dB over its target, at the first resonance, which had moved down 8 %.

In a search of OSTI, Crossref, arXiv, NTRS, Google Patents and the open web, with partial Semantic Scholar coverage (2026-11-11 to 15), we found no published multi-shaker SRS control law that follows a structure whose response changes with level, and no network used in any shock control loop. One paper on closed-loop multi-shaker SRS control (Zheng, Ma, Wei and Chen 2023) could be read only as an abstract. This article reports what we tried that did not work.

2. The Test Model

2.1 The Structures

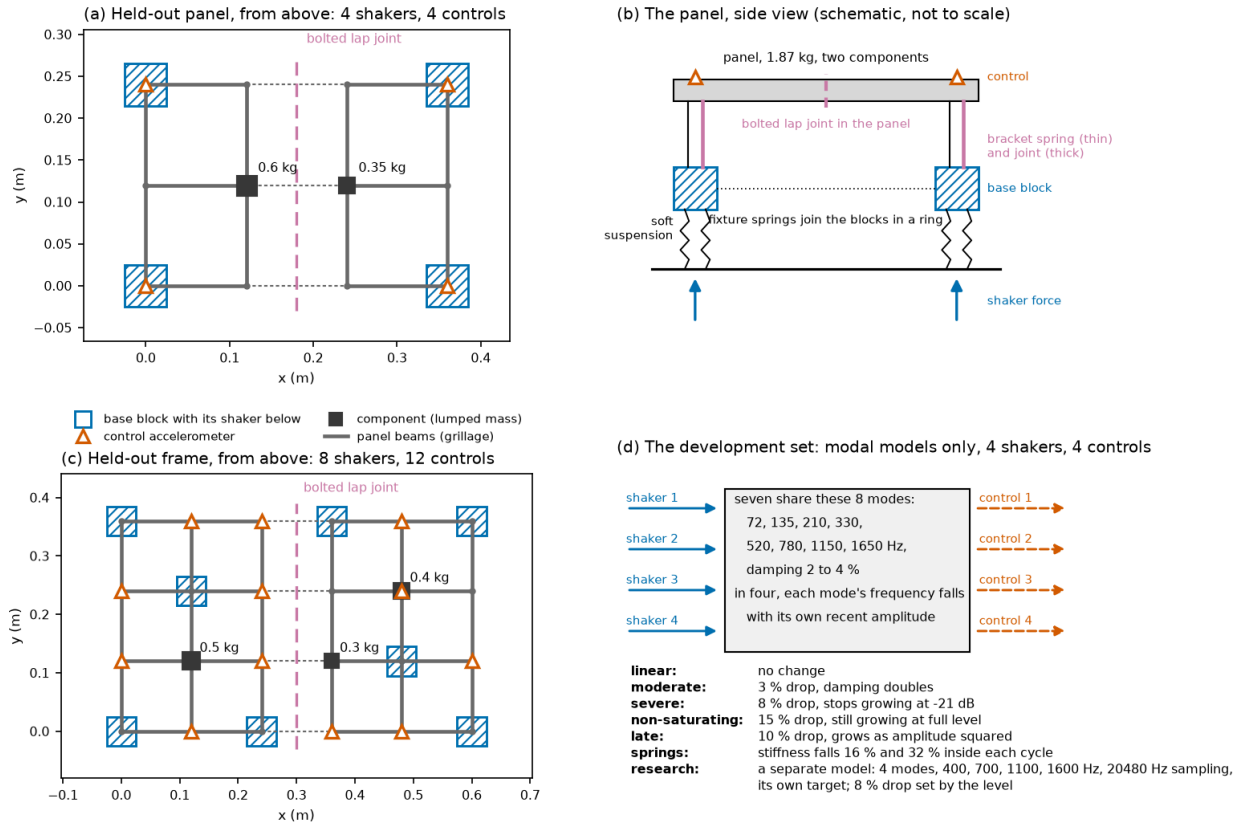


Figure 1. The simulated structures. (a) to (c) are drawn from the held-out structures' own definitions. (b) is a schematic and not to scale. (d) The development structures exist only as modal models, so they are drawn as

a block diagram; the box describes the seven that share one model, and the research model is listed below it. Data: `data/fig1_geometry.csv`, `data/fig1_development_modes.csv`, `data/fig1_development_structures.csv`; code: `code/fig1_setup.py`.

Two sets of simulated structures were used (Figure 1).

The development set is the eight structures the methods were built and tuned on. All are modal models of a setup with four shakers and four control accelerometers, with no geometry. Seven of them share one model: eight modes from 72 to 1650 Hz with 2 to 4 % damping, sampled at 12,800 Hz in records of 16,384 samples (1.28 s). In four of these, each mode's frequency falls with that mode's own recent amplitude, which is how this simulator represents a jointed structure:

- **moderate:** the frequency falls by up to 3 % and the damping doubles;
- **severe:** up to 8 %, with no damping rise; the drop stops growing once the mode reaches a set amplitude, which on the first mode happens near −21 dB;
- **non-saturating:** up to 15 %, still growing at full level;
- **late:** up to 10 %, growing as the square of the amplitude, so almost nothing shows below −12 dB.

Two **springs** soften another way: the stiffness falls with displacement inside every cycle, by up to 16 % and 32 % (about 8 % and 16 % in frequency). The seventh is **linear**. The target of these seven is a "twist": a flat 50 g SRS from 100 Hz to 2 kHz, read in third octaves from 50 Hz, with two of the four controls inverted, so that the article must twist rather than move as one. The eighth, the **research** model, is separate: four modes from 400 to 1600 Hz with 2 % damping, sampled at 20,480 Hz (records of 0.8 s), whose frequencies drop by up to 8 % with the level of the pulse; it has its own target, read in twelfth octaves from 100 Hz, and a fixed identification.

The held-out set was built, mostly in physical coordinates, and kept back from the design of the methods of Sections 5 to 7 and of the probe check in Table 1: their developers did not see these structures or any result on them until each method's setting was written down, and each method was run on them from its saved setting, independently of its development. The hold rule of Section 4 is different; see there. The main structure (Figure 1a, b) is a stiff ribbed panel, 0.36 m by 0.24 m and 1.87 kg, modeled as a grid of beams with two components on it. It is carried at its four corners by four base blocks; each block sits on a soft suspension and is driven by its own shaker, and fixture springs join the blocks in a ring. Each corner support is a bracket spring beside a bolted joint, and a bolted lap joint crosses the panel. The controls are the panel's four corners. The nine modes in the test band lie between 138 and 1873 Hz, with 0.3 to 0.7 % damping.

The joints are Iwan joints (Segalman 2002, in the form written by Porter, Balaji, Little and Brake 2022): many small springs in parallel, each with its own slip force. As the load rises, more of them slip, so the joint gets softer and dissipates more energy, and it does so with memory, as real bolted joints do. Variants of the panel:

- **realistic joints:** slip forces set so the full-level pulse loads each joint to 0.6 to 0.85 of the force at which it slips completely; the modes drop by 0.1 to 1.7 % and their damping rises 1.1 to 2.5 times;
- **onset:** as realistic, with one support joint slipping completely at full level;
- **low preload:** every joint slipping completely from about −30 dB, as a badly loosened joint would;
- **scatter:** low preload, with every joint's slip force redrawn by 8 % from pulse to pulse and a slow loss of preload;
- **hardening:** a contact in every joint that closes and stiffens it;
- **softening then hardening:** realistic joints plus the closing contacts;
- **with shakers:** the realistic panel driven through modeled shakers (coil, armature and flexure) and a voltage clip; and the development set's severe structure behind a shaker and a clip.

Frames (Figure 1c): a 0.6 m by 0.36 m panel on eight driven blocks with twelve controls, linear or with realistic or low-preload joints, under the twist target or under a target derived from the frame's own response so that it can be met.

2.2 The Test Sequence

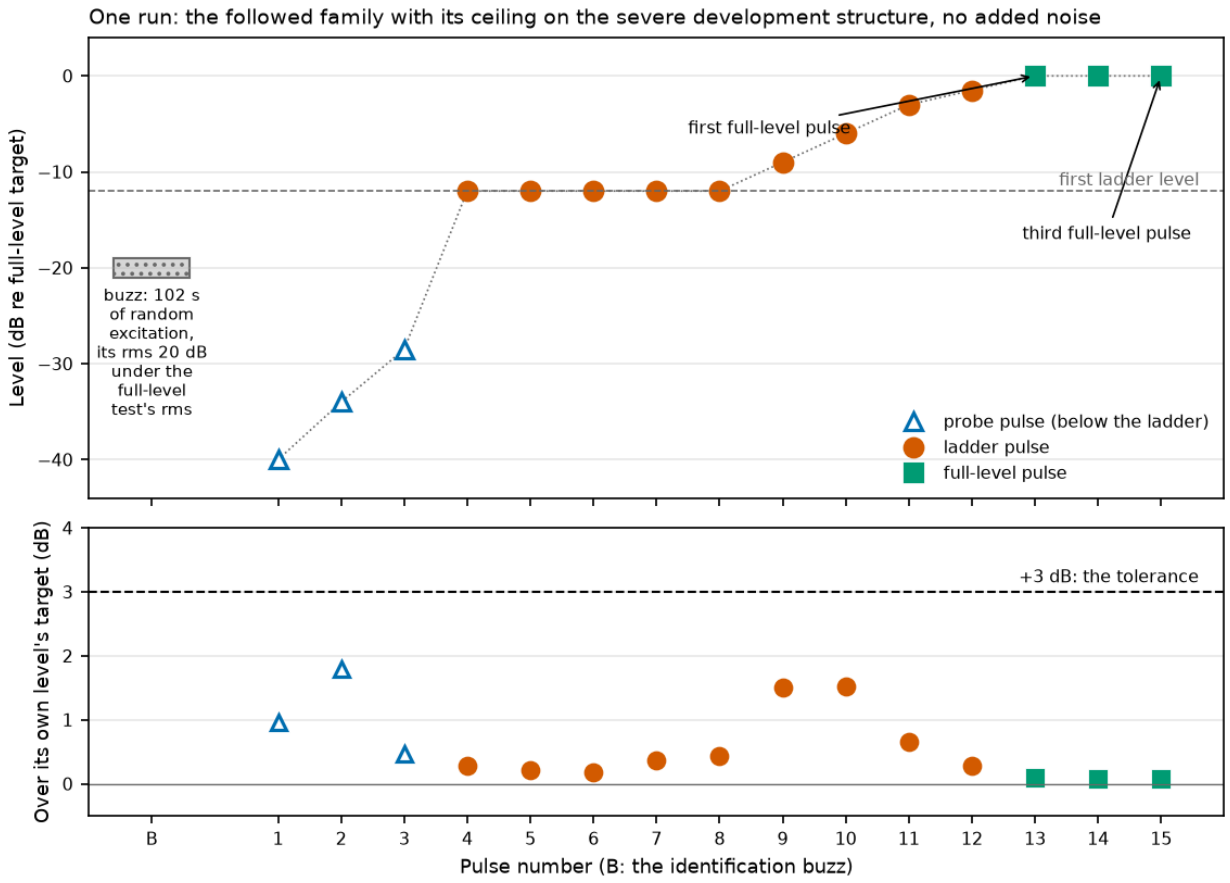


Figure 2. The test sequence of one run: the followed family with its 12 % ceiling (Section 4) on the severe development structure, no added noise. The level of the first ladder pulses repeats until every SRS point of a pulse is within 3 dB; the early −12 dB pulses were under target at some points. Data: `data/fig2_test_sequence.csv`; code: `code/fig2_test_sequence.py`.

Every run follows the same sequence (Figure 2).

1. **The identification buzz.** Independent random voltages on every shaker at once, 5 Hz to 2.5 kHz, at a level that puts the controls 20 dB below the full-level test's rms. For four shakers it lasts 102 s (40 averages of two 1.28 s records). The model is estimated from it.
2. **Probe pulses.** The law's own drive fired well below the ladder, at −40 and −34 dB re the full-level target, to see how the article responds before the real climb. The ladder may insert a level between them and the first ladder level when the step is large (Figure 2, pulse 3), and some laws add more.
3. **The ladder.** Pulses from −12 dB (the first ladder level) up through −9, −6, −3 and −1.5 dB to full level. A level is left only when every SRS point of a pulse is within 3 dB; at most six pulses are allowed at a level, and a step that looks too large may be split. Some development structures run a fixed ladder with one pulse a level.

4. **Full-level pulses**, at least three.

Two words are used throughout. The **over-test** of a pulse ("over") is the largest amount, in decibels, by which any control's SRS at any frequency point exceeds the target of the level the pulse was fired at. Its **absolute level** is the same measured against the full-level target: a pulse at −12 dB that is 14 dB over its own target puts a control 2 dB above the full-level target.

Measurement noise is added to every measured record, the same size at every level, as a fraction of the rms of the full-level target waveform: 0, 1e-4, 1e-3, 3e-3 and 1e-2. At 1e-2 a probe at −40 dB is as small as the noise.

2.3 How Hard the Structures Are

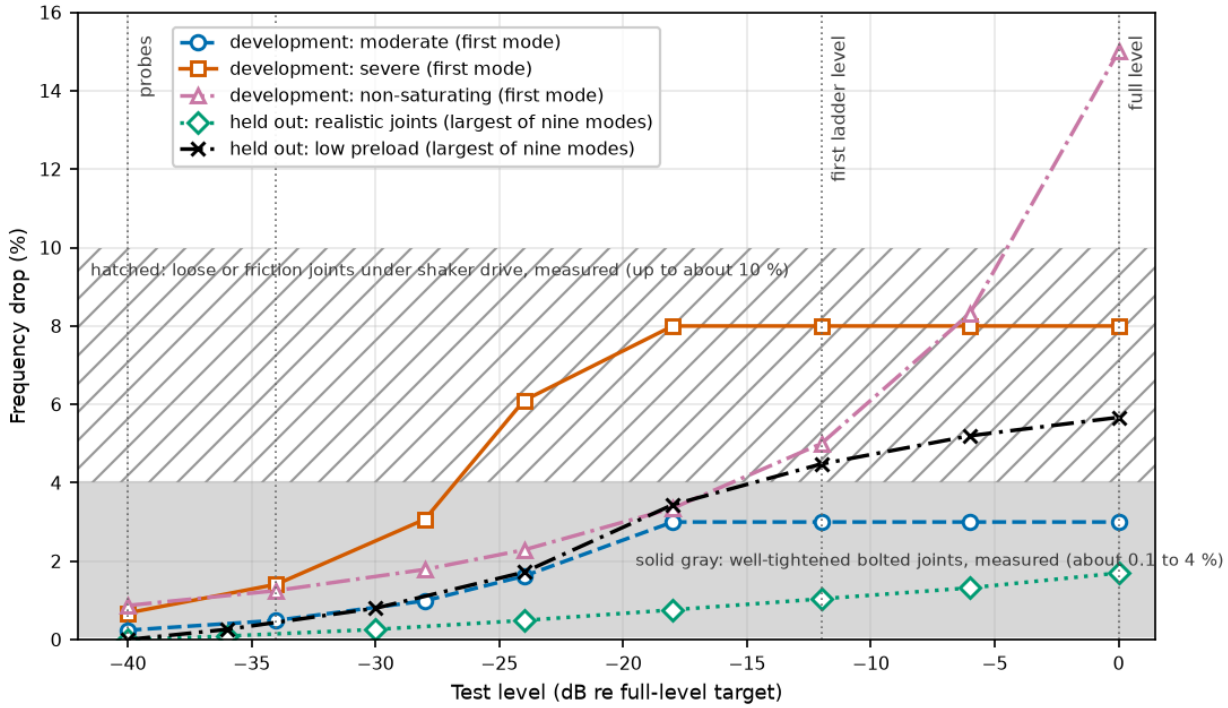


Figure 3. How far the structures soften at each test level. Development structures: the first mode's drop at the peak of a pulse at that level. Held-out panel: the largest drop of its nine modes, from the free decay of a burst of that mode at the amplitude the pulse reaches, relative to −40 dB. The bands are the ranges of total drops measured on bolted and friction joints below full slip (Section 2.3), not curves against level. Data: `data/fig3_softening_by_level.csv`; code: `code/fig3_softening_by_level.py`.

How much do real joints soften? Well-tightened bolted joints, struck with a hammer, lose about 0.1 to 4 % of a mode's frequency below full slip (Roettgen 2016; Kuether and Brake 2016; Wall, Allen and Kuether 2022). Looser joints and friction interfaces lose more: about 6 % on a lightly torqued two-beam structure (Deaner and co-authors 2015), about 8 to 10 % on a metal slug press-fitted in foam cups under shaker sine (Kuether and Najera 2017, from the measurements of Jacobs-O'Malley and Hofer 2017), about 10 % on a single bolt under shaker random (Khan and co-authors 2021, from presentation slides), and about 28 % on a beam in full slip (Bonney and co-authors 2016). Damping usually rises as the frequency falls, in one tabulated case elevenfold (Roettgen 2016), and one structure's first mode hardened by about 2 % (Wall, Allen and Kuether 2022). How a joint is driven matters, and few measurements use shock. On the press-fit interface, shaker shock pulses of 8.5 to 35 g left the frequency unchanged, while sine sweeps of 1 to 10 g lowered it 8 to 10 %, though the shock tests

used different foam cups and some preload was lost during the sine tests (Jacobs-O'Malley and Hofer 2017). At pyroshock level, a resonant plate with loosely bolted damping bars moved 1.6 to 12.7 % down on five modes and 20 % up on one between a hammer test and shock (Schoenherr, Soine and Witt 2022).

The development set's moderate structure softens 3 %, as much as a well-tightened bolted joint; its other jointed models and its springs soften 8 to 16 %, as much as the loose and friction joints or more (Figure 3), most with no damping rise. In my experience softening has generally shown up only in the last 6 dB, with a caution: that was on very high-force shocks that simulate pyrotechnic events such as separation or ignition, and the small modal shakers on today's multi-shaker frames cannot deliver shocks of that force.

Results are therefore given in two tiers. The **realistic tier** is the linear and moderate development structures and three held-out panels: realistic joints, realistic joints with shakers, and softening then hardening. The **stress tier** is the rest: the severe, non-saturating, late, research and spring development structures, and the held-out onset, hardening, low-preload, scatter and severe-with-shaker panels. Figure 3 also shows the first lesson of this article: at the probe levels the jointed development structures have moved 0.2 to 1.4 %, and they move 3 to 15 % by full level.

3. How the Ideas Were Judged

The strict criteria. A run passes only if all seven hold:

1. every probe and the first ladder pulse is at most 3 dB over its own level's target, and no identification pulse puts a control more than 3 dB over the first ladder level's target;
2. no pulse below full level is more than 3 dB over the full-level target;
3. the first full-level pulse is at most 3 dB over;
4. the third full-level pulse is within 3 dB everywhere;
5. the run completes, and its last pulse is within 3 dB everywhere;
6. no pulse of any kind puts a control more than 3 dB over the full-level target;
7. no full-level pulse is more than 3 dB over.

They are strict and binary: a probe 0.2 dB past its own level's limit fails a run in which the article never comes near the full-level target. An earlier article in this journal (Cross and Zwink 2026) proposes judging a controller by whether it is safe and whether it reaches what the setup can reach. Where it changes the verdict, results are therefore also given under two looser readings:

- **reading A:** the run ends at full level with every control inside 3 dB;
- **reading B:** nothing, at any level, puts any control more than 3 dB above the full-level target.

What a method had to be. One setting for every structure, with no number the operator enters about the structure; no harm on a linear structure; and a computation between pulses of at most 2 s at 16 controls, 8 shakers and 16,384 samples, on one processor core.

How strong the evidence is. The work ran in rounds. In the last round every method's setting was written down before it met a held-out structure, every method ran beside the baselines of Section 4 on the same structures, seeds and noise, and each result reported here as strong was rerun independently from the saved setting (How This Was Made). Pass rates come from five or twenty random seeds where stated (a seed changes the noise in every pulse, and in the identification of every development structure but the research model, whose identification is fixed). Held-out results are one seed. Earlier rounds used one to three seeds, the developers' own models, and few or no held-out structures; Section 10 says so per idea.

4. The Laws Every Idea Was Compared With

Three laws serve as baselines here. The companion article describes them fully.

The frozen law. The drive is computed once from the model, to put each control on its target waveform. After each pulse, each control's target is scaled by a smooth gain per SRS point, from the ratio of target to measured SRS, and the drive is recomputed from the same model. The model is never updated. With one shaker this reduces to the single-shaker magnitude correction. A correction of this kind on the targets of several drives, with a model measured once, is published in a patent application (Zheng, Chen and Wei 2022). On the development set it passes the linear and late structures only.

The followed family with its ceiling. Instead of one model, the law uses a family: copies of the measured model with every resonance lowered by none, half and all of an assumed softening, and fires the drive that is least wrong on average over the copies. The assumed softening is read from the probes and from each level's pulses (the downward stretch of the measured model, per octave band, that best explains the pulse), carried up to the next level as a power of the response amplitude no larger than one, with a margin, and never more than a ceiling of 12 % that the operator sets. It passes seven of the eight development structures, all but the 32 % spring, where its −34 dB probe is 3.2 dB over.

The hold rule. The followed family with no ceiling, a check before each probe that lowers its drive when the model predicts it over, an allowance for the rounding of the fitted softening when its growth with level is extrapolated, and one rule: once a pulse at a level is inside the 3 dB tolerance, fire the same drive again. Over twenty seeds it passes 160, 159, 158 and 139 of 160 development runs at measurement noise 0, 1e-3, 3e-3 and 1e-2, against 140, 140, 125 and 118 for the followed family with its ceiling. Unlike the methods of Sections 5 to 7, the hold rule was written after held-out results of the law it comes from (Section 7) were known; its setting was written down before it was itself run on the held-out structures.

Two more baselines ran beside every method: the frozen law climbing from −40 dB in 3 dB steps, and a family with a fixed 8 % softening. They change no conclusion here.

5. A Neural Network in the Control Law

The idea. A network trained beforehand on many simulated structures, then adapted to the article from the test's own data, to tell the law how far the article will soften at the next level. A network that maps any target straight to a drive for any article was ruled out before anything was built: the article's own dynamics are not among its inputs, and nothing could show it safe on an article it was not trained on. The form that does make sense is a base network adapted on the fly to the buzz and the first low pulses.

Why it was reasonable. A network trained on a stream of simulated systems can identify a new system from a short record of its input and output without retraining (Forgione, Pura and Piga 2023). In feedforward control, networks have worked as corrections on a physics model, kept close to it where data is missing (Bolderman and co-authors 2024). Near this problem, a network has driven one shaker to reproduce a force history (Peláez-Rodríguez, Magdaleno, Salcedo-Sanz and Lorenzana 2023), and one turns an SRS into a time history, open loop (Watts, Jeon, Newton and Bowering 2026). We found none inside a multi-shaker or SRS control loop.

What was built. A small network that reads, band by band, what the buzz and up to four low pulses (at −40, −34, −28 and −22 dB) showed, and predicts how far each octave band of the measured model should be stretched down in frequency at the next level. It has 14,276 parameters and takes 20 microseconds per prediction on one core; training took 23 s and adaptation needs no weight update. It was trained on 700 simulated jointed four-shaker structures (3 to 8 modes from 60 to 1800 Hz, 0.8 to 4 % damping, frequency drops of 0 to 20 %); the springs, the research model and every held-out mechanism were never shown to it. It was tried in three roles: as a predictor; as the law's model (the drive is computed from the stretched model, with the stretch limited

to 0 to 40 %); and as the setter of the followed family's ceiling (its predicted 90th percentile of the next level's softening).

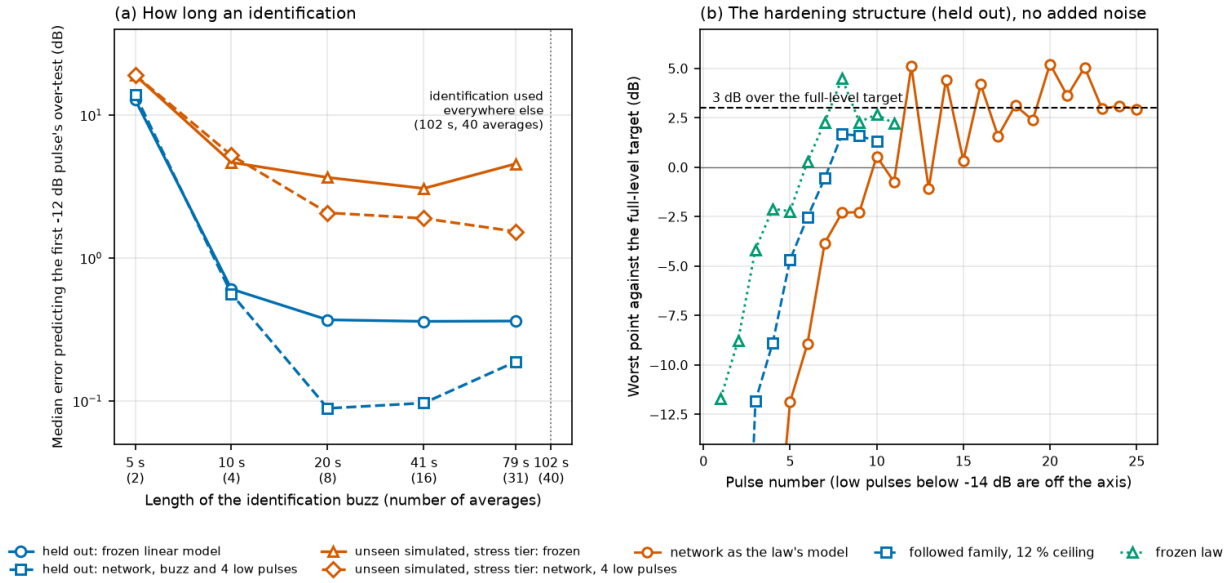


Figure 4. The network. (a) How long an identification it needs: the median error predicting the first −12 dB pulse's over-test, on the nine held-out four-shaker structures and on 41 unseen simulated stress-tier structures. (b) The network as the law's model on the held-out hardening panel, with the two baselines, no added noise; pulses below −14 dB are off the axis. Data: `data/fig4a_identification_length.csv`, `data/fig4b_hardening_pulses.csv`; code: `code/fig4_networks.py`.

What happened.

- **As a predictor.** From the buzz alone it did no better than the frozen linear model: median errors predicting the next level's over-test of 1.78 against 1.65 dB on the realistic tier and 2.34 against 2.47 dB on the stress tier, over 200 unseen simulated structures. The buzz carries the linear structure and nothing of the change with level. Given four low pulses it beat the simple power-law extrapolation of the same pulses in the median (1.40 against 1.68 dB realistic, 1.89 against 2.66 stress), and on the nine held-out structures in the tail as well: 95th percentile 2.3 against 16.4 dB over 55 predictions. On the realistic held-out tier every predictor, the frozen model included, was already within 1.6 dB at the 95th percentile.
- **How long an identification** (Figure 4a). Two averages (5 s) were unusable, with a 12.8 dB median error on the held-out structures; four (10 s) gave 0.5 to 0.8 dB; from eight (20 s) on the error was 0.1 to 0.4 dB with no further trend. The unseen simulated structures have larger errors but the same plateau.
- **As the law's model.** On the development set (one seed, no noise) it passed 4 of 8, against 7 for the followed family. Three of its four misses there were small (an adaptation pulse 1.3 dB past its limit; a first ladder pulse and a first full-level pulse 3.6 dB over, 0.6 dB past the limit), and every run ended at full level inside tolerance; the fourth was an adaptation pulse at −22 dB that put a control 2.5 dB above the full-level target on the 32 % spring. On the held-out structures (Figure 5) it passed no more realistic-tier runs than the followed family (4 of 6), and it failed to finish the softening-then-hardening panel, which every plain law finished; 2 of 10 stress-tier runs passed, against 5 for the followed family. On the hardening panel (Figure 4b) it fired pulses up to 5.1 dB above the full-level target before full level and full-level pulses up to 5.2 dB over; the followed family stayed within 1.7 dB. With a mode at 2,200 Hz (it was trained up to 1,800 Hz) its worst first pulse of a level was 25.6 dB over, against 9.8 dB for the followed family on the same twelve simulated structures; the

same law with a power-law extrapolation in place of the network did as badly (25.1 dB), so the fault there is firing one predicted model, not the network's training range.

- **As the ceiling.** It passed 4 of 8 development structures, against 7 for the same family with no ceiling at all. On the severe structure it narrowed the family and a pulse at −9 dB was 14.4 dB over, 5.4 dB above the full-level target; on the non-saturating structure 16.7 dB over, 7.7 dB above it. Held out, it did the same on the severe structure behind a shaker and clip (14.2 dB over at −9 dB, 5.2 dB above the full-level target), which the followed family passes. Its 90 % prediction covered only 71 % of unseen cases of its own mechanism. A three-parameter regression on the same inputs also did worse than no ceiling (5 of 8). The limit placed on the predicted softening bound for the network on the 32 % spring and for the regression on four structures, and halving it moved the severe structure's worst pulse by 12.5 dB: the number doing the work was a statement about the structure, chosen by the designer instead of the operator.

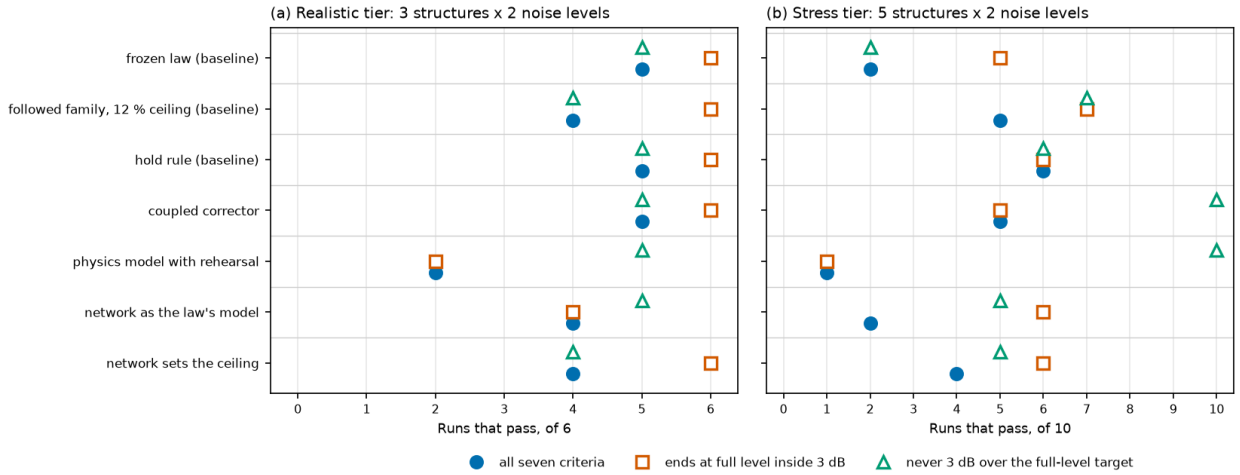


Figure 5. Runs that pass on the held-out four-shaker structures (one seed, measurement noise 0 and 1e-3), by law, tier and reading. The realistic tier is the realistic, realistic-with-shakers and softening-then-hardening panels; the stress tier is the onset, hardening, low-preload, scatter and severe-with-shaker panels. The linear panel is not shown: every law passes it at both noise levels except the physics model, which fails it at 1e-3. Data: `data/fig5_held_out_runs.csv`; code: `code/fig5_held_out.py`.

Why it did not work. A law that fires the drive computed from a predicted model is only as safe as its lowest prediction, and nothing bounds a prediction that is too low. On the hardening panel the stretch can only lower resonances, so the model was wrong in a direction it cannot represent. Speed was never the issue: the computation that matters is the model fit and the drive, not the network.

Under the readings. By reading B the failures that count are over-tests above the full-level target: 5.2 dB on the hardening panel, 5.2 to 7.7 dB in the ceiling role. Reading A, which looks only at the end of a run, does not see them: those runs end inside tolerance. By reading A the network as the law's model still finishes fewer realistic held-out runs than the plain laws (4 of 6 against 6 of 6, Figure 5a). On the development set, by reading A, it and the followed family both end in tolerance on all eight structures.

What survives. A pre-trained network fed four low pulses is a modestly better predictor of the next level than a power-law extrapolation, mostly on stress-tier structures. The identification-length result applies to any method: on these four-shaker models, about eight averages were enough.

6. A Physics Model That Changes Inside the Pulse

The idea. Fit, from the test's own pulses, a modal model in which each mode's frequency and damping are curves of that mode's own amplitude, so that the model is nonlinear inside the pulse. Before each pulse, **rehearse** the drive: simulate it through the model over a set of possible continuations of each curve beyond the largest amplitude seen, and adjust it until even the worst continuation stays inside 3 dB. Choose the size of the next level step from how far the last prediction missed. No ceiling, no operator number.

Why it was reasonable. Amplitude curves of this kind, read mode by mode from free decays after hammer impacts, describe the frequency of measured jointed structures well (Roettgen 2016; Kuether and Brake 2016). Two-step iterative learning control corrects a model and then optimizes the next input on it (Volckaert, Diehl and Swevers 2011). We found neither used to fit such curves from a controlled test's own pulses and compute the next drive.

What happened. On the development set it passed seven of eight structures with no ceiling, and stopped safely on the 32 % spring at −5.5 dB after 18 pulses; that held over five seeds. It passed the severe structure at noise 3e-3, where the followed family with or without its ceiling fails (first pulse 4.65 dB over), on both seeds tried. It also showed something useful about identification: a mode's frequency curve read from a pulse did not depend on how far the pulse was off target (within 1 percentage point between a drive and the same drive 4 dB low).

On the held-out structures it failed (Figure 5). Of six realistic-tier runs it passed 2, and of ten stress-tier runs 1. It failed two structures the frozen law passes: the realistic panel with shakers, where it stopped at −7.5 dB, and the onset panel without noise, where it never left −16 dB (with noise 1e-3 it passed that panel, its one stress-tier pass). On the linear held-out panel it needed 44 pulses where the frozen law needs 8, and at noise 1e-3 it failed that linear panel with a probe 3.75 dB over its own level. A computation between pulses took 2.4 to 9.8 s on average at four shakers on the development set and up to 110 s on the held-out structures, against a 2 s budget; at 16 controls and 8 shakers it was estimated at 15 to 25 s.

Why it did not work. The amplitude memory the model uses is the same running-envelope form as the simulator's own joint law, so the development set could not test the model's form. On the held-out structures, built in physical coordinates, the fitted model explained too little from the second pulse on, even on the linear panel; a safety gate then dropped the law to 1 dB steps, and it crawled or stopped. Why the gate tripped on a linear structure was not traced.

Under the readings. Its failures are mostly slowness and stops: by reading B it never put a control 3 dB above the full-level target on any stress-tier held-out run (10 of 10), and did so once on the realistic tier, by 0.2 dB. By reading A, which asks that the test finish, it reached full level in tolerance on 2 of 6 realistic and 1 of 10 stress-tier runs, against 6 and 6 for the hold rule. It does worse than the plain baselines under any reading that asks the test to finish, and it cannot meet the time budget.

What survives. The identification finding (one seed, on the development set, not rerun independently), and rehearsal as a way to choose a safe step: without rehearsal the same law never left its probes.

7. A Coupled Corrector That Was Not the Main Reason Its Law Passed

The idea. Between pulses at a level, the plain law adjusts each control's gains from that control's own SRS error. The corrector instead takes one Levenberg-Marquardt step (a damped Gauss-Newton step) on the SRS errors of every control and frequency point together. It uses the sensitivity of each SRS peak to each drive, taken at the oscillator's peak time through the model stretched to the level's softening, a trust region from pulse to pulse, and a dead band: no step when every point is within 1.5 dB.

Why it was reasonable. A family wider than the structure's softening leaves an error that per-control gains cannot remove because the controls are coupled: on an exact linear model the same too-wide family converged within 1 dB in seven pulses when the coupling was removed, and held one control 6.5 to 8 dB over for 30 pulses when it was not. A step that sees the coupling removed that stall: with a family two and three times too wide, the severe structure completed (third full-level pulse within 1.06 and 0.79 dB) where per-control gains did not converge. Published SRS iteration for one shaker and several responses uses a scalar sensitivity (Heitman, Babuska, Cap and Reid 2016), and published correction for several drives scales each target by its own SRS error (Zheng, Chen and Wei 2022); we found no published step that uses the sensitivity of each SRS peak to each of several drives.

What happened. As part of a law with no ceiling (with the probe check and the rounding bound of Section 4), it passed all eight development structures and 198 of 200 runs over five seeds and noise up to 1e-2 on the first six criteria (195 when every full-level pulse is judged). Three things undid it.

- **It was not the main reason.** The same law with the corrector's step removed, its dead band and replay kept, passed 8 of 8 without noise and 24 of 24 at noise 3e-3, with a better worst first pulse on the severe structure (1.48 against 2.76 dB over). Of 315 corrector steps, 260 measured a worse worst error than the per-control update would have from the same state, by a median of 0.53 dB and at most 3.02 dB. The step did matter in one place: without it the same law failed the 32 % spring under noise on two or three of five seeds (at 1e-3 and 1e-2), with full-level pulses up to 10.6 dB over, judged ones among them. Holding a pulse once it is inside 3 dB rather than 1.5 dB (the hold rule, Section 4) removed those failures without the step.
- **It fired over-tests that some judging does not see** (Figure 6). On the 32 % spring at noise 1e-3, three of five seeds fired full-level pulses 5.97, 7.36 and 8.87 dB over target, between the first, third and last full-level pulses. Judged on those three pulses alone, all five runs pass. The hold rule, which is this law with the corrector removed and any pulse inside 3 dB fired again, never exceeded 2.37 dB there, and the followed family with its ceiling 1.61 dB. This is why the strict criteria judge every full-level pulse (criterion 7).
- **Held out,** it lost the hardening panel at noise 1e-3 (it did not converge at −3 dB; the same law without the corrector passes), and on the softening-then-hardening panel at 1e-3 it fired a full-level pulse 3.7 dB over. Its computation at 16 controls and 8 shakers took 3.98 s on average and up to 7.18 s on one core; even with the sensitivity computation removed the largest call took 2.30 s.

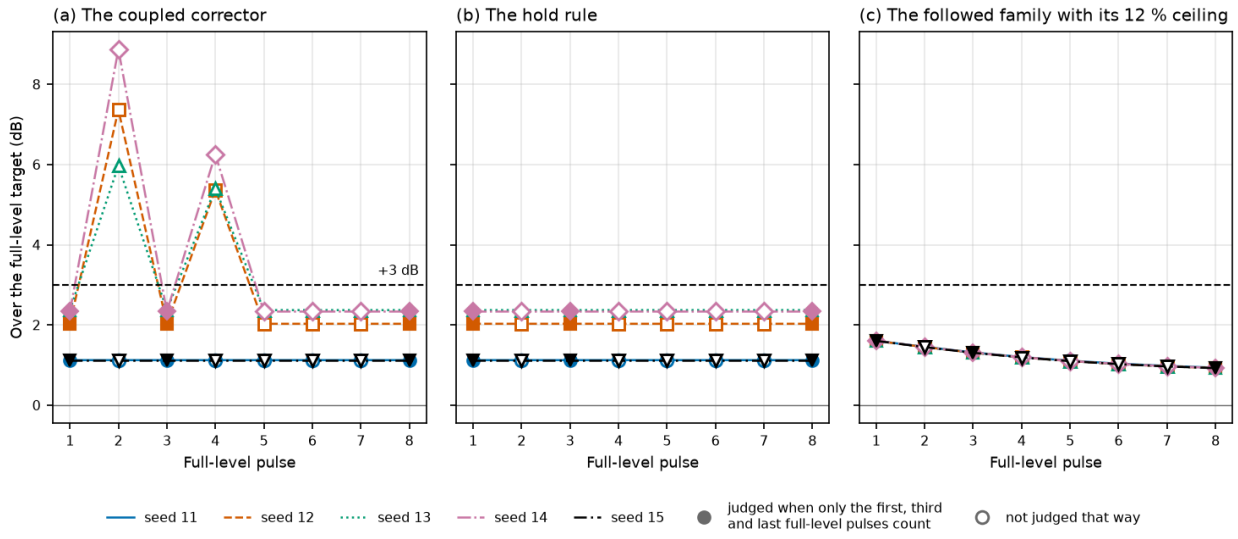


Figure 6. Every full-level pulse on the 32 % spring at measurement noise 1e-3, five seeds, three laws. Filled markers are the pulses a test judged on its first, third and last full-level pulses would look at. Data: `data/fig6_full_level_pulses.csv`; code: `code/fig6_unjudged_pulses.py`.

Why it did not work. A first-order step is only as good as the model's sensitivity, and near full level on the 32 % spring the structure responded to a change of drive nearly twice as strongly as the model predicts: in an earlier form of the corrector, a 0.25 dB step moved the worst point by 0.48 dB and a 2.1 dB step by 3.3 dB. And a step that reduces the sum of squared errors does not reduce the worst one.

Under the readings. The development-set over-tests are 6 to 9 dB above the full-level target and fail reading B; by reading A, which looks only at the end of the run, these runs pass. On the stress-tier held-out panels the corrector never put a control 3 dB above the full-level target (10 of 10 by reading B), where the hold rule did on 4 of 10; it got there by not converging on the low-preload and scatter panels, which no law here controls.

What survives. The corrector does remove the stall of a family deliberately made too wide (one seed, not rerun independently), and it is what passes the softening-then-hardening panel without noise (without it the same law's first full-level pulse is 3.11 dB over). And the lesson about judging: a test judged on some of its full-level pulses can pass while the article sees others 9 dB over.

8. A Step Rule in Place of the Softening Ceiling

The idea. The followed family's 12 % ceiling is a number about the structure that the operator enters. Instead, choose the size of each step up the ladder from how well the law predicted the last pulse, as numerical path-following methods shrink a step when the last prediction was poor. That needs no number about the structure.

What happened. On the development set (one seed), three forms built on the followed family:

- predicting each step through the family's fitted member: 7 of 8, with the 32 % spring's third full-level pulse 18.6 dB off target and its last 7.0 dB off;
- adding the same check to the family with its measured power: 7 of 8, and the severe structure at noise 3e-3 failed with a first pulse 4.65 dB over;
- predicting through the widest member the joint model allows: every structure tried stopped before full level (0 of 4).

The physics-model law of Section 6, whose step rule has a nonlinear model behind it, stopped the 32 % spring at −5.5 dB. An earlier attempt to measure both the extrapolation power and the margin from the run's own pulses, a quick set of low-level mapping pulses in place of the operator's number, failed both springs: on the 16 % spring a first pulse of a level 4.72 dB over and the first full-level pulse 3.45 dB over; on the 32 % spring a third full-level pulse 3.82 dB off and a last pulse 9.21 dB off.

Why it did not work. Without a model that is nonlinear inside the pulse, the only predictions of the next level are linear models with a fixed amount of softening. The fitted one cannot see a softening that has not happened yet, so it is too bold; the widest one assumes the worst for every structure, so it stops everything. Measured ratios do not help either: the per-band softening fitted to a pulse depends on how far that pulse was off target as well as on the level (the per-mode curve of Section 6, indexed by the mode's own amplitude, does not), so ratios taken from successive corrected pulses do not measure the structure; and the total softening a joint will reach cannot be read from levels below full slip.

Under the readings. A third full-level pulse 18.6 dB off and a last pulse 7.0 dB off fail the strict criteria and reading A; whether the 18.6 dB was over or under the target was not saved, so reading B cannot be judged. Stopping every structure fails any reading that asks the test to finish.

What replaced the ceiling. Removing the ceiling outright changes only one development structure, the 32 % spring. What carried that structure in the end was a check of the probes and the rule that holds a pulse once it is in tolerance (the hold rule, Section 4).

9. Two Cures for Failures Under Heavy Noise

The problem they were meant to cure. Over twenty seeds, the hold rule passes 139 of 160 development runs at noise 1e-2. About one run in eight fails, mostly on criterion 1, with first ladder pulses up to 14.3 dB over their own target (2.3 dB above the full-level target), and no run stops safely. The cause, from one failing run read in full and a split of the two noise sources: on the moderate structure the softening at the −40 and −34 dB probes is smaller than a record that noisy can show, the fit returns almost none, and the law then jumps 22 dB to the first ladder level on that estimate; on the severe and non-saturating structures the noise in the identification alone, with clean pulses, caused the failures.

Cure 1: average each probe 16 times (12 dB less noise on the probes). Passes at 1e-2 fell from 139 to 134 of 160; the moderate structure fell from 13 to 4 of 20. With cleaner probes the fit returns a small softening with confidence, and the family it sets is too narrow for the 22 dB jump. Some of the single-probe passes had been noise read as softening.

Cure 2: climb in 6 dB steps (extra levels at −28, −22 and −16 dB between the probes and the first ladder level, with no check on them), for about three more low-level pulses a run. On the realistic tier it removed every failure: 40 of 40 at every noise level (Figure 7c). Overall it passed fewer runs under the strict criteria at every noise level (120 of 160 at no noise, against 160), because on two stress structures an inserted step was over its own level's target: on the severe structure the −16 dB step read 3.5 to 3.7 dB over without noise and up to 5.6 dB at 1e-2 (never above −10.5 dB against the full-level target), and on the 32 % spring the −16 dB step read 7.4 dB over and the first −12 dB pulse 7.2 dB over. Giving the inserted steps the existing probe check did not work: in three of four runs tried, the check lowered successive steps until the ladder would not accept them.

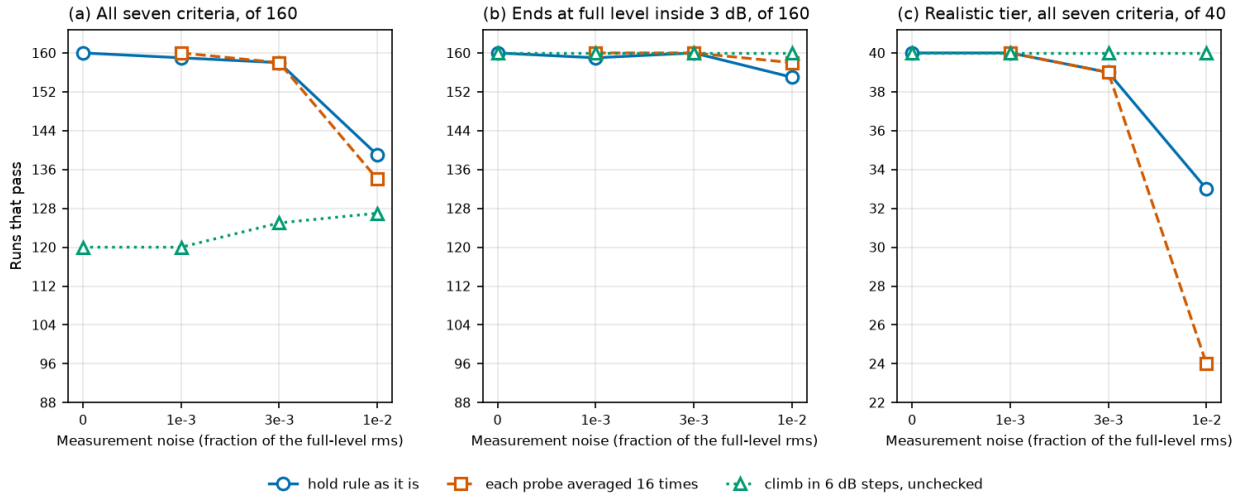


Figure 7. Two attempted cures for the hold rule's failures under heavy measurement noise, eight development structures and twenty seeds at each noise level. Data: `data/fig7_noise_cures.csv`; code: `code/fig7_noise_cures.py`.

Under the readings (Figure 7b), both cures end in tolerance more often than the hold rule at 1e-2 (158 and 160 of 160, against 155). By reading B, at 1e-2 neither cure ever puts a control 3 dB above the full-level target, and the hold rule does once (a full-level pulse 12.6 dB over on the research model). What the strict criteria count against the cures is low-level pulses over their own level's target.

The verdict. Averaging the probes is not a cure; it makes the realistic tier worse even by the strict criteria it was meant to fix. The climb is a cure for the realistic tier and only partly one overall: as built it trades a large first-pulse over-test for smaller ones on its inserted steps, and those steps need a check. The one the probes have scales the drive under the level and stalled three of four runs; the level-lowering check of Table 1 was not tried on them. These results are on the development set only, twenty seeds, and were not run on the held-out structures or rerun independently.

10. Other Ideas, With Weaker or Narrower Evidence

Table 1. Further ideas that did not work. "Development set, one seed" means the developers' own models, no held-out structures and no independent rerun. Where an idea was suggested from practice or by the author, the Acknowledgments say so.

Idea	What happened	Evidence
Identify the model from four copies of the drive, fired 6 dB under the first ladder level with each shaker's polarity coded, so the joints are loaded as the test loads them	The first pulse was much better (0.85 dB over on the severe structure, against 15.4 dB from the buzz), but the coded pulses themselves put controls up to 6.9 dB above the full-level target on the development set and 11.5 dB above it on the held-out panels. Checking each pulse beforehand helped where the article softens little: with the check, the pulses met the exposure limit on four development structures and the whole run passed on two; on the severe, non-saturating and spring structures the first coded pulse measured 14 to 24 dB over its prediction. As built, it fails the strict exposure criterion on every structure; by reading B it puts the article more than 3 dB above the full-level target on two development structures and on seven of the nine held-out four-shaker panels; by reading A most runs end in tolerance.	Strong: all development and held-out structures, run independently
A noise-aware probe check that lowers a probe's level, judges it there, and can stop the test before firing	It passed 193 of 200 development runs against the hold rule's 198. Of its 13 stops on the held-out structures, 6 came after its own first probe had already over-tested, and one was false (a linear frame the frozen law completes). Its stop before firing was right on 3 of 3 runs at noise 3e-3, where every law that fired over-tested. The stop is worth keeping; the check is not.	Strong: five seeds, held out, run independently
Per-control correction on frames with more controls than shakers, from a drive designed for the SRS	Diverged at full level on 11 of 12 linear frames, the last pulse 1.8 to 6.6 dB off and rising, where a correction through the model passed all 12. With more	Development frames only, simulated in steady state rather than from rest, one seed

Idea	What happened	Evidence
	controls than shakers the controls cannot be corrected one at a time.	
Sparse penalties (the lasso and its grouped form, which set small changes exactly to zero) on each pulse's change of the model, so only the modes or bands that moved are updated	No help without noise. Under noise, 0.1 to 2.5 dB worse in two model forms and 1 to 7 dB worse in a third, up to 0.2 dB better in two cases. The trend starts at the lowest levels, where noise is largest, and a threshold set in noise units removes exactly the small early changes that grow later.	Development models whose modes are set by the level, one seed
Regularization of the first pulse of a level, weighted by the model's uncertainty	The first full-level pulse was 0.9 to 2.2 dB over, against 0.1 to 0.6 dB for the same law without it, and the next pulse jumped 8.5 to 11.5 dB. Holding the update back where the model is uncertain holds it back exactly where the structure moved.	As above
Limit how deep the inverse's notches can go, by adding damping to the model	The severe structure's first −12 dB pulse was still 12.3 dB over at best (15.4 dB without), and the linear structure's first pulse rose to 1.6 to 3.5 dB over.	Development set, one seed
Identify with random noise at a higher level, up to the test's own rms for 41 s	The severe structure's first −12 dB pulse was still 13.4 dB over. Random noise does not load the joints as a pulse does.	Development set, one seed
Narrow the family of models as evidence comes in, within a level or across levels	Every run of the evidence-weighted family (36 of 36) ended at the first or second level, on a law fault or without converging; narrowing 0.7 a step faulted on the severe structure and diverged on the non-saturating one; resetting the gains when narrowing still faulted. A drive from a sharper model moves more than the previous gains allow.	Development set, one to three seeds
Follow the softening by stretching the measured model's frequency axis (a frequency warp) on the jointed structures	Diverged on the severe structure, faulted on the moderate one, completed only on the linear structure. It had worked on models whose modes are set by the level and constant inside a pulse.	Development set, three seeds
Set the softening once from the probes and keep it	The research model's first full-level pulse was 7.9 dB over: its identification error hid its softening from the probes. The severe structure was marginal: a level's first pulse 3.06 dB over and a third full-level pulse 2.93 dB off.	Development set, three seeds

Idea	What happened	Evidence
Control a rigid table under an article: a matrix per frequency band updated from one pulse	Worse in every case; with control on the article itself, errors of 11 to 23 dB. Inside a band one pulse's drive has effectively one direction, so one pulse cannot update a matrix.	A separate rigid-table model, one seed
The belief that the combined SRS of several inputs lies between the largest single contribution and their sum	Wrong below: contributions cancel. Removing one input raised the response at one point by 4.4 to 5.3 dB. Only the upper bound (the sum) holds.	Rigid-table model, a direct check

11. What the Failures Have in Common

Low-level pulses do not show the softening that matters. At the probe levels the jointed development structures had moved 0.2 to 1.4 %; by full level, 3 to 15 % (Figure 3). Under noise even that small signal is lost. The network, the step rule, the softening set once from the probes and the averaged probes all extrapolated from there, and many of the strict failures in this work were the first pulse after a large step on such an extrapolation.

A law that fires a predicted drive is only as safe as its lowest prediction. The network as the law's model, the network as ceiling, the step rule through the fitted member and the checked coded pulses all failed this way. A law that holds what it has measured, or averages over a family of possibilities, did not.

A model that contains the simulator's own law proves little. The physics model passed the development set because its form matched how those structures were simulated. Structures built a different way, and kept back from the design, are what showed it.

Judging some pulses hides over-tests. The corrector's 9 dB full-level pulses passed a judgment of the first, third and last. A test article sees every pulse.

One or five seeds can mislead. The physics model's worst first pulse on the severe structure was 0.59 dB over at one seed and about 4.0 dB at two others. The hold rule passed 38 of 40 runs at noise 1e-2 over five seeds and 139 of 160 over twenty.

Most strict failures are low-level pulses over their own target. Under the looser readings several of these ideas are closer to the plain laws than the strict counts suggest, and the slower climb of Section 9 does better than the law it was meant to fix. The ideas that fail reading B as well as the strict criteria are the ones that put the article well above its full-level target: the network on the hardening panel and as ceiling, the corrector between judged pulses, and the coded identification pulses. Reading A, which looks only at the end of a run, misses most of them.

Added mechanisms rarely beat the plain family. Every idea of Sections 5 to 7 ran beside the followed family on the same structures, seeds and noise, and those of Sections 8 and 9 beside the law they were meant to improve. On the held-out structures only the corrector won a run for a reason of its own (the softening-then-hardening panel without noise), and it lost the hardening panel under noise for the same reason. The law that did best there, the hold rule, differs from the plain family by a probe check, a rounding bound and a rule to stop changing a drive that works.

12. Limits

- **Simulation only.** Nothing here has been on a shaker. The held-out joints are a model too (Iwan joints), and the model leaves out normal-direction joint dynamics, contact opening, heating and wear.

- **Most of the development set softens more than well-tightened joints.** Its moderate structure softens 3 %, within the about 4 % or less measured on well-tightened bolted joints; its other jointed models and springs soften 8 to 16 %, against about 6 to 10 % measured on loose or friction joints under shaker drive, most with no damping rise, and their softening follows a running envelope rather than a slipping joint. Few joint measurements use shock: on one press-fit interface shaker shock left the frequency unchanged where sine lowered it (with different foam cups), and a bolted resonant plate at pyroshock level moved 1.6 to 12.7 % down on five modes and 20 % up on one.
- **Mostly four shakers.** Frames appear only in Table 1; the larger open problems (dense modes, modes one or two lines wide, 16 controls on 8 shakers) are not covered here.
- **Seeds.** Twenty seeds for the noise studies, five for the development-set comparisons of the corrector and the probe check, one for the held-out comparisons, one to three for the earlier ideas in Table 1.
- **The yardstick.** The strict criteria are ours. Where a looser reading changes a verdict, it is stated; a different standard could change more.
- **The network** was trained on one simulated population. A larger or different one might narrow its predictions on its own mechanism; it cannot be expected to bound a prediction outside it.

What would change these conclusions: a jointed test article on shakers, run with the plain laws and with any of these ideas; and held-out tests of the earlier ideas in Table 1.

How This Was Made

AI tools were used in this work. The simulations, the control laws, the independent reruns, the figure code and the drafting of the text were done with Claude Code, an AI coding assistant running Anthropic's Claude Opus models, under Kevin Cross's direction; the networks were trained with PyTorch. The author is responsible for every word, number and line of code.

Acknowledgments

Four of the ideas tested were mine: trying sparse methods (the penalties of Table 1), a network pre-trained on structural dynamics and adapted on the test's own data (Section 5), identification from low-level coded pulses, from shock test practice (Table 1), and setting the softening allowance from a quick set of low-level mapping pulses instead of an operator's number (Section 8). The others were developed in the course of the work, from the published sources below.

Published work of Sandia National Laboratories is used throughout and credited to Sandia National Laboratories: single-shaker SRS control and response limiting (Coleman and Skousen of Sandia, with Smallwood), SRS iteration for one shaker and several responses (Heitman, Babuska and Cap, Sandia, with Reid of the Massachusetts Institute of Technology), the four-parameter Iwan joint model (Segalman), and the measurement of instantaneous frequency and damping from ring-down data (Kuether and Brake). Measured joint behavior comes from published work of the University of Wisconsin-Madison and Sandia National Laboratories and their co-authors (Roettgen; Wall, Allen and Kuether; Deaner, Allen and co-authors; Bonney and co-authors; Jacobs-O'Malley and Hofer; Kuether and Najera; Khan and co-authors; Schoenherr, Soine and Witt). The networks for SRS we found are the work of Watts, Jeon and Newton of Los Alamos National Laboratory and their co-authors.

References

- Bolderman, M., Butler, H., Koekebakker, S., van Horssen, E., Kamidi, R., Spaan-Burke, T., Strijbosch, N., Lazar, M. (2024). Physics-guided neural networks for feedforward control with input-to-state-stability guarantees. Control Engineering Practice 145, 105851. DOI 10.1016/j.conengprac.2024.105851.

- Bonney, M. S., Robertson, B. A., Mignolet, M., Schempp, F., Brake, M. R. W. (2016). Experimental Determination of Frictional Interface Models. In Dynamics of Coupled Structures, Volume 4 (IMAC XXXIV), 473-490. DOI 10.1007/978-3-319-29763-7_47. SAND2015-9817C.
- Coleman, R., Smallwood, D., Skousen, T. (2016). Response Limited Shaker Shock Testing. 86th Shock and Vibration Symposium. SAND2016-1335C, OSTI 1344695.
- Cross, K., Zwink, B. (2026). What Is a Good Multi-Shaker Test? A Proposed Practice for Judging MIMO Vibration and Shock Tests on Flexible Articles. Hertzforge Journal 1, 2026-001. DOI 10.5281/zenodo.24235533.
- Deaner, B. J., Allen, M. S., Starr, M. J., Segalman, D. J., Sumali, H. (2015). Application of Viscous and Iwan Modal Damping Models to Experimental Measurements From Bolted Structures. Journal of Vibration and Acoustics 137(2), 021012. DOI 10.1115/1.4029074.
- Forgione, M., Pura, F., Piga, D. (2023). From System Models to Class Models: An In-Context Learning Paradigm. IEEE Control Systems Letters 7, 3513-3518. DOI 10.1109/LCSYS.2023.3335036.
- Heitman, C., Babuska, V., Cap, J. S., Reid, J. B. (2016). Development of a Single Input Multiple Output (SIMO) Input Derivation Algorithm for Oscillatory Decaying Shocks. 86th Shock and Vibration Symposium. SAND2016-3995C, OSTI 1365337.
- Jacobs-O'Malley, L. D., Hofer, J. H. (2017). Nonlinear Feature Extraction and Energy Dissipation of Foam/Metal Interfaces. SAND2017-4293, Sandia National Laboratories. DOI 10.2172/1595879.
- Khan, M. Y., Hunter, P., Pacini, B. R., Roettgen, D. R., Schoenherr, T. F. (2021). Evaluation of Joint Modeling Techniques Using Calibration and Fatigue Assessment of a Bolted Structure. Presentation slides, IMAC XXXIX, SAND2020-13875C, OSTI 1836908, DOI 10.2172/1836908. The paper is in Nonlinear Structures and Systems, Volume 1, 33-52, DOI 10.1007/978-3-030-77135-5_4.
- Kuether, R. J., Brake, M. R. W. (2016). Instantaneous Frequency and Damping from Transient Ring-Down Data. In Dynamics of Coupled Structures, Volume 4 (IMAC XXXIV), 253-263. DOI 10.1007/978-3-319-29763-7_24. SAND2015-9173C.
- Kuether, R. J., Najera, D. A. (2017). Modeling Nonlinear Energy Dissipation of the Ministack Assembly. SAND2017-10517, Sandia National Laboratories. OSTI 1596201, DOI 10.2172/1596201.
- Peláez-Rodríguez, C., Magdaleno, Á., Salcedo-Sanz, S., Lorenzana, A. (2023). Human-induced force reconstruction using a non-linear electrodynamic shaker applying an iterative neural network algorithm. Bulletin of the Polish Academy of Sciences: Technical Sciences 71(3), e144615. DOI 10.24425/bpasts.2023.144615.
- Porter, J. H., Balaji, N. N., Little, C. R., Brake, M. R. W. (2022). A quantitative assessment of the model form error of friction models across different interface representations for jointed structures. Mechanical Systems and Signal Processing 163, 108163. DOI 10.1016/j.ymssp.2021.108163.
- Roettgen, D. R. (2016). Experimental Dynamic Substructuring Using Nonlinear Modal Joint Models. PhD dissertation, University of Wisconsin-Madison.
- Schoenherr, T. F., Soine, D. E., Witt, B. L. (2022). Experimental Modal Analysis of a Resonant Plate During a Mid-Field Pyroshock Replication Test. In Sensors and Instrumentation, Aircraft/Aerospace and Dynamic Environments Testing, Volume 7 (IMAC XL), 41-52. DOI 10.1007/978-3-031-05415-0_4.
- Segalman, D. J. (2002). A Four-Parameter Iwan Model for Lap-Type Joints. SAND2002-3828, OSTI 805877.
- Volckaert, M., Diehl, M., Swevers, J. (2011). Iterative learning control for nonlinear systems with input constraints and discontinuously changing dynamics. Proceedings of the 2011 American Control Conference, 3035-3040. DOI 10.1109/ACC.2011.5991346.
- Wall, M., Allen, M. S., Kuether, R. J. (2022). Observations of modal coupling due to bolted joints in an experimental benchmark structure. Mechanical Systems and Signal Processing 162, 107968. DOI 10.1016/j.ymssp.2021.107968.
- Watts, A., Jeon, A., Newton, D., Bowering, R. (2026). Inverse Reconstruction of Shock Time Series from Shock Response Spectrum Curves using Machine Learning. arXiv 2603.03229v3.

- Zheng, R., Chen, H., Wei, X. (2022). Shock response spectrum reproduction control method for multi-axis shock vibration testing. Chinese patent application CN115130506A, Nanjing University of Aeronautics and Astronautics; deemed withdrawn 2025.
- Zheng, R., Ma, Y., Wei, X., Chen, H. (2023). Multi-shaker shock response spectra replication control. Journal of Vibration and Control 29(15-16), 3744-3755 (online 2022). DOI 10.1177/10775463221104577.

Version History

Version 1, 2026-11-15. DOI 10.5281/zenodo.24254182. First published.

Reviews. This article is reviewed in public. Its reviews, its review status and any later version are at https://hertzforge.com/journal/2026-004/

Copyright 2026 Kevin Cross. Text, figures and data CC BY 4.0; code MIT.

Methods are published for research. Anyone who applies one to a test is responsible for protecting the equipment and the article under test.