

Closing the Loop Behind the Cochlea

Monitoring and device architecture for slow-growth retrocochlear hearing loss

Matthew "Bear" Schulz

Independent researcher

Belmont, California, United States

ORCID: [0009-0000-0956-5743](https://orcid.org/0009-0000-0956-5743)

Licence: Creative Commons Attribution 4.0 International (CC BY 4.0)

Preprint · 27 August 2026

Companion papers: *Ultrasonic Cochlear Excitation — Device Scoping*; *The Auditory Pathway as a General-Purpose Data Channel*

Keywords: vestibular schwannoma · retrocochlear hearing loss · otoacoustic emissions · medial olivocochlear reflex · closed-loop control · cochlear implant telemetry · focused ultrasound · holonomy

Summary

Some hearing loss happens because the ear itself is damaged. This paper is about a different kind: a slow-growing, usually benign tumour presses on the nerve carrying sound from a healthy ear to the brain. **The ear works. The wire is being squeezed.**

Most hearing technology exists to replace a broken ear, so it does not help here. Roughly one person in a thousand develops such a tumour, many are watched rather than treated, and the hardest question in that management is *when to act* — too early means surgery nobody needed, too late means hearing and balance that cannot be recovered. That decision is currently made on scans months apart, a hearing test that blends the ear and the nerve into one number, and a patient's recollection of months they have quietly adapted to.

The central observation is that the ear is already wired as a loop. Nerve fibres carry sound inward to the brainstem, and the brainstem reaches back down the same bundle to adjust the ear's own amplifier. That reaching-back is measurable from outside the head, in minutes, with equipment audiology clinics already own — and squeezing the nerve should weaken it. The disease opens a loop the body already runs, and the loop's state can be read.

A second observation follows from the first. The faint sound a healthy ear emits turns out to be **insensitive to damage on the nerve side** — animal work shows it essentially unchanged even when most inner hair cells are gone. That is usually reported as a limitation. Here it is the point: a measurement that ignores the nerve is exactly what you want when you are trying to tell the ear apart from the nerve, and the standard hearing test cannot.

Third, hearing implants are open-loop. They carry the hardware to check whether a stimulus produced its intended neural response, use it once at fitting, and then never ask again. That is reasonable for an ear that is not changing and wrong for a nerve that is.

The paper decomposes hearing assistance into eight functional stages, identifies which stage fails for each of four populations, and proposes four measurements — efferent loop integrity, emission phase rather than amplitude, paired tests of the two hair cell populations, and vestibular function tracked alongside hearing. Five of the capabilities discussed **already exist and are not novel**; they are simply not assembled into a surveillance protocol for this population. An appendix records seven approaches evaluated and not pursued, each with the number that closed it.

None of this restores hearing. It aims to make the decisions around it better — and in a disease managed by watching, the decision is the treatment.

1. An engineering orphan

Roughly one person in a thousand develops a vestibular schwannoma. Most are small, many grow slowly or not at all, and standard practice for a substantial share is observation with interval imaging and repeat audiometry rather than immediate surgery or radiation. That produces a population living for years with a hearing loss that is changing, on one side, from a cause that is not the cochlea.

Their device options were designed for someone else. Hearing aids amplify into a pathway that cannot carry more signal. Cochlear implants replace a transduction stage that has not failed and depend on a pathway that has. Contralateral routing devices move sound to the good ear and discard the geometry that made it informative. Auditory brainstem implants address the case where the nerve is gone entirely and deliver limited speech understanding.

This is not a failure of the industry. It is a consequence of scale: the sensory-loss population is orders of magnitude larger, the engineering there is correspondingly mature, and it is prior-art dense enough that newcomers add little. The retrocochlear case is where questions have gone unasked, because almost nobody was asking them.

2. Eight functional stages

A hearing device of any kind performs eight functions in sequence. Naming them separately is what makes the two populations distinguishable, because they fail at different stages.

Table 1.

#	stage	what it does
1	Acquisition	captures sound from the world
2	Analysis	splits it into frequency channels
3	Mapping	assigns each channel to a delivery site
4	Encoding	decides what about each channel is transmitted — envelope only, or fine structure too
5	Delivery	gets energy to the site
6	Transduction	converts that energy into neural firing
7	Calibration	sets every parameter for this individual
8	Feedback	reports whether any of it worked

Stages one to six are a signal chain. Stage seven is a configuration step. Stage eight, in current practice, is close to absent.

Table 2. Where the chain breaks.

stage	sensory (cochlear) loss	retrocochlear loss
1 Acquisition	intact	intact
2 Analysis	intact	intact
3 Mapping	intact in principle	intact
4 Encoding	intact	intact
5 Delivery	intact	intact
6 Transduction	FAILED — hair cells lost	intact
— the pathway	intact	FAILED — compression, ischaemia
7 Calibration	one-off, adequate	inadequate — the target moves
8 Feedback	absent, tolerable	absent, and it is the problem

The cochlear implant is a correct and elegant answer to the left column: it replaces stage six. In the right column it replaces a stage that had not failed.

Two rows carry the rest of this paper. **Calibration** is specified as a fitting procedure, done once and revisited when a patient complains — which assumes a stable target, and slow growth means the target is not stable by definition. **Feedback** is absent from the device, and, as the next section argues, present in the patient.

4. The loop the ear already has

The cochlea is not an open-loop sensor. Auditory nerve fibres project to multipolar neurons of the posteroventral cochlear nucleus, which project to medial olivocochlear neurons of the superior olivary complex, which project back bilaterally to the outer hair cells. By hyperpolarising those cells the efferents inhibit their electromotility and reduce the gain of the cochlear amplifier. That is a feedback loop with a brainstem controller, and it has an externally observable output: otoacoustic emissions, whose amplitude falls when the loop is driven.

The standard assay is contralateral suppression. Present noise to the opposite ear, measure the reduction in emission amplitude in the test ear, and the size of that reduction reports the functional integrity of the efferent pathway. Suppression is reduced in ears whose vestibular nerve has been sectioned, which severs the olivocochlear bundle, and reduced suppression has been reported in acoustic neuroma.

4.1 What a compressed nerve does to that loop

Compression of the eighth nerve reduces the afferent traffic that drives the loop. Less afferent drive means less efferent output, and therefore less suppression of outer hair cell gain — the cochlear amplifier runs with its brake partly released.

The animal evidence is more specific than that argument alone, and it is worth setting out carefully because it points somewhere better.

Selective destruction of inner hair cells in chinchillas, sparing outer hair cells, has been studied repeatedly. The consistent finding across that work is that **distortion-product emissions are essentially unaffected**. Losses of 80–100% of inner hair cells with outer hair cells retained do not reduce emission amplitude or the cochlear microphonic, while completely abolishing the compound action potential; a dose-response study reached the same conclusion; and a review of the series states plainly that eliminating the inner hair cell and type-I neuron system has no effect on distortion-product emissions or the cochlear microphonic.

Separately, one study using *click-evoked* emissions — a different measurement — reported amplification from regions with extensive inner hair cell damage.

Read together, this is a better result for monitoring than a rise would have been. The property that matters is not that emissions increase but that distortion-product emissions are **insensitive to loss on the neural side of the hair cell**. That insensitivity is usually reported as a limitation of the test. Here it is precisely the specification: a measurement that ignores the nerve is what is needed when the clinical problem is telling the ear apart from the nerve, and a pure-tone audiogram cannot.

Two consequences follow directly.

Divergence is the signal. Emissions holding steady while the audiogram worsens localises the loss behind the cochlea. Emissions declining alongside it says something is reaching the outer hair cells as well — which in this anatomy most plausibly means the blood supply travelling through the same canal. The two tests are not two measurements of the same thing that ought to agree; their *disagreement* is the diagnostic content.

And the audiogram is less sensitive than expected. The same animal work found thresholds in quiet remaining surprisingly normal until inner hair cell losses exceeded 80%, with the remaining fibres retaining sharp tuning and low thresholds. If a threshold measure can miss most of the neural population, a monitoring protocol built only on thresholds is monitoring the wrong quantity late.

What this does not support. The efferent argument at the head of this section predicts weaker suppression, not larger emissions, and the emission-amplitude literature above is consistent with no change rather than a rise. Any protocol should therefore read *suppression* — the assay of §4 — as the loop measure, and emission amplitude as a stable cochlear reference against which the audiogram is compared. A claim that emissions rise as the nerve fails is not supported by the distortion-product evidence and is not made here.

4.2 The question the literature has not settled

Whether human efferent suppression reliably weakens in schwannoma is unresolved. One line of work reports suppression absent or reduced in acoustic neuroma. A dedicated pilot in nine patients with unilateral schwannoma, comparing tumour and non-tumour ears against matched controls, found no statistically significant difference, and concluded that future studies should include a larger population across different grades of hearing loss.

Proposal 1 — run that study, longitudinally, in the observation cohort. The existing work is cross-sectional and small. What the mechanism predicts is not a difference between patients at one moment but a *trend within a patient over time*: efferent suppression weakening as afferent traffic declines, in step with the lesion. That is a within-subject design in a population already returning for interval imaging, using equipment audiology clinics already own.

How it fails. If efferent suppression in this cohort proves dominated by measurement variability, by middle-ear status, or by hearing threshold rather than by lesion progression, the trend carries no information. The pilot's null result is a warning that this is possible, and the longitudinal design exists precisely to distinguish a real weak effect from an absent one.

5. The loop the device does not have

Every modern cochlear implant can evoke and record a compound action potential: the array stimulates and, a fraction of a millisecond later, records the synchronous neural response. The capability ships under a trade name in every major system, automated threshold-seeking software is fielded, and closed-loop fitting is an active research direction — a demonstrated proof of concept records cortical responses from the implant alone, aimed at automatic and adaptive fitting.

All of that targets **fitting a stable ear more conveniently**. In ordinary operation the device stimulates from a stored map and never asks whether the intended response occurred.

For a stable deaf ear, that is a reasonable engineering decision. The plant does not change and measurement costs power. Every other safety-relevant control system built in the last forty years is closed-loop because plants drift; this one is not, because this plant does not.

For a nerve under slowly increasing compression, the plant does drift, and the device is blind to it. Threshold, response latency and the recruitment of fibres at a given stimulus level all change as the lesion grows, and every one is measurable with hardware already inside the device.

Proposal 2 — closed-loop stimulation with drift tracking, and the residual as a growth marker. Sample the evoked response on a duty cycle and store the trend. Use it twice: adapt stimulation to maintain the intended neural response as the pathway changes, and surface the trend clinically.

The quantity a control system watches is the **residual** — the gap between the response commanded and the response achieved — together with its derivatives. A stable residual says the pathway is stable. Steady drift says steady change. A residual whose *rate of change is itself increasing* says the driving process has accelerated, and that second derivative is what a team following a slow lesion wants. Growth rate is the decision variable in observation management, and it is currently estimated from serial imaging with substantial uncertainty on small lesions.

5.1 Why this matters more than it first appears

Between appointments months apart, the instrument of record for symptoms is the question "how have you been?" That instrument is weak here for three compounding reasons. Change is gradual, so no single day differs noticeably from the one before. The brain adapts, so a patient recalibrates to a slowly worsening baseline and reports it as normal. And ordinary day-to-day variation in hearing is large enough to hide a slow trend inside it.

A continuous objective record does not replace the patient's account; it gives patient and clinician something to check it against. A drift the patient has normalised becomes visible as a line on a chart, and a complaint that turns out not to correspond to any change in the record is equally informative.

How it fails. If evoked-response parameters are dominated by electrode-tissue impedance drift, array position shift, or ordinary physiological variability rather than by pathway change, the trend carries no clinical information. That is directly testable in an existing implanted population before any device is modified.

What makes it possible. Evoked responses are microvolts against millivolt noise and are recovered by averaging — a thousand sweeps buys about 30 dB, but only if trigger phase is stable across them. Averages compared across months therefore require a phase reference stable across months. An implanted timing reference is not a refinement of this proposal; it is what makes a longitudinal residual comparable to itself.

6. Making the measurement survive the clinic

Emissions report on the cochlea independently of the nerve, which is exactly the separation this population needs. The difficulty is that emission **amplitude** is fragile in serial use: reliability degrades whenever the probe is refitted between sessions, absolute values differ substantially between commercial systems, and a patient followed for years across clinics and instruments accumulates variance that can mask real change or manufacture false alarms. The field's response has been improved stimulus calibration, which reduces the problem without removing its cause.

Proposal 3 — track phase gradient rather than amplitude. Emission phase, swept across frequency, has a gradient, and that gradient is a derivative computed *within a single probe insertion*. The insertion-dependent transfer function that corrupts absolute level largely cancels in the derivative. Amplitude asks how loud the emission is, which depends on everything between the cochlea and the microphone. Phase gradient asks how fast the phase turns as frequency advances, which depends principally on the cochlear mechanics generating it.

Phase-gradient analysis itself is established — the separation of emission sources into wave-fixed and place-fixed components is done on exactly this basis, and the argument that clinical practice discards phase has been published with swept-tone evidence and a recommendation that phase be integrated into screening. That work argues from **sensitivity**. The argument here is **robustness to probe re-insertion in longitudinal use**, which is the confound that dominates surveillance rather than screening.

How it fails. If phase gradient shows test-retest variability across deliberate probe refits comparable to or worse than calibrated amplitude, the argument collapses. Repeated swept measurements with deliberate repositioning between them, both metrics computed from the same recordings, settles it — no new hardware and no patients beyond a normal-hearing panel.

6.1 Why a phase gradient is the robust quantity

The reason is structural and worth naming, because it generalises. The quantity being read is **holonomy** — what a system accumulates when carried around a closed path, as distinct from how large anything is along the way. The stimulus travels inward to the generation site, the emission is created, and it travels back out to the microphone; what a phase gradient measures is the phase accumulated around that round trip per unit of frequency.

Holonomy is magnitude-blind by construction. A winding rate is a property of a path, and no uniform rescaling of the signal can change it. That is precisely the right invariance when the dominant nuisance — how deep the probe sits — acts approximately as a uniform gain. Holonomy is not being imported here as a framework; it is already the operative principle of emission source separation, and it hides in plain sight wherever a field has learned to distrust its own amplitudes.

7. Acoustic delivery: what changes when the energy arrives as sound

Every barrier in Appendix A.4 comes from making sound by heating tissue. If the energy arrives as sound already, none of them applies — no absorption-to-pressure conversion, no unavoidable DC bias, no thermal dose. That is a different mechanism with its own literature, and it is not closed by anything above.

7.1 What is already established

Humans perceive **bone-conducted ultrasound** well above the roughly 24 kHz limit of air conduction. Amplitude-modulated ultrasound carries speech, some profoundly deaf individuals perceive it, and work using cisplatin-induced cochlear change found that *direct* ultrasonic stimulation of the cochlea contributes to perception rather than demodulation alone, with the cochlear basal turn implicated. Devices built on this principle have been patented. None of this is speculative and none of it is claimed here.

Separately, **focused ultrasound neuromodulation** acts mechanically rather than thermally, operates at 0.2–1.5 MHz, and reaches deep structures transcranially.

7.2 The property that matters most, and it is not the one usually cited

Measurements of bone-conducted stimuli across 250 Hz to 64 kHz show attenuation rising with frequency, and the consequence stated in that work is decisive for this population: at audiometric frequencies, stimulating one side of the head stimulates the other almost equally, while **in the ultrasonic range the contralateral side is progressively isolated**.

That is the crosstalk problem, solved by moving band. Transcranial crosstalk is what makes bone-conduction devices deliver a signal to the wrong cochlea, degrade the intact ear's own spectral cues, and require adaptive cancellation to undo the damage. It is a property of the audio band, not of bone conduction. An ultrasonic carrier is naturally lateralised.

7.3 The spatial argument, which attacks the field's oldest ceiling

Electrical stimulus spreads widely through perilymph because the intracochlear fluids are electrolytes, and the resulting channel interaction constrains frequency selectivity and speech recognition. The spread distance is not a universal constant: it is characterised as a space constant fitted per patient from evoked-response measurements, and it is that distribution, not a single figure, which sets how many independent sites an array can address.

Acoustic focusing is diffraction-limited instead:

Table 3.

frequency	wavelength	focal spot	independent sites over 20 mm
0.25 MHz	6.00 mm	6.0 mm	3
0.50 MHz	3.00 mm	3.0 mm	7
1.00 MHz	1.50 mm	1.5 mm	13
1.50 MHz	1.00 mm	1.0 mm	20

The comparison is therefore conditional, and honestly so. Against a space constant of 3–4 mm a 1.0–1.5 mm focus gives a four- to sixfold gain in independent sites; against 1.5 mm it gives a factor of two, which would not justify a device programme; against 1.0 mm the argument is dead. Establishing that distribution from evoked-response data implant manufacturers already collect is cheaper than any other step here and gates all of them.

The second open question is biological: whether spiral ganglion neurons respond to focused ultrasound at intensities within safety limits. **Neither question is about geometry, and neither is settled by this paper.**

7.4 Where the hardware already exists

Transcranial delivery through the skull vault is costly: at roughly 10, 20 and 40 dB/cm for 0.25, 0.5 and 1.0 MHz through 7 mm of bone, transmitted intensity falls from about 20% to 4% to 0.16%, which is why clinical systems use large phased arrays with aberration correction.

But the cochlea is not behind the vault. It sits within the temporal bone, and mastoid or ear-canal coupling avoids the vault entirely — which is the geometry a bone-anchored hearing device already uses, with an established surgical pathway, established osseointegration, and an existing regulatory route. The transducer siting problem for this application is largely solved hardware operating in the wrong frequency band.

7.5 The confound that is the mechanism

The largest methodological problem in focused ultrasound neuromodulation is that focused ultrasound makes a sound. Cortical activation follows an indirect auditory route; contralateral bias demonstrates it originates at the ear nearest the focus; and deafened animal models are now standard specifically to control for it.

For a hearing application that confound is the effect being sought. **But the same work states the boundary precisely: transecting the auditory nerves, or removing cochlear fluids, eliminates it.** Ultrasound-evoked hearing runs *through* the cochlea and the auditory nerve.

7.6 What this does and does not offer this population

It is not a therapy for retrocochlear loss. It does not bypass a compressed nerve, because it depends on that nerve. Any claim otherwise contradicts the evidence in §7.5.

Proposal 4 — ultrasound as a pathway probe. What a millimetre-scale, non-invasively steerable acoustic focus offers is the ability to excite a chosen location and ask whether the pathway still carries it. Combined with the evoked responses of §5 and the emissions of §4, that turns a single global question — *is hearing worse* — into a spatially resolved one: *which parts of the pathway still conduct*.

How it fails. If ultrasonic excitation at safe intensity cannot be localised finely enough in the intact human ear to distinguish adjacent cochlear regions, the probe has no more resolution than existing tests and adds nothing.

What a device would have to be. A falsification test says whether an idea is wrong; it does not say what would be built if it were right. The functional placement of an ultrasonic aperture, the acoustic output envelope imposed by IEC 60601-2-37 and FDA Track 3 limits, the standards that would apply, and the staged evidence required before any human exposure are set out in a companion scoping document. Two points from it belong here. **The binding constraint is the time-average intensity limit rather than the mechanical index**, because a hearing device must run continuously: at the pulse-average ceiling the permitted duty cycle is 0.38%, so continuous audio delivery sits far below the peak pressures the mechanical index would allow. And **an acoustic aperture is bidirectional**, so the same hardware that stimulates can image the structure it is stimulating — which inverts the proposition. A cochlear implant is a delivery device carrying telemetry; an ultrasonic aperture is an imaging device that may also stimulate, and the diagnostic indication is both the cheaper regulatory path and the one that makes no therapeutic claim.

7.7 Two hard limits on the delivery route, if anyone pursues it

Bone-conducted ultrasound has a dynamic range near 30 dB, against roughly 120 dB for air conduction — and speech alone spans about 35 dB of level variation, so there is no headroom left after speech. And its perceived pitch sits near 8–16 kHz **regardless of carrier frequency**, so the carrier cannot code pitch and all frequency information must ride in the modulation envelope. That is the same constraint cochlear implants already live under, which is worth knowing before anyone treats the ultrasonic carrier as a richer channel than it is.

7.8 Where this rhymes beyond this population

Three adjacent applications share the structure and are noted without being developed. The eight-channel ceiling belongs to every cochlear implant recipient, not only this cohort, so §7.3 is a general argument if the biology cooperates. Non-invasive stimulation would address people who do not qualify for or do not want implantation. And a spatially steerable probe that maps where conduction fails is a diagnostic tool for any retrocochlear or neuropathic hearing loss, including auditory neuropathy, where the same measurement ambiguity applies.

8. Where the wall is, for each population

The proposals above were tested against one hard case. That is efficient rather than narrow, because a constraint that closes an approach also names the population where the approach belongs. Four populations, four different walls.

Table 4.

population	what has failed	the wall	what that permits
Sensory (cochlear) loss — the large population	hair cells	electrical stimulus spreads through perilymph, so channel interaction limits selectivity and performance has plateaued near eight effective channels	anything that beats conduction-limited spread: tighter focusing, distributed local stimulation, mechanical rather than electrical delivery
Conductive and mixed loss	the middle-ear mechanical chain	a replacement strut transmits whatever its geometry gives; the transfer function is not designed	a graded passive resonator array — a prosthesis with a specified transfer function rather than an incidental one
Retrocochlear loss — this paper	the pathway, with the cochlea intact	every delivery route depends on the nerve it delivers into, so none of them transfers	measurement rather than delivery: the efferent loop, emission phase, paired hair-cell tests, vestibular function
Normal hearing, elective use	nothing	surgical risk accepted for convenience, and a device that cannot be removed	assistive cases where the benefit is capability rather than convenience — and governance written before a market exists

Three consequences follow that are not obvious from any single row of that table, and each one points somewhere different.

The walls are different in kind, not merely in place. Sensory loss is bounded by conduction physics, conductive loss by mechanical design, retrocochlear loss by a biological dependency, and elective use by an ethical asymmetry. Only the first two are engineering problems in the ordinary sense.

A closure in one row is a redirection in another. Every architecture examined in the companion scoping paper failed for this population and failed *specifically* — implanted elements burden the surveillance imaging that defines this population's care; a passive resonator array duplicates a cochlea that here is intact; ultrasonic excitation depends on the compressed nerve. Each of those constraints is absent from at least one other row.

And the last row inverts the usual reading of risk. Elective use carries the weakest benefit-to-risk ratio and the strongest misuse profile — but a large elective market is also how a surgical safety envelope gets characterised, funded and made routine, and small populations historically inherit mature technology they could never have paid to develop. Both readings are true. This paper states the asymmetry in both directions and does not resolve it; that judgement belongs to patients, clinicians, regulators and the people who would live with the device. The companion paper on the general-purpose channel treats it at length.

Table 5. What a clinic could do today, and what needs work.

capability	status	what it requires
separate cochlear from neural decline over time	available now	serial emissions on the imaging schedule; equipment already owned
measure efferent loop integrity	available now	contralateral suppression, standard protocol
track balance alongside hearing	available now	vestibular evoked myogenic potentials; comprehensive schwannoma batteries are already published
identify which nerve branch is involved	available now	bone-conducted VEMP, demonstrated for nerve of origin
preserve spatial cues under a bone device	available now	adaptive crosstalk cancellation with an error sensor near the cochlea, published with a verification procedure
robust longitudinal emission tracking	needs work	phase-gradient retest study (Proposal 3)
efferent suppression as a progression marker	needs work	longitudinal within-subject study (Proposal 1)
implant that tracks a changing nerve	needs work	duty-cycled telemetry, stable timing reference, residual trending (Proposal 2)

The first five rows are the immediate contribution of this paper and none of them is novel: they are existing capabilities not currently assembled into a surveillance protocol for this population. **Assembling them is available today and requires no invention.**

10. What this means for a patient

Nothing here restores hearing, and nothing here slows a tumour. It would be dishonest to imply otherwise, and any account of benefit has to start there.

What it offers is information, and in a disease managed by watching, information is what the treatment decision is made of. Acting too early means surgery or radiation that was not needed, with its own risks to hearing, balance and facial nerve. Acting too late means losing hearing and balance that cannot be recovered. That decision is currently made on interval imaging, an audiogram that blends the ear and the nerve into a single number, and a patient's recollection of months they have adapted to.

Better information changes four things concretely. It can show whether decline is coming from the cochlea or the pathway, which are different trajectories with different implications. It can supply a functional trend between scans rather than only at them. It can make gradual change visible to a patient who has normalised it. And it can bring balance into the same picture as hearing, when unsteadiness and visual instability are frequently what patients report as most disabling and are downstream of the same millimetre of anatomy.

In fifty words: this does not give hearing back. It aims to tell you sooner and more precisely what is changing and why, so the decision about whether and when to intervene is made on evidence rather than on a blended number and a recollection — and that decision is what determines how much hearing and balance are kept.

Table 6. Falsification summary.

#	proposal	what would kill it
1	efferent suppression as a longitudinal progression marker	suppression dominated by measurement variability, middle-ear status or threshold rather than by lesion progression
2	closed-loop stimulation with residual trending	evoked-response parameters dominated by impedance or position drift
3	phase gradient over amplitude in serial use	phase-gradient retest variability across probe refits $\geq$ calibrated amplitude
4	focused ultrasound as a pathway probe	safe-intensity excitation cannot be localised finely enough to distinguish adjacent cochlear regions

Proposals 1 to 3 are answerable with existing equipment and need no device modification or implanted patients. Proposal 4 requires a bench and animal path before any human work.

12. Prior-art position

A targeted rather than systematic search was performed. Nothing in this paper is a new physical effect and none is claimed.

The efferent loop, its anatomy and the contralateral suppression assay are established. Reduced suppression in acoustic neuroma has been reported, and a dedicated schwannoma pilot found no significant difference; the discrepancy is the open question Proposal 1 addresses, and the call for a larger study is the literature's own.

Evoked-response telemetry ships in every major implant system, and closed-loop fitting is an active research direction aimed at stable ears. The application to chronic drift tracking in a changing pathway, and the use of the residual and its derivatives as a progression marker, were not found in the searched literature.

Emission phase-gradient analysis is established and its clinical under-use has been argued in print from sensitivity; the robustness-to-re-insertion argument for longitudinal surveillance was not found.

Bone conduction devices already serve this population and that literature must be acknowledged. Bone conduction implants for hearing rehabilitation in skull-base tumour patients are an established subject, including bone-anchored devices placed after retrosigmoid approach for vestibular schwannoma, implantation in patients with prior craniotomy, a systematic review and meta-analysis of quality-of-life outcomes with bone conduction devices in single-sided sensorineural deafness, and direct comparison of speech-in-noise and localisation benefit between contralateral-routing aids and bone-anchored implants. None of that work addresses the monitoring questions this paper raises, but it establishes that the population is served surgically and that any measurement proposal enters an existing clinical pathway rather than an empty one.

Comprehensive audio-vestibular batteries for schwannoma follow-up are published, including bone-conducted VEMP used to identify a tumour's nerve of origin. Adaptive crosstalk cancellation for bone-conduction devices is published with a verification procedure. Both appear in §9 as available capabilities rather than as proposals, on that basis.

13. Limitations

No patient data underlies this paper and no device has been built. The staging in §9 is constructed from published natural history rather than from a cohort. Proposals are ordered by how cheaply they can be tested, not by expected benefit. The population is small, which is why the engineering is absent and also why any trial would be slow, multi-centre and expensive — a further reason to begin with the two proposals answerable from equipment clinics already own.

Appendix A — Approaches evaluated and not pursued

Recorded because a documented dead end is worth more to the next reader than another proposal. Each entry gives the approach, why it is attractive, and the quantity that closes it.

A.1 Using bone-conduction crosstalk to build an artificial direction cue. In a bone-conduction arrangement the intact cochlea receives two signals — one through the outer ear carrying the pinna's direction-dependent spectral notches, one by bone carrying none — and the interference between them is direction- dependent. *Closes because:* modelling of bilateral bone-conduction crosstalk finds the interference impairs rather than assists localisation, and a study of nineteen single-sided listeners on contralateral bone devices found no localisation benefit, attributing it to masking of the monaural spectral cues the listener has learned to use. The productive direction is the inverse — cancelling the crosstalk to preserve those cues — which is published and appears in §9.

A.2 A master timing oscillator carried around the skull as a shared substrate. Cranial bone carries longitudinal waves near 2800 m/s and a head is roughly 0.58 m around, so the skull spans about 0.05 wavelengths at 250 Hz, 0.21 at 1 kHz and 0.83 at 4 kHz. *Closes because:* below roughly 5 kHz the skull is a fraction of one wavelength and moves as a rigid body plus a few low modes, so there is no spatial phase pattern to carry a reference around. The deeper obstacle is independent of the skull: a cochlea integrates everything arriving at each place into one displacement before any neuron sees it, so a second injected channel becomes a spectral pattern at the same place rather than a comparable signal. Binaural hearing works because there are two integrators and a circuit built to difference them.

A.3 The saccule as a second sound channel. The saccule responds to sound and has its own nerve branch, making it a genuine second receptor. *Closes because:* saccular acoustic thresholds sit near 120 dB SPL air-conducted — roughly 60 dB above conversational speech and at the pain threshold. It remains excellent as a diagnostic channel, which is why it appears in §9.

A.4 Direct radiofrequency stimulation. The microwave auditory effect is real: pulsed radiofrequency deposits energy, thermoelastic expansion launches a pressure wave, and the cochlea hears it. The natural next thought is that clever waveform design should reduce the heat while keeping the sound.

Closes because heating and sound are the same quantity in two units. The thermoelastic pressure is $p = \Gamma \cdot A$, where Γ is the Grüneisen parameter and **A is absorbed energy density, not total energy** — and that same density is what raises tissue temperature. Spreading the deposition over ten times the volume therefore divides the peak absorption by ten and the acoustic pressure by ten with it, a 20 dB loss. A phased array genuinely reduces *surface* hot spots by distributing entry paths, but it cannot reduce density at the point where sound is made, because that density is the sound. There is no beam geometry that separates them.

And the waveform is not the obstacle either, though the reason is instructive. The launched wave follows the *rate* of deposition, so synthesising a signal requires the heating history to track its integral. Heat is unipolar — energy can be added, not removed — so that integral must be non-decreasing, which forces the signal non-negative. A unipolar source cannot produce a bipolar wave. The standard escape is to bias: deposit a DC heating level and modulate around it, with the bias at least equal to the signal amplitude. Arbitrary waveforms then become achievable, and the DC heat becomes mandatory. The waveform is free; the bias is not.

The dose that results. Threshold perception is reported near 16 mJ/kg per pulse. Synthesising at Nyquist across a 1 kHz bandwidth requires 2,000 such events per second — about 32 W/kg, against a 2 W/kg limit for localised head exposure in the general public. Sixteen times over, at barely audible loudness; 4 kHz bandwidth gives 128 W/kg, sixty-four times over. Stress confinement, by contrast, is *not* the barrier: 4 kHz needs pulses under 125 μ s, which is straightforward.

Nor does coded transmission rescue it. Pulse compression buys real gain — a 0.1 s code across 4 kHz gives a time-bandwidth product of 400, or 26 dB — but that gain is only realisable by a receiver correlating against the known code. The cochlea is a bank of bandpass filters running continuous spectral analysis, not a correlator, and has no access to the code. There is no coding gain to recover at the ear.

What survives. Every constraint above arises from converting electromagnetic energy into sound by heating. If the energy arrives as sound already, none of them applies. That route is §7.

A.5 Cellular-automaton modelling of the head shadow. *Closes because:* head shadow is acoustic diffraction around a semi-rigid obstacle, with rigid-sphere solutions, boundary-element methods on measured geometry, and direct head-related-transfer-function measurement already validated against measurement. A discrete approximation scheme substitutes a weaker solver into a solved problem.

A.6 Choosing a comb-filter notch ratio for maximal inharmonicity. Where two delayed copies of a signal sum, the resulting comb can in principle be spaced to avoid confusion with the source's own harmonic structure, and irrational ratios that resist rational approximation are the natural candidates. *Closes because:* harmonics lie at integer multiples, so in log frequency they fall with gaps that shrink monotonically — 1.000, 0.585, 0.415, 0.322 — and no constant-ratio comb can track them for any ratio. A log-spaced comb is already poorly aligned with harmonic structure by construction, and the ratio buys nothing there. *What survives:* the real constraint on the ratio is to avoid placing artificial notches in the 6–10 kHz band that already carries monaural elevation and front-back cues.

A.7 Comb geometry, retained for anyone working on bone-conduction crosstalk. A pure delay places notches at linear frequency spacing while the cochlea maps frequency to place logarithmically, so the two never align except at isolated points. Whether notches are perceptually resolved or fuse into coloration depends on their spacing against auditory filter bandwidth: a 0.5 ms delay stays resolved to about 18 kHz, 1 ms to 9 kHz, 2 ms to 4.4 kHz, 5 ms to 1.6 kHz and 10 ms to 700 Hz. Measured transcranial delays are frequency-dependent — near 0.06 ms at 250 Hz, about 1.1 ms near 800 Hz, about 0.6 ms at 4 kHz — so natural crosstalk is already dispersive rather than a clean single-delay comb. On the Greenwood map over 500 Hz to 8 kHz a 1 ms linear comb places successive notches 6.77 mm apart apically and 1.01 mm apart basally, a 6.7-fold variation in place spacing, while a constant-ratio comb at 1.25 varies by 1.23-fold. A linear comb therefore concentrates its effect in one cochlear region; a ratio-spaced comb spreads it evenly and presents the same spatial pattern wherever signal energy sits.

Related work by the same author

The methodological framing used here — reading a system by what it accumulates around a closed path rather than by magnitude — is developed in *Loop Walking: A Field Guide to Knowing Whether You Found Anything* (concept DOI 10.5281/zenodo.21811588) and its translation companion *The Crossover Atlas: One object, many walls* (concept DOI 10.5281/zenodo.22064125). Neither is required to read this paper. §6.1 records where that framing turned out to be already present in cochlear measurement practice rather than imported into it.

References

Retrieved and verified during preparation. Entries marked **[partial]** were identified by title and source but their complete bibliographic record was not obtained; they are listed with enough detail to be found. No author list, volume or year appearing below without that marker was inferred.

Inner and outer hair cell loss, and what emissions do and do not report

1. Trautwein, P.G.; Hofstetter, P.; Wang, J.; Salvi, R.; Nostrand, A. Selective inner hair cell loss does not alter distortion product otoacoustic emissions. *Hearing Research* **1996**, 96, 71–82. doi:10.1016/0378-5955(96)00040-8
2. Hofstetter, P.; Ding, D.; Powers, N.; Salvi, R.J. Quantitative relationship of carboplatin dose to magnitude of inner and outer hair cell loss and the reduction in distortion product otoacoustic emission amplitude in chinchillas. *Hearing Research* **1997**, 112, 199–215. doi:10.1016/S0378-5955(97)00123-8
3. Salvi, R.; Sun, W.; Ding, D.; Chen, G.-D.; Lobarinas, E.; Wang, J.; Radziwon, K.; Auerbach, B.D. Inner hair cell loss disrupts hearing and cochlear function leading to sensory deprivation and enhanced central auditory gain. *Frontiers in Neuroscience* **2017**, 10, 621. doi:10.3389/fnins.2016.00621
4. Lobarinas, E.; Salvi, R.; Ding, D. Selective inner hair cell dysfunction in chinchillas impairs hearing-in-noise in the absence of outer hair cell loss. *Journal of the Association for Research in Otolaryngology* **2016**, 17(2), 89–101. doi:10.1007/s10162-015-0550-8
5. Wake, M.; Anderson, J.; Takeno, S.; Mount, R.J.; Harrison, R.V. Otoacoustic emission amplification after inner hair cell damage. *Acta Oto-Laryngologica* **1996**, 116, 374–381. doi:10.3109/00016489609137860
6. Qiu, C.; Salvi, R.; Ding, D.; Burkard, R. Inner hair cell loss leads to enhanced response amplitudes in auditory cortex of unanesthetized chinchillas: evidence for increased system gain. *Hearing Research* **2000**, 139, 153–171.

Emissions in vestibular schwannoma

7. Ferri, G.G.; Modugno, G.C.; Calbucci, F.; Ceroni, A.R.; Pirodda, A. Hearing loss in vestibular schwannomas: analysis of cochlear function by means of distortion-product otoacoustic emissions. *Auris Nasus Larynx* **2009**, 36(6), 644–648. doi:10.1016/j.anl.2009.02.006 (*the 183-patient series*)
8. Telischi, F. An objective method of analyzing cochlear versus noncochlear patterns of distortion-product otoacoustic emissions in patients with acoustic neuromas. *Laryngoscope* **2000**, 110(4), 553–562. doi:10.1097/00005537-200004000-00007
9. Telischi, F.F.; Roth, J.; Stagner, B.B.; Lonsbury-Martin, B.L.; Balkany, T.J. Patterns of evoked otoacoustic emissions associated with acoustic neuromas. *Laryngoscope* **1995**, 105, 675–682.
10. Kagoya, R.; Shinogami, M.; Kohno, M.; Yamasoba, T. Distortion-product otoacoustic emission tests evaluate cochlear function and differentiate cochlear and vestibular schwannoma. *Otolaryngology-Head and Neck Surgery* **2013**, 148(2), 267–271. doi:10.1177/0194599812469502
11. Perlman, H.B.; Kimura, R.; Fernandez, C. Experiments on temporary obstruction of the internal auditory artery. *Laryngoscope* **1959**, 69, 591–613. (*the vascular mechanism invoked in §6*)

Dead regions and frequency lowering

12. Salorio-Corbetto, M.; Baer, T.; Moore, B.C.J. Comparison of frequency transposition and frequency compression for people with extensive dead regions in the cochlea. *Trends in Hearing* **2019**, 23, 2331216518822206. doi:10.1177/2331216518822206
13. Vinay, Moore, B.C.J. Prevalence of dead regions in subjects with sensorineural hearing loss. *Ear and Hearing* **2007**, 28(2), 231–241. doi:10.1097/AUD.0b013e31803126e2
14. Pepler, A.; Munro, K.J.; Lewis, K.; Kluk, K. Prevalence of cochlear dead regions in new referrals and existing adult hearing aid users. *Ear and Hearing* **2014**, 35(3), e99–e109. doi:10.1097/AUD.0000000000000011
15. Baer, T.; Moore, B.C.J.; Kluk, K. **2002** — origin of the finding that amplification above about 1.7 times the edge frequency of a dead region gives no benefit. **[partial — journal and pages not retrieved]**

Bone conduction, transcranial transmission, and this population

16. Stenfelt, S. Transcranial attenuation of bone-conducted sound when stimulation is at the mastoid and at the bone conduction hearing aid position. *Otology & Neurotology* **2012**, 33, 105–114.
17. Rigato, C.; Reinfeldt, S.; Håkansson, B.; Fredén Jansson, K.-J.; Renvall, E.; Eeg-Olofsson, M. Effect of transducer attachment on vibration transmission and transcranial attenuation for direct drive bone conduction stimulation. *Hearing Research* **2019**, 381, 107763. doi:10.1016/j.heares.2019.06.006
18. Placke, L.; Appelbaum, E.; Patel, A.; Sweeney, A. Bone conduction implants for hearing rehabilitation in skull base tumor patients. *Journal of Neurological Surgery Part B: Skull Base* **2019**, 80(2), 139–148. doi:10.1055/s-0039-1677690
19. Hampton, T.; Milinis, K.; Whitehall, E.; Sharma, S. Association of bone conduction devices for single-sided sensorineural deafness with quality of life: a systematic review and meta-analysis. *JAMA Otolaryngology-Head & Neck Surgery* **2022**, 148(1), 35–42. doi:10.1001/jamaoto.2021.2769
20. Bone-conduction crosstalk and localisation: a 2025 wave-interference model reporting impaired localisation under bilateral bone conduction, and a 2018 study of nineteen single-sided listeners on contralateral devices reporting no localisation benefit. **[partial — full records not retrieved]**

Cochlear implant telemetry

21. Evoked compound action potential telemetry, fielded across major implant systems under manufacturer trade names, and a 2021 proof of concept for intracorporeal cortical response telemetry aimed at adaptive fitting. **[partial — manufacturer literature and primary study not retrieved]**

Companion documents and methodological framing

22. Schulz, M. Loop Walking: A Field Guide to Knowing Whether You Found Anything. Concept DOI 10.5281/zenodo.21811588
 23. Schulz, M. The Crossover Atlas: One object, many walls — a translation companion to the Field Guide. Concept DOI 10.5281/zenodo.22064125
-

Declaration

This page is a standing declaration attached to every preprint in this series. It states what was and was not done, so a reader can judge the work rather than infer its rigour from its tone.

Publication status

This is a preprint. It has not been peer reviewed. It is a checkpoint on a road, published so the reasoning is available and checkable while the work is still in progress, not a finished result presented as settled.

The literature pass is limited and incomplete. Searches were targeted and natural-language, run against general web indices. They did **not** include structured database searching with controlled vocabulary, systematic patent searching by classification code, clinical trial registries, or non-English literature. Patent literature in particular is where ideas of this kind predominantly sit and is where this method is weakest.

A screening pass is not a clearance. Where a claim is described as unclaimed or not found, the correct reading is *not found by these searches*, not *absent from the art*. Anything intended for filing, commercial use, or a novelty claim requires a professional prior-art search first.

Both successes and failures are reported. Approaches evaluated and not pursued are recorded in an appendix with the quantity that closed each one, on the view that a documented dead end is worth more to a reader than an additional proposal.

Authorship and assistance

Prepared by a part-time, unaffiliated research effort, published free to the world. An AI language model (Anthropic Claude) assisted with structural drafting, literature recall, arithmetic and unit checking, and targeted literature search. No automated tool generated any observation from a patient, a device, or an instrument. **The author is responsible for the correctness of all content and citations.**

Provenance of quantities

Every load-bearing number is tagged by how it was obtained. This exists because the characteristic failure of speculative engineering is the inherited constant: a figure taken from a reference propagates through documents unchecked, because each reader assumes an earlier reader verified it.

Two classes are separated, and they warrant different levels of concern. A number is **decision-relevant** if changing it by a factor of two changes whether the work should proceed. It is **design-relevant** if it changes how something is built but not whether. An unverified design-relevant number is untidy; an unverified decision-relevant number can invalidate the work.

Table 7.

tag	meaning
SOURCED	traced to a named primary or authoritative source; count of independent sources given
COMPUTED	derived in this work from SOURCED inputs; the calculation is reproducible from the stated relation
UNVERIFIED	recalled or widely repeated, not checked against a source in preparing this document
ASSUMED	chosen as representative for the purpose of an estimate; not a measurement

Where a quantity is SOURCED to a single source and carries safety consequence, that is noted, because one source is not corroboration.

Table 8. Provenance table — this document.

quantity	tag	source / note
medial olivocochlear efferent anatomy and contralateral suppression assay	SOURCED ×2	efferent pathway reviews; schwannoma suppression studies
reduced suppression reported in acoustic neuroma	SOURCED ×1	contradicted by a 9-patient pilot finding no significant difference — the discrepancy is \$4.2

DPOAE unaffected by selective inner hair cell loss (80–100%)	SOURCED ×3	chinchilla carboplatin series; dose-response study; review of the series
thresholds in quiet near-normal until IHC loss exceeds 80%	SOURCED ×1	same review
click-evoked emissions amplified after IHC damage	SOURCED ×1	single study, different emission type. Not relied upon; §4.1 states why
eCAP telemetry in every major implant system; closed-loop fitting research	SOURCED ×2	manufacturer literature; 2021 intracorporeal cortical telemetry proof of concept
emission amplitude reliability degrades on probe refit	SOURCED ×2	test-retest and inter-system calibration studies
phase-gradient source separation established	SOURCED ×2	wave-fixed / place-fixed separation literature; swept-tone phase study recommending clinical use
bone-conduction crosstalk impairs localisation	SOURCED ×2	2025 interference model; 19-listener contralateral-device study
transcranial attenuation of bone-conducted sound, measured at mastoid and hearing-aid positions	SOURCED ×2	Stenfelt, <i>Otol Neurotol</i> 2012;33:105–114; Rigato et al., <i>Hear Res</i> 2019;381:107763 (direct-drive attachment)
that attenuation <i>continues rising into the ultrasonic range</i> , isolating the contralateral side	UNVERIFIED	the two sources above are audiometric-band measurements and do not establish the ultrasonic extension. The claim is retained only where marked and should not be relied upon
bone conduction devices in skull-base tumour and SSD populations	SOURCED ×4	skull-base rehabilitation review; post-retrosigmoid Baha series; prior-craniotomy implantation series; QoL systematic review and meta-analysis
BCU dynamic range ~30 dB; carrier-independent pitch	SOURCED ×1	bone-conducted ultrasound perception study
ultrasound neuromodulation auditory confound; effect abolished by nerve transection	SOURCED ×3	multiple rodent studies and a deafened-model control paper
electrical space constant, numeric value	UNVERIFIED — DECISION-RELEVANT	per-patient fitted parameter. §7.3 states the comparison as a function of it rather than assuming a value
~8 effective channels	UNVERIFIED	widely repeated; primary study not checked
Greenwood map constants; cochlear length	UNVERIFIED	standard values, not re-checked here
skull wave speed 2800 m/s; head circumference 0.58 m	ASSUMED	used only in Appendix A.2, where the conclusion survives a large error
saccular acoustic threshold ~120 dB	UNVERIFIED	VEMP thresholds are quoted in several unit conventions; the comparison is order-of-magnitude only
Frey-effect threshold 16 mJ/kg per pulse	UNVERIFIED	recalled from the microwave-auditory literature; Appendix A.4's conclusion holds across a wide range of this value
comb resolvability, Greenwood place spacing, radiation pressure, duty-cycle budget	COMPUTED	from the relations stated at each point