


Does It Matter Which Fibre You Lose?

Marta Campi^{1,2}, Elie Partouche², Paul Avan², Tobias Goehring¹

¹ Deep Hearing Lab, University of Zurich, University Hospital Zurich

² CeRIAH, Institut de l'Audition, Université Paris Cité, Institut Pasteur,
Fondation Pour l'Audition, IHU reConnect, Paris, France



Forum Acusticum
Graz, September 2026

USZ Universitäts
Spital Zürich




**University of
Zurich** ^{UZH}



Speech-In-Noise Difficulty with a Normal Audiogram

- **About 1 in 8 adults** report speech-in-noise difficulty with a normal audiogram^[1]
- Hearing thresholds **within 2 dB of each other**
- Speech-in-noise thresholds **spread over 20 dB**




Guest et al. 2018, *Hear Res* 364:142–151, Figs. 2A and 3


One proposed lesion: **cochlear synaptopathy** (*hidden hearing loss* names what the audiogram misses) ^[2,3]

Cochlear Synaptopathy

Loss of synapses between inner hair cells and auditory-nerve fibres



Simulated Auditory Nerve Responses




stimulus: /aa/ · 113 ms · 150 characteristic frequencies · 65 dB SPL

- Healthy auditory nerve: low, medium and high-SR fibres all responding
- Low- and medium-SR fibres are preferentially vulnerable ^[1,2]
- They are high-threshold: the audiogram, tested near threshold, does not see them ^[3]


How Much Response Survives?

21 of 22 studies measure one number



Would that tell us what the listener has lost?

Two ways of degrading speech



How well each speech feature survives

ABR wave I changes by 0–10 % when every low- and medium-rate fibre is gone [1,2,3]

Same amount of information lost. Different features survive[4]

How much response remains does not tell us what information has been lost

[1] Bramhall et al. 2019, Table 1, and five later studies · [2] Bourien et al. 2014 · [3] Budak et al. 2021

What Does Fibre Loss Do to the Neural Representation?

Different phonemes evoke different patterns of auditory-nerve activity^[1]. We measure **how far apart** those patterns stay.

CODE CHANGE

Rescaling all distances between phonemes shrink by a common factor

Rearrangement phoneme responses move relative to each other^[2]

OUTCOME

What a decoder can extract

Do phoneme distinctions survive?^[3]

Research Questions (RQs)

RQ1 Structure

Does fibre loss rescale the neural representation, or rearrange it?

RQ2 Robustness & Behaviour


Does the surviving structure persist across stages and listening conditions, and what does it cost in words?

RQ3 Recoverability

Can a decoder trained on the intact code read the lesioned code and can retraining recover what remains?

Stimuli, Neural Simulations & Experimental Design

SPEECH STIMULI 171 TIMIT sentences^[1] · 28 talkers · 10 070 phoneme tokens in 49 classes 65 dB SPL · quiet and at +10, +5, 0 dB SNR (WHAM!)^[2]



EXPERIMENTAL DESIGN

UNIFORM LOSS the synaptopathy pattern

Low- and medium-SR fibres removed at every CF

6 % 19 % 37.5 %

2/3/10 0/3/10 0/0/10

L/M/H = fibres remaining at each CF
healthy = 3/3/10

FIBRE-TYPE CONTROL matched control

Same number removed, different fibre type

3 medium 3 high 3 M + 3 H 6 high

3/0/10 3/3/7 3/0/7 3/3/4

matched to uniform loss by fibre count

FREQUENCY-SPECIFIC LOSS constructed

Loss confined to one band, then deepened

125 Hz 10 kHz

1 – 3.5 kHz

in band: 3/3/10 → 0/0/10 → 0/0/5 → 0/0/0
first is the control · also at 4–8 and 1.5–2.5 kHz

CALIBRATION not a lesion

Test whether our methods detect artificial structural loss

Silencing
a proportion of channels set to zero

Pooling
channels averaged — resolution lost

applied after simulation · no fibres removed


● low ● medium ● high ○ removed

[1] Garofolo et al. 1993 · [2] Wichern et al. 2019

Methods

Representation geometry

Gromov–Wasserstein (GW) ^[1]



Computed at AN and IC stages


What changes is which pairs are close, not how strong the response is

RQ1 · level, or arrangement

RQ2 · and does the answer hold in noise, at 85 dB, at the midbrain

Decoder


Linear phoneme classifier



RQ3 · the mapping moved, or the distinctions are gone

Synthesis + recognition

NeuroVoc ^[2] + Whisper ^[3]





Computed at the AN stage only

RQ2 · whether the surviving code still supports words

RQ1 Structure

Auditory nerve, 65 dB SPL, quiet · 49 phoneme classes
GW distance between phoneme geometries


L/M/H = fibres remaining at each CF · healthy = 3/3/10




● measured ● with overall firing matched

Take the overall firing out and every **uniform loss/fibre-type control** becomes indistinguishable from chance.
In the **frequency specific-loss** the lesion rearranges

RQ2 Robustness & Behaviour



Only the severest loss is detectable, and even that is a change of scale




The code is rescaled, not rearranged. Recognition degrades, mostly with the listening condition

RQ3 Recoverability

Decoder: a kernel ridge readout of 49 phoneme classes from the auditory nerve. What changes is what it was trained on


UNIFORM LOSS

same proportion removed at every CF



FREQUENCY-SPECIFIC LOSS

fibres removed inside one frequency band



The decoder fails, but the neural code still carries the information

Summary & Conclusions

Uniform fibre loss, up to half the auditory-nerve fibres, **rescales the code without rearranging it.**

Rearrangement emerges only when fibre loss is **frequency-specific.**

The shape of the code is **unchanged across stage, noise and level.** **The recognition cost is not.**

Even where the decoder trained on the healthy ear fails, **the distinctions remain recoverable after retraining.**

In synaptopathy, how much neural response is lost is a poor proxy for how much speech information is lost

Thank you!

marta.campi@uzh.ch

Dr Elie Partouche



Prof Paul Avan




Prof Tobias Goehring



Does the lesion pick on particular sounds?

effect size per phoneme category, with the overall firing matched · rows ordered by the nerve



Weakening the synapse instead of removing the fibre

Synaptopathy is a loss of synapses. This study removes whole fibres. **Does that choice hide a different regime?**

Backup · single tone at 2 kHz, 65 dB SPL, one characteristic frequency, one seed — not the corpus

Two ways to lower the rate

Removing fibres



8 of 16 removed
fewer terms in the sum

Halving release sites



n ■■■■ → ■■

all 16 still there
it just cannot fire as fast

Why this cannot change the shape

$$r = s / (s D + 1)$$

$$D = t_{\text{abs}} + t_{\text{rd}} / n + t_{\text{rel}}$$

D is the dead time after each spike. As the drive s grows, r approaches $1/D$: the rate has a ceiling, and the ceiling is set by n .

$n = 4$: $D = 4.70 \text{ ms} \rightarrow \text{ceiling } 213 \text{ s}^{-1}$

$n = 2$: $D = 8.20 \text{ ms} \rightarrow \text{ceiling } 122 \text{ s}^{-1}$

ratio of ceilings = 0.573

A ceiling is a multiplicative constant.

What the simulation gives

	mean	onset	sust.	sust./on set
$n = 4$	141.7	148.6	141.0	0.9483
$n = 2$	88.1	92.3	87.4	0.9469
ratio	0.622	0.621	0.620	0.9985

spikes per second · onset is the peak over the first 20 ms, sustained the mean over the final fifth

All three quantities fall by the same factor. The ratio of sustained to onset moves by one and a half parts in a thousand.

Weakening the synapse lands in the same regime as removing the fibre: it rescales, it does not rearrange.