

A Mathematical Model for Temporary Threshold Shift (TTS)

Abstract

A significant portion of the population experiences Permanent Threshold Shift (PTS), often due to prolonged exposure to loud environments. This condition is largely preventable if individuals understand the relationship between sound intensity and exposure duration. Standards from NIOSH and OSHA provide specific guidelines for safe exposure levels:

- **OSHA Standard:** A 5 dB increase halves the allowed exposure time (e.g., 2 hours at 100 dBA).
- **NIOSH Standard:** A 3 dB increase halves the allowed exposure time (e.g., 15 minutes at 100 dBA).

Noise dosimeters typically monitor Sound Pressure Level (SPL) and compute a total "dose." However, these devices often assume the ears are fully recovered by the next day, which is an inaccurate assumption. Research indicates that ear recovery can take up to 4000 minutes. Without knowing the actual recovery state, users risk stacking stress on already stressed ears, significantly increasing the risk of PTS.

This project designs a model that monitors ambient SPL to calculate Temporary Threshold Shift (TTS) dynamically, including the recovery phases. This approach allows for a more accurate assessment of hearing fatigue than standard daily resets.

The author of this document, tried to design a model that monitors ambient sound pressure levels (SPL) and calculates the Temporary Threshold Shift (TTS) that results from it. This also includes recovery. Basically, if the SPL goes up, the TTS goes up, and vice versa.

Patuzzi (1998) Research Foundation

This paper contains the research done by Patuzzi R. et al. What Patuzzi did mostly (in the experiments in the paper) is expose his ears at 100dB SPL, 4kHz for a certain amount of time, while tracking the manner in which his TTS rose, and then track the manner in which it fell when he was recovering.

The paper concludes that the manner of rising and falling of TTS is similar across experiments, which means that it must be Mathematically predictable.

The rest of this document will use images, examples and references from that paper.

The TTS Model

We have already established that the curves are similar across experiments. For visual context, here is one diagram and the annotations straight from the paper.




Fig. 1. A: When the mean thresholds over the last 2 min of the 5 min silences are plotted against cumulative exposure duration, the growth of threshold elevation (dB) is seen. Two examples are shown here (open squares for a 3400 min exposure; closed squares for a total of 600 min exposure). Note that the onset time course of the TTS is reproducible, and occurs in two phases, with an initial plateau at about 100 min, followed by a slower rise over 2000 min or so, reaching a second plateau or asymptote for time longer than 2000 min. In this case the asymptotic TTS or ATS is 48 dB. B: When the exposure tone is turned off for a prolonged period (in this case days), the threshold recovers towards normal in a complex manner. If the initial 5 min of recovery is ignored, the data are well described if we assume a multi-exponential reactivation of MET channels (see text). Note that for recovery times more than 300 min, the recovery time course could be described as 'linear with the log of time', as suggested by Ward. C: The data of A on a linear time scale. D: The data of B on a linear time scale. E: The data in A and C on an expanded time scale showing the first 100 min of the growth in TTS. F: The first 100 min of recovery on a linear time scale.

Before we proceed it is important that we talk about the "Synaptic phase" as Patuzzi calls it. It is the recovery phase that takes place within the first 5 minutes of recovery. It can "eat up approximately 20 dB of TTS. Apart from mentioning that fact, the paper doesn't mention any other way that synaptic recovery occurs. Say one is exposed for some time and their TTS rises to 20 dB, does the synaptic phase eat up all of it? Meaning that they can recover the 20dB TTS within the first 5 minutes? Since I did not have information about that, I decided to leave it out of the model entirely.

Methodology and Logic

Intuition: When the SPL falls, TTS falls and vice versa. It is a continuous process and it is something that should be computed dynamically, not with a fixed formula.

For simplicity, I chose to use equations of rate of change of TTS vs TTS. Reason being you eliminate the need for a timer (which would need a reset on every regime change). Some would question its simplicity, but when one uses TTS vs time formulas, the results are approximately the same.

A lookup table for the rate of change of TTS vs. TTS was constructed using data from Patuzzi's research. The process involved several steps:

- Find the formula for the curve in any of the 2 figures. I extracted the points using Web-Plot-Digitizer then used Sci-Py (A library in python for scientific computations) to obtain the formula of the curve.
Onset equation: $TTS(t) = 15.83 * (1 - \exp(-t / 6.44)) + 9.05 * (1 - \exp(-t / 354.1)) + 21.99 * (1 - \exp(-t / 1074.2))$
Recovery equation: $TTS(t) = 11.00 * \exp(-t / 17.93) + 37.62 * \exp(-t / 1334.8)$

The formulas were used to generate a curve to determine its accuracy visually;

<See  #1>
 <See  #2>

- Differentiate the curve to obtain the rate of change of TTS vs time.

$$dTTS_{onset}/dt = (15.83/6.44) \cdot \exp(-t/6.44) + (9.05/354.1) \cdot \exp(-t/354.1) + (21.99/1074.2) \cdot \exp(-t/1074.2)$$
- Using the curves obtained from step 1 and 2, build a lookup table of rate of change of TTS_rate vs TTS, using linear interpolation method.

Note: the rate of onset is always going to be positive, and the rate of recovery is always negative. But that lookup would only be helpful if the SPL always sat at a constant 100dB SPL. Hence, my first set of

assumptions:

- a. There is a boundary where neither onset nor recovery is going to happen. Cross the boundary upwards, onset begins, cross it downwards, recovery begins. I could not figure out what Patuzzi meant by "recovery in silence" thus I made an estimate: 45dB SPL average. And for the boundary, I picked one that seemed reasonable enough: 50dB. The accuracy of this boundary should not matter so much because recovery and onset taking place around this boundary are negligible.
- b. For the onset curve, the higher the dB SPL, the higher the rates are going to be, and vice versa. Since at the boundary (50dB) the rate of onset is zero, I assumed, as the SPL drops from 100dB to 50dB, the rate vs TTS curve slowly obtains a flat shape going towards the horizontal axis.
- c. For the recovery curve, the lower the SPL, the lower the recovery rates and vice versa. As the SPL approaches 50dB, the rates approach zero. I capped the lower boundary of the recovery rates at 20dB since I do not want the model recording an infinite recovery rate when the SPL approaches zero. I picked 20dB because it seems sensible— not backed by any data.

Two linear scaling functions are defined, one for each regime:

$$h(\text{SPL}) = (\text{SPL} - 50) / 50, \text{ for } \text{SPL} \geq 50$$

$$g(\text{SPL}) = (50 - \text{SPL}) / 5, \text{ for } \text{SPL} < 50, \text{ capped at } g = 6 \text{ for } \text{SPL} \leq 20$$

At each timestep, the rate obtained from the lookup table is multiplied by the appropriate scaling function:

$$dTTS/dt \text{ (onset)} = TTS_rate_lookup(TTS) * h(\text{SPL})$$

$$dTTS/dt \text{ (recovery)} = TTS_rate_lookup(TTS) * g(\text{SPL})$$

One more thing: as indicated by the Patuzzi R. paper, the shape of the curves varies with the initial TTS. A recovery from 20dB to zero looks "flatter" and takes way less time than a recovery from 60dB. Here are some of the graphs from the paper that illustrates this.

<See image #3>

Fig. 8. The amount of threshold elevation and the recovery time course vary with the total or cumulative exposure duration. A: With 3400

min of exposure, the initial TTS is large (about 60 dB), and thresholds recover in three clear phases: an initial phase over the first 5 min or

so, a second phase between 5 and 30 min, and a final phase lasting to about 3000 min. The longer-lasting phases in this case are well described by a multi-exponential reactivation of MET channels, with two exponential components with time constants of 30 min and 800 min (solid curve). Note also the 'bounce' in the first 2 min of recovery. B: If the cumulative exposure is only 600 min, the initial TTS is less than in A (about 50 dB), and the second phase of the recovery becomes less distinct, while the long-lasting phase remains. C: If the exposure is only 100min, the second phase of recovery is absent, and the recovery is well described by a single exponential reactivation of MET channels with a time constant of 240 min. D: If the pooled data of Fig. 5B (relative thresholds in the first 5 min of recovery) are plotted in a similar format (black dashes), with a logarithmic time axis, we see that the initial phase of recovery over the first 5 min is common to A, B and C. The open triangles represent un-

normalised threshold recovery after a single 15 min exposure with no prior exposures, while the open circles represent un-normalised data after a 30 min exposure (data not available after 100 min of recovery). The solid curves are multi-exponential fits to the data (ignoring the first 5 min; see Table 1), except for the normalised data (black dashes) where the curve is a cubic polynomial (see Fig. 5B).

As for the onset curve, we do not know if this is the same since the onset experiments from the paper all began from an initial TTS of 0dB. Hence another assumption: The onset rate profile remains the same regardless of what the initial TTS was.

So, how do we scale the recovery curves according to the initial TTS? Patuzzi mentions that the individual components in the multi-exponential curve scale differently.

To capture this, both amplitudes are scaled by L_{anchor} (the L at the moment recovery begins) but with different exponents:

$$A_1(L_{\text{anchor}}) = 11.00 \cdot (L_{\text{anchor}} / 48.62)^1 \quad A_2(L_{\text{anchor}}) = 37.62 \cdot (L_{\text{anchor}} / 48.62)^3$$

The fast amplitude ($\tau \approx 18$ min) scales linearly. The slow amplitude ($\tau \approx 1335$ min) scales cubically. This is what is referred to in the rest of the document as the **(1, 3) cubed-slow rule**.

Reasoning: This problem was presented to Anthropic's Sonnet 4.6 and after "running some tests" it suggested that this was the most appropriate way of scaling the amplitudes in order to best match the paper's curves. It was instructed to use Python for any arithmetic and simulation since it does not have an ALU.

Now, those experiments done by Patuzzi were using a 4KHz monotone. Real world sound contains a wide spectrum of frequencies. Since the cochlear is split into sections each recording its own range of frequencies, I found it best if we could calculate the TTS for all those component regions. I made an assumption that all frequency components affect the ear in the same manner that the 4KHz tests by Patuzzi did. This means one needs to do a fourier transform on a spectrum of sound to extract it's component frequencies. Most frameworks (used in programming languages) like Kiss-FFT give you the spectrum grouped into bins. The one implementing this needs to make a decision about which regions they care about most so that they can tweak the math to provide bins with those regions. If the regions are bigger than the bins, pick the highest value in that range of bins, and use that to calculate dB SPL. I chose to use dB SPL for the TTS calculations since it represents the raw energy of the sound wave.

Also, since frequencies also affect parts of the cochlear that are not the main parts of that frequency, it is necessary that one includes an algorithm for simulating how much a certain frequency affects other parts of the cochlear too. I did not search for this because I was too burned out.

The pseudocode therefore becomes:

INPUTS AT EACH STEP:

```

SPL_CURRENT : CURRENT SPL (dB)

DT : TIME STEP (MIN)

STATE VARIABLES:

TTS : CURRENT TTS (dB)

SPL_ANCHOR : SPL AT LAST REGIME CHANGE

TTS_ANCHOR : L AT LAST REGIME CHANGE

PREV_REGIME : 'ONSET' OR 'RECOVERY'

REC_TABLE : RATE-VS-L LOOKUP FOR CURRENT L_ANCHOR

EACH STEP:

1. DETECT REGIME CHANGE
IF |SPL_CURRENT - SPL_ANCHOR| > 5:
    SPL_ANCHOR <- SPL_CURRENT
    TTS_ANCHOR <- TTS
    REC_TABLE <- REBUILD_RECOVERY_TABLE(TTS_ANCHOR)
    NEW_REGIME <- 'RECOVERY' IF SPL_ANCHOR < 50 ELSE
'ONSET'    PREV_REGIME <- NEW_REGIME

2. COMPUTE RATE
BUILD THE LOOKUP, IF NOT ALREADY BUILT
IF W_ANCHOR >= 50:
    RATE <- ONSET_RATE_AT_TTS(TTS) * H(W_ANCHOR)

3. INTEGRATE
TTS <- MAX(0, TTS + RATE * DT)

```

Tests

The simulation algorithm used is the same as the pseudocode above.

Three tests were run against the original Patuzzi data.

Test 1 – Continuous 100 dB SPL exposure

The model is held at SPL = 100 dB for 4000 minutes and TTS is tracked.

Quantity	Paper	Model
TTS at $t = 100$ min	22 dB	20.01 dB
Final TTS (asymptote)	47 dB	46.34 dB

This test essentially confirms that the fit to Figure 1A was correct.

Test 2 – Steady-state L vs W

The model is held at constant W from 40 to 100 dB until L reaches equilibrium.

SPL (dB)	Equilibrium TTS (dB)
40	0.0
50	20.0 (boundary – see note)
55	36.6
60	42.3
70	45.9
80	46.6
90	46.8
100	46.9

Note on SPL = 50: at the boundary, both onset and recovery rates are exactly zero, so TTS stays at whatever value it had. The model is initialised at TTS = 20 for this test, so equilibrium is 20. This is a known artefact of the sharp boundary – the equilibrium TTS at the boundary is undefined.

Test 2 – Recovery from various TTS_0

The model is initialised at TTS_0 and exposed to SPL = 45 dB (recovery), with the synaptic phase active. Time to recover to TTS = 1 dB is recorded.

TTS₀ (dB)	Paper estimate (min)	Model (min)
15	≈ 150	0.5
22	≈ 1000	664
30	≈ 1000	2870
38	≈ 2500	3792
60	≈ 3000	4921

The middle-range TTS₀ values (22, 30) are within a factor of 3 of paper, which is reasonable for a model fitted to a single recovery dataset. The extremes are off:

- TTS₀ = 15 collapses to ~0 because the synaptic phase alone subtracts about 20 dB. This is a known limitation, flagged in section 8.
- TTS₀ ≥ 38 is slower than the paper – the slow component of the fit ($\tau \approx 1335$ min) has a long tail that the (1, 3) rule does not fully tame.

Test 3 – Intermittent exposure

The model is exposed to a repeating 8 seconds at 100 dB / 2 seconds at 45 dB cycle for 2000 minutes.

12

Configuration	Paper estimate	Model
80% duty cycle at 100 dB	≈ 30 dB	0.0 dB (with synaptic), 17.3

	plateau	dB (without)
--	---------	--------------

The paper reports that intermittent exposure produces a much lower asymptotic TTS than continuous exposure (30 dB at 80% duty cycle vs. 47 dB at 100%). The model reproduces the **direction** of this effect (intermittency lowers asymptote), but under-predicts the magnitude – and the synaptic phase, retriggering every recovery cycle, drives TTS all the way to zero.

Known Limitations and Open Problems

L1 – Intermittency under-prediction

The model does not yet reproduce the paper's intermittency data quantitatively. There are two possible interacting causes:

- Synaptic phase.

In the Patuzzi paper, synaptic recovery happened as the first phase of recovery. But it is not specified whether the synaptic process happens every time when the ear enters recovery trajectory. If we assumed it did, it causes excessive recovery for longer exposure periods

for a very short recovery period. How this process changes with regime change, I do not know. We know for sure there is a synaptic phase during recovery. But during intermittency, it does not seem to be there or rather, it does not seem to affect TTS values significantly. (See Patuzzi R. (1998) fig. 9). I therefore nuked it from the model entirely.

- Current TTS not affecting onset rate.

The model's onset depends only on current TTS (via the rate-vs-TTS lookup) and W . We do not scale the amplitudes like we did with the recovery rate curve. It does not "remember" that TTS was

recently shaken loose by a brief recovery. The paper's data, however, seems to imply that the slow components of onset are particularly sensitive to brief interruptions and do not fully re-engage after a short break. This will be revisited once additional data or expert consultation is available.

L2 – Recovery slower than paper for high L0

For $TTS_0 \geq 38$, recovery times are longer than the paper's estimates by 30– 60%. But again, it might be a blessing in disguise: the model is less likely to tell the user has fully recovered before they actually do so.

L3 – Synaptic phase eats short recoveries

For $TTS_0 < 20$ dB, the synaptic phase (which subtracts ≈ 20 dB) drives TTS to zero almost immediately. This is unrealistic – the paper's short-exposure data shows synaptic recovery proportional to (not capped by) the total TTS. The current implementation is "synaptic always subtracts ~ 20 dB", which over-corrects for small TTS_0 . Again, we do not have the data for when and how synaptic recovery affects recovery profiles.

L4 – Single-frequency model

The Patuzzi data is for 5 kHz exposure and 7 kHz TTS measurement. The model uses the same concept (that were obtained from data picked at 4kHz monotone) for all the sticks (see more about stick number in section 9.1). Different frequencies regions in the cochlear might be affected differently from other regions when they receive the same SPL at their corresponding frequencies.

L5 – No upper cap on $h(SPL)$

The onset scaling $h(W) = (W - 50) / 50$ grows without bound. At $W = 130$, $h = 1.6$, meaning the model predicts 60% more onset than the data ever measured. This is reasonable as an extrapolation (louder sounds do cause more damage) but the model has not been validated for $W > 110$.

L6 – Equilibrium TTS at boundary is undefined. At $SPL = 50$ dB, both rates are exactly zero. TTS stays at whatever value it had. This value changes between different humans as documented experiments suggest.

9. Multi-Frequency Extension

The single-frequency model can be extended to handle real sounds by treating the cochlea as a series of independent "sticks", each tracking a different frequency band.

9.1 Independent band state

Each frequency band tracks its own TTS independently. This means that when a sound source produces energy across multiple frequency bands

simultaneously, each band responds only to its own SPL value, with no knowledge of what is happening in neighbouring bands. In reality, the cochlea does not behave this way – a tone at a given frequency does not excite only its corresponding cochlear location. Due to the mechanics of the travelling wave along the basilar membrane, energy spreads asymmetrically to neighbouring locations, more so toward higher frequencies than lower ones, and this spread widens at higher SPLs. The current model ignores this cross-frequency excitation entirely, meaning it will underestimate TTS at frequency bands adjacent to a dominant sound source, particularly on the high-frequency side. This is a known limitation and is flagged for a future revision.

9.2 Cross-frequency spread

In a real cochlea, a tone at frequency f does not stimulate only its own cochlear location. Energy spreads asymmetrically – wider toward higher frequencies (toward the cochlear base) than toward lower frequencies. The shape and width of this spread are level-dependent, with broader spreads at higher SPLs.

9.3 Frequency-dependent damage susceptibility

The Patuzzi data is for one exposure frequency. Real ears have frequency-dependent sensitivity – the middle ear transmits energy most efficiently around 1–4 kHz, so equal dB SPL at 100 Hz delivers less energy to the cochlea than at 2 kHz. A naive per-stick application of the model will overpredict damage at very low and very high frequencies. This is a known simplification of the initial implementation. This issue may be resolved in the future

Conclusion:

The model performs well on the original paper's continuous-exposure and large-magnitude recovery data. It is known to under-predict intermittency effects and to over-correct for very small TTS values; both issues are

documented as bookmarks and will be addressed in subsequent work.

The intended use is in a consumer application that monitors a user's

sound


environment and informs them about their current hearing fatigue and recovery state – closing the gap left by standard dosimeters, which report

dose without reporting recovery time.


The weaknesses caused by assumptions made in the model will be fixed

Images

1.



2.



3.

