Vinyl Intermodulation Distortion (IMD) Analyzer

Technical Reference & User Guide

`IMDImproved.py` — version 2.0

ACD & Claude Pilotéen

March 16, 2026

Contents

1 Purpose

This tool measures intermodulation distortion (IMD) from two-tone test tracks on the CBS Laboratories STR-112 test record (or compatible records with tone pairs).

Both L and R channels of every WAV file are always analysed independently. The report contains separate sections for each channel, and four plots are generated (one spectrum plot and one bar chart for each direction, with L and R channel panels).

The two recordings serve different cartridge alignment purposes:

- `-fileVertical` — vertically modulated track. Used to optimise VTA (Vertical Tracking Angle) or SRA (Stylus Rake Angle). The vertical modulation excites the stylus in the vertical plane; incorrect VTA/SRA generates excessive IMD in this axis.
- `-fileHorizontal` — laterally (horizontally) modulated track. Used to optimise LTA (Lateral Tracking Angle) or cartridge zenith rotation in the headshell. The lateral modulation excites the stylus in the horizontal plane; incorrect zenith generates excessive IMD in this axis.

The analysis is identical for both files — only the physical meaning of the result differs.

Results are presented as:

- A detailed text report (console + `.txt` file).
- Two spectrum plots (one per direction, each with L and R channel panels).
- Two bar-chart plots (one per direction, each with side-by-side L vs R IMD product comparison).

By default, all reported levels are normalized so that the tallest test tone (f_L or f_H) sits at 0 dB. This removes gain-chain dependence and makes results comparable across different recordings.

2 Version

The script exposes a `__version__` variable at module level (currently `"2.0"`). This can be queried programmatically:

```python
from IMDImproved import __version__
print(__version__)    # "2.0"
```

The version is incremented when meaningful changes are made to the analysis logic, output format, or options file interface.

Version 2.0 changes (relative to 1.0):

- `f_low` and `f_high` now accept 1, 2, or 4 comma-separated values to support per-channel frequencies (Section 9.4).
- Tone discovery uses a two-stage approach: guided (closest peak) then blind (Section 9.5).
- Both L and R channels are always analysed — a tone-location error on one channel does not skip the other (Section 10).
- Per-direction plot frequency range: `plot_freq_max_vertical` and `plot_freq_max_horizontal` (Section 11).
- Error reporting includes per-channel rerun suggestions with the 4-value syntax (Section 16).
- Non-numeric values in the options file are caught and reported via the error file rather than causing an unhandled crash.

3 Background Theory

3.1 Intermodulation Distortion

When two tones at frequencies f_L and f_H ($f_\mathrm{L} < f_\mathrm{H}$) pass through a non-linear system (such as a phono cartridge), the output contains not only the original tones but also new spectral components at frequencies that are integer linear combinations of f_L and f_H [1]. These are called *intermodulation products*.

For a system with a polynomial transfer characteristic of order n [3]:

$$y(t) = a_1\,x(t) + a_2\,x^2(t) + a_3\,x^3(t) + \cdots + a_n\,x^n(t) \tag{1}$$

an input $x(t) = A_L\cos(2\pi f_\mathrm{L} t) + A_H\cos(2\pi f_\mathrm{H} t)$ produces output components at:

$$f_\mathrm{IMD} = |m f_\mathrm{H} \pm k f_\mathrm{L}|, \qquad m, k \in \mathbb{Z}_{\geq 0}, \quad m + k \leq n \tag{2}$$

The *order* of a product is defined as $m + k$. The most important products for cartridge alignment assessment are the 2nd-order pair:

$$f_\mathrm{diff} = f_\mathrm{H} - f_\mathrm{L} \qquad \text{(difference tone)} \tag{3}$$

$$f_\mathrm{sum} = f_\mathrm{H} + f_\mathrm{L} \qquad \text{(sum tone)} \tag{4}$$

3.2 Total IMD

The total IMD is defined as the RSS (root sum of squares) of all individual product amplitudes, expressed as a percentage of the pilot tone (f_H) amplitude [1]:

$$\mathrm{IMD}_\mathrm{total} = \sqrt{\sum_{i=1}^{P} r_i^2} \times 100\% \tag{5}$$

where

$$r_i = 10^{\,D_i/20} \tag{6}$$

is the linear amplitude ratio of the i-th product relative to the pilot tone, D_i is the product level in dBc (dB relative to carrier, i.e. relative to f_H), and P is the total number of measured products.

The corresponding total level in dBc is:

$$\mathrm{IMD}_\mathrm{dBc} = 20\log_{10}\left(\sqrt{\sum_{i=1}^{P} r_i^2}\right) \tag{7}$$

3.3 The CBS STR-112 Two-Tone Test

The CBS Laboratories STR-112 contains two groups of intermodulation test tracks:

Table 1: CBS STR-112 IMD test track groups.

Group	**Pilot (f_H)**	**Test tone (f_L) levels**
Vertical (3 tracks)	4000 Hz at −6 dB	+6, +9, +12 dB above standard
Horizontal (5 tracks)	4000 Hz at −6 dB	+6, +9, +12, +15, +18 dB above standard

The f_L frequency is either 200 Hz or 400 Hz. The high-frequency pilot tone is intentionally recorded at a low level (−6 dB) so that the intermodulation products it generates are not masked by its own amplitude. The low-frequency test tone is recorded at progressively higher levels to increasingly stress the cartridge/stylus system.

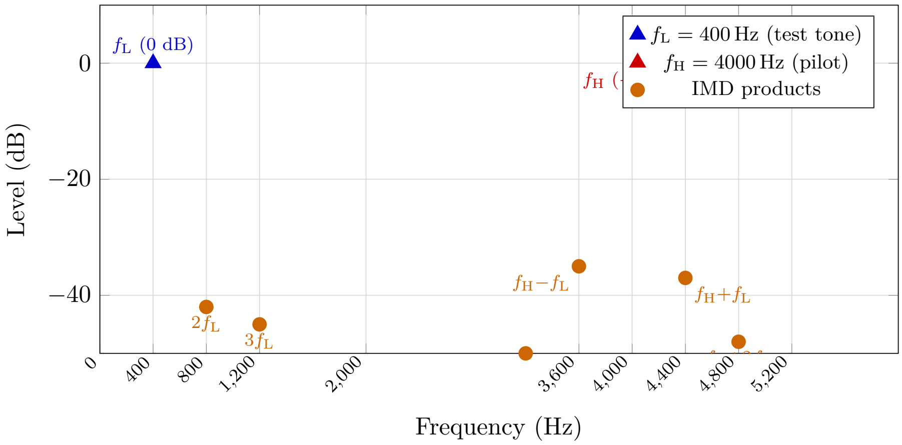

Figure 1: Schematic two-tone IMD spectrum for $f_L = 400$ Hz, $f_H = 4000$ Hz. The test tone (f_L) at 0 dB and the pilot (f_H) at −6 dB are the intended signals. All other peaks are distortion products generated by the non-linear system.

3.4 Why STFT for IMD Measurement?

The test tracks on the CBS STR-112 contain a steady two-tone signal over several seconds. While a single long FFT could in principle capture the entire spectrum, using the Short-Time Fourier Transform (STFT) [3] provides several advantages:

- **Robustness:** Individual STFT windows where the pilot tone falls below the noise floor (e.g. at the start/end of the track or during lead-in groove) are automatically discarded.
- **Statistical aggregation:** The median dBc across all valid windows is computed for each IMD product, which is robust to outliers caused by surface noise, ticks, or pops.
- **Consistency checking:** The per-window standard deviation of each product level serves as a quality indicator.

3.5 The Hann Window and Coherent Gain Correction

Each STFT block is multiplied by a periodic Hann window before the FFT [4]. The Hann window is defined as:

$$w[n] = \frac{1}{2}\left(1 - \cos\left(\frac{2\pi n}{N}\right)\right), \qquad n = 0, 1, \ldots, N-1 \tag{8}$$

where N is the window length in samples and the window is *periodic* (`sym=False` in SciPy).

The corrected single-sided amplitude spectrum is:

$$A[k] = \frac{2\,|\mathrm{FFT}[k]|}{N \cdot G_{\mathrm{coh}}} \tag{9}$$

where the coherent gain is [4]:

$$G_{\mathrm{coh}} = \frac{1}{N}\sum_{n=0}^{N-1} w[n] = 0.5 \quad \text{(for the Hann window)} \tag{10}$$

The DC bin ($k = 0$) and the Nyquist bin ($k = N/2$ for even N) are *not* doubled.

3.6 dBc: Decibels Relative to Carrier

All IMD product levels are expressed in dBc (decibels relative to the carrier) [1], where the "carrier" is the pilot tone f_{H}:

$$D_{\text{product}} = L_{\text{product}} - L_{f_{\mathrm{H}}} \qquad [\text{dBc}] \tag{11}$$

where L_{product} and $L_{f_{\mathrm{H}}}$ are absolute levels in dB. Since all IMD products are weaker than the pilot tone, $D_{\text{product}} < 0\,\text{dBc}$. The more negative the better.

This ratio-based measure is unaffected by the level normalization setting (Section 3.7).

3.7 Level Normalization

By default, all reported *absolute* dB levels (f_{H}, f_{L}, spectrum plot) are shifted so that the tallest test tone (f_{L} or f_{H}, whichever is louder) sits at 0 dB:

$$L_{\text{norm}} = L_{\text{raw}} - L_{\text{ref}}, \qquad L_{\text{ref}} = \max(L_{f_{\mathrm{L}}},\ L_{f_{\mathrm{H}}}) \tag{12}$$

This removes dependence on the recording gain chain and makes results comparable across different recordings and equipment setups. The normalization can be disabled via `normalize_levels = false` in the options file.

Importantly, dBc values and the total IMD percentage are *unaffected* by this setting because they are ratios.

3.8 Speed Error Estimation

The turntable speed error is estimated from the detected pilot tone frequency relative to the expected frequency:

$$\varepsilon = \frac{f_{\mathrm{H,detected}} - f_{\mathrm{H,expected}}}{f_{\mathrm{H,expected}}} \times 100\% \tag{13}$$

A positive value means the turntable is running fast; negative means slow. This value is used to adjust the frequency search tolerance for the STFT analysis. Old records with significant speed error (up to ±5%) are accommodated by a wide initial search tolerance (`freq_tolerance_wide`, default 200 Hz).

3.9 Track Level Auto-Detection

The CBS STR-112 records the test tone f_{L} at several known levels relative to the standard recording level [7]. The analyser auto-detects which track level variant was recorded by comparing the measured $f_{\mathrm{L}}/f_{\mathrm{H}}$ ratio to the known level sets:

$$\Delta_{\text{measured}} = L_{f_{\mathrm{L}}} - L_{f_{\mathrm{H}}} \tag{14}$$

The expected difference for each known level ℓ is:

$$\Delta_{\text{expected}}(\ell) = \ell - L_{f_{\mathrm{H}},\,\text{nominal}} \tag{15}$$

where $L_{f_{\mathrm{H}},\,\text{nominal}} = -6\,\text{dB}$. The track level is matched to the known level that minimises $|\Delta_{\text{measured}} - \Delta_{\text{expected}}|$. A match distance greater than 3 dB is flagged as "uncertain".

4 IMD Products Measured

The analyser measures the following intermodulation products (default up to 5th order). Table 2 uses $f_L = 400$ Hz, $f_H = 4000$ Hz as an example.

Table 2: IMD products measured (example: $f_L = 400$ Hz, $f_H = 4000$ Hz, up to 5th order).

Label	**Frequency**	**Order**	**Notes**
$f_H - 1f_L$	3600 Hz	2	Difference tone
$f_H + 1f_L$	4400 Hz	2	Sum tone
$f_H - 2f_L$	3200 Hz	3	
$f_H + 2f_L$	4800 Hz	3	
$f_H - 3f_L$	2800 Hz	4	
$f_H + 3f_L$	5200 Hz	4	
$f_H - 4f_L$	2400 Hz	5	
$f_H + 4f_L$	5600 Hz	5	
$2f_L$	800 Hz	2	2nd harmonic of f_L
$3f_L$	1200 Hz	3	3rd harmonic of f_L
$4f_L$	1600 Hz	4	4th harmonic of f_L
$5f_L$	2000 Hz	5	5th harmonic of f_L
$2f_H$	8000 Hz	2	2nd harmonic of f_H
$3f_H$	12000 Hz	3	3rd harmonic of f_H

The 2nd-order difference and sum products ($f_H \pm f_L$) are the most important for cartridge alignment assessment [1]. Products with frequency ≤ 0 or above the Nyquist frequency [3] are excluded.

The maximum order can be configured via `max_imd_order` in the options file (allowed range: 2–10, default: 5).

5 Algorithm

Algorithm 1 describes the complete STFT-based IMD measurement procedure for a single WAV file and a single channel.

Algorithm 1 STFT-based IMD measurement for one WAV file and one channel.

Require: WAV file, channel (L or R), direction (vertical or horizontal)
Ensure: IMD product levels (dBc) and total IMD percentage

Read WAV, convert to float64 in $[-1, 1]$
Check L/R channel identity; warn if channels differ
Extract selected channel signal $\mathbf{s}$
▷ — Two-stage tone-pair discovery —
$\mathbf{b} \leftarrow$ first $\sim 2\,\mathrm{s}$ of $\mathbf{s}$
Compute spectrum of $\mathbf{b}$ (Hann window + coherent gain)
Estimate noise floor (median of spectrum in dB)
Find all prominent peaks $\geq$ `tone_prominence_db` above noise floor
if no prominent peaks found **then**
    raise AnalysisError
end if
▷ Stage 1: Guided — pick peak closest to expected frequency
for each expected tone (f_{L}, f_{H}) for this direction and channel **do**
    Search prominent peaks within ±`freq_tolerance_wide`
    Select the peak *closest in frequency* (not loudest)
end for
if both guided peaks found **then**
    Accept unconditionally; warn if deviation > `tone_acceptance_tolerance`
    goto STFT analysis
end if
▷ Stage 2: Blind — fallback to two loudest peaks
if fewer than 2 prominent peaks **then**
    raise AnalysisError
end if
Assign the two strongest peaks as f_{L} (lower) and f_{H} (higher)
Check deviations from expected `f_low` and `f_high`
if either deviation > `tone_acceptance_tolerance` **then**
    Log per-channel suggestion; **raise** ToneLocationError
end if
▷ — STFT analysis —
$\varepsilon \leftarrow (f_{\mathrm{H,detected}} - f_{\mathrm{H,expected}})/f_{\mathrm{H,expected}}$ ▷ Speed error
Compute adjusted STFT tolerance from speed error
$N_w \leftarrow \max(\lfloor T_w \cdot f_s \rfloor,\ 2048)$ ▷ $T_w =$ `stft_window_sec`
$H \leftarrow \max(\lfloor (1-O) \cdot N_w \rfloor,\ 1)$ ▷ $O =$ `stft_overlap`
Build list of IMD product definitions from f_{L}, f_{H}
$p \leftarrow 0$
while $p + N_w \leq$ total samples **do**
    Compute spectrum of $\mathbf{s}[p : p + N_w]$
    Find f_{H} peak within ± adjusted tolerance
    if f_{H} not found or f_{H} level < `min_pilot_db` **then**
        $p \leftarrow p + H$; **continue**
    end if
    Find f_{L} peak within ± adjusted tolerance
    for each IMD product definition $(f_{\mathrm{exp}}, \mathrm{label}, \mathrm{order})$ **do**
        Find peak near f_{exp} within ± tolerance
        Record dBc $= L_{\mathrm{product}} - L_{f_{\mathrm{H}}}$
    end for
    $p \leftarrow p + H$
end while
▷ — Aggregation —
Compute median dBc for each product across all valid windows
Compute total IMD (Eq. 5)
Auto-detect track level variant (Section 3.9)
Optionally normalize levels (Section 3.7)

Both L and R channels are analysed independently for each WAV file, producing separate result sets. A tone-location error on one channel does not prevent analysis of the other channel.

6 Signal Chain and Measurement Setup

Figure 2: Signal chain from vinyl playback to IMD analysis.

6.1 Recording Procedure

1. Connect turntable → phono preamp → ADC (audio interface).
2. Play one of the CBS STR-112 *vertical modulation* IMD tracks and record to a stereo WAV file (`-fileVertical`).
3. Play one of the CBS STR-112 *horizontal (lateral) modulation* IMD tracks and record to a stereo WAV file (`-fileHorizontal`).
4. Ensure recording levels are not clipping.
5. Sample rates of 44.1 kHz, 48 kHz, or 96 kHz all work well.

Each WAV file should contain a single test track (or a concatenation of same-level tracks). The CBS STR-112 IMD tracks are mono signals recorded identically on both channels; the analyser verifies this and warns if the channels differ.

6.2 Recommended Test Records

- CBS Laboratories STR-112 [7]
- CBS STR-100 (earlier version with similar tone pairs)
- Any compatible record using a 200/400 Hz + 4000 Hz two-tone test signal

7 Analysis Mode

The options-file key `mode` controls which analyses are performed and which file arguments are required:

Table 3: Analysis mode values.

Value	**Required files**	**Description**
`BOTH`	`-fileVertical` and `-fileHorizontal`	Both analyses (default).
`VERTICAL`	`-fileVertical` only	VTA/SRA optimisation only.
`HORIZONTAL`	`-fileHorizontal` only	LTA/zenith optimisation only.

8 Usage

8.1 Command Line

```
python IMDImproved.py
       -fileVertical = <vertical.wav>
       -fileHorizontal = <horizontal.wav>
       [-equipment = <equipment.txt>]
       [-options = <imd_analysis_options.txt>]
```

8.2 Arguments

Table 4: Command-line arguments.

Argument	Required	Description
`-fileVertical = <file>`	Depends on mode	WAV from a vertical-modulation track (VTA/SRA).
`-fileHorizontal = <file>`	Depends on mode	WAV from a lateral-modulation track (LTA/zenith).
`-equipment = <file>`	No	Equipment info file (shown on all plots).
`-options = <file>`	No	Options file (overrides built-in defaults).

Ordering rules:

- File arguments may appear in any order.
- When `mode = BOTH`, both `-fileVertical` and `-fileHorizontal` are required but neither needs to come first.
- When `mode = VERTICAL` or `mode = HORIZONTAL`, only the relevant file argument is required; the other is ignored if given.
- `-equipment` and `-options` may appear in any order.
- Spaces around = are allowed and ignored.
- Both `-fileHorizontal` and `-fileHorisontal` spellings are accepted.

8.3 Examples

```
# Both directions (default mode = BOTH)
python IMDImproved.py
       -fileVertical = vert_track.wav
       -fileHorizontal = lat_track.wav

# With equipment info and custom options
python IMDImproved.py
       -fileVertical = vert_track.wav
       -fileHorizontal = lat_track.wav
       -equipment = EquipmentMeasured.txt
       -options = imd_analysis_options.txt

# Vertical only (mode = VERTICAL in options file)
python IMDImproved.py
       -fileVertical = vert_track.wav
       -options = imd_analysis_options.txt

# Horizontal only (mode = HORIZONTAL in options file)
python IMDImproved.py
       -fileHorizontal = lat_track.wav
       -options = imd_analysis_options.txt
```

9 Options File

The `-options` argument specifies a plain text file containing analysis parameters. If `-options` is not given on the command line, all built-in defaults are used.

9.1 File Format

- Lines starting with `#` are comments and are ignored.
- Blank lines are ignored.
- Format: `key = value` (spaces around `=` are allowed).
- Keys are case-insensitive.
- All options are individually optional — if not present, the built-in default is used.
- **Inline comments are not supported.** Do not place `# text` after a value on the same line. Use a separate comment line instead.

9.2 Available Options

Table 5 lists all available options.

Table 5: Options file keys, defaults, and descriptions.

Key	Default	Description
`mode`	`BOTH`	Analysis mode
`f_high`	4000	Pilot tone freq (Hz); 1, 2, or 4 values
`fh_nominal_db`	−6	Nominal pilot level (dB)
`f_low`	200	Test tone freq (Hz); 1, 2, or 4 values
`tone_prominence_db`	20	Min tone-to-noise separation (dB)
`tone_acceptance_tolerance`	50	Max deviation from expected (Hz)
`track_level_vertical`	auto	Override vertical track level (dB)
`track_level_horizontal`	auto	Override horizontal track level (dB)
`freq_tolerance`	30	STFT peak search tolerance (Hz)
`freq_tolerance_wide`	200	Initial scan wide tolerance (Hz)
`stft_window_sec`	0.2	STFT window length (seconds)
`stft_overlap`	0.75	STFT overlap fraction
`min_pilot_db`	−40.0	Min pilot level to accept window
`db_floor`	−100.0	Lowest dB value reported
`channel`	L	(legacy, overridden internally)
`channel_identity_threshold`	0.001	L/R identity check threshold
`normalize_levels`	`true`	Normalize tallest tone to 0 dB
`max_imd_order`	5	Maximum product order (2–10)
`plot_freq_max_vertical`	12000	Upper freq for vertical plots (Hz)
`plot_freq_max_horizontal`	12000	Upper freq for horizontal plots (Hz)
`plot_freq_max`	—	Legacy: sets both directions (Hz)

9.3 Example Options File

A default options file named `imd_analysis_options.txt` is provided with all options documented:

```
# IMD Analyzer Options File
# Format:  key = value

mode = BOTH
f_high = 4000
fh_nominal_db = -6
f_low = 200
tone_prominence_db = 20
tone_acceptance_tolerance = 50
freq_tolerance = 30
```

```
freq_tolerance_wide = 200
stft_window_sec = 0.2
stft_overlap = 0.75
min_pilot_db = -40.0
db_floor = -100.0
channel = L
channel_identity_threshold = 0.001
normalize_levels = true
max_imd_order = 5
plot_freq_max_vertical = 12000
plot_freq_max_horizontal = 12000
```

9.4 The `f_low` and `f_high` Options

These options accept 1, 2, or 4 comma-separated values to support per-direction and per-channel expected frequencies:

Table 6: `f_low` and `f_high` value formats.

Count	Format	Meaning
1	`value`	Used for all directions and both channels.
2	`vert, horiz`	Vertical and horizontal; L and R get the same value.
4	`vertL, vertR, horizL, horizR`	Per-direction, per-channel.

Examples:

```
f_low = 200                    # same for everything
f_low = 60, 200                # vert=60, horiz=200 (L=R)
f_low = 60, 60, 200, 200      # vertL=60, vertR=60, horizL=200, horizR=200
```

The 4-value syntax is needed when the L and R channels of a WAV file have different prominent frequencies (e.g. on old worn records where channel crosstalk or groove damage shifts a peak). The error file includes 4-value suggestions when L and R differ (Section 16).

9.5 Tone Discovery

The analyser uses a two-stage approach:

Stage 1: Guided

For each expected frequency (f_L, f_H), the script searches all prominent peaks within ±`freq_tolerance_wide` and selects the one *closest in frequency* to the expected value (not the loudest). This ensures that when the user specifies frequencies (especially after a rerun suggestion), the correct peak is selected even if a stronger peak exists nearby at a different frequency.

If both peaks are found, they are accepted unconditionally. If the deviation from the expected frequency exceeds `tone_acceptance_tolerance`, a warning is printed but the analysis proceeds.

Stage 2: Blind (fallback)

If the guided search fails to find one or both peaks, the script falls back to picking the two loudest prominent peaks. If those are far from the expected frequencies, a `ToneLocationError` is raised with per-channel rerun suggestions.

9.6 Track Level Overrides

The `track_level_vertical` and `track_level_horizontal` options allow manually specifying the f_{L} level variant (in dB above the standard recording level) when auto-detection is unreliable. Leave them commented out for auto-detection (default).

10 Channel Analysis

Both L and R channels are always analysed independently. The CBS STR-112 IMD test tracks are mono signals recorded identically on both channels. The analyser verifies this by computing:

$$\mathrm{RMS}_{\mathrm{diff}} = \sqrt{\frac{1}{N}\sum_{n=0}^{N-1}\left(s_L[n] - s_R[n]\right)^2} \tag{16}$$

If $\mathrm{RMS}_{\mathrm{diff}}$ exceeds the `channel_identity_threshold` (default 0.001, i.e. 0.1% of full scale), a warning is issued. The `channel` option in the options file is accepted for backward compatibility but is overridden internally — both channels are always processed.

Important (v2.0): A tone-location error on one channel (e.g. L) does *not* prevent analysis of the other channel (R). This is necessary because old or worn records may have different prominent frequencies on L and R. Each channel's expected frequencies are looked up independently from the `f_low` and `f_high` settings.

11 Per-Direction Plot Frequency Range

The spectrum plot frequency range can be set independently for the vertical and horizontal analyses:

```
plot_freq_max_vertical   = 4000
plot_freq_max_horizontal = 12000
```

This is useful when the vertical file uses low-frequency tones (e.g. $f_{\mathrm{H}} = 949$ Hz) where a 12000 Hz plot range would leave most of the spectrum empty, while the horizontal file uses higher tones (e.g. $f_{\mathrm{H}} = 4000$ Hz) that benefit from the full range.

The per-direction value also controls the upper frequency limit for tone discovery — peaks beyond the plot range are ignored.

The legacy key `plot_freq_max` sets both directions to the same value. If per-direction keys are also present, they take precedence.

12 Equipment Information File

The `-equipment` argument specifies a plain text file whose contents are displayed on all plot images (left-aligned below the chart area). The file is plain text, free-form. If not given, no equipment info is shown and no error is produced.

The text is rendered in monospace font (fontsize 7.5), colour `#444444`, left-aligned below the chart panel.

13 File Date Information

The analysis run date and the modification dates of all input files (WAV, options, equipment) are recorded in two places:

- The text report — in a `FILE DATES` section near the end, after the `SUMMARY`.

- All plot images — right-aligned below the chart area, opposite the equipment info block. The format is:

```
Analysis run:  2026-03-01 19:32:28
Vertical WAV:  vert_track.wav  [2026-02-28 14:05]
Horizontal WAV:  lat_track.wav  [2026-02-28 14:12]
Options:  imd_analysis_options.txt  [2026-02-25 10:00]
Equipment:  EquipmentMeasured.txt  [2026-02-20 09:30]
```

Only files actually provided on the command line are listed. If a file's modification date cannot be determined, "(unknown)" is shown.

The date text is rendered in monospace font (fontsize 7.5), colour `#444444`, right-aligned below the chart area. The figure height is dynamically increased to accommodate the taller of the two footer text blocks.

14 Input File Requirements

Each file should be a stereo (2-channel) WAV [6] containing a single two-tone IMD test track. Supported sample formats:

Table 7: Supported WAV sample formats.

Format	**NumPy dtype**
16-bit PCM	`int16`
24/32-bit PCM	`int32`
32-bit float	`float32`
64-bit float	`float64`

Integer formats are normalised to $[-1, 1]$ by dividing by the maximum integer value (2^{15} for 16-bit, 2^{31} for 32-bit). Float formats are used as-is.

15 Output Files

Up to five files are written to the current working directory on success:

Table 8: Output files.

Filename	**Contents**
`imd_report_<base>.txt`	Full text report (identical to console output).
`imd_plot_<base>_vertical_spectrum.png`	Spectrum plot for the vertical direction (L + R panels).
`imd_plot_<base>_horizontal_spectrum.png`	Spectrum plot for the horizontal direction (L + R panels).
`imd_plot_<base>_vertical_LR_comparison.png`	Bar chart for the vertical direction (L vs R). Suffix is `_products` if only one channel succeeded.
`imd_plot_<base>_horizontal_LR_comparison.png`	Bar chart for the horizontal direction (L vs R). Suffix is `_products` if only one channel succeeded.

Here `<base>` is derived from the `-fileVertical` filename (or `-fileHorizontal` if only that file is provided). When `mode = VERTICAL` or `mode = HORIZONTAL`, only the relevant direction's plots are generated.

16 Error Reporting

16.1 Usage Errors

If the script is invoked incorrectly (missing arguments, unknown arguments, files not found, invalid option values), the error message is written to:

`imd_usage_error_report.txt`

in the current working directory, in addition to being printed to the console.

Non-numeric values in the options file (e.g. from inline comments such as `f_high = 4000 # pilot`) are caught and reported via `_die()` with a clear message, rather than causing an unhandled `ValueError` crash.

16.2 Analysis Errors

If individual file analyses fail (e.g. no tones found, pilot tone too weak), the error details are written to:

`imd_analysis_errors.txt`

The remaining analyses continue and produce partial results. This file is deleted on a fully successful run. Each error entry records the WAV filename, channel, and a diagnostic message explaining why the analysis failed.

16.3 Rerun Suggestions

When tone-location errors occur, the error file includes a **SUGGESTED OPTIONS-FILE VALUES FOR RERUN** section. If L and R channels have different prominent frequencies, the suggestions include:

- Per-channel values (for two separate runs, applying one channel's frequencies to both L and R).
- The 4-value syntax (for a single run with independent per-channel frequencies).

Example output when L and R differ:

```
NOTE: L and R channels have different prominent
frequencies.  You may need two separate runs,
or use the 4-value syntax (see below).

Vertical file, channel L:
  f_low  = 60
  f_high = 949

Vertical file, channel R:
  f_low  = 60
  f_high = 1039

Or use the 4-value syntax for a single run
(verticalL, verticalR, horizontalL, horizontalR):

  f_low  = 60, 60, 200, 200
  f_high = 949, 1039, 4002, 4002
```

17 Text Report Contents

17.1 Header

- Date and time
- Test record identification
- Pilot tone frequencies and nominal level (per direction, per channel)
- Test tone frequencies (per direction, per channel)
- Tone prominence threshold
- Acceptance tolerance
- Vertical and horizontal track level sets
- Analysis mode and channels
- Normalize levels setting
- STFT window size and overlap
- Frequency tolerance (normal and wide)
- Maximum IMD order

17.2 Per-Direction, Per-Channel Section

Each direction (vertical, horizontal) has separate L and R sections. The vertical sections are headed:

"`VERTICAL MODULATION - VTA / SRA Assessment [LEFT channel]`"

The horizontal sections are headed:

"`HORIZONTAL (LATERAL) MODULATION - LTA / Zenith Assessment [LEFT channel]`"

Each section contains:

- Source file path
- Guidance on what to adjust (VTA/SRA or LTA/zenith)
- Sample rate, duration, channel identity status
- Analysis channel
- Number of valid STFT windows
- Level normalization offset (if enabled)
- Detected f_H and f_L frequencies and levels
- Speed error
- Track level (auto-detected or from options file)
- IMD product table: label, frequency, level (dBc), order, window count
- Total IMD percentage and dBc

17.3 Summary

- One line per direction/channel with total IMD
- Guidance: "Lower IMD values are better"
- Adjustment advice for each direction analysed
- Warnings for uncertain track levels or channel mismatches

17.4 File Dates

- Analysis run timestamp
- Modification dates of all input files

18 Plot Descriptions

18.1 Spectrum Plots (one per direction)

Each spectrum plot contains one or two panels (L and R channels):

- **Title:** "IMD Spectrum Analysis [VERTICAL/HORIZONTAL] [normalized]"
- **Upper panel:** L channel.
- **Lower panel:** R channel.

Table 9: Spectrum plot visual elements.

Element	Style
Average spectrum	gray, linewidth = 0.5, $\alpha = 0.7$
f_H marker	red dashed vertical, triangle marker
f_L marker	blue dashed vertical, triangle marker
IMD product markers	colour-coded dotted verticals, triangle markers
Total IMD annotation	wheat-background text box, upper right

Axes:

- Horizontal: frequency, 0 to `plot_freq_max_vertical` or `plot_freq_max_horizontal` (default 12000 Hz each). Ticks at: 200, 500, 1k, 2k, 4k, 6k, 8k, 10k, 12k (ticks beyond the range are automatically excluded).
- Vertical: magnitude in dB (normalized or raw).
- Grid: major ticks only, $\alpha = 0.3$.

Footer below the chart area:

- Left side: equipment information (if provided).
- Right side: file dates and analysis run timestamp.

18.2 Bar Charts (one per direction)

Each bar chart shows the IMD product levels side by side for easy comparison between L and R channels:

Table 10: Bar chart visual elements.

Element	Style
LEFT channel bars	royalblue, $\alpha = 0.8$
RIGHT channel bars	crimson, $\alpha = 0.8$
Summary box	lightyellow background, monospace font

Axes:

- Horizontal: IMD product labels (rotated 45°).
- Vertical: level in dBc relative to pilot tone.
- Grid: y-axis only, $\alpha = 0.3$.

When only one channel is available, a single set of bars is shown and the chart title reflects the single channel. The filename suffix changes from `_LR_comparison` to `_products`.

Footer: identical to the spectrum plots.

19 Factors Affecting Vinyl IMD

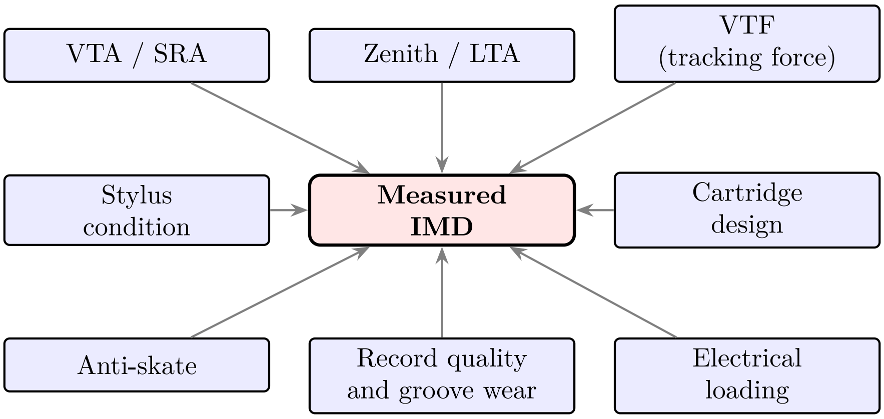

Figure 3: Factors contributing to measured IMD. VTA/SRA primarily affects the vertical result; zenith/LTA primarily affects the horizontal result.

- **VTA/SRA** — the vertical tracking angle or stylus rake angle determines how the stylus contacts the groove walls in the vertical plane. Incorrect VTA/SRA generates excessive IMD in the vertical modulation test [2].
- **Zenith/LTA** — the lateral tracking angle or cartridge zenith rotation affects horizontal-plane tracking. Incorrect zenith generates excessive IMD in the horizontal modulation test.
- **VTF (Tracking Force)** — insufficient or excessive downforce causes mistracking and increased distortion.
- **Stylus condition** — worn or damaged styli cannot trace the groove accurately, increasing all forms of distortion.
- **Cartridge design** — the generator linearity and compliance characteristics vary by cartridge type.
- **Anti-skate** — incorrect anti-skate force causes asymmetric tracking behaviour.
- **Record quality** — groove wear and pressing defects contribute to measured distortion.
- **Electrical loading** — impedance and capacitance mismatch can affect the cartridge's mechanical behaviour through electromagnetic damping.

20 Processing Pipeline Overview

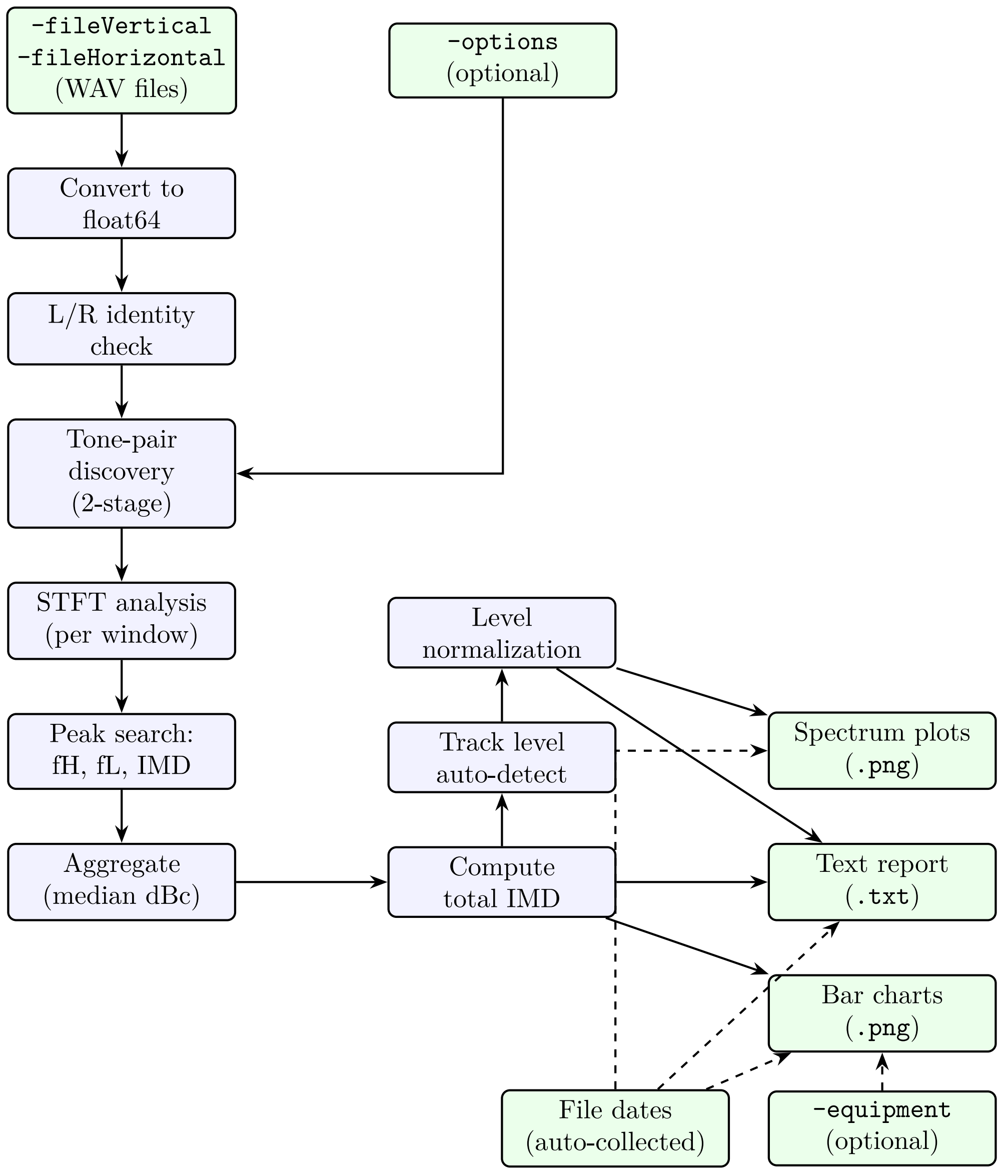

Figure 4: Complete processing pipeline. Each WAV file is processed on both L and R channels independently. The optional equipment file and auto-collected file dates feed into all output files.

21 Internal Constants

All constants listed in Table 11 are built-in defaults that can be overridden by the options file.

Table 11: Internal constants (defaults, overridable via options file).

Constant / Option key	**Default**	**Description**
`mode`	`BOTH`	Analysis mode
`f_high`	4000	Pilot tone freq (Hz); 1/2/4 values
`fh_nominal_db`	−6 dB	Nominal pilot level
`f_low`	200	Test tone freq (Hz); 1/2/4 values
`tone_prominence_db`	20 dB	Min tone-to-noise separation
`tone_acceptance_tolerance`	50 Hz	Max deviation from expected
`freq_tolerance`	30 Hz	STFT peak search tolerance
`freq_tolerance_wide`	200 Hz	Initial scan wide tolerance
`stft_window_sec`	0.2 s	Window length (200 ms)
`stft_overlap`	0.75	Overlap fraction (75%)
`min_pilot_db`	−40.0 dB	Min pilot level per window
`db_floor`	−100.0 dB	Lowest dB value reported
`max_imd_order`	5	Maximum product order
`plot_freq_max_vertical`	12000 Hz	Upper freq for vert plots
`plot_freq_max_horizontal`	12000 Hz	Upper freq for horiz plots
`normalize_levels`	`true`	Normalize tallest tone to 0 dB
`channel_identity_threshold`	0.001	L/R identity check threshold

22 Dependencies

Table 12: Python dependencies.

Package	**Modules used**
Python ≥ 3.8	`sys`, `os`, `datetime`
`numpy`	Array operations, FFT
`scipy`	`scipy.io.wavfile`, `scipy.signal.windows`
`matplotlib`	`matplotlib.pyplot`, `matplotlib.ticker`

Install with:

```
pip install numpy scipy matplotlib
```

A Derivation: Total IMD from Individual Product Levels

The total IMD combines all product amplitudes by root-sum-of-squares in the linear domain [1]. Given P products with levels $D_1, D_2, \ldots, D_P$ in dBc:

1. Convert each to a linear ratio: $r_i = 10^{D_i/20}$.
2. Compute the RSS: $r_{\text{total}} = \sqrt{\sum_{i=1}^{P} r_i^2}$.
3. Express as a percentage: $\text{IMD\%} = r_{\text{total}} \times 100$.
4. Express in dBc: $\text{IMD}_{\text{dBc}} = 20 \log_{10}(r_{\text{total}})$.

This is equivalent to computing the total *power* of all distortion products relative to the pilot tone power, then taking the square root to return to amplitude. It is the standard method used by audio test equipment [1].

B Derivation: Coherent Gain of the Hann Window

The coherent gain of a window is [4]:

$$G_{\text{coh}} = \frac{\sum_{n=0}^{N-1} w[n]}{N} \tag{17}$$

For the periodic Hann window $w[n] = \frac{1}{2}(1 - \cos(2\pi n/N))$, the sum evaluates to $N/2$, giving $G_{\text{coh}} = 0.5$.

Without this correction, a full-scale sine would appear at −6 dB instead of 0 dB. The analyser divides the raw FFT magnitudes by $N \cdot G_{\text{coh}}$ (and multiplies by 2 for single-sided) to restore correct absolute levels (Eq. 9).

The DC bin ($k = 0$) and the Nyquist bin ($k = N/2$, for even N) are exempt from the factor-of-2 doubling because they have no negative-frequency mirror.

C IMD Product Generation Algorithm

The list of IMD products is generated algorithmically from f_L, f_H, and the maximum order $N_{\max}$:

Algorithm 2 IMD product list generation.

Require: f_L, f_H, $N_{\max}$
Ensure: List of (label, frequency, order)
for $n = 1$ **to** $N_{\max} - 1$ **do**
  $f_- \leftarrow f_H - n \cdot f_L$
  $f_+ \leftarrow f_H + n \cdot f_L$
  order $\leftarrow n + 1$
  if $f_- > 0$ **then**
    Add (“fH−nfL”, f_-, order)
  end if
  Add (“fH+nfL”, f_+, order)
end for
for $n = 2$ **to** $N_{\max}$ **do**
  Add (“nfL”, $n \cdot f_L$, n) ▷ Harmonics of f_L
end for
for $n = 2$ **to** 3 **do**
  Add (“nfH”, $n \cdot f_H$, n) ▷ Harmonics of f_H
end for
return product list

D Command-Line Argument Parsing Details

The argument parser handles flexible spacing around `=` by processing the shell-tokenised `sys.argv` list, identically to the cross-talk analyser. The supported forms are:

Table 13: Supported argument spacing variants.

Command-line form	Shell tokens
`-fileVertical=track.wav`	1 token
`-fileVertical =track.wav`	2 tokens
`-fileVertical= track.wav`	2 tokens
`-fileVertical = track.wav`	3 tokens

Argument keys are normalised to lowercase with hyphens and underscores removed before matching. Both `-fileHorizontal` and `-fileHorisontal` are accepted.

E Nomenclature, Abbreviations, and Key Quantities

This section defines the technical terms, abbreviations, symbols, and physical quantities used throughout this document. It is intended as a reference for readers who are not familiar with vinyl playback technology or audio signal analysis.

E.1 Abbreviations

Table 14: Abbreviations used in this document.

Abbreviation	Meaning
ADC	Analogue-to-Digital Converter. A device that converts a continuous (analogue) audio signal from the phono preamp into a discrete digital representation (samples) stored in a WAV file [3].
CBS	Columbia Broadcasting System. CBS Laboratories, the research division of CBS, produced the STR-112 test record used by this analyser.
DAW	Digital Audio Workstation. Software used to record the analogue signal as a WAV file.
DFT / FFT	Discrete Fourier Transform / Fast Fourier Transform. The FFT is an efficient algorithm for computing the DFT, which decomposes a discrete-time signal into its frequency components [3, 5].
IMD	Intermodulation Distortion. Non-linear distortion that generates new spectral components at frequencies that are sums and differences of the input frequencies. Defined in IEC 60268 [1].
LTA	Lateral Tracking Angle. The angle between the cantilever and the tangent to the groove in the horizontal (lateral) plane. When the cartridge is rotated about its vertical axis (i.e. zenith is misaligned), the LTA deviates from the ideal and generates excess lateral-plane IMD.
PCM	Pulse Code Modulation. The standard method of representing sampled audio as integer values [3].
RIAA	Recording Industry Association of America. The RIAA equalisation curve defines the standard pre-emphasis and de-emphasis applied to vinyl records and phono preamps.
RMS	Root Mean Square. A statistical measure of the magnitude of a varying quantity, used here to quantify the difference between L and R channels.
RSS	Root Sum of Squares. The method used to combine individual IMD product amplitudes into a single total IMD figure [1].
SRA	Stylus Rake Angle. The angle between the stylus contact edges and the record surface, measured in the vertical plane. SRA directly affects vertical-plane tracking accuracy and is closely related to VTA.
STFT	Short-Time Fourier Transform. A technique that computes the DFT of successive overlapping segments (windows) of a signal, providing a time–frequency representation [3, 4].
VTA	Vertical Tracking Angle. The angle between the cantilever's axis of motion and the record surface, measured in the vertical plane. The playback VTA should match the recording cutter head's vertical angle (typically 20°–22°) to minimise IMD [2].
VTF	Vertical Tracking Force. The downward force (in grams or millinewtons) applied by the tonearm to keep the stylus in the groove. Incorrect VTF causes mistracking and increases all forms of distortion.
WAV	Waveform Audio File Format. A standard uncompressed audio container format using PCM encoding [6].

E.2 Symbols and Physical Quantities

Table 15: Symbols and physical quantities.

Symbol	Definition
f_H	High-frequency *pilot tone* (Hz). On the CBS STR-112 this is 4000 Hz recorded at −6 dB relative to the standard recording level [7]. All IMD product levels are expressed relative to this tone (in dBc).
f_L	Low-frequency *test tone* (Hz). On the CBS STR-112 this is either 200 Hz or 400 Hz, recorded at various elevated levels (+6 to +18 dB) to progressively stress the cartridge [7].
$A[k]$	Corrected single-sided amplitude spectrum at frequency bin k (Eq. 9). Obtained by windowing with a Hann window and applying coherent-gain correction [4, 3].
D_i	Level of the i-th IMD product in dBc (decibels relative to the pilot tone f_H). Defined in Section 3.6.
G_{coh}	Coherent gain of the analysis window (Eq. 10). For the Hann window, $G_{coh} = 0.5$ [4].
N	Window length in samples. Determined by the `stft_window_sec` option and the sample rate f_s.
f_s	Sample rate of the WAV file (Hz). Common values: 44 100, 48 000, 96 000 Hz.
H	Hop size (stride) in samples between successive STFT windows. $H = N \times (1 - O)$, where O is the overlap fraction.
$w[n]$	Hann window function applied to each STFT block before computing the FFT. Defined in Section 3.5 [4].
ε	Turntable speed error (%), estimated from the detected f_H frequency vs. its expected value. See Section 3.8.

E.3 Units

Table 16: Units of measurement.

Unit	Definition
dB (decibel)	A logarithmic unit for expressing the ratio of two amplitudes or powers. For amplitudes: $L = 20 \log_{10}(A/A_{ref})$. In this document, absolute dB levels are referenced to digital full-scale (0 dBFS) unless stated otherwise [3].
dBc	Decibels relative to carrier. Here the "carrier" is the pilot tone f_H. A product at −40 dBc is 40 dB below the pilot tone. See Section 3.6 and [1].
dBFS	Decibels relative to digital full-scale. 0 dBFS corresponds to the maximum amplitude representable in the digital format (e.g. ±32,768 for 16-bit PCM).
Hz (hertz)	Unit of frequency; one cycle per second.
% (percent)	The total IMD is expressed as a percentage of the pilot tone amplitude: IMD% $= r_{total} \times 100$. See Section 3.2.

E.4 Audio and Vinyl Playback Concepts

The following concepts are central to understanding the measurements performed by this tool.

Phono cartridge
An electromechanical transducer mounted at the end of the tonearm. The stylus (needle) rides in the record groove and its vibrations are converted to an electrical signal by a magnetic or piezoelectric generator.

Groove modulation (vertical vs. lateral)
A stereo vinyl groove carries two independent signals encoded as motion of the groove walls at ±45° to the record surface. Pure *vertical modulation* moves the stylus up and down; pure *lateral modulation* moves it side to side. The CBS STR-112 provides separate tracks with each modulation type to isolate vertical-plane and horizontal-plane distortion mechanisms [7].

Two-tone IMD test
A standard method for measuring intermodulation distortion. Two tones at frequencies f_{L} and f_{H} ($f_{\mathrm{L}} \ll f_{\mathrm{H}}$) are applied simultaneously. Any non-linearity in the system under test generates new spectral components at $mf_{\mathrm{H}} \pm kf_{\mathrm{L}}$. The levels of these products, relative to the pilot tone, quantify the system's non-linearity. The CBS STR-112 uses a two-tone method ($f_{\mathrm{L}} = 200$ or $400\,\mathrm{Hz}$, $f_{\mathrm{H}} = 4000\,\mathrm{Hz}$) as described in [1].

Zenith (cartridge rotation)
The rotational alignment of the phono cartridge body about its vertical axis in the headshell. When zenith is misaligned, the stylus contact edges are not tangent to the groove walls in the lateral plane, producing excess lateral-plane IMD.

Anti-skate
A compensating force applied by the tonearm to counteract the inward skating force that results from stylus friction in the groove. Incorrect anti-skate biases the stylus toward one groove wall, causing asymmetric distortion between left and right channels.

Hann window
A raised-cosine window function applied to each STFT block before the FFT to reduce spectral leakage. It provides good frequency resolution with moderate side-lobe suppression. See Section 3.5 and [4].

Coherent gain
The amplitude scaling factor introduced by a window function. Multiplying a signal by a window reduces its apparent amplitude; the coherent gain correction restores correct absolute levels. See Section 3.5 and Appendix B, and [4].

Nyquist frequency
Half the sample rate ($f_s/2$). No frequency component above the Nyquist frequency can be represented without aliasing in a sampled signal [3]. IMD products above this limit are excluded from the analysis.

References

[1] International Electrotechnical Commission, "IEC 60268-3: Sound system equipment – Part 3: Amplifiers," IEC, Geneva, Switzerland, 2018. *Defines the standard method for measuring intermodulation distortion using two-tone (SMPTE/DIN) and multi-tone signals, including the RSS combination of product amplitudes and the definition of dBc.*

[2] International Electrotechnical Commission, "IEC 60098: Analogue audio disk records and reproducing equipment," IEC, Geneva, Switzerland, 1987. *Defines the standard vertical tracking angle (20°) for stereo disk records and the geometry of the ±45° stereo groove.*

[3] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 3rd ed. Upper Saddle River, NJ: Pearson, 2010. *Standard reference for the DFT, FFT, STFT, windowing, spectral analysis, and the Nyquist sampling theorem. Provides the mathematical foundations for the spectral analysis techniques used in this tool.*

[4] F. J. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform," *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51–83, Jan. 1978. *Definitive survey of window functions (including Hann), their spectral properties, coherent gain, and scalloping loss. The coherent gain correction used in this analyser follows the methodology described here.*

[5] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, Apr. 1965. *The original paper describing the Fast Fourier Transform (FFT) algorithm, which reduces the DFT computation from* $O(N^2)$ *to* $O(N \log N)$. *The FFT is the core of the spectral analysis performed in every STFT window.*

[6] Microsoft Corporation and IBM Corporation, "Multimedia Programming Interface and Data Specifications 1.0," Microsoft, Redmond, WA, 1991. *Defines the RIFF/WAVE file format (WAV) used as the input format for this analyser. Covers PCM encoding at 8, 16, 24, and 32 bits per sample, as well as IEEE 754 floating-point variants.*

[7] CBS Laboratories, "STR-112: Standard Test Record – Broadcast and Recording," CBS Inc., Stamford, CT, 1978. *The test record containing intermodulation distortion test tracks used as the primary reference for this analyser.*