

White Paper on the Harmonic Series Metrics

Version 1 - January 2023

Document number IARO23-1

Change History

Version 1 is derived from IARO21-3, which dealt with peak harmonic prominence below 5 Hz. This new paper extends this concept to include peak and total harmonic prominence as well as harmonic sound-pressure-level (SPL) metrics.

Preamble

Harmonic series are rare in nature. They are far more commonly associated from human activity. This paper looks at several measures that can be calculated from harmonic series, more specifically those that can be calculated from the frequency spectrogram of a recording.

Two separate classes of metrics are considered; those that deal with the SPL of the series and those that deal with the prominence of the series above the sound background.

The definition of prominence for these metrics comes from the Matlabⁱ function ‘findpeaks,’ which returns a list of peaks from (in this case) the 1/36th-octave, narrow-band-filter frequency spectrum of a sound file. The prominence of these peaks is defined as part of this functionⁱⁱ and reproduced in the appendix.

Definition

A qualifying harmonic series is a harmonic series within a 1/36th-octave, narrow-band-filter, frequency spectrum in dBZ units (also known as dBLin).

The Peak Harmonic Prominence, P_{peak} , metric for a qualifying harmonic series is the maximum prominence of all the harmonic peaks, including the fundamental.

The Peak Harmonic Prominence (0.5–5 Hz), P_{harm} , metric for a qualifying harmonic series is defined as the maximum prominence of all harmonic peaks between 0.5 Hz and 5 Hz (including the fundamental).

The Total Harmonic Prominence, P_{tot} , metric for a qualifying harmonic series is the energy-summed total of the prominences of all the harmonic peaks, including the fundamental.

The Peak Harmonic SPL, H_{peak} , metric for a qualifying harmonic series is the maximum SPL of all the harmonic peaks, including the fundamental.

The Total Harmonic SPL, H_{tot} , metric for a qualifying harmonic series is the energy-summed total of the SPLs of all the harmonic peaks, including the fundamental.

The metric for a recording is the maximum metric of all qualifying harmonic series.

Discussion

Frequency Domain vs Time Domain

The major harmonic series of interest are those resulting from a train of pulses, as seen in wind-turbine acoustic signature (WTAS), for example. The interest here is the size of the pulses. This could be measured from pulse heights in the time domain but for the effects of the background noise. The heights of the pulses are often distorted.

In the frequency domain, however, the peaks representing the components at the harmonic frequencies are much easier to measure since they are averaged over the entire integration period and isolated by their frequencies from much of the underlying noise. The total energy in the harmonic series can be established by performing an energy-sum of all the peaks in the harmonic series.

The down side to this approach is that it ignores all the phase information, essentially finding the area of the pulse in the time domain, and not its height.

Sound Pressure Level

The sound pressure level (SPL) of a sound is used to establish its energy content. This can be used to infer the effect of this energy on some object. For example, it can be used to identify the damage that is caused to the human hearing system due to excessively large energy content.

The assertion is that the energy content, in itself, is the primary cause of any effect. This may not be correct. Taking the hearing damage example again, it might be considered that it is not the energy content per se but the maximum bending of the hair cells in the cochlea that is the primary cause, in which case a metric based on the maximum amplitude of the time-series signal would be more appropriate.

Prominence

Traditional measures relating to acoustics rely on the absolute value of sound pressure level (SPL) as they are more interested in such concepts as hearing thresholds and noise-induced hearing loss. When the focus switches to concepts such as annoyance, or 'special audible characteristics' it is the SPL above the background noise that becomes important, i.e., the prominence. This is because in human hearing, provided an acoustic feature is audible, it is the prominence that determines how clearly it can be distinguished.




Figure 1: Prominence of a series of harmonic peak (blue) above the background (green).

Identifying the prominence above a background normally involves measuring the SPL when only the background is present and then when the feature of interest is present. The assumption that the background remains constant is invalid unless factors such as wind speed or rainfall can be accounted for.

When a narrow-band frequency spectrum (or Fourier transform) is taken of a sound signal, more traditional signal-processing techniques can be applied to find the prominence of specific frequencies above the rest of the background noise. The 'findpeaks' function in Matlab is used to find the peaks (see Appendix 1). The prominences returned by this function are significantly too high in some cases and so a further procedure is needed to then determine their prominence (see Addendum 1).

Fight-or-flight

The fight-or-flight response occurs when a stimulus is interpreted by the human or animal hindbrain as a threat. It causes a sudden release of cortisol that prepares the body for instant and strenuous action. Repeated or prolonged triggering of this response can lead to a deterioration in the health of the organism.

Entrainment

Entrainment in the biological sense refers to the predilection of the body to synchronisesynchronize some of its functions to an external rhythm. In particular, the heart rate can become synchronisesynchronised with an external rhythm that is a reasonably close match or the harmonic of a reasonably close match.

Such entrainment has a negative effect upon the body and can lead to heart disorders.

Summation

When measures related to SPL are to be summed or averaged, the summation is normally performed on the energy of the measures. This is in accord with primary interest of acoustics on

concepts such as hearing thresholds and noise induced hearing loss, which are related to the energy. In these it is the energy of the sound which creates the damage to the cochlea that causes the latter.

For all other matters the summation and averaging is customarily done using the deciBel scale as this relates to the perception of sound, i.e., that doublings of energy are perceived as linear increases in deciBel levels. Thus every increase of 6 dB doubles the energy but is perceived as an incremental increase in loudness.

Since the concept of prominence is most closely associated with perception, and many responses are due to perception, this is the approach taken here.

Note that the maximum deciBel level can also be relevant.

Frequency analysis

The magnitude of the prominences will vary with the width of the narrow-band filters used. Specifically, the height of the peaks will remain the same but the level of the background between them will fall as the width of the filters is reduced, below about 12th octave.

Between 1/36th-octave and 1/48th-octave filters, however, the largest prominences will increase by roughly 5%, with a similar difference in the P_{harm} measure (see the following two figures). The application of 1/36th-octave filters has been a standard for analysis by the IARO members for some years, since this bandwidth forms a suitable balance between adequate frequency resolution at higher frequencies and adequate temporal resolution at lower frequencies. The prominence analysis used for the P_{harm} measure is always performed on a spectrum created using 1/36th-octave filters.




Figure 2: Harmonic series of peaks analysed with narrow-band filters of various widths.




Figure 3: The measured size of the prominences in Figure 2 for various widths of narrow-band filters.

Implementation

All metrics

The metrics defined in this white paper can be constructed from a recording by the following operations.

For all of the metrics defined here;

1. Construct a Class 0 filter bank of narrowband filters according to either ANSI[®] S1.11-2004 or IEC 61260:1995, with 1/36-octave bandwidths. The filter bank should extend from 0.5 Hz to at least 100 Hz.
2. The recording is filtered using the filter bank.
3. The filtered series are converted to an RMS average for each band over the length of the recording and then to decibels to produce a frequency spectrum.
4. The peaks in the frequency spectrum are identified using the 'findpeaks' function in Matlab, returning the position of the peaks and their SPL values.
5. The prominences are then determined for each peak by locating the nearest valley on each side of a peak and using the higher one as the background level. (see Addendum 1) Only peaks more than 1 dB above the background are considered.
6. A harmonic analysis is carried out on these peaks to identify any harmonic series. This is achieved by starting with a peak and seeking peaks at 2, 3, 4 etc. times the frequency of the peak. The tolerance on the frequencies is slightly larger than the width of the narrow-band filter responsible, i.e., as the frequencies increase the tolerance becomes larger.
7. Harmonic series with fundamental frequencies that are exact multiples of any other harmonic series are eliminated.

Peak harmonic prominence, P_{peak}

1. Include any harmonic series that has a) one prominence greater than 3 dB and four or more peaks, or b) two prominences greater than 3 dB and three or more peaks. The second-largest prominence must be at least half the size of the largest prominence.
2. The P_{peak} value for each qualifying harmonic series is the size of the largest of the included prominences.
3. The maximum P_{peak} value from all the qualifying harmonic series is recorded as the P_{peak} for the recording.
4. For a specific noise source that has a harmonic series, P_{peak} is the P_{peak} value for that harmonic series. For instance, the P_{peak} for a wind turbine will be the P_{peak} of the harmonic series that has the blade-pass frequency as its fundamental frequency, referred to as the blade-pass P_{peak} .

Peak harmonic prominence (0.5–5 Hz), P_{harm}

1. All peaks that lie outside of the 0.5–5 Hz range in the spectrogram are removed from consideration.
2. Include any harmonic series that has a) one prominence greater than 3 dB and four or more peaks, or b) two prominences greater than 3 dB and three or more peaks. The second-largest prominence must be at least half the size of the largest prominence.
3. The P_{harm} value for each qualifying harmonic series is the size of the largest of the included prominences.

4. The maximum P_{harm} value from all the qualifying harmonic series is recorded as the P_{harm} for the recording.
5. For a specific noise source that has a harmonic series, P_{harm} is the P_{harm} value for that harmonic series. For instance, the P_{harm} for a wind turbine will be the P_{harm} of the harmonic series that has the blade-pass frequency as its fundamental frequency, referred to as the blade-pass P_{harm} .

Total harmonic prominence, P_{tot}

1. Include any harmonic series that has a) one prominence greater than 3 dB and four or more peaks, or b) two prominences greater than 3 dB and three or more peaks. The second-largest prominence must be at least half the size of the largest prominence.
2. The P_{tot} value for each qualifying harmonic series is the energy-sum of all the prominences.
3. The maximum P_{tot} value from all the qualifying harmonic series is recorded as the P_{tot} for the recording.
4. For a specific noise source that has a harmonic series, P_{tot} is the P_{tot} value for that harmonic series. For instance, the P_{tot} for a wind turbine will be the P_{tot} of the harmonic series that has the blade-pass frequency as its fundamental frequency, referred to as the blade-pass P_{tot} .

Peak harmonic SPL, H_{peak}

1. Include any harmonic series that has a) one prominence greater than 3 dB and four or more peaks, or b) two prominences greater than 3 dB and three or more peaks. The second-largest prominence must be at least half the size of the largest prominence. (Note: The peaks must still be distinguished from the surrounding noise, even if the prominence levels are not used in the final calculation.)
2. The H_{peak} value for each qualifying harmonic series is the maximum SPL value of all the peaks.
3. The maximum H_{peak} value from all the qualifying harmonic series is recorded as the H_{peak} for the recording.
4. For a specific noise source that has a harmonic series, H_{peak} is the H_{peak} value for that harmonic series. For instance, the H_{peak} for a wind turbine will be the H_{peak} of the harmonic series that has the blade-pass frequency as its fundamental frequency, referred to as the blade-pass H_{peak} .

Total harmonic SPL, H_{tot}

1. Include any harmonic series that has a) one prominence greater than 3 dB and four or more peaks, or b) two prominences greater than 3 dB and three or more peaks. The second-largest prominence must be at least half the size of the largest prominence. (Note: The peaks must still be distinguished from the surrounding noise, even if the prominence levels are not used in the final calculation.)
2. The H_{tot} value for each qualifying harmonic series is the energy-sum of SPL values of all the peaks.
3. The maximum H_{tot} value from all the qualifying harmonic series is recorded as the H_{tot} for the recording.
4. For a specific noise source that has a harmonic series, H_{tot} is the H_{tot} value for that harmonic series. For instance, the H_{tot} for a wind turbine will be the H_{tot} of the harmonic series that

has the blade-pass frequency as its fundamental frequency, referred to as the blade-pass H_{tot} .

The ACHE software suite can be used to create these metrics using the output of its harmonic analysis. This stores the identified harmonic series, along with their prominences, in the /Derived/Data directory in files whose filenames end in "Harmonic Data - xxxx.csv," where "xxxx" is the name of the channel, usually "Blue" or "Red."

Fundamental-Frequency Plots

Important harmonic series can be identified by plotting one of these metrics for each 10-minute recording against its fundamental frequency as well as the frequency histogram of the fundamental frequency (i.e., how often it occurs). (See the following figure for P_{harm} .)

The strong vertical lines in the top plot and the matching tall lines in the lower plot indicate those fundamental frequencies with a) large P_{harm} values and b) frequent occurrence. The discrete nature of the narrow-band filters can be seen in the upper plot and the twinned lines in both indicate that the fundamental frequencies varied slightly.

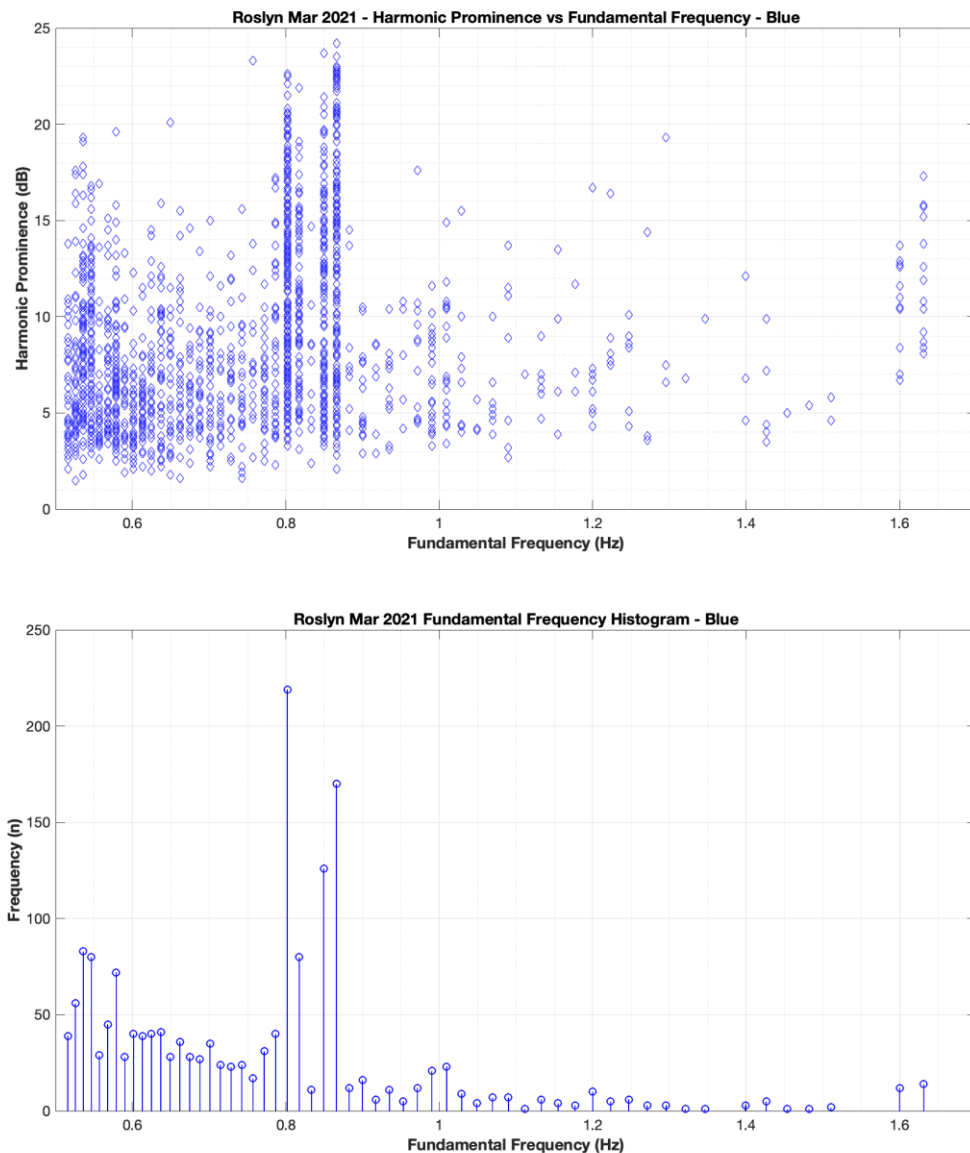


Figure 4: P_{harm} of multiple 10-minute recordings is plotted against the fundamental frequency of the P_{harm} harmonic series (top) and the frequency histogram of the fundamental frequencies (bottom).

By observation of the twin plots the two most important fundamental frequencies are seen to be roughly 0.8 Hz and 0.85 Hz. A weighted average could be calculated for each, e.g., between 0.80 and 0.81 Hz for the lower fundamental frequency. In reality, a value halfway in between is sufficient to pick up instances of both. The twinned peaks, can be narrowed down to the fundamental frequencies of 0.81 and 0.86 Hz, respectively.

Time-of-day plot

Plotting any of these metrics as a time-of-day vs day plot can reveal times of the day when prominences are particularly small or large. The following diagram is a P_{harm} time-of-day plot. Here each horizontal bar is the P_{harm} value of a 10-minute recording, the colour showing the value, as indicated by the legend on the right. Each vertical column is one day of the recording.

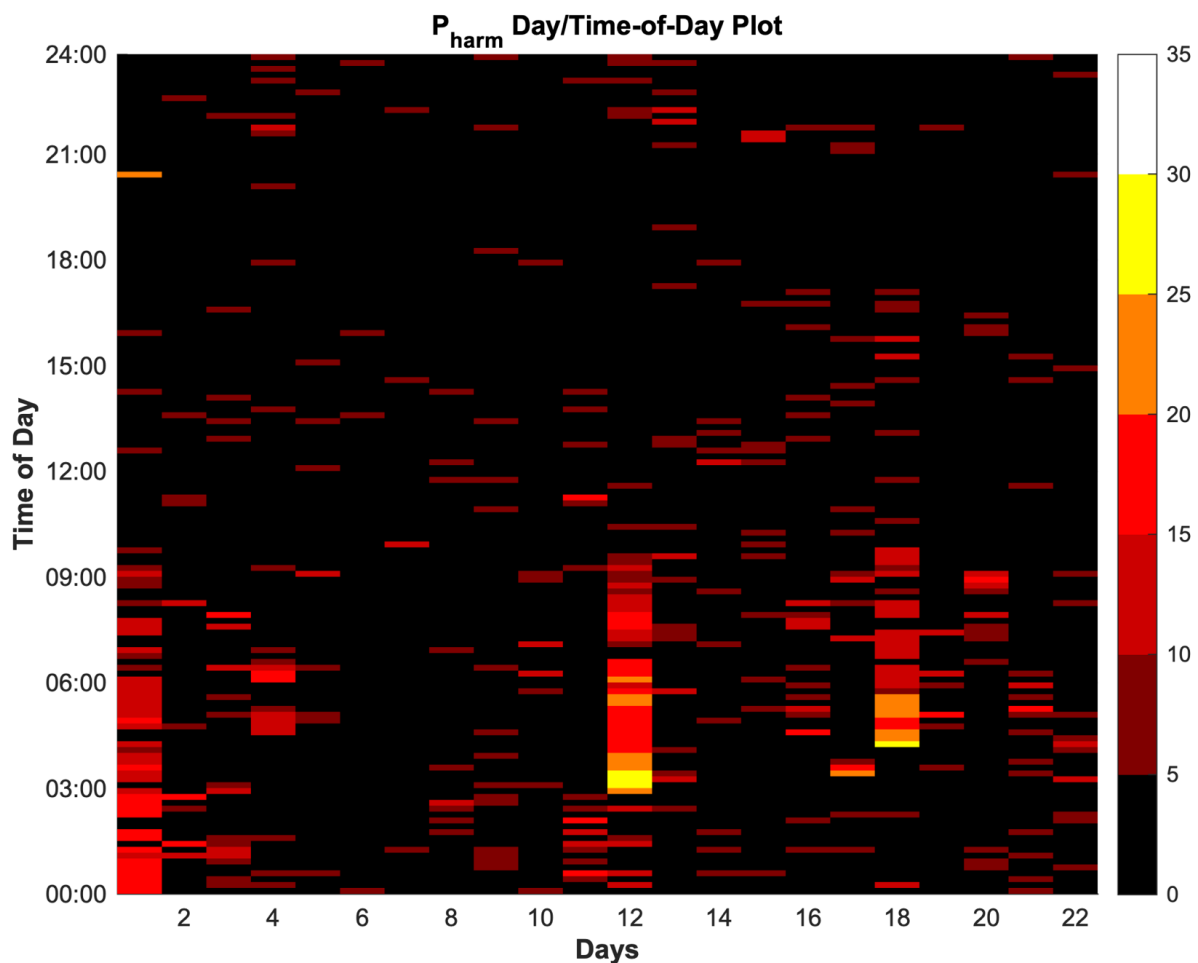


Figure 4: A harmonic prominence time-of-day plot with days horizontally, time of day vertically and the harmonic prominence value in colour according to the colour legend on the right.

Wind Rose

The information obtained for any of these metrics, from a number of successive recordings, can be compactly represented on a wind rose, analogous to a normal wind-speed wind rose but representing any of these metrics, instead. This can be constructed using the following operations;

1. Recordings are analysed to calculate the chosen metric. (These may be ‘regularly spaced,’ referring to the fact that, due to the onerous computational requirements of narrow-band

filtering, it may be more effective to only process one or two 10-minute recordings per hour, usually the 0–10-minute and 30–40-minute recordings.)

2. Each of these values is matched with the concurrent wind speed and direction at the site or nearby weather station.
3. Bins are created for the combination of wind direction and metric level, e.g., bins for the wind direction in 5° increments and the metric in 5-decibel increments.
4. The values are placed into their appropriate bins and tallied.
5. The prominence wind rose is then plotted as a stacked polar histogram where the axial direction denotes wind direction and the radial direction denotes the number of 10-minute recordings. Thus, the length of each 5-deciBel sector denotes the number of 10-minute recordings, while the colour denotes the P_{harm} level.

Such a P_{harm} wind rose is shown below.




Figure 6: A harmonic prominence wind rose where the direction of a sector indicates the wind direction, the length of each section of the sector indicates the number of 10-minute intervals and the colour indicates the level of the harmonic prominence as shown in the legend.

Appendix - Matlab prominence definition


Prominence

The *prominence* of a peak measures how much the peak stands out due to its intrinsic height and its location relative to other peaks. A low isolated peak can be more prominent than one that is higher but is an otherwise unremarkable member of a tall range.

To measure the prominence of a peak:

1. Place a marker on the peak.
2. Extend a horizontal line from the peak to the left and right until the line does one of the following:
 - Crosses the signal because there is a higher peak
 - Reaches the left or right end of the signal
3. ?
4. Find the minimum of the signal in each of the two intervals defined in Step 2. This point is either a valley or one of the signal endpoints.
5. The higher of the two interval minima specifies the reference level. The height of the peak above this level is its prominence.

findpeaks makes no assumption about the behavior of the signal beyond its endpoints, whatever their height. This is reflected in Steps 2 and 4 and often affects the value of the reference level. Consider for example the peaks of this signal:




Peak Number	Left Interval Lies Between Peak and	Right Interval Lies Between Peak and	Lowest Point on the Left Interval	Lowest Point on the Right Interval	Reference Level (Highest Minimum)
1	Left end	Crossing due to peak 2	Left endpoint	a	a
2	Left end	Right end	Left endpoint	h	Left endpoint
3	Crossing due to peak 2	Crossing due to peak 4	b	c	c
4	Crossing due to peak 2	Crossing due to peak 6	b	d	b
5	Crossing due to peak 4	Crossing due to peak 6	d	e	e
6	Crossing due to peak 2	Right end	d	h	d

7	Crossing due to peak 6	Crossing due to peak 8	f	g	g
8	Crossing due to peak 6	Right end	f	h	f
9	Crossing due to peak 8	Crossing due to right endpoint	h	i	i

Addendum 1 - Changes to Finding Peaks

Issues with Matlab findpeaks function

In July of 2021 it was noticed that the `findpeaks` function in Matlab overestimates the highest prominence in a periodogram by as much as 5 dB. This can be seen in the following figure, where the periodogram is in blue, the peaks are shown in red and the baseline is shown in black. The prominence is the distance between the red line and the black line at each peak.



The highest peak is second from the left and clearly shows that the baseline has dropped well below the base of the peak. This is a consequence of the algorithm used by the `findpeaks` function, which considers the baseline of the highest peak to be the larger of the minimum value to the left and to the right of the peak.

The P_{harm} of this 10-minute recording is given as 30.8 dB by the `findpeaks` function. From inspection it should be closer to 24 or 25 dB.

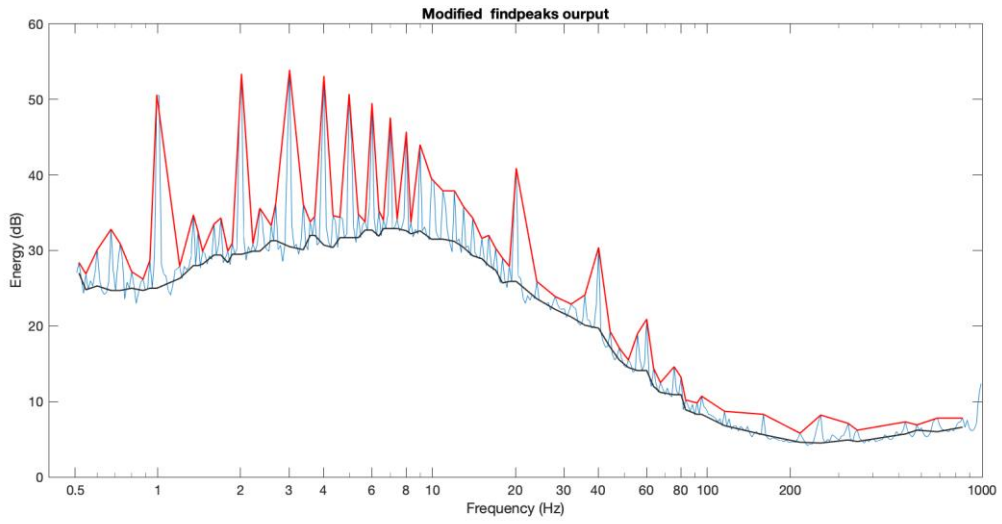
This behaviour is undesirable in the harmonic prominence measure as it will significantly overestimate the value of the harmonic prominence for these peaks.

Changes to the `findpeaks` function

A more limited, but more consistent, method of determining the baseline of the periodograms was developed.

In this method the valleys to the left and right of a peak are determined by looking for the first instance where the slope from the peak goes from negative to positive. Like the `findpeaks` function, the higher of these two is used as the baseline. The prominence is then the height of the peak minus the height of the baseline.

By comparison with the figure, above, the new baseline is seen below. The black baseline now more closely follows the base of the peaks. The P_{harm} for this recording is given as 24.6 dB, in line with a visual estimate.



ⁱ Matlab 9.7, The Mathworks Inc (Natick, MA, USA)

ⁱⁱ Retrieved from

https://au.mathworks.com/help/signal/ref/findpeaks.html?s_tid=srchtitle#bufbbs1-p, 18 Sept. 2020.