1. Introduction

Figure 1: Soundskrit Boom Headset Demo Kit (Mantis)

The Boom Headset Demo Kit (Mantis) is designed to evaluate the performance of Soundskrit microphones and software in a headset scenario. Our most powerful headset configuration combines an omnidirectional and dipole microphone to enable several beamforming and AI algorithms. There are three positions on the Mantis board with an omnidirectional microphone and a SKR0610 dipole microphone. One near the mouth to simulate a boom headset, one further back to simulate a reduced boom length, and one near the ear as found in boomless headphones.

Soundskrit provides a variety of audio processing algorithms for the Mantis kit, categorized by boom location, with each boom location corresponding to a specific algorithm family. These algorithms include linear, adaptive, and nonlinear beamformers, as well as AI-based denoising models, all of which can be configured and evaluated using the Soundskrit Demo Kit Interface software, either with or without a connected Mantis demo kit. Developed specifically for demo kit evaluation, Soundskrit Demo Kit Interface platform enables customers to quickly test and compare different algorithms on the demo kit within seconds.

What's In the Box	
Mantis Board	The Mantis development board with omni + dipole microphone pairs in 3 locations.
Soundskrit PARDI audio interface	Multichannel audio interface to connect microphones over USB
Molex cable	To connect Mantis microphone board with audio interface
USB-A to USB-C Cable	A cable to connect the board to your PC.
Mounting Putty	Putty to mount the board to an existing headset for evaluation.

Table of Contents

1. Introduction 1
2. Mantis Hardware Introduction 3
3. Mantis Software Introduction 3
4. Mantis Algorithm Introduction 4
   4.1. Beamformers 5
   4.2. AI Denoising 5
   4.3. Algorithm Performance 6
5. Setup Guide 6
   5.1 Mount Mantis Board to a Headset 6
   5.2 Installation 7
   5.3 Audio Performance Evaluation 8
      5.3.1 Option 1: Load built-in audio samples 8
      5.3.2 Option 2: Load an External .wav File 9
      5.3.3 Option 3: Record Live with the Mantis 9
6. Troubleshooting 11
7. Additional Support 11
8. Revision History 11

2. Mantis Hardware Introduction

The Mantis board has a Molex connector and 6 MEMS microphones. The “front” of the board is the bare side of the PCB, or the side opposite to which the microphones are mounted. When using the board, face the bare side of the PCB towards the mouth. Each algorithm uses one omnidirectional microphone and one dipole microphone. These microphone pairs are positioned in three different locations so that you can compare the performance difference between a headset with a long boom, a short boom, and one without a boom. These pairs will be referred to as the long boom pair, short boom pair, and boomless pair respectively. Generally, the long boom and short boom pairs provide similar performance, offering the best vocal quality and highest noise rejection. The boomless pair trades off performance to embed the microphones in the earcup rather than using a boom.

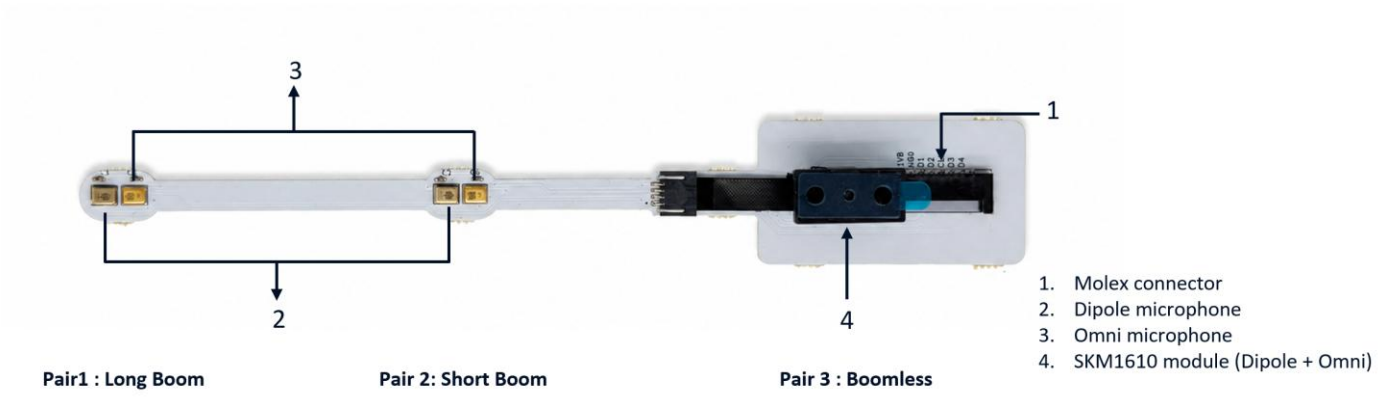

Figure 2: Mantis Board Microphone Pairs

The PARDI audio interface has a USB-C connector to connect it to a PC and a Molex connector to connect the Mantis board. The cable comes pre-connected and only needs to be reattached if it becomes disconnected.

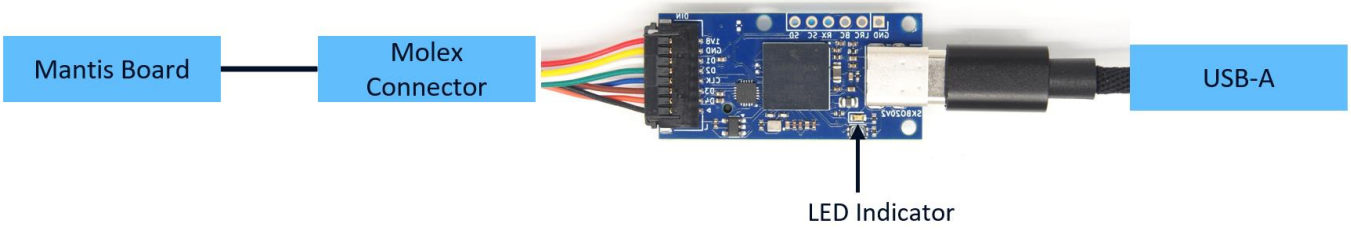

Figure 3: Mantis PARDI Board Interface

3. Mantis Software Introduction

The Soundskrit Demo Kit Interface is a platform used for evaluating performance of the Mantis demo kit. It enables users to quickly test the different algorithms provided by Soundskrit, while also supporting multiple headset boom locations simultaneously.

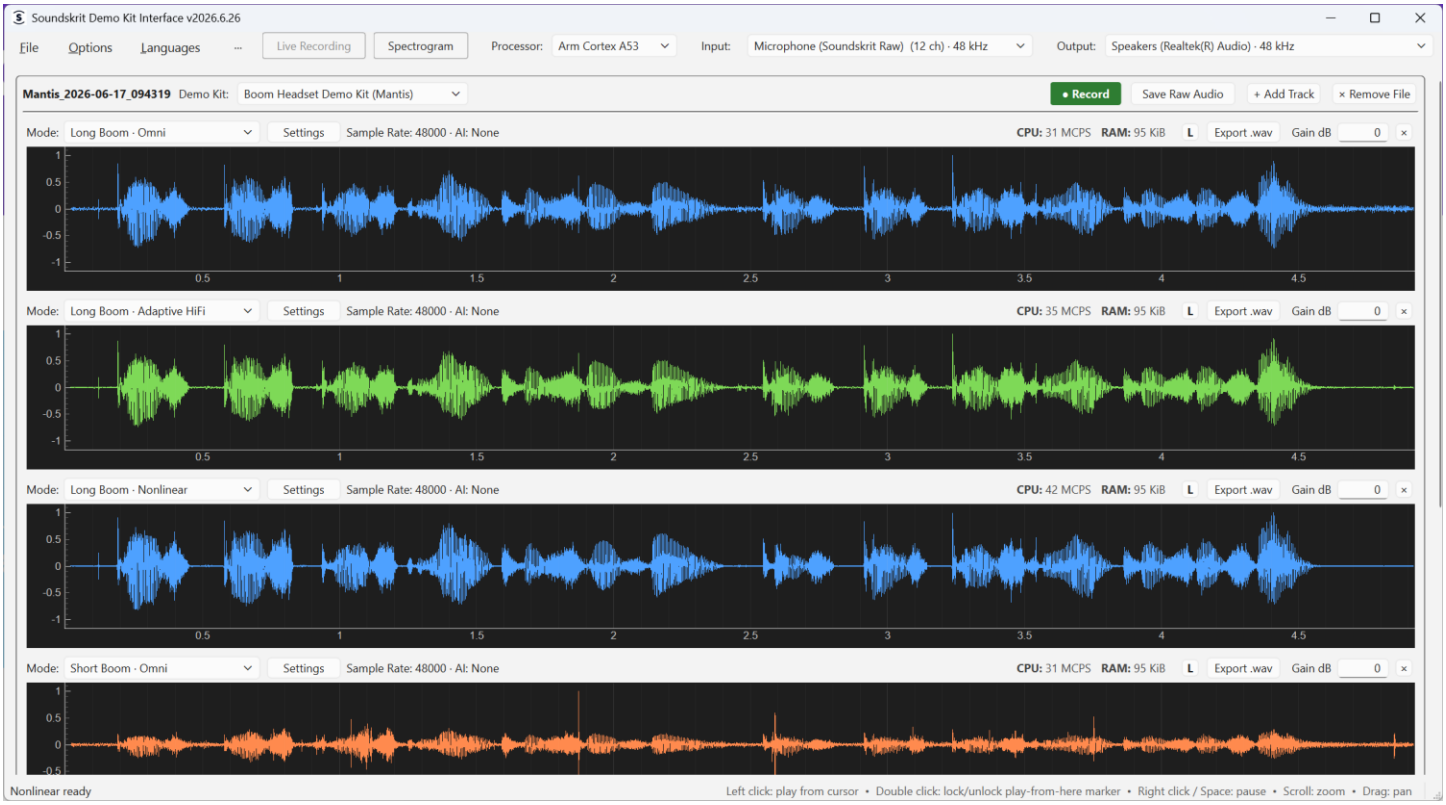

Figure 4: Soundskrit Demo Kit Interface

The software allows users to capture a recording via Live Recording mode and process the raw signals through various algorithms. This mode lets users compare algorithm performance side by side through playback. To facilitate quick evaluations, sample audio files are also provided in the software, allowing users to test the algorithms even without connecting to the Mantis demo kit.

The Demo Kit Interface also provides relevant computational metrics across all algorithm configurations, including CPU usage and RAM usage for various processor cores to help customers understand the algorithm requirements in embedded implementation.

For more detailed introduction of Soundskrit Demo Kit Interface, please refer to **Soundskrit Demo Kit Interface User Manual** from the Downloads page on Soundskrit's website.

4. Mantis Algorithm Introduction

To help customers understand the performance differences between algorithms and identify the most suitable solution for processing audio signals in their projects, Soundskrit provides various algorithm families, including beamformers and AI-based models, across three boom locations. These algorithms are optimized for isolating voice pickup and suppressing background noise. They are specifically tuned for headset applications, with a typical microphone-to-mouth distance of approximately 5 cm for the long-boom configuration and may not perform optimally at distances greater than 10 cm. Below is a detailed overview of these algorithms.

4.1. Beamformers

Soundskrit provides several beamforming algorithms for users to evaluate the performance in terms of voice capture and noise suppression and further decide which one works best for their application. The table below shows the beamformers that can be evaluated with the Mantis.

Beamformer Descriptions*	
Omni	Outputs the equalized omnidirectional microphone signal, used as a reference for performance comparison.
Dipole	Outputs the equalized dipole microphone signal, showcasing the pure hardware directionality.
Linear	A fixed, distortionless beamformer toward the speaker's mouth preserving natural voice quality.
Adaptive HiFi	A beamformer that dynamically switches polar patterns to minimize background noise, tuned to remove plosives and maintain minimal distortion.
Adaptive Focus (Long boom only)	A beamformer that dynamically switches polar patterns to minimize background noise, tuned to better suppress interfering speech.
Nonlinear	An enhanced beamformer, providing stronger background noise suppression, while introducing some distortion and artifacts.
Adaptive AI (Long boom only)	Combines an adaptive beamformer with an AI denoiser running in parallel to achieve stronger noise suppression while maintaining low CPU usage.

** Please refer to the computational requirement shown in the GUI.*

4.2. AI Denoising

The GUI comes equipped with AI denoising models that are divided into three categories: Small, Medium, and X-Large. These models can be cascaded on top of a chosen beamformer. The table below shows the detailed description for each AI model.

AI Denoiser Description*	
Small	Small AI model for noticeable background noise suppression, with some distortion and artifacts.
Medium	Medium AI model for enhanced background noise suppression, with less distortion and artifacts
X-Large	Extra-large AI model for premium background noise suppression, with minimal distortion and artifacts

** Please refer to the computational requirements shown in the GUI.*

4.3. Algorithm Performance

The table below summarizes the measured performance of each algorithm across the three boom locations. Scores were obtained using the 3QUEST methodology, which separately rates speech quality and background-noise suppression under realistic noise conditions. For every metric, a higher value is better.

Performance is measured using three metrics, all scored so that higher values are better:

- S-MOS (Speech Mean Opinion Score) rates how natural and undistorted the voice sounds on a scale of 1 to 5
- N-MOS (Noise Mean Opinion Score) rates how well background noise is suppressed on a scale of 1 to 5, with higher scores meaning the residual noise is less audible
- SNR (Signal-to-Noise Ratio), expressed in decibels, captures how much louder the desired speech signal is compared to the level of unwanted background noise.

Algorithm	Long Boom			Short Boom			Boomless		
	S-MOS	N-MOS	SNR	S-MOS	N-MOS	SNR	S-MOS	N-MOS	SNR
Omni	4.62	3.03	17.40	4.45	2.52	10.15	4.06	2.26	5.64
Linear	4.64	3.37	25.84	4.14	2.60	12.91	4.06	2.40	10.17
Nonlinear	4.30	4.59	40.80	4.05	4.35	33.36	3.74	4.15	29.07
Adaptive HiFi	4.65	3.50	28.02	4.21	2.74	14.32	4.19	2.55	11.74
Adaptive HiFi + Medium AI	4.64	4.43	38.86	4.33	3.82	28.01	4.14	3.72	24.17
Adaptive Focus	4.39	3.53	27.93	—	—	—	—	—	—
Adaptive AI	4.39	4.48	38.16	—	—	—	—	—	—

5. Setup Guide

5.1 Mount Mantis Board to a Headset

There are two methods to mount the Mantis board to a headset and take recordings, referring to the figures below. Please note that the back side of the Mantis board should face the headset and the long boom microphone should be positioned approximately 5 cm away from the mouth.

Figure 5: Mantis board with The Cable Facing Down

Figure 6: Mantis Doard with Wable Routed over The Headband

5.2 Installation

Download and launch the installer file from Downloads on Soundskrit website then follow the installation steps. A prompt may appear requiring approval for an unknown publisher.

Connect the Mantis to your PC using the included USB-C to USB-A cable.

To start the software, search from the Windows start menu for ‘Soundskrit Demo Kit Interface’ or use the desktop shortcut generated by the installer.

Make sure the Mantis is correctly connected to the PC and that the PARDI board is in **raw mode**. If the PARDI board is not in raw mode (this may apply to units shipped before June 2026), a window will appear, prompting you to change the PARDI board to raw mode.

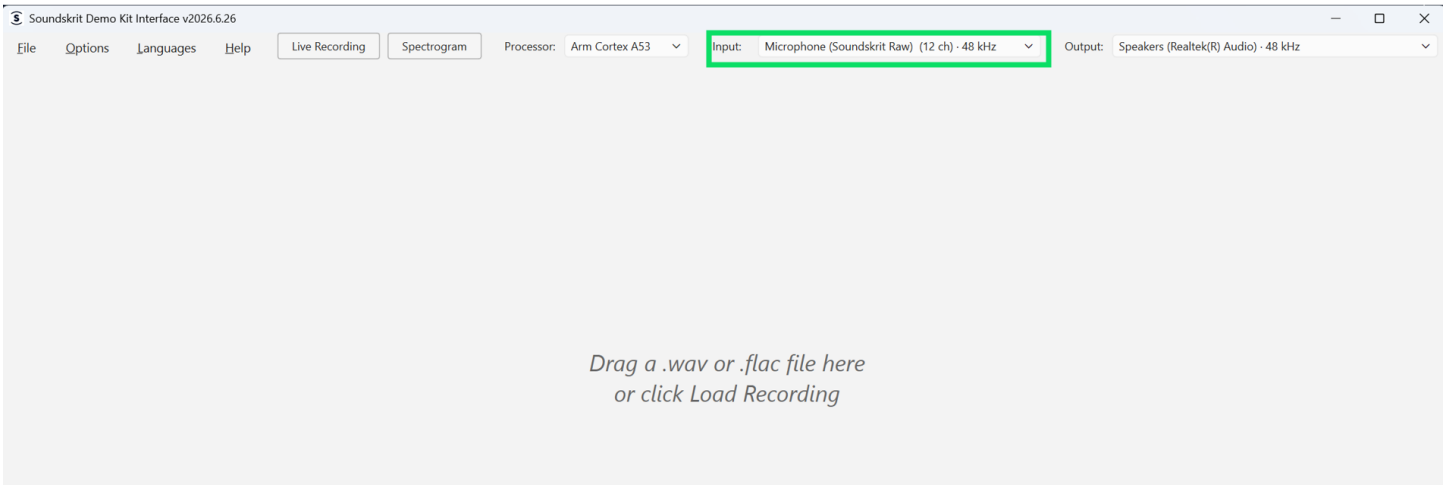

Figure 7: Raw Mode Detected

5.3 Audio Performance Evaluation

There are **three ways** to get audio for audio performance evaluation.

5.3.1 Option 1: Load built-in audio samples

Users can find pre-recorded audio samples and load them in the Soundskrit Demo Kit Interface for quick performance comparison without the Mantis kit connected. To find the audio samples, click **File**in the top-left corner, and choose **Samples**. In the new window displayed, you can choose the samples in Section drop-down menu.

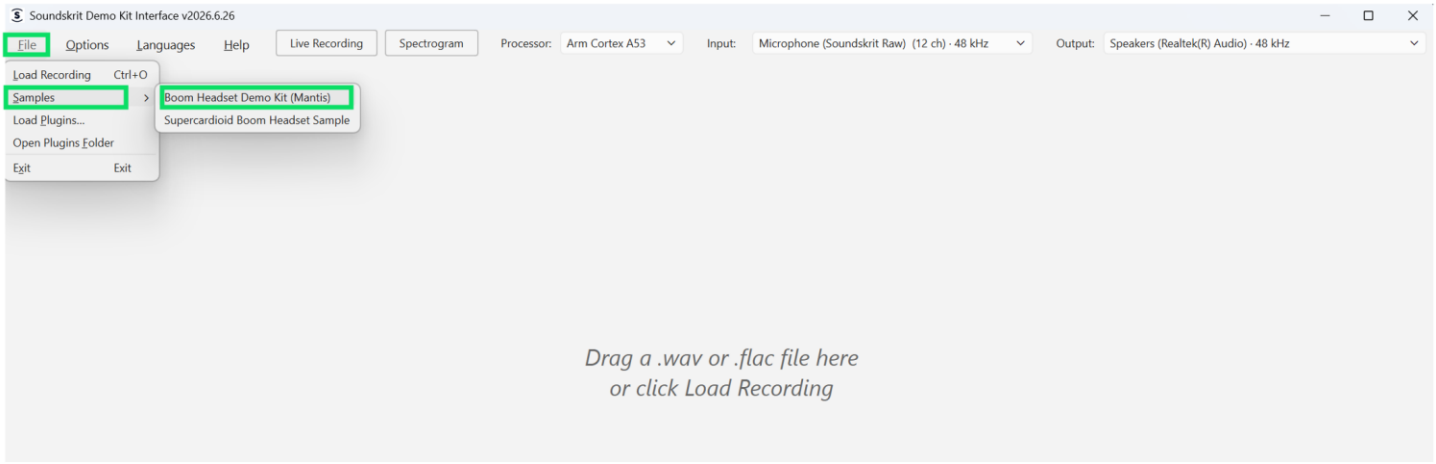

Figure 8: Audio Samples in Soundskrit Demo Kit Interface

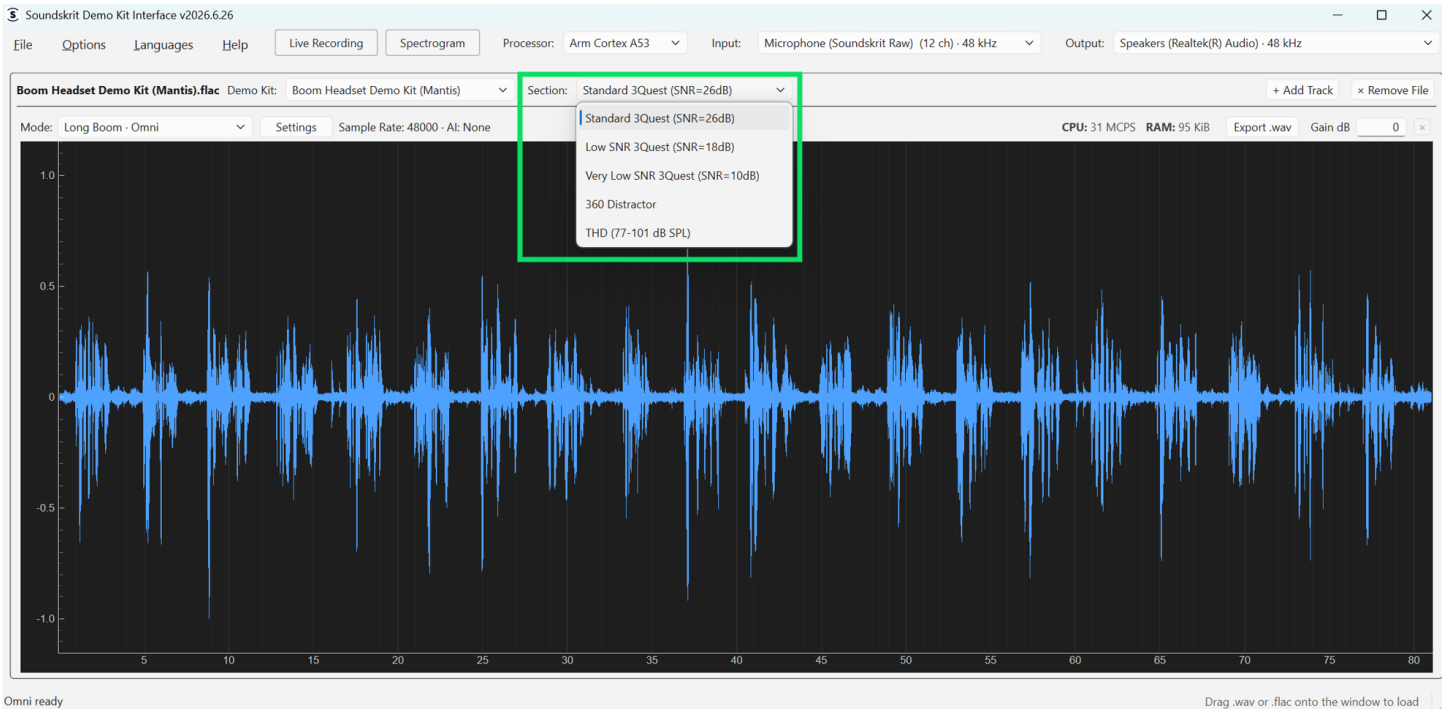

Figure 9: Different Audio Samples for Mantis

3QUEST samples assess speech quality and noise suppression in typical noise. In a 3QUEST test, a talker speaks a set of standard sentences while realistic background noise is played around the device. Three versions are provided at decreasing signal-to-noise ratios for speech level to background noise: 26dB (standard 3Quest setup), 18dB, and 10dB SNR. The SNR is computed at the HATS mouth reference point (25mm in front of the mouth).

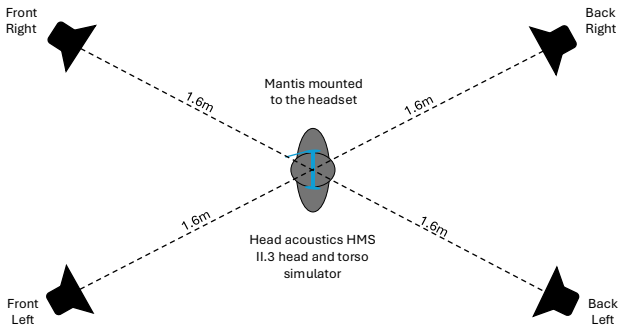

Figure 10: 3Quest Testing Setup

THD samples test for distortion at high speech levels. In this test, a talker's voice is played at progressively increasing levels, from 77 to 101 dB SPL, to simulate a person speaking louder and louder. The recording contains speech rising in level with no added background noise, letting you check whether the algorithms introduce any distortion as the voice gets louder.

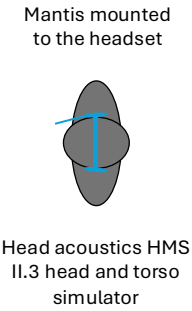

Figure 11: THD Samples Testing Setup

The **360 Distractor** sample evaluates how well competing talkers in the environment are removed. Throughout the test, a constant background noise at 26dB SNR is present while an interfering speaker (the distractor) is played through loudspeakers placed at five positions around the listener. The distractor speech is played in the following order: front center, front left, front right, back left, and back right. For each position, the sample contains three segments back-to-back: the target speaker alone, the distractor alone, and then the two together.

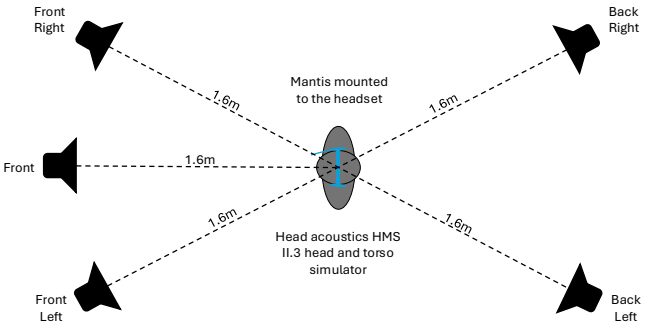

Figure 12: 360 Distractor Testing Setup

5.3.2 Option 2: Load an External .wav File

Users can upload external .wav file for audio processing. The .wav file must be a raw signal recording captured from the Mantis demo kit.

Please note, if the external audio file is not recorded by Mantis demo kit, the audio file will only be marked as ‘External’, and no demo kit can be chosen for it. Under this circumstance, external audio file could be set as benchmark reference to compare the audio performance between Soundskrit Mantis Kit and other audio solutions.

5.3.3 Option 3: Record Live with the Mantis

Use Live Recording (see image below) to capture audio with the Mantis and process it in real-time.

1. Click **Live Recording** to open the recording window. By default, six tracks are configured to compare performance across boom locations and algorithms in parallel (see Tracks Map below).
2. Click **Record** to start and **Stop** to end. The audio is distributed across the tracks with preset algorithms applied. All audio tracks are normalized to facilitate easy comparisons. Recording again replaces the previous take.
3. To compare more algorithms, click **Add Track** and set the algorithm **Mode** and other parameters in **Settings**, such as AI models for each new track.
4. Click **Save Raw Audio** to keep all six raw channels (re-loadable later for processing), or **Export .wav** to save a single processed track.

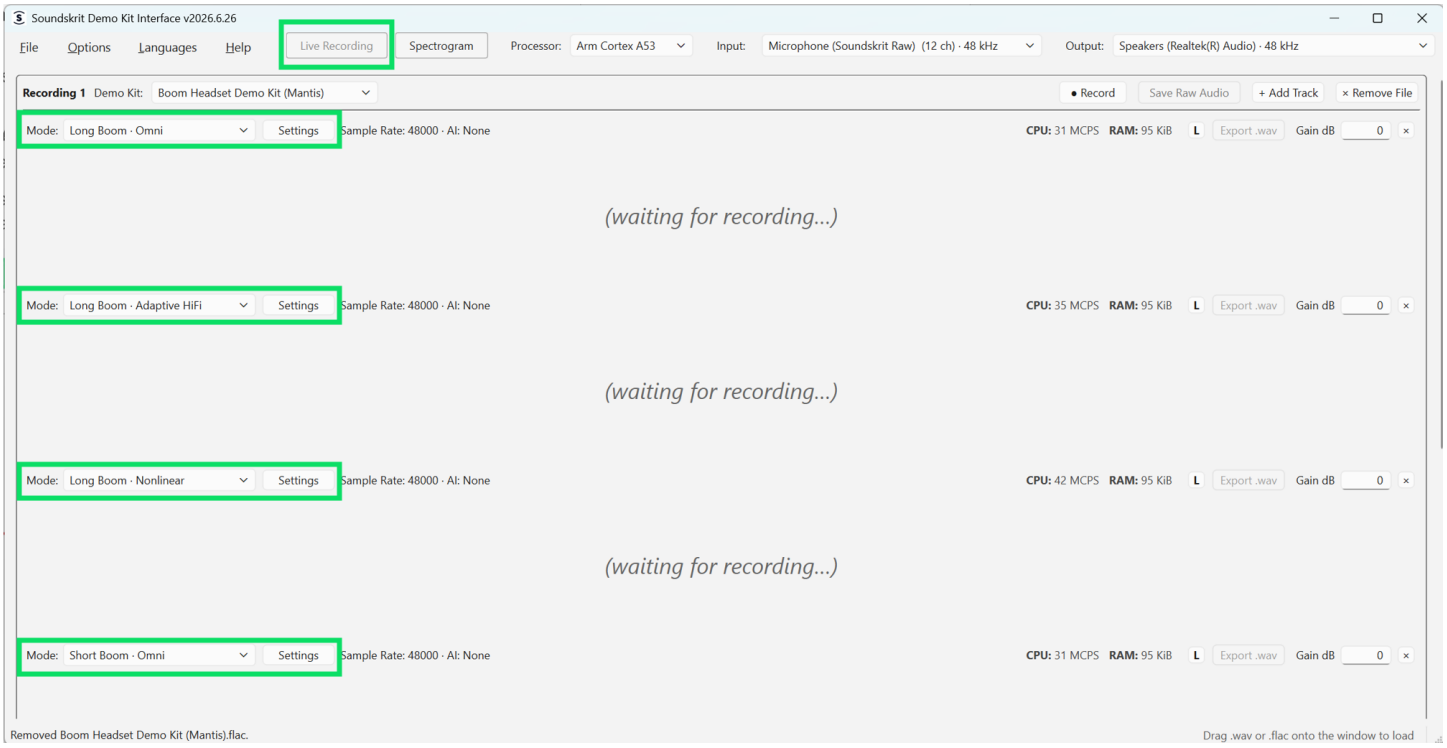

Figure 13: Automatic Tracks Setup in Live Recording

Tracks Map	
Track 1	Long boom, Omni (reference audio)
Track 2	Long boom, Adaptive HiFi
Track 3	Long boom, Nonlinear
Track 4	Short boom, Omni (reference audio)
Track 5	Short boom, Adaptive HiFi
Track 6	Short boom, Nonlinear

6. Troubleshooting

This section provides solutions to common issues encountered during installation and operation of the demo kit.

Problem	Possible Cause	Solution
Soundskrit Demo Kit Interface keeps showing **Soundskrit Raw not detected**	Mantis is not correctly connected to PC	Verify the connection between PARDI board and mantis, as well as between PARDI board and PC. A green LED indicator on PARDI board should be flashing if the connection is correct.
Cannot locate boom position in track after loading external audio file recorded by Mantis	'Demo Kit' is not chosen as Mantis	Set Mantis as the demo kit.

7. Additional Support

For further information on Soundskrit's products, visit our website at http://www.soundskrit.ca where you can find more application notes, datasheets, and purchasing information. If you have any questions or need technical support, please reach out to applications@soundskrit.ca.

8. Revision History

Revision Label	Revision Date	Sections Revised
-	Jun 2026	Initial release

Figure 18: Add benchmark recording in Mantis GUI

Soundskrit developed the first high-performance directional MEMS microphone on the market, leveraging years of research in bio-inspired MEMS based on how spiders and other insects in nature hear. In combination with Soundskrit's in-house audio processing algorithms, directional microphones can be used to capture and isolate any sound in an environment with a fraction of the size, power, and computation of traditional omnidirectional-based microphone arrays.

Soundskrit was founded in 2019 and is headquartered in Montreal, Quebec with an R&D facility in Ann Arbor, Michigan.

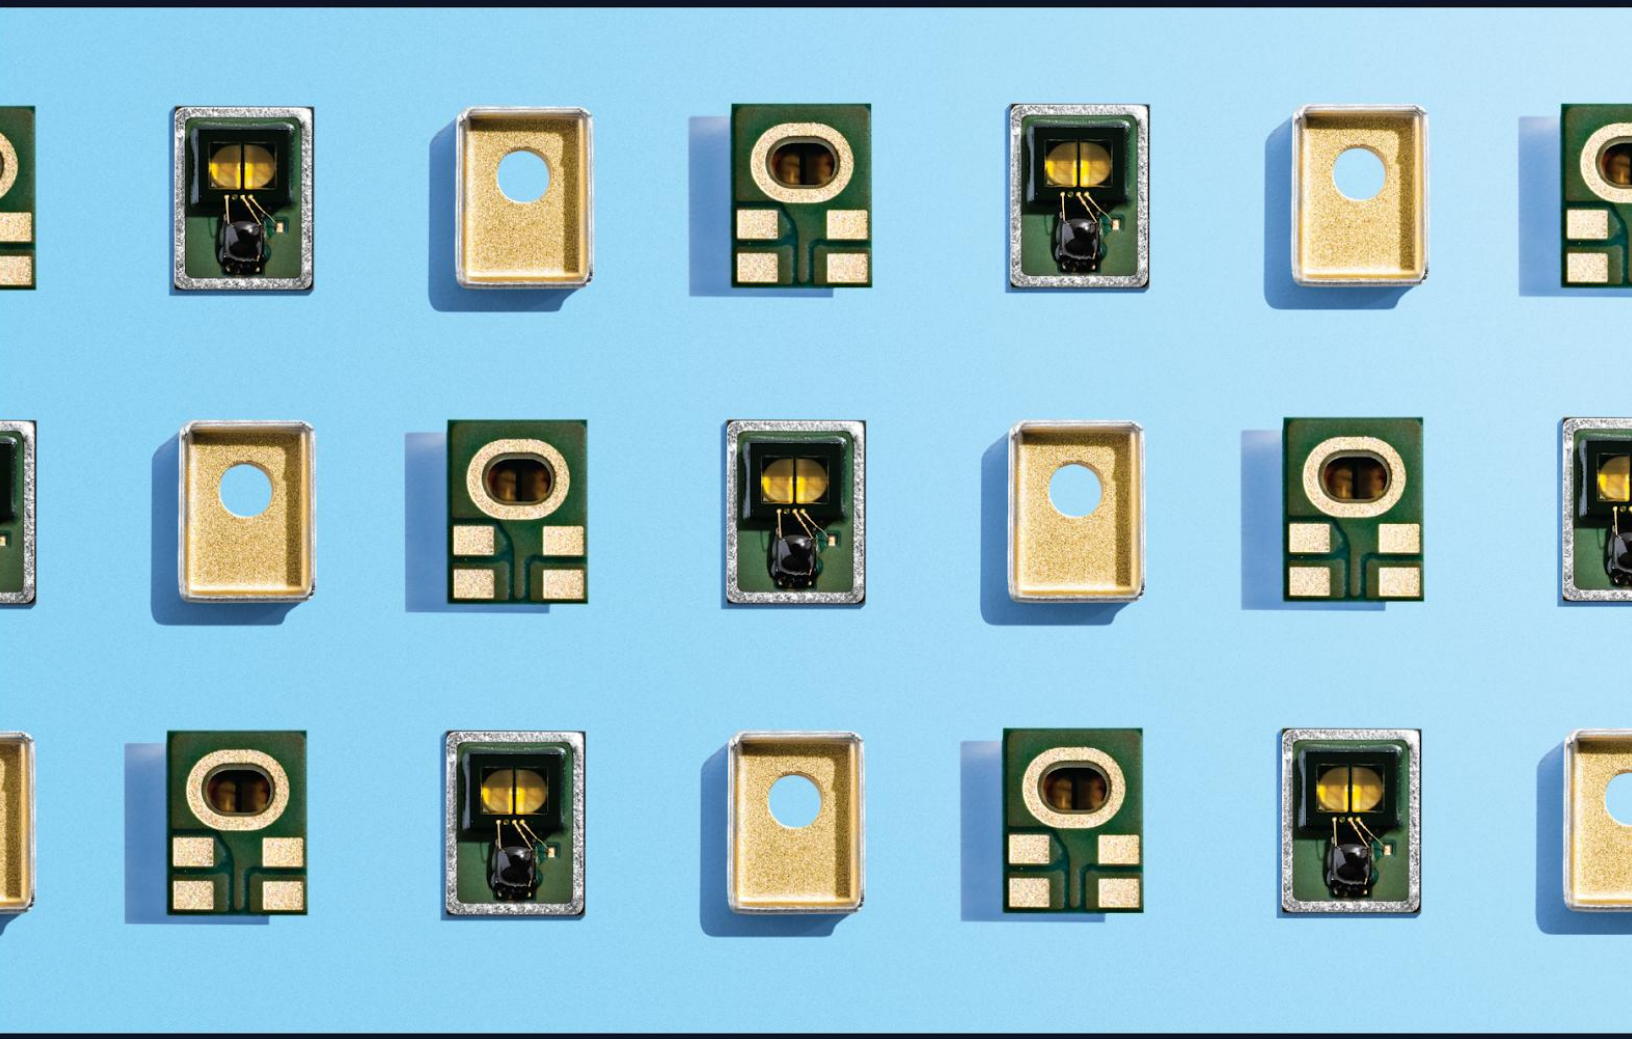

soundskrit
Hear the impossible.