redi™ ACOUSTICS
RESEARCH EDUCATION DEVELOPMENT INITIATIVE
www.rediacoustics.com

Room Optimization of Cuboid Spaces (ROCS)

User Guide

v1.0

Contents

List of Figures

1 Introduciton

This guide provides practical, user-focused documentation on working with the ROCS application. It offers step-by-step instructions for interacting with the app, interpreting results, and leveraging optimization features in real-world scenarios.

2 What is ROCS?

The Room Optimization for Cuboid Spaces (ROCS) is a tool designed to help you find the best setup for your room. You can input the dimensions of your room and the position of your speakers and listener, and the tool will provide you with a preview of the room model and the frequency response at the listening position. You can also run an optimization to find the best setup for your room, including the best speaker and listener positions.

3 Coordinate System

The coordinate system used in ROCS is a Cartesian coordinate system, illustrated by the X, Y, and Z axes below. The origin point is located at the center of the intersection between the rear wall and the floor, i.e., the point (x0, y0, z0) is located at the center of the rear wall and the floor. The coordinate system is designed so the listener is always centered in the room width and facing the front wall.

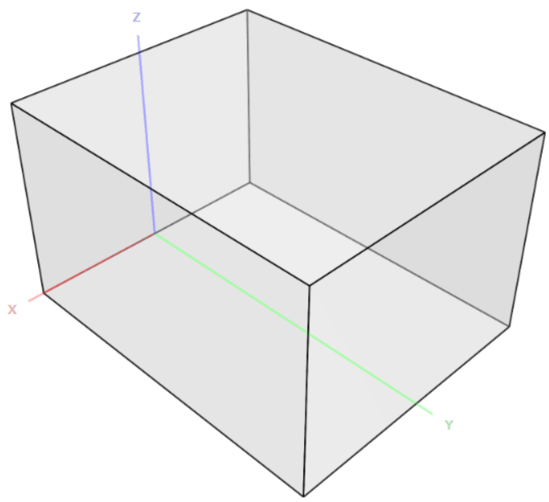

Figure 1: Cartesian coordinate system.

4 Input Parameters

The ROCS tool allows users to define various input parameters for room acoustic modeling and optimization. These inputs form the basis for calculating optimal speaker and listener placements, room dimensions, and the associated room response. A brief description of the core input sections is provided below, with detailed explanations of all parameters to follow in subsequent sections.

- Input Values: Users provide the room's physical characteristics, such as dimensions (length, width, height), speaker placements, and listener locations. These inputs allow for precise modeling based on the room's constraints.

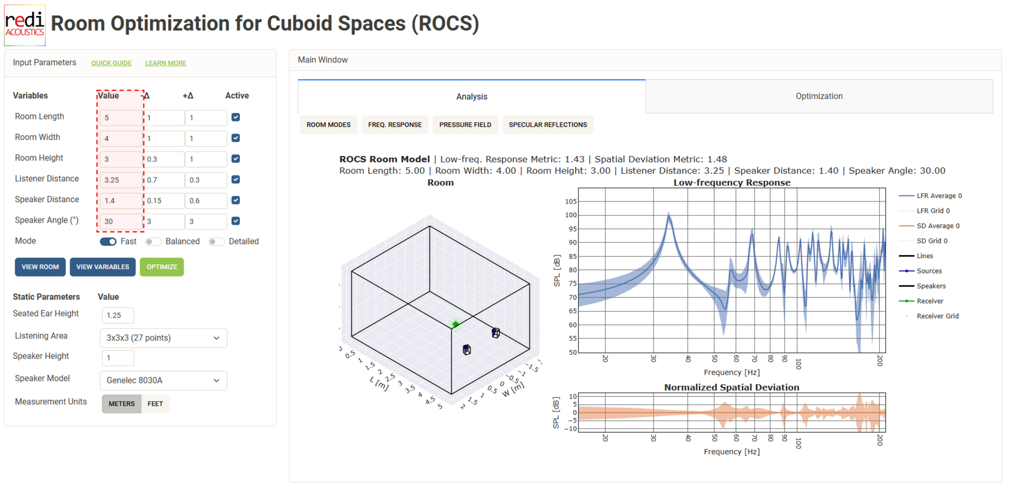

Figure 2: Input values column.

- Variable Ranges: Each input parameter has an associated range, enabling users to experiment with different configurations within reasonable boundaries for optimal results. The ranges can be displayed using the **View Variables** button in the input section.

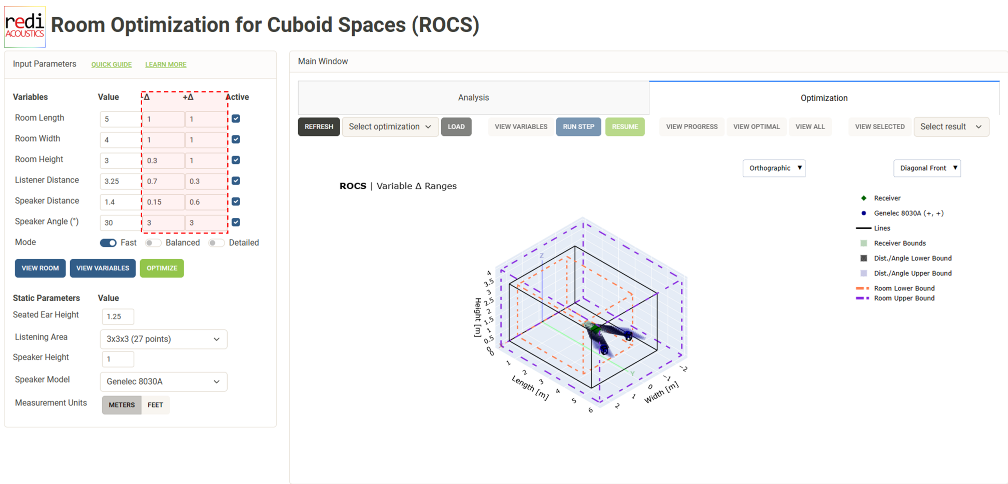

Figure 3: Variable ranges column.

- Active Variables: These refer to the parameters that users choose to actively modify during optimization. For instance, if a room is already constructed, its dimensions may remain fixed. In such cases, users can deactivate the room dimensions and only adjust the positions of the receiver and sources. This approach allows for the optimization of these positions within the constraints of the existing room structure.

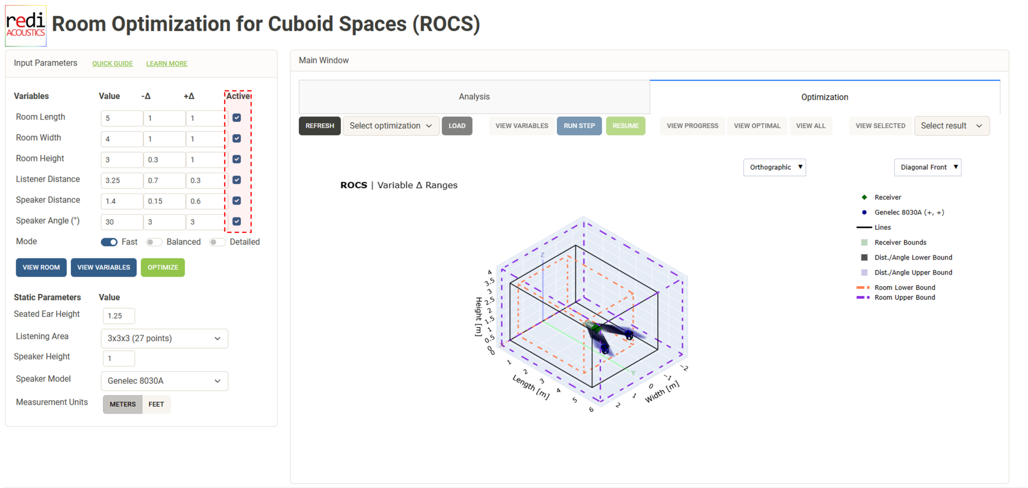

Figure 4: Active variables column.

- Static Parameters: Certain parameters remain fixed during the optimization process, often determined by user preferences or architectural constraints. For instance, the heights of the receiver and speakers are set based on user height and furniture arrangement. Additionally, specific speaker models and receiver grid configurations are chosen beforehand and do not vary during optimization. This allows the algorithm to concentrate on adjustable variables while ensuring consistency in the established aspects of the room setup.

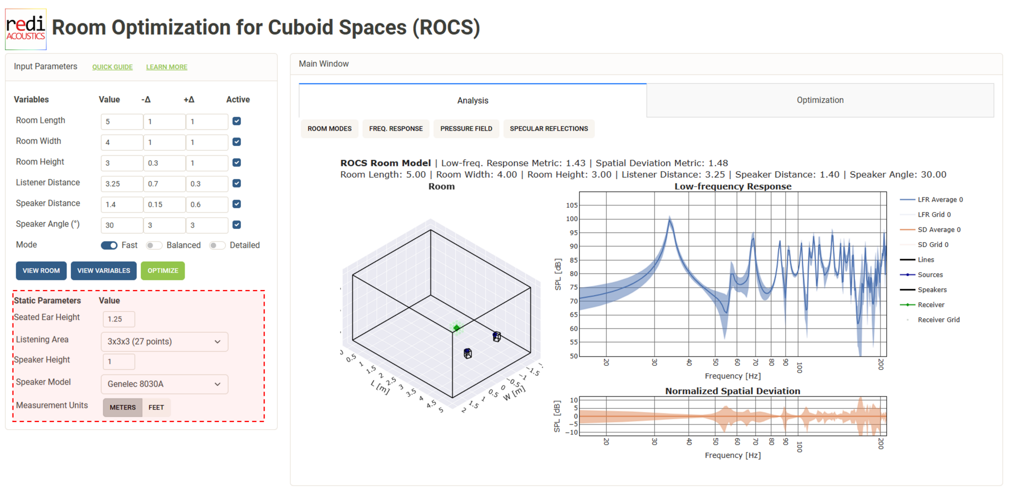

Figure 5: Static parameters section.

All measurements can be switched from meters to feet at the bottom of the Static parameters section.

4.1 Room Length

The room length represents the front-to-back distance of the room, i.e., the distance between the rear and the front walls. In ROCS, the room length is the orientation aligned with the listener, even if that is not the largest room dimension.

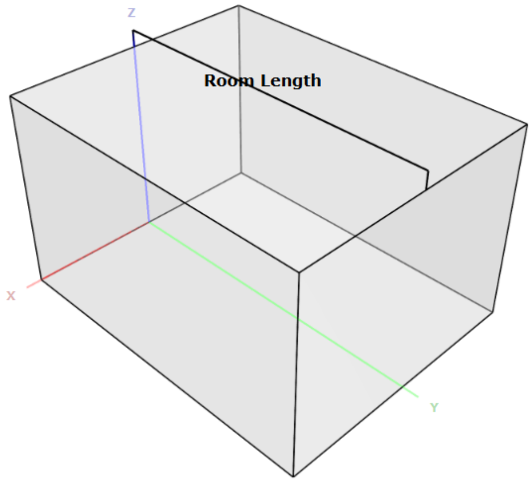

Figure 6: Room length definition.

4.2 Room Width

The room width represents the left-to-right distance of the room, i.e., the distance between the left and right walls.

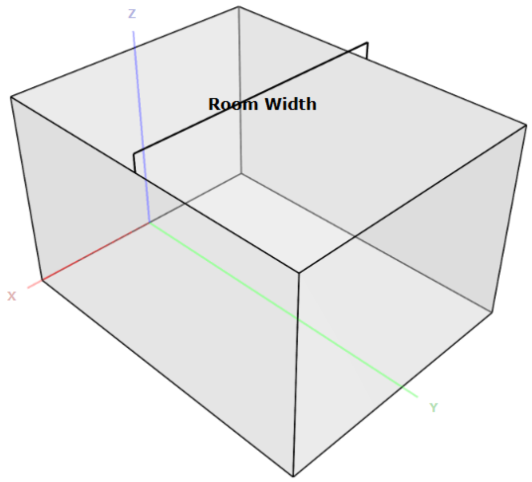

Figure 7: Room width definition.

4.3 Room Height

The room height represents the floor-to-ceiling distance of the room, i.e., the distance between the floor and the ceiling.

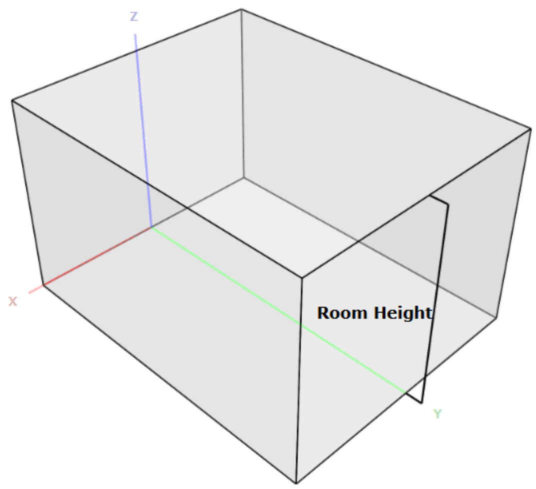

Figure 8: Room height definition.

4.4 Listener Distance

The listener distance is the distance between the listener and the rear wall. For example, if a listener is seated in the middle of a room with a length of 5 meters, the listener distance would be 2.5 meters.

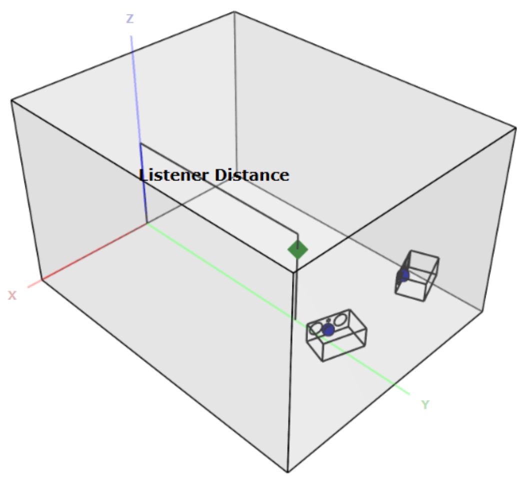

Figure 9: Listener distance definition.

4.5 Speaker Distance

The stereo speaker pair and the listener are positioned in a triangle configuration. The listener position is the lower triangle corner, and the speaker positions are the upper triangle corners. The speaker distance, shown in orange, is the distance between one speaker and the listener, i.e., the length of the left or right triangle sides.

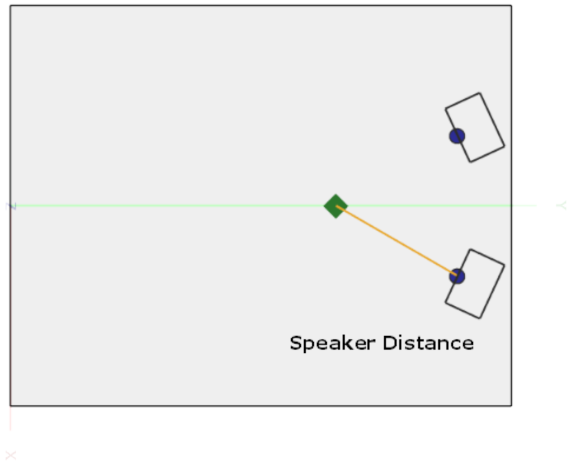

Figure 10: Speaker distance definition.

4.6 Speaker Angle

The speaker angle, shown in purple, is the angle of the speaker relative to the listener. The angle is measured from the listener to the speaker position, with the listener position as the origin and the speaker position as the endpoint. In other words, a centered speaker has an angle of 0°, and the usual stereo pair has an angle between 20° and 35°, depending on the application.

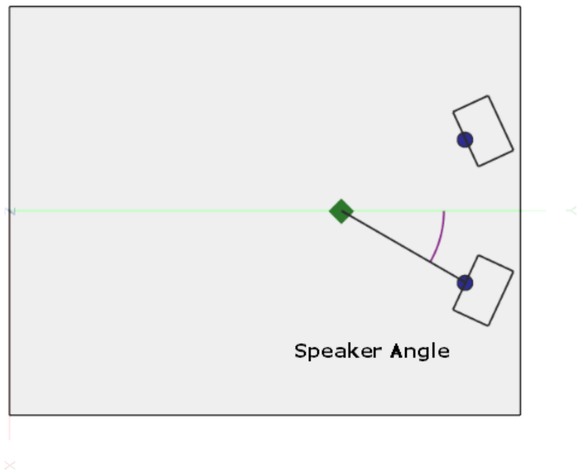

Figure 11: Speaker angle definition.

4.7 Seated Ear Height

The seated ear height is the height of the listener's ears when seated. This parameter is used to calculate the listener's head position, and along with the Listener Distance, it determines the listener's position in the room and sets the center point of the Listening Area.

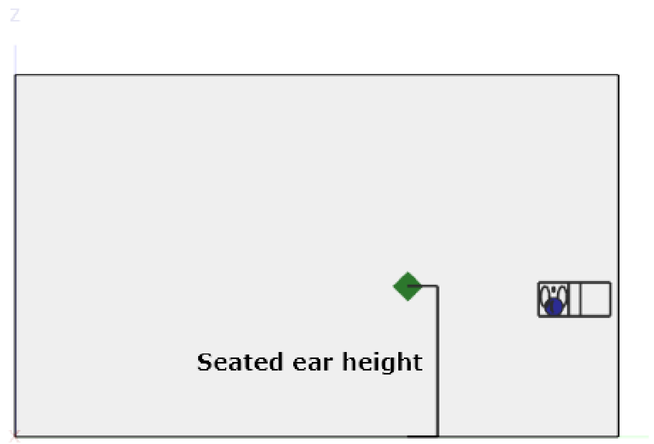

Figure 12: Seated ear height definition.

4.8 Listening Area

The listening area determines the grid size used for spatial averaging. An example of a 7x5x3 grid can be seen in Figure 13. The grid is used to calculate the frequency response at several points surrounding the listener, which are then used to evaluate how the room response changes with the listener's position. By nature, larger listening areas will present more spatial variation in the frequency response, while smaller ones will present less spatial variation. Smaller rooms tend to be suited for smaller listening areas, while bigger rooms (especially with larger furniture) might require a larger listening area.

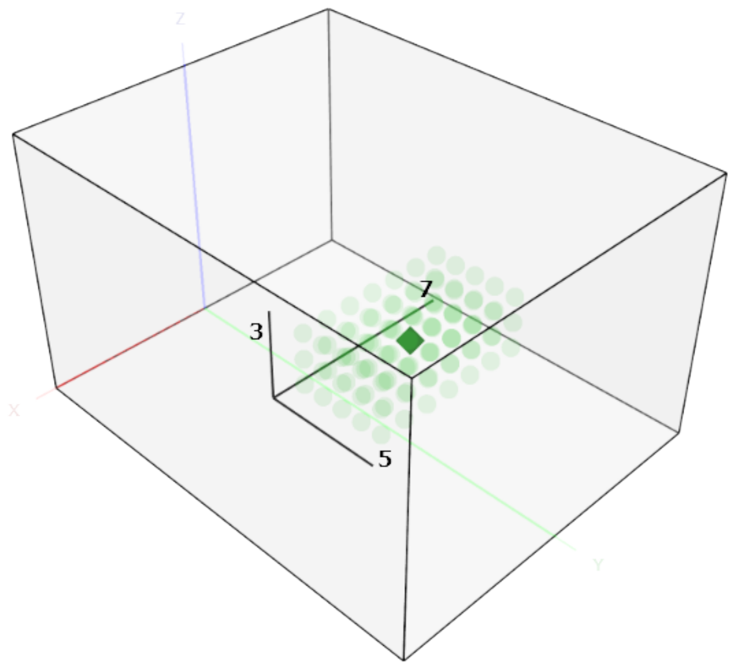

Figure 13: Listening area definition.

4.9 Speaker Height

The speaker height is the distance between the floor and the base of the speaker cabinet. This parameter is used to place the speaker and should be adjusted to each speaker model.

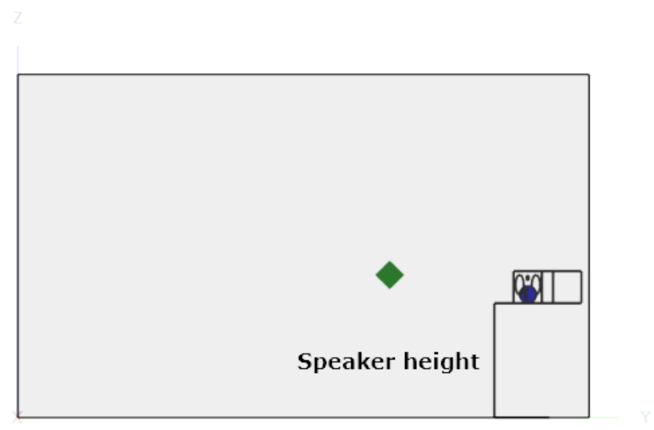

Figure 14: Speaker height definition.

4.10 Speaker Model

The speaker manufacturer's brand and model. This parameter is used to load the speaker model and correctly place the speaker drivers for all the available models.

4.11 Mode

The mode selection determines the level of detail used in the optimization. The fast mode is the quickest but will take the longest to converge. The detailed mode is the slowest, but it will likely converge faster. The balanced mode compromises the two and is recommended for most critical listening spaces. More details of the optimization modes will be given in Section 6.

4.12 Construction Type

Users can choose between four construction types: Gypsum, Concrete, Masonry, and Gypsum-Plywood-Gypsum. Each type possesses unique rigidity and absorption characteristics that influence the acoustic performance of the room.

5 Analysis

The Analysis tab allows the user to view all the relevant acoustic parameters of the room defined in the input parameters section. This set of visualization tools can be used to evaluate an existing room as well as to analyze further a specific solution given by the optimization algorithm. Detailed descriptions of each analysis option will be provided in the following sections.

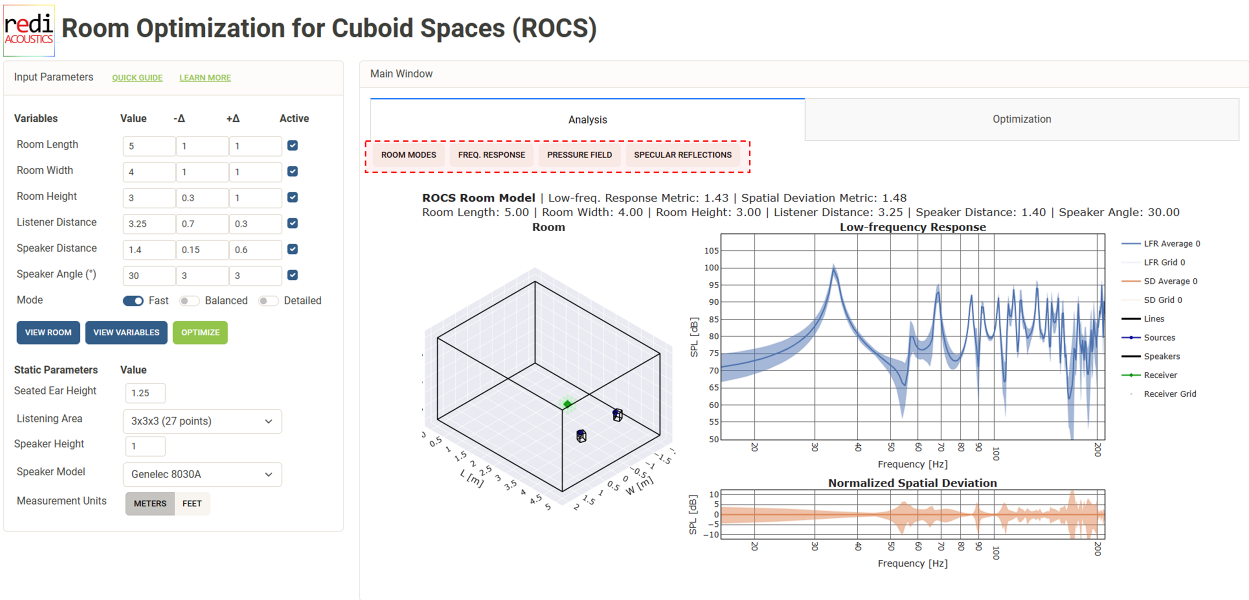

Figure 15: Analysis tab options.

5.1 Room Modes

This analysis calculates the analytical natural resonances or 'modes' of the room across the frequency spectrum. Room modes are frequencies at which standing waves occur, leading to areas of acoustic pressure buildup or nulls. Understanding room modes is essential for identifying potential issues with low-frequency response and making informed decisions about room dimensions, speaker placement, and acoustic treatment.

5.2 Frequency Response and Spatial Deviation

The **Frequency Response** analysis provides a graphical representation of the room's acoustic response across the low-frequency range. It reveals how evenly the room reproduces audio frequencies, highlighting any peaks or dips that may affect sound quality. This analysis is crucial for evaluating the acoustic performance of your room and identifying frequencies that may require treatment to achieve a balanced listening environment.

Spatial Deviation is an objective metric designed to assess the uniformity of the frequency response within a specified listening area. Centered around a critical listening position - often the primary position in a room where optimal sound reproduction is desired - this metric offers a nuanced understanding of how the acoustic properties vary across different points within the listening area.

In the Frequency Response tab, the spatial deviation is illustrated in light blue in the low-frequency response graph and in light orange in the normalized spatial deviation graph. The frequency response in the latter graph is normalized to zero, allowing spatial deviation to be isolated and visualized more clearly.

Lower values of this metric imply a more uniform auditory experience in the room, making it ideal for tasks such as audio mixing, monitoring, or focused listening. In contrast, higher values indicate a heterogeneous sound field, which may require additional acoustic treatment to attain the desired level of uniformity.

5.3 Pressure Field

This calculation maps the distribution of acoustic pressure within the room at low frequencies, illustrating areas of high and low sound pressure. The **Pressure Field** analysis helps visualize how sound waves interact with the room's geometry and surfaces, offering insights into the spatial variation of sound levels and guiding the placement of acoustic treatments and listening positions.

5.4 Specular Reflection

The **Specular Reflections** analysis identifies the paths that sound waves take as they reflect off room surfaces before reaching the listening position. This tool is invaluable for pinpointing early reflections, which are critical to sound clarity and imaging. Effectively managing these reflections is crucial for accurate sound reproduction. Knowing the location of the primary interfering reflections enables the precise positioning of targeted acoustic treatment, improving the auditory integrity of the space.

6 Optimization

The ROCS software employs a multi-objective genetic algorithm to optimize room parameters. This technique effectively addresses problems involving multiple conflicting objectives by evolving a population of candidate solutions over successive generations. Each candidate, referred to as an individual, is evaluated based on predefined objective functions.

In the context of acoustic optimization, each individual represents a unique room configuration. The objectives are to enhance low-frequency response and minimize spatial deviation. The population size defines the number of candidate configurations being evolved in each generation, influencing the algorithm's search space and diversity. The number of generations determines how long the algorithm continues to evolve solutions, impacting the depth of the search for optimal designs. By adjusting these parameters, the algorithm systematically improves room acoustics over time.

The population evolves over generations through genetic operations:

- Initial Population: The algorithm begins by creating different room designs based on the defined input parameters and their specified ranges.
- Evaluation: Each configuration is evaluated based on the specific acoustic objectives.

- Survival: Configurations are ranked according to their performance, determining which individuals will continue to the next generation.
- Selection: Individuals are selected for mating based on their acoustic performance to promote improved configurations.
- Crossover: Selected individuals are combined to create new room designs using relevant problem information like room geometry and listener and speaker positions.
- Mutation: Random variations are introduced to newly created offspring to maintain diversity and explore new potential solutions.

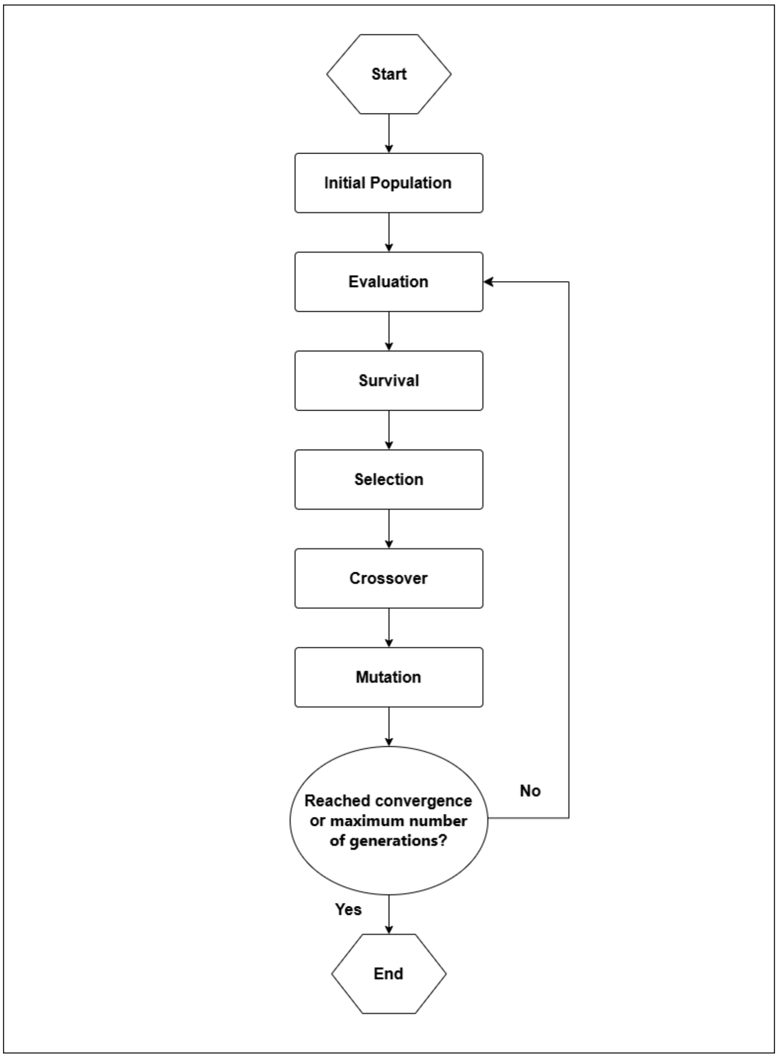

Figure 16: Genetic optimization algorithm diagram.

The user can choose from three optimization modes: **Fast**, **Balanced**, and **Detailed**. Each mode varies in the number of individuals and the maximum number of generations allowed during the optimization process, as outlined in Table 1. This variation effectively balances result quality with computation time; an increased number of individuals and generations leads to longer processing times. By selecting the appropriate mode, users can customize the optimization process to meet their specific needs and time constraints.

To prevent unnecessary computation, once the best solution has been identified, the algorithm features a convergence stop mechanism. This mechanism stops the process if the optimization parameters remain unchanged for ten generations, suggesting that an optimal solution is likely reached before all generations are completed.

In a multi-objective optimization like the one used in ROCS, the algorithm returns multiple optimal solutions rather than a single "best" result. Each solution represents a different balance between different acoustic objectives. For instance, one solution might achieve a better low-frequency response but show slightly higher spatial deviation, while another balances both metrics differently. All are considered optimal because none is superior across all metrics. These solutions give flexibility, offering users a range of choices depending on their specific acoustic needs.

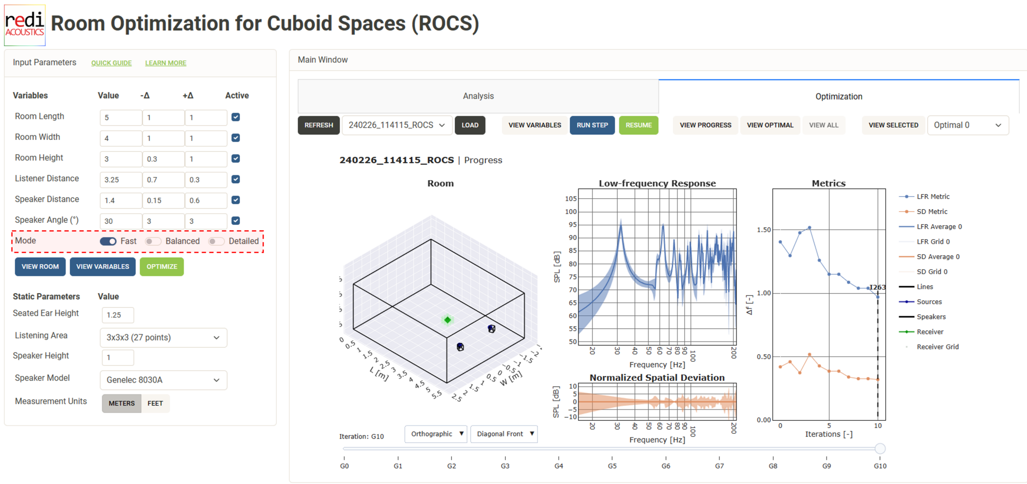

Figure 17: Mode selection.

Mode	Number of individuals	Maximum number of generations
Fast	25	10
Balanced	50	100
Detailed	100	150

Table 1: Optimization modes.

The user can initiate the optimization process by clicking the Optimize button in the input parameter section. This action will automatically switch to the Optimization tab, where various analysis and display tools become available.

Here, a 3D visualization of the room illustrates the placement of speakers, listeners, and room boundaries relative to the chosen room in the slider, offering a clear view of the layout. View options enable users to adjust how the layout is displayed. To the right, the Low-Frequency Response (LFR) and Normalized Spatial Deviation (SD) graphs provide key insights into sound behavior and distribution. The Metrics graph tracks changes in both LFR and SD metrics across optimization iterations, allowing users to monitor improvements. The individual number generated by the optimization of the currently displayed rooms is shown in the Metrics graph, and users can also select that individual from the drop-down menu using the **View Selected** function. Note that the rooms shown in the Metrics graph and available for selection via the slider are dependent on the chosen view option: **View progress**, **View optimal**, **View all**, or **View selected**. A detailed description of the computation management options and display tools is provided in the following sections.

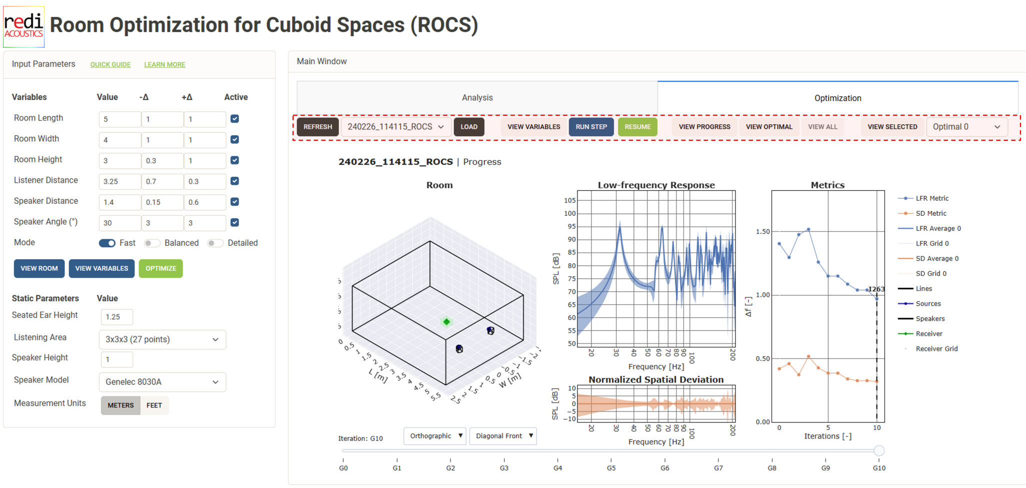

Figure 18: Optimization tab options.

6.1 Refresh

Updates the list of saved optimization in the drop-down menu.

6.2 Select Optimization

This feature allows for the selection of various saved optimization sessions. Users can easily select their desired session from a drop-down menu, which lists all previously computed processes, each with different starting values, constraints, and modes. This selection enables efficient navigation through past optimization efforts.

6.3 Load

Once an optimization session is selected, users can click the load button to retrieve that specific session. This functionality ensures users can resume work seamlessly, avoiding starting from scratch and maintaining continuity in their optimization processes.

6.4 View Variables

This section displays the room parameters and their range used in the optimization. Users can review these variables to understand the optimization landscape and make informed adjustments.

6.5 Run Step

By clicking this option, users can execute a single step (Generation) in the optimization process. This allows for a more controlled approach, enabling users to monitor progress and analyze changes iteratively.

6.6 Resume

Users can quickly resume optimization after loading, canceling, or completing a previous session. When resuming from a previously completed run, the algorithm will compute an additional set of generations according to the selected mode. This feature allows for flexibility in managing the optimization workload without losing progress.

6.7 View Progress

This tool displays the optimization process's progress, illustrating how the algorithm evolves across generations. Users can interact with a slider to view changes in room configurations at each stage, enabling them to assess improvements and make informed decisions throughout the optimization journey. Each generation displays one of the optimal room configurations, while the accompanying metrics graph represents the average performance of the generation. This provides insight into how the generations are improving as a whole.

6.8 View Optimal

This option allows users to explore the top-performing room configurations identified by the algorithm. Each solution strikes a different balance between acoustic objectives. Users can review these results to find the configuration that best meets their specific needs, providing flexibility in selecting the one that aligns best with their preferences.

6.9 View All

Users can access a comprehensive list of all room configurations generated during the optimization process. This feature allows for detailed analysis of the performance of each configuration.

The optimization process may generate and evaluate thousands of individuals depending on the selected mode and the number of generations computed. To maintain performance, this option becomes unavailable when the number of individuals exceeds a manageable threshold for display.

6.10 View Selected

The View Selected option displays any of the tested individuals or optimal configurations available from the drop-down menu. Users can choose specific configurations to review, enabling a focused examination of particular designs. This feature allows for easy access to both general and optimal results, making it straightforward to compare and analyze different room designs based on user-defined criteria.

6.11 Information Display

The information display provides important insights into the optimization process through various internal performance metrics:

- EPS and indicator are internal variables that reflect the algorithm's effectiveness in converging toward optimal solutions.
- n_nds shows the number of optimal individuals identified, while n_eval indicates the total number of individuals evaluated.
- n_gen represents the number of generations completed, providing context on the duration of the optimization process.
- Run time gives the total time taken for the optimization, helping assess efficiency.
- The convergence log indicates whether the algorithm has reached a stable solution, signaling when further iterations may no longer yield improvements.

Together, these metrics offer a comprehensive view of the algorithm's performance and progress throughout the optimization.