

APU Loudness Compressor



User Manual
v3.0.3

Table of Contents

Table of Contents	2
1. Introduction	4
2. Parameters	4
2.1. Compression	4
2.1.1. Presets	5
2.1.2. Loudness type	5
2.1.3. Momentary window	5
2.1.4. RMS window	6
2.1.5. Output type	6
2.1.6. Visualization	6
2.1.7. Raster type	6
2.1.8. Sidechain	6
2.1.9. Dynamics type	7
2.1.10. Ballistics type	7
2.1.11. Low-Level type	7
2.1.12. Source range	7
2.1.13. Learn source range	8
2.1.14. Source range presets	8
2.1.15. Target range	8
2.1.16. Learn target range	8
2.1.17. Target range presets	8
2.1.18. Threshold	8
2.1.19. Ratio	9
2.1.20. Max range	9
2.1.21. Attack/Release	9
2.1.22. Linear Attack/Release	9
2.1.23. Look-ahead	9
2.1.24. Look-ahead mode	10
2.1.25. Knee width	10
2.1.26. Dry/Wet	10
2.1.27. Output gain	10
2.2. Loudness	11
2.2.1. Adaptive startup	11
2.2.2. Adaptive edges	11
2.2.3. Adaptive type	11
2.2.4. Adaptive threshold	11
2.3. Limiter	12
2.3.1. Limiter mode	12
2.3.2. Limiter threshold	12
2.3.3. Limiter ceiling	12
2.3.4. Limiter threshold link	12
2.3.5. Limiter lookahead	12
2.4. Visualization	13
2.4.1. Auto range	13
2.4.2. Visual range presets	13
2.4.3. Loudness range	13
2.4.4. Layout options	14
2.4.5. Layout source mode	14
2.4.6. Layout target mode	14
2.4.7. Layout output mode	14
2.4.8. Delta field type	15
2.4.9. Super-sampling mode	15
2.4.10. Bits per pixel	15
2.4.11. Breath residual	15
2.4.12. Bucket size	15
2.4.13. Histogram hold	15
2.4.14. Delta X/Y/T	15
2.4.15. Delta field	15
2.4.16. Shader params	16
2.4.17. Show thresholds	16
2.4.18. Panel params	16
2.4.19. History length	16
2.4.20. Show xfer func	16
2.4.21. Show peak meter	16
2.4.22. Show axis	16
2.4.23. Show axis lines	17
2.4.24. Show param context	17
2.5. Theme	17
2.5.1. Theme presets	18
2.5.2. Source color	18
2.5.3. Normal color	18
2.5.4. Target color	18

2.5.5. Output color	18
2.6. Settings	19
2.6.1. BPM	19
2.6.2. Host BPM	19
2.6.3. Dither	19
2.6.4. Velocity sensitive knobs	19
2.6.5. Load preset visuals	19
2.6.6. Load preset ranges	20
2.6.7. Dark mode	20
2.6.8. Offline HD	20
2.6.9. Blocksize	20
2.6.10. Look-ahead budget	20
2.6.11. Delay compensation	20
2.7. Update	20
2.8. About	20
3. Tips & Tricks	20
3.1. Compression/expansion	21
3.2. Automatic gain control	21
3.3. Gain riding	21
4. Glossary	21
4.1. Range sliders	21
4.2. History	21
4.3. Histogram	22
4.4. Source	22
4.5. Target	22
4.6. Output	22
5. Credits	22
5.1. MIT License (libebur128)	23
5.2. MIT License (melatonin_blur)	23

1. Introduction

[APU Loudness Compressor](#) is a modern loudness compressor / expander plug-in released by [APU Software, LLC](#). This plug-in is designed to improve several aspects of dynamic range control, both in terms of sound quality and user controls. Here you can read about the different features and parameters available with the plug-in, and get a general sense of how it differs from other compressors and why.

2. Parameters

This software supports a variety of parameters, most of which can be adjusted in real-time. This section provides an overview of each parameter, with each subsection detailing the parameters for the associated tab in the user interface. From within the user interface, you can also hover the mouse over a slider, combo-box, or checkbox to see a popup description of the parameter.

2.1. Compression



The compression tab contains configuration and visualization for loudness compression/expansion. This section details each of the individual parameters and components.

2.1.1. Presets

The compressor's presets combo-box contains a collection of basic presets, one for each combination of [Loudness type](#) and latency configuration. The low-latency presets introduce minimal delay, while the look-ahead presets introduce more delay in exchange for the ability to preemptively respond to the source signal.

See [look-ahead](#) for an understanding of how delay compensation works in the plug-in.

When switching between presets, your [source range](#) and [target range](#) won't be modified. This is done to make it easier to flip through different compressor presets after you've configured the range settings. Note that peak modes have their own units scale and as a result they use a separate source and target range; the sliders for these settings are swapped into view automatically when switching between loudness modes.

You can configure whether or not to [load preset visuals](#) when switching between presets, and you can switch between loading dark or light themes through the [dark mode](#) parameter. These settings are visual only, they don't have any impact on the audio behavior of the compressor.

Currently, the following presets are available ("Full Reset" restores all default parameter values):

- Full Reset
- Mixing & Mastering - Focus (low-latency)
- Mixing & Mastering - Present (low-latency)
- Mixing & Mastering - Deep (low-latency)
- Mixing & Mastering - Traditional (low-latency)
- Mixing & Mastering - Maximizer TP (low-latency)
- Mixing & Mastering - Maximizer PK (low-latency)
- Mixing & Mastering - Focus (look-ahead)
- Mixing & Mastering - Present (look-ahead)
- Mixing & Mastering - Deep (look-ahead)
- Mixing & Mastering - Traditional (look-ahead)
- Mixing & Mastering - Maximizer TP (look-ahead)
- Mixing & Mastering - Maximizer PK (look-ahead)

2.1.2. Loudness type

The loudness type combo-box and channel-link button allow you to configure which type of loudness measurement is used by the compressor. Each loudness type supports channel-split or channel-linked processing. Channel-split treats each channel independently, while channel-linked integrates across all channels.

Three of the supported loudness types use a popular modern measurement standard for perceived loudness called [LUFS](#). These loudness types are momentary (400ms window), short-term (3s window) and Integrated (infinite window).

There are also two [peak](#) loudness types (True Peak and Peak) as well as traditional [RMS](#). True Peak is an improvement on [peak](#) which takes into consideration waveform behavior between samples, which allows it to properly respond to inter-sample peaks. Peak mode loudness types use a [blocksize](#) window.

Integrated loudness type window can be reset by pressing the reset button.

2.1.3. Momentary window

The momentary window parameter allows you to adjust the duration of the momentary window in milliseconds. Lower window durations respond more quickly to changes in loudness, while higher window durations provide a more stable loudness measurement. The momentary window is used by the Momentary [Loudness type](#).

2.1.4. RMS window

The RMS window parameter allows you to adjust the duration of the RMS window in milliseconds. Lower window durations respond more quickly to changes in loudness, while higher window durations provide a more stable loudness measurement. The RMS window is used by the RMS [loudness type](#).

2.1.5. Output type

The output type options allows you to configure which type of samples to output.

- Default output type corresponds to the compressor's output
- Bypass output type corresponds to the compressor's source
- Delta output type corresponds to the compressor's output minus source

Output gain and dry/wet are bypassed for Delta output type

2.1.6. Visualization

The visualization component displays a continuously evolving real-time view of your [source](#), [target](#), and [output](#) loudness over time. There are currently two main types of visualizations: [history](#) and [histogram](#). [History](#) displays a rolling window of the most recent history of loudness samples. [Histogram](#) divides this history up into buckets and displays the relative proportion of each bucket in real-time.

You can hover the mouse over a bucket to see the loudness range the bucket corresponds to.

The visualization component also displays the [source range](#) and [target range](#), allowing you to precisely target these ranges based on the dynamics of the audio. Detailed visualization settings can be adjusted from the [visualization](#) tab, with some features also available by right clicking on the visualization component.

You can double-click the visualization component to resize and hide everything else.

2.1.7. Raster type

The raster type parameter allows you to specify the type of raster visualization. You can select between [history](#) and [histogram](#) views, with either loudness or delta values. Loudness values match the currently selected [loudness type](#) while delta values take the difference between [source](#) and [target](#) or [output](#) loudness.

Delta view is useful for determining how much loudness is being changed, and how much residual there is between [target](#) and [output](#) loudness.

You can also visualize [target](#) vs [output](#) residual via [breath residual](#) option.

2.1.8. Sidechain

The sidechain toggle allows you to enable/disable sidechain processing. When enabled, the source signal will come from the plug-in's sidechain channel(s). This allows you to drive compression/expansion using an external source, which can be used for example to [duck](#) the sidechain signal.

Since this compressor is capable of both compression and expansion, the reverse of ducking is also possible (source amplifies in response to sidechain signal above threshold).

2.1.9. Dynamics type

The dynamics type combo-box allows you to change the dynamics algorithm type.

- Source:Target
 - transforms from source loudness range to target loudness range
- Compression (downward)
 - reduces the loudness of audio above [threshold](#) by [ratio](#).
- Compression (upward)
 - increases the loudness of sounds below [threshold](#) by [ratio](#).
- Expansion (downward)
 - reduces the loudness of audio below [threshold](#) by [ratio](#).
- Expansion (upward)
 - increases the loudness of sounds above [threshold](#) by [ratio](#).
- Soft Limiting
 - limits the loudness of audio above [threshold](#).

2.1.10. Ballistics type

The ballistics type combo-box allows you to change the compressor's ballistics behavior.

- Natural (dithered) applies light dithering to natural ballistics deltas.
- Natural (direct) applies normal smooth attack and release ballistics.
- Inertial (dithered) applies light dithering to inertial ballistics deltas.
- Inertial (direct) applies attack and release ballistics with inertia.

Inertial ballistics are more responsive to sudden changes in loudness, but may sound less natural.

Dithering can help reduce quantization noise, but may be less accurate

2.1.11. Low-Level type

The low-level type combo-box allows you to change the compressor's behavior for loudness levels below the [source range](#).

- Makeup applies the difference between [target](#) min and [source](#) min to all loudness levels.
- Bypass leaves loudness levels below [source](#) min untouched (no makeup).
- Gate applies silence to loudness levels below [source](#) min.

The [attack/release](#) and [linear attack/release](#) ballistics apply to all low-level types

2.1.12. Source range

The source range parameter allows you to specify a min and max loudness value for the [source](#) signal. Loudness values within the source range are mapped onto the [target range](#) to determine the amount of gain reduction or amplification. The resulting target is then followed by the compressor using the configured [attack/release](#) and/or [linear attack/release](#) parameter settings.

See [range sliders](#) for information about range sliders in general.

2.1.13. Learn source range

The learn source range parameter allows you to enable or disable learning for [source range](#).

While active, [source range](#) will be adjusted to accommodate the min and max input loudness. This is useful for quickly setting the source range to match the input loudness. Once the source range is set, you can disable learning to prevent the source range from changing.

Learn accounts for source distribution asymmetries, making it great for hitting a target range accurately.

2.1.14. Source range presets

Source range presets are provided for convenience to quickly reset the range or select the full range.

Currently, the following source range presets are available:

- APU Default (-20/10)
- Full Range (≥ -70)

2.1.15. Target range

The target range parameter allows you to specify a min and max loudness value for the [target](#) signal. Loudness values within the source range are mapped onto the [target range](#) to determine the amount of gain reduction or amplification. The resulting target is then followed by the compressor using the configured [attack/release](#) and/or [linear attack/release](#) parameter settings.

See [range sliders](#) for information about range sliders in general.

2.1.16. Learn target range

The learn target range parameter allows you to enable or disable learning for [target range](#).

While active, [target range](#) will be adjusted to accommodate the min and max input loudness. This is useful for quickly setting the target range to match the input loudness. Once the target range is set, you can disable learning to prevent the target range from changing.

2.1.17. Target range presets

Target range presets are provided for convenience as an example set of industry related loudness ranges. The average loudness and limiter peak are provided in the preset names in parenthesis. These values were transcribed from the [RTW](#) delivery standards page.

Currently, the following target range presets are available:

- APU Default (-20/10)
- Spotify Loud (-11, -2)
- Spotify/Tidal/YouTube (-14, -1)
- Deezer (-15, -1)
- Apple (-16, -1)
- AES Streaming (-18, -1)
- EBU R128 Broadcast (-23, -1)
- US TV ATSC A/85 Broadcast (-24, -1)
- Netflix Online Streaming (-27, -1)

2.1.18. Threshold

The threshold parameter allows you to specify the loudness level where the compressor becomes fully activated. The precise meaning of threshold is dependent on the current [dynamics type](#).

This parameter is unused for source:target [dynamics type](#).

2.1.19. Ratio

The ratio parameter allows you to specify the compression/expansion ratio. The ratio is applied to signals above or below [threshold](#) (depending on upward/downward compression/expansion).

This parameter is unused for source:target [dynamics type](#).

2.1.20. Max range

The max range parameter allows you to specify the max gain reduction or gain amplification for the associated [dynamics type](#). Note this parameter has a positive and a negative version, which is automatically displayed based on the current [dynamics type](#).

This parameter is unused for source:target and soft limiting [dynamics type](#).

2.1.21. Attack/Release

The attack and release parameters allow you to specify the amount of time it takes for a given amount of gain reduction to be applied to the [source](#) signal. Attack refers to the “attack” of applying gain reduction, while release refers to the “release” of this gain reduction. More generally, release refers to gain amplification since it isn’t always necessary for the compressor to “attack” prior to “release” and the two are not always correlated.

Attack and release are traditional ballistic parameters, similar to what you will find on most compressors. See the [Linear attack/release](#) parameters for more direct attack/release behavior. In addition to specifying the attack and release in milliseconds, you can specify them in sample, BPM-relative units, or ratio of [look-ahead](#) duration.

2.1.22. Linear Attack/Release

The linear attack and release parameters allow you to directly specify the rate of change in gain reduction applied to the [source](#) signal. Attack refers to the “attack” of applying gain reduction, while release refers to the “release” of this gain reduction. More generally, release refers to gain amplification since it isn’t always necessary for the compressor to “attack” prior to “release” and the two are not always correlated.

Linear attack and release parameters are different from traditional ballistic parameters. The actual rate of change in gain reduction with the traditional controls is dependent on both the attack or release time and the residual between target gain reduction and current gain reduction. This means the alterations made to the source signal are a complicated combination of the audio’s dynamics. For linear attack and release, the actual rate of change is constant.

This setting can be configured to apply directly, completely overriding the [attack/release](#) ballistics, or as a limit to their rate. The former can be a way to get extremely fast response, while the latter can serve as a barrier to prevent the compressor from attacking or releasing too rapidly during extreme fluctuations in [source](#) loudness.

2.1.23. Look-ahead

The look-ahead parameter allows you to configure how far in advance the compressor looks for the [source](#) signal. This can be used to align the compressor’s response to account for the loudness window time and/or [attack/release](#) times, or simply as a creative effect.

Look-ahead requires that the [look-ahead budget](#) parameter is configured via the [settings](#) tab. This method of look-ahead configuration is used to support real-time adjustments of look-ahead without causing stuttering which would otherwise interfere with A/B testing.

Note that [look-ahead budget](#) adds latency to the compressor in order to function.

2.1.24. Look-ahead mode

The look-ahead mode parameter allows you to specify the look-ahead algorithm.

- OFF completely disables look-ahead
- Offset reacts to the newest sample in the look-ahead window
- Average reacts to the average of the entire look-ahead window
- Peak reacts to the newest sample in the look-ahead window (attack only)
- Floor reacts to the newest sample in the look-ahead window (release only)

The look-ahead window duration can be configured using the [look-ahead](#) parameter.

2.1.25. Knee width

The knee width parameter allows you to specify the width of the transition regions at [source range](#) and [target range](#) min and max boundaries. Inside these regions, the compressor's gain reduction is smoothly transitioned from on to off and vice versa. Adjusting the knee width effectively makes this transition more or less sudden.

The lower knee transition begins at source min and ends at source min minus knee width. The upper transition begins at source max and ends at source max plus knee width.

Note that the target range knee width is implicitly scaled in proportion to the amount of compression/expansion.

2.1.26. Dry/Wet

The dry/wet parameter allows you to configure the percentage of compression to mix into the dry signal.

This feature can be used to apply [parallel compression](#).

2.1.27. Output gain

The output gain parameter allows you to apply additional gain to the output of the compressor.

2.2. Loudness



The loudness tab contains parameters relating to loudness measurements. You can configure adaptive window behavior here, which can be used to reduce transient artifacts during compression.

2.2.1. Adaptive startup

The adaptive startup parameter enables or disables adaptive startup behavior for loudness measurements. When enabled, the loudness window begins small and grows as the window fills. This can be useful for reducing transient artifacts at the start of playback.

2.2.2. Adaptive edges

The adaptive edges parameter enables or disables adaptive edge detection for loudness measurements. When enabled, loudness values crossing the configured [adaptive threshold](#), or silence, will be detected and the loudness window will be reset. This causes the [adaptive startup](#) behavior to apply to every transition to/from silence. This can be useful for reducing transient artifacts during sudden changes in loudness.

2.2.3. Adaptive type

The adaptive type parameter allows [adaptive startup](#) and [adaptive edges](#) to be configured for Silence or [adaptive threshold](#) modes. Silence mode will only reset the loudness window when the input signal is silent. Adaptive threshold mode will reset the loudness window when the input signal crosses the configured [adaptive threshold](#).

2.2.4. Adaptive threshold

The adaptive threshold parameter allows you to configure the threshold used by [adaptive edges](#) in [adaptive threshold](#) mode. This threshold is specified in the same units as the current [loudness type](#).

2.3. Limiter



The limiter tab contains parameters relating to the built-in limiter. The limiter scans the [limiter lookahead](#) region of blocks for peak or true peak loudness values over the [limiter ceiling](#). When these peaks are encountered, the current gain reduction trajectory is adjusted to accommodate the peak.

2.3.1. Limiter mode

The limiter mode parameter lets you decide which peaks, if any, will be used for detection. The following options are available:

- None: No limiting
- Peak: Digital peaks are detected
- True Peak: Intra-sample peaks are detected according to the LUFS True Peak specification.

2.3.2. Limiter threshold

The limiter threshold parameter determines the loudness at which the limiter starts to apply gain reduction. Loudness values above threshold are raised to the [limiter ceiling](#).

2.3.3. Limiter ceiling

The limiter ceiling parameter determines the maximum peak loudness value that the limiter will allow. When the limiter detects a peak above this value, it will pre-emptively begin to apply gain reduction. Note that make-up gain is not automatically applied.

2.3.4. Limiter threshold link

The limiter threshold link parameter allows you to link the [limiter threshold](#) and [limiter ceiling](#) parameters. When linked, the ceiling will automatically adjust to match the threshold.

2.3.5. Limiter lookahead

The limiter look-ahead parameter determines the amount of time to look into the future when detecting peaks. Shorter lookahead times will result in more aggressive gain reduction, with the results becoming more gritty and distorted. Longer lookahead times will result in more transparent gain reduction, but will push the average level of the audio further down.

2.4. Visualization



The visualization tab contains parameters related to real-time configuration of the compressor's visualizations. This section describes these various settings.

2.4.1. Auto range

The auto range parameter enables or disables the visualization's auto range finding capabilities. This feature operates by analyzing the continuous histogram at each frame to determine a reasonable range for that moment in time. This range is then followed and adjusted smoothly over time.

2.4.2. Visual range presets

Visual range presets are provided for convenience as an example set of industry related loudness ranges.

Currently, the following visual range presets are available:

- Auto Range
- Standard Reference (≥ -60)
- Full Range (≥ -70)
- Custom Range ...

2.4.3. Loudness range

The loudness range parameter allows the visualization loudness range to be set manually. In order for this range to be enabled and used, the [auto range](#) parameter must be disabled.

Note that there are separate loudness ranges for windowed loudness types and peak loudness types.

See [range sliders](#) for information about range sliders in general.

2.4.4. Layout options

The options menu allows you to configure the layout of the visualization. The following options are available:

- [layout source mode](#) : Specifies the layout mode for the [source](#) signal.
- [layout target mode](#) : Specifies the layout mode for the [target](#) signal.
- [layout output mode](#) : Specifies the layout mode for the [output](#) signal.
- [raster type](#) : Specifies the type of raster visualization.
- [show thresholds](#) : Specifies whether or not to show thresholds.
- [show xfer func](#) : Specifies whether or not to show transfer function.
- [show peak meter](#) : Specifies whether or not to show peak meter.
- [show axis](#) : Specifies whether or not to show the axis.
- [show axis lines](#) : Specifies whether or not to show axis lines.
- [show param context](#) : Specifies whether or not to show parameter context while certain parameters are changing.

Each layout mode has an AUTO option which automatically selects the best layout mode based on the current context.

2.4.5. Layout source mode

The layout source mode parameter allows you to configure the layout of the visualization's source panel.

Currently, the following layout source modes are available:

- None (No source panel is displayed)
- Auto (Automatically selects the best layout mode based on the current context)
- Source only (Draws just the [source](#) signal)
- Source over Target (Draws the [target](#) signal first, then the [source](#) signal on top)
- Source over Output (Draws the [output](#) signal first, then the [source](#) signal on top)

2.4.6. Layout target mode

The layout target mode parameter allows you to configure the layout of the visualization's target panel.

Currently, the following layout target modes are available:

- None (No target panel is displayed)
- Auto (Automatically selects the best layout mode based on the current context)
- Target only (Draws just the [target](#) signal)
- Target over Source (Draws the [source](#) signal first, then the [target](#) signal on top)
- Target over Output (Draws the [output](#) signal first, then the [target](#) signal on top)

2.4.7. Layout output mode

The layout output mode parameter allows you to configure the layout of the visualization's output panel.

Currently, the following layout output modes are available:

- None (No output panel is displayed)
- Auto (Automatically selects the best layout mode based on the current context)
- Output only (Draws just the [output](#) signal)
- Output over Source (Draws the [source](#) signal first, then the [output](#) signal on top)
- Output over Target (Draws the [target](#) signal first, then the [output](#) signal on top)

2.4.8. Delta field type

The delta field type parameter allows you to specify the type of delta field used by the visualization.

The delta field is basically a mapping from each source pixel on the screen to a destination pixel. Pixel shaders are used to iteratively apply this delta field using interpolation and some light dithering effects. This feature is purely for aesthetic purposes and has no effect on the audio.

[Raster type](#) effects are drawn into the delta field on each frame.

2.4.9. Super-sampling mode

The super-sampling mode parameter determines the amount of super-sampling to use when rendering visualizations. The default render dimensions match the default plug-in window size. You can enable super-sampling to multiply this resolution by up to 4x.

2.4.10. Bits per pixel

The bits-per-pixel parameter determines the number of bits per pixel to use when rendering visualizations. The default bpp is the best performing. You can increase bpp to improve visual quality.

2.4.11. Breath residual

The breath residual parameter can be used to help visualize compressor behavior. For each frame of the visualization, the current residual between [target](#) and [output](#) loudness is calculated and used to multiply the magnitude of the visualization delta field.

This feature helps you find sections of audio where the compressor is having a difficult time keeping up with the amount of [target](#) gain. This can help guide the configuration of the [attack/release](#) and [linear attack/release](#) parameters. The direction of the multiplication factor is also set to reflect whether the residual occurred in the attack or release direction, so it's possible to get a sense of the overall balance between these residuals.

2.4.12. Bucket size

The bucket size parameter determines the size in pixels of visualization buckets (rectangles).

2.4.13. Histogram hold

The histogram hold parameter enables or disables “hold” mode for the [histogram](#) raster types. While “hold” mode is enabled, the histogram will continuously accumulate measurements indefinitely. This can be used to capture a full duration view of your source audio. This allows you to easily adjust ranges to target to a specific region of the full dynamic range.

The histogram will still continuously accumulate measurements while “hold” mode is disabled, but for each new measurement the oldest measurement in history is replaced. This results in a continuously evolving histogram of duration specified by the [history length](#) parameter.

2.4.14. Delta X/Y/T

The delta X/Y/T parameters are provided to the visualization [delta field](#) in order to manipulate the delta field in real-time. The exact behavior of each parameter depends on the active [delta field type](#).

2.4.15. Delta field

The delta field parameter allows you to enable or disable the visualization delta field.

The delta field is basically a mapping from each source pixel on the screen to a destination pixel. Pixel shaders are used to iteratively apply this delta field using interpolation and some light dithering effects. This feature is purely for aesthetic purposes and has no effect on the audio.

[Raster type](#) effects are drawn into the delta field on each frame.

2.4.16. Shader params

The shader parameters control various aspects of visualization rendering. The following options are available:

- Shader param 1: Delta field fade rate. Lower values leave longer trails.
- Shader param 2: Curve fill transparency. Lower values make the curve more transparent, leaving only the edge(s) visible.
- Shader param 3: Curve edge thickness. Lower values make the edge thinner.

2.4.17. Show thresholds

The show thresholds parameter allows you to enable or disable threshold drawing in the visualization.

If this option is disabled, thresholds will still be drawn while changing the associated parameters.

Thresholds include [source range](#), [target range](#), [threshold](#), and [knee width](#).

2.4.18. Panel params

The panel parameters control various aspects of panel rendering. The following options are available:

- Panel param 1: Panel alpha (transparency), 0.0 to remove panel completely
- Panel param 2: Panel hue (rotates through normal, source, target, output).
- Panel param 3: Panel hue intensity

2.4.19. History length

The history length parameter allows you to change the duration of visualization history. This history is a rolling-window where each new loudness sample pushes out the oldest sample, maintaining a continuous history length with the duration you've configured here. Note that this setting applies both to [history](#) and [histogram](#) raster types.

During [histogram hold](#) the histogram has an effectively infinite history length.

2.4.20. Show xfer func

The show transfer function parameter allows you to enable or disable transfer function drawing in the visualization.

If this option is enabled, the transfer function will be drawn next to the configured [raster type](#).

The transfer function is rendered using the current [Loudness range](#).

2.4.21. Show peak meter

The show peak meter parameter allows you to enable or disable peak meter and gain reduction drawing next to output gain.

Positive gain reduction is drawn from bottom-up, negative gain reduction is drawn from top-down.

2.4.22. Show axis

The show axis parameter allows you to enable or disable axis drawing in the visualization.

The axis is scaled relative to the [loudness range](#), which can be manual or [auto range](#).

2.4.23. Show axis lines

The show axis lines parameter allows you to enable or disable axis line drawing in the visualization.

axis lines are drawn at each [Show axis](#).

2.4.24. Show param context

The show param context parameter allows you to enable or disable param context text bubbles during parameter changes.

2.5. Theme



The theme tab contains parameters related to color themes. Here you can select between the bank of color theme presets, or configure the individual colors yourself.

2.5.1. Theme presets

The theme presets combo-box allows you to switch between a variety of theme presets. Each theme has a light and dark variation. Switching between theme presets will load values into [source color](#), [normal color](#), [target color](#) and [output color](#).

Currently, the following themes are available:

- APU Default
 - APU Red
- APU Orange
- APU Yellow
- APU Chartreuse
- APU Green
- APU Spring
- APU Cyan
- APU Azure
- APU Blue
- APU Violet
- APU Magenta
- APU Rose

2.5.2. Source color

The source color parameters control the red, green, and blue components of the [source](#) color.

The source color is used in a variety of contexts, from interactive widgets to visualization elements. This color signifies that an element relates to the [source](#) signal in some way. This color is expected to contrast against the [normal color](#) to some extent.

2.5.3. Normal color

The normal color parameters control the red, green, and blue components of the “normal” color.

The normal color is used in a variety of contexts, from interactive widgets to visualization elements. This color signifies that an element is essentially neutral, not related to [source](#), [target](#) or [output](#) signal. This color is expected to contrast against the [source color](#), [target color](#), and [output color](#) to some degree.

2.5.4. Target color

The target color parameters control the red, green, and blue components of the [target](#) color.

The target color is used in a variety of contexts, from interactive widgets to visualization elements. This color signifies that an element relates to the [target](#) signal in some way. This color is expected to contrast against the [normal color](#) to some extent.

2.5.5. Output color

The output color parameters control the red, green, and blue components of the [output](#) color.

The output color is used in a variety of contexts, from interactive widgets to visualization elements. This color signifies that an element relates to the [output](#) signal in some way. This color is expected to contrast against the [normal color](#) to some extent.

2.6. Settings



The settings tab contains various additional parameters. These parameters are broken down between General and Latency parameters. Since Latency parameters impact [delay compensation](#), changes to these parameters are deferred until you click the Apply button. It is generally not advisable to automate the parameters in the Latency section.

2.6.1. BPM

The bpm option allows you to set the BPM used by tempo-relative parameters.

2.6.2. Host BPM

The host bpm option enables usage of the host's BPM for tempo-relative parameters. When disabled, the [bpm](#) parameter is used instead.

This parameter is not available (nor applicable) to the standalone application.

2.6.3. Dither

The dither parameter determines the strength of ballistics dithering. This parameter applies only to dithering [ballistics type](#) modes.

2.6.4. Velocity sensitive knobs

If enabled, this will turn on velocity-sensitive dragging, so that the faster the mouse moves, the bigger the movement to the knobs. This helps when making accurate small-scale adjustments.

This parameter is saved at user scope, so it will be remembered between sessions.

2.6.5. Load preset visuals

The load preset visuals option allows you to enable/disable the loading of visualization and [theme](#) settings when loading [presets](#). All other settings will still be loaded; this setting only has an impact on the visual appearance of the plug-in and does not impact the audio.

2.6.6. Load preset ranges

The load preset ranges option allows you to enable/disable the loading of [source range](#) and/or [target range](#) when loading [presets](#). All other settings will still be loaded.

2.6.7. Dark mode

The dark mode toggle enables/disables dark mode. When enabled, theme colors have their brightness inverted.

2.6.8. Offline HD

The Offline HD combo-box allows you to configure the Offline HD feature. This feature, when enabled, performs super-sampling when the compressor is operated in “offline” mode. Offline mode typically applies to DAW features including Freeze and Export, but not real-time operation. The exact scenarios where this feature depends on the host application.

When enabled and active, super-sampling 2x is effectively the same as halving the block size. Likewise, 3x and 4x effectively divide the block size by 3 or 4 respectively. This has the end result of increasing the time resolution of loudness measurements. Although more loudness samples is generally better, it does come with a performance cost and so extremely low block sizes are currently only enabled for offline mode.

2.6.9. Blocksize

The block size parameter determines the time resolution of the compressor’s [source](#) loudness measurements. Generally speaking, lower block sizes will give more accurate results. However, lower block sizes also require more CPU resources, so it is necessary to find a balance. You can squeeze improved quality and/or performance out of the compressor by tuning this setting based on your available CPU resources.

Fast attack and/or release times may benefit from similarly low block size.

BPM units for block size are evaluated once at the time you press Apply.

2.6.10. Look-ahead budget

The look-ahead budget parameter determines the latency budget which is available for the [look-ahead](#) parameter. Once you have configured a look-ahead budget, you can adjust the [look-ahead](#) parameter in real-time within this range without introducing artifacts.

2.6.11. Delay compensation

The delay compensation parameter determines whether or not the plug-in will report latency to the host. Delay compensation is used by hosts to keep audio synchronized across channels.

The delay compensation option is not available (nor applicable) to the standalone application.

2.7. Update

The update tab allows you to check for the latest product versions. Just click “Check for updates” to see the latest version numbers. If you’re not running the latest version, you can click “Download” to open the download page in your default browser.

2.8. About

The about tab contains basic information about the plug-in. This is also where you can activate or deactivate your product keys and check license status.

3. Tips & Tricks

This section provides some general tips and strategies for using this compressor in a variety of contexts.

3.1. Compression/expansion

Compression is applied when the [target range](#) is smaller than the [source range](#) and the [source](#) signal is within the source range. If the [target range](#) is larger than the [source range](#), expansion is applied instead.

Signals outside of the source range have neither compressor nor expansion applied, but are amplified according to the difference between source and target min values. Manipulating the [source range](#) and/or [target range](#) allows you to compress, expand, and/or apply make-up gain.

One effective process for compressing or expanding with make-up gain is to:

- Configure your [source range](#) using the [histogram](#) and/or [history](#) visualizations for guidance
 - Target the particular range of loudness you want to compress or expand
- Configure your [target range](#) using [target range presets](#) and/or manual configuration.
 - Target the loudness range you want the [source range](#) to be mapped onto

3.2. Automatic gain control

This compressor is well suited for automatic gain control since it has support for multiple [LUFS](#) loudness types. For example, this configuration will provide very soft automatic gain control:

- Select one of the “Deep” [presets](#) which utilizes short-term LUFS loudness type
- Configure your [source range](#) so the minimum is above the noise floor and max is above the max source signal
- Configure the [target range](#) to be relatively narrow, centered at your target average loudness
- Lower the [attack/release](#) and/or [linear attack/release](#) settings as needed for transparency or responsiveness
- Fine tune [look-ahead](#) time if using a look-ahead preset

3.3. Gain riding

The compressor’s [source range](#) and [target range](#) can be dragged while maintaining range by clicking and dragging the slider area between the two range thumbs. The actual gain applied is always filtered through your attack/release settings.

4. Glossary

This section defines some of the concepts used within the software.

4.1. Range sliders

Range sliders are used throughout the plug-in in order to specify the upper and lower boundaries of a range. These ranges can be controlled via mouse in various ways.

- Click and drag the lower thumb to adjust the minimum value.
- Click and drag the upper thumb to adjust the maximum value.
- Click and drag the region between slider thumbs to move both values. This allows you to adjust the average loudness without expanding/contracting the dynamic range.
- Ctrl + click and drag the region between slider thumbs to expand/contract range without changing the average (center) value. Drag the mouse up and down, left and right are ignored.
- Shift + click and drag the region between slider thumbs to combine both. Drag the mouse up and down to expand/contract range. Drag the mouse left and right to move both values.

4.2. History

One of the primary views into your audio that this software provides is real-time history. The history view maintains a recent history of loudness measurements, continuously displaying them in [FIFO](#) order. This view helps you understand how the signal is changing over time, in real-time. [History length](#) and [bucket size](#) can be configured dynamically without forcing the history to reset.

4.3. Histogram

One of the primary views into your audio that this software provides is a real-time histogram. Histograms in general provide a quick and intuitive way to understand the relative frequency of different measurements. This is very useful when judging the overall dynamic range of the audio. The histogram provided by this software is capable of changing [history length](#), [bucket size](#) and size continuously.

4.4. Source

The term “source” is used throughout the plug-in to identify the plug-in’s input source signal. This signal is represented in the user interface by the current theme’s [source color](#).

4.5. Target

The term “target” is used throughout the plug-in to identify the “ideal” target output loudness. The actual loudness achieved depends on [attack/release](#) and [linear attack/release](#) parameter settings.

4.6. Output

The term “output” is used throughout the plug-in to identify the actual output loudness. This can differ from [target](#) loudness due to non-instantaneous [attack/release](#) and/or [linear attack/release](#) parameter settings.

5. Credits

This software was developed by [APU Software, LLC](#) and is available as VST (windows x64/x86, macOS universal), Audio Unit (macOS universal), or Standalone Application (windows x64/x86, macOS universal). The software libraries below are utilized for portions of the software:

- [JUCE](#) (cross-platform audio and user interface framework)
- [Boost](#) (header-only algorithms)
- [libebur128](#) (loudness measurements)
- [melatonin_blur](#) (blur effects)

Demo video song credits:

- Annasara - A Lifetime Rolls By, licensed via [Audioio](#)



5.1. MIT License (libebur128)

Copyright (c) 2011 Jan Kokemüller

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

5.2. MIT License (melatonin_blur)

Copyright (c) 2023 Sudara Williams

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.