RPS SOFTWARE

3DCombine

User's Guide — concise edition

> **Automatic 2D → 3D video conversion for VR headsets and anaglyph glasses.**
> A reference for experienced 3D users. It covers the conversion workflow, the output formats, the settings that affect the result, and the filters.

3DCombine v6.231 (64 bit)
Windows 10 1903 or later · macOS (video features)
Assumes you already work in stereo: disparity, depth maps, SBS and over-under packing are used without explanation.

Contents

1.The interface 3
- Menus and the button ribbon 3
- Global settings 4
- Preview 3D mode 4
- What the Lite version leaves out 5

2.Automatic 2D → 3D video conversion 6
- Choosing a method 6
- Step by step: source to headset 6
- Choosing an output format for your device 7
- Saving the video 7
- Batch processing 7

3.The manual four-step video workflow 8

4.Stereo output formats 9
- Format reference 9
- Choosing a container, codec and quality 9
- Audio 10

5.Image 2D → 3D conversion 11

6.Working with existing stereo material 12

7.Filters reference 13

8.Troubleshooting 14

1 The interface

The window has a menu bar, a button ribbon below it, and two large image panes. Filters are added to a filter list and applied to the open media.

1.1 Menus and the button ribbon

The File menu

Item	Purpose
`New`	Close all open files and reset the editor.
`Open Image(s)`	Sub-menus: *3D* (a stereo pair), *2D Left*, *2D Right / Depthmap*. Multiple files open as a batch.
`Open Video`	Sub-menus: *3D*, *2D Left*, *2D Right*, *Dual* (two stereo streams in one file).
`Save 2D Image(s)`	Save the left or right pane as a flat image — useful for exporting a depth map.
`Save 3D`	Enter Preview 3D mode (same as the ribbon button).
`Batch Process Videos`	Run the current filter list over a queue of video files. Professional only.
`Project`	Save or load a `.prj` file containing the filter list and its settings.
`Exit`	Close the application.

The Guides menu

This is the fastest route to a result. For video conversion:

Entry	What it does
Video 2D to 3D	
`Automatic`	Converts a 2D video to 3D using image analysis. Runs *Intelligent Moving* at the default depth. No prompts — pick the file, choose a format, save.
`Customise Output`	Same conversion, but you choose the method, whether to output a depth map instead of a converted image, and the depth strength. See §2.1.
`Manual Step 1` … `Manual Step 4`	Four separate entries: generate keyframes, turn them into a video depth map, optionally clean it up, and convert it to stereo. See chapter 3.
Video Editing	
`Combine left and right`	Combine two separate videos (left and right) into one stereo video.
`Repair Anaglyph`	Convert a colour anaglyph video back to a full-colour stereo pair.
`Align`	Correct a stereo pair for vertical misalignment and exposure.
`Vuze`	Stitch a Vuze video into a panorama.
Image 2D to 3D	
`Automatic`	Convert 2D images to 3D using the neural depth model at the default depth strength. No prompts.
`With Depthmap`	Convert 2D images to 3D using a depth map you supply.
`From Portrait Mode Photo`	Use the depth data stored in a phone photo taken in portrait mode.
`For Looking Glass`	Tile the result for a Looking Glass or other tiled display.
Image Editing	
`Combine left and right`	Combine two separate images into a stereo pair.
`Repair Anaglyph`	Convert a colour anaglyph image back to a full-colour stereo pair. Professional only.
`Align`	Correct a stereo pair for misalignment and exposure.

> **Prompt user**
> The `Guides → Prompt user` item is on by default. Each step of a guide shows a short message explaining what to do. Switch it off if you know the workflow; the same information appears in the status bar instead.

Other menus

- `Edit` — *Details* shows the dimensions of the open images and video (frame rate, size, audio), and *Undo* reverts the last filter.
- `Settings` — the two global settings described in §1.2.

- `Help` — the About box (which also reports whether hardware video compression and GPU acceleration are available) and a link to a short survey. Two submenus carry the web links: *Manual Pages* has the online 3D formats page and the batch mode and filters page; *On the Web* has the homepage, Facebook and YouTube.

The button ribbon

The ribbon is a set of toolbars you can show and hide individually. Click the double-line handle on its left edge to undock it and move it anywhere, including off the main window. The toolbars are:

Toolbar	Contains
Simple	Everyday operations: flip, mirror, swap, greyscale, crop to same size, alpha channels.
Tweak	Manual adjustments: align, rotate, resize, crop, pad, smart pad, convergence, colour filter, gamma.
Automatic	Filters that make their own decisions: align, match, auto-scale, crop borders, de-ghost, find edges.
Adjust	Interactive sliders for convergence, rotation, cropping, scaling, brightness and gamma, applied to the displayed image. Keeping the result adds the adjustment to the filter list.
VR	Panorama and headset preparation: Cardboard, Oculus, Fisheye, and the Vuze tools.
Depth	Depth map creation and the 2D → 3D filters. The *Video 2D->3D* and *Image 2D->3D* buttons sit at the front of this bar.

The ribbon also carries the media navigation controls — rewind, step back, step forward, skip, and play — which operate on the open movie. The *Filters* button shows the list of filters applied to the current media, with buttons to edit, remove or clear them.

1.2 Global settings

Two settings in the `Settings` menu affect every automatic 2D → 3D conversion, for images and videos alike. Both are remembered between sessions, so you set them once.

Background Fill Quality

Controls how the background revealed by the warp is filled in.

Setting	Method	Character and cost
Quick	Native fill	Fastest. Fills from the revealed surroundings, favouring background pixels. Soft and slightly blurred where the gap is large. The default on a fresh install.
Balanced	OpenCV Telea inpainting	Moderately slower. Produces noticeably more structure in the filled background.
High Quality	LaMa neural inpainting	Slowest. The best background quality, with the fewest obvious clones.

High Quality is spent selectively rather than on every frame: in a video the neural filler is used on the first frame of a scene, and on later frames only when more than about 5% of the frame is still blank. Otherwise the faster Balanced method is used. A still image has no previous frame to fall back on, so it always gets the full neural treatment. This is why High Quality costs less than a naive reading of the table suggests, and why it is worth enabling for a final render.

> **Practical advice**
> For a first look at how a conversion will go, leave Fill at Quick and use the Quick depth model. When you are happy with the structure, switch to High Quality and Detailed for the final render. The two stages cost very different amounts of time.

Model

Controls which neural network estimates the depth map.

Setting	Network	Character and cost
Quick	Monocular depth	Fast. Coarser depth boundaries. The default on a fresh install.
Detailed	Depth Anything 3 Mono Large	Substantially slower — it is a far larger network. Noticeably better object separation and depth ordering.

Both networks work at a fixed resolution — the source is resized before it reaches the model — so the depth estimation costs the same whatever the size of your file. The background fill is done at full resolution, so a 4K source still costs more to fill than a 720p one.

If the Detailed model cannot be loaded, 3DCombine says so once and falls back to Quick so the conversion still runs. The same applies to the High Quality fill model, which falls back to Balanced.

1.3 Preview 3D mode

The *Preview 3D* button (also `File → Save 3D`) switches the window into 3D display mode. It does three things at once:

1. It presents the left and right images as a stereo pair, so you can judge the result.
2. It fills the output format drop-down list with the available stereo layouts. **The format you select here is the format you save.**
3. It reveals the *Save Image*, *Fullscreen* and *Save Video* buttons. The video button is absent in Lite, and *Fullscreen* is a Windows-only feature.

Changing the format re-renders the preview immediately. Clicking *Preview 3D* again leaves the mode and returns to editing — any filters you have applied remain in the filter list.

On Windows, *Fullscreen* shows the 3D image full screen; the left and right arrow keys or the mouse wheel switch between the eyes, and `S` swaps them.

> **Previewing a video**
> Once a video is open, the navigation buttons step through it. Because the conversion is applied per frame, playback proceeds at the speed the conversion runs. The *skip* button is the practical way to sample a long film — jump ahead, look, jump again — without waiting for every frame in between.

1.4 What the Lite version leaves out

3DCombine starts in the full Professional mode, and the trial period gives you all of it. When the trial expires and no licence key has been entered, the application drops back to Lite, which is a still-image editor: the video features are switched off and the 2D → 3D guides are removed rather than merely disabled. If you are looking for `Open Video` and it is not there, that is the reason.

Feature	Lite	Professional
Open, combine, align and adjust still images	Yes	Yes
Open, convert and save video	No	Yes
Batch process videos	No	Yes
Save projects (filter lists)	No	Yes
`Guides → Video 2D to 3D`	No	Yes
`Guides → Video Editing`	No	Yes
`Guides → Image 2D to 3D`	No	Yes
Repair Anaglyph	No	Yes
Open more than one image at a time	No	Yes
The VR and Depth bars on the ribbon	No	Yes
`Guides → Image Editing`	Yes	Yes

Two consequences are worth spelling out. In Lite the ribbon carries only the *Simple*, *Tweak*, *Automatic* and *Adjust* buttons, so the depth filters and the panorama tools are not offered at all — and with the `Image 2D to 3D` guide missing, there is no automatic 2D → 3D image conversion either. *Repair Colour Anaglyphs* is also absent from the *Automatic* bar, and the `Image Editing` guide offers *Combine left and right* and *Align* but not *Repair Anaglyph*.

The About box names the version you are running — *Lite* or *Professional* — and in Lite it points you at the upgrade.

2 Automatic 2D → 3D video conversion

This is the workflow the rest of the guide is built around: take a flat 2D video, convert it to stereo, and save it in a form a VR headset or a pair of anaglyph glasses can play.

2.1 Choosing a method

When you choose `Guides → Video 2D to 3D → Customise Output`, or press the *Video 2D->3D* button on the Depth toolbar, 3DCombine asks you three questions: the method, the output, and the depth strength. The first is the important one.

Method	Use it when
Intelligent Static	The scene is mostly static — a talking head, an interview, a locked-off camera. Produces the smoothest and most stable result of the three.
Intelligent Moving	General footage with camera movement. **This is the default and the best all-round choice.**
Frameshifting	When you want zero artifacts and will accept that the effect comes and goes with motion. No depth map and no neural network are used, so nothing is ever invented. Static shots show no depth; vertical motion shows little.
Intelligent Dynamic	A fallback when the other two wobble or misbehave. The depth map follows the frame exactly, so it is more accurate frame to frame but less smooth — watch for shimmer, which is the price of not blending.

The three Intelligent methods all use the same depth network, chosen in `Settings → Model`. They differ only in how the depth map is held steady from one frame to the next.

The output

Converted Image (the default) produces the finished stereo picture. **Depthmap** produces a depth map movie instead — the right-hand image of every frame is the depth map, unwarped. Use it when you intend to correct the depth maps by hand; see chapter 3.

The 3D effect magnitude

The depth strength, default **1.00**, ranges from 0 to 100 in steps of 0.05. It scales how far apart the two eyes are placed.

Start at 1.00. Raising it increases the depth effect *and* the amount of background that must be invented, so artifacts grow with it. Values much above 2 will strain most viewers; much below 0.5 will be hard to see.

2.2 Step by step: source to headset

The full path from a source file to a file your headset can play.

1. **Prepare the source.**
   3DCombine accepts any container FFmpeg can read, but a clean source converts better. If your file has letterbox bars, remove them first — the *Crop Borders* filter on the Automatic toolbar detects and removes black bars at the top and bottom.

   A higher bitrate gives the converter more real information to work with, and costs nothing in processing time — the depth network works at a fixed resolution.

2. **Set the global settings.**
   Choose *Background Fill Quality* and *Model* in `Settings`. Quick for both for a trial run; consider High Quality fill and the Detailed model for the final render.

3. **Start the conversion.**
   Choose `Guides → Video 2D to 3D → Automatic` for the default path, or *Customise Output* to pick the method yourself. Either way, select your 2D source file when prompted.

4. **Preview and tune the depth.**
   3DCombine enters Preview 3D mode. Use the navigation buttons and the *skip* button to sample several points in the film — a wide shot, a close-up, a fast pan. Press *Preview 3D* again to leave the mode. The depth strength belongs to the conversion itself, so to change it you change the setting on the *Video 2D->3D* filter and run the conversion again.

   Look at object edges. A smeared or torn background is the fill's fault; if the depth itself feels wrong, the model is at fault.

5. **Choose the output format.**
   Still in Preview 3D mode, pick the layout from the drop-down list. See §2.3 for which to pick.

6. **Save.**
   Click *Save Video*. Choose the container in the file dialog, then the codec, audio and quality in the Video Compression Options dialog. See §2.4.

2.3 Choosing an output format for your device

This is the decision that determines whether the result plays correctly, so it is worth getting right before you spend hours on a conversion. The full format list, with suffixes and MP4 tagging, is in §4.1.

Target	Format	Why
VR headset	Parallel	Side by side, full resolution. The safest general choice, and the one most players expect.
	Above/Below Full	Over-under, full resolution. The traditional VR projection layout; use it if the headset expects over-under.
	Above/Below	Over-under, half height per eye. Only for a display that expects the half-height packing.
Anaglyph glasses	Colour Anaglyph	Red/cyan — the standard, and the only one that works with ordinary cardboard glasses. The left eye is greyscaled, so colour is lost permanently.
	Grey Anaglyph	Both eyes greyscale. The most predictable result, and best for monochrome material.
3D TV or projector	SBS/Sensio	Half side-by-side, both eyes squeezed into one frame. The usual choice at 1920×1080 or 1280×720.
	Parallel	Full side-by-side at full resolution — for two projectors, for authoring a 3D disc, or for editing before a final encode.
	Interlaced / DLP 3D	Field-interlaced or checkerboard, for older shutter hardware and DLP projectors.

Two points matter for headsets in particular.

- **Aspect ratio.** Headset video is usually displayed through a lens, so the practical requirement is a source whose *per-eye* aspect ratio matches the device. If your device expects a square per-eye image, use *Tweak → Smart Pad* to letterbox the source into the right shape before converting; *Auto Crop* crops to a chosen aspect ratio, given as a fraction (6×4 would be 1.5).
- **Frame size.** 3DCombine does not enforce any resolution for headset output. The encoder pads the width up to a multiple of 16 and the height up to a multiple of 8, so avoid odd dimensions. 1920×1080 and 1280×720 are the usual choices.

Anaglyph cannot be played on a normal monitor without glasses. If you might want a full-colour version later, keep the Parallel or Above/Below master and derive the anaglyph from it rather than converting to anaglyph directly.

2.4 Saving the video

Saving is a two-dialog sequence. Both appear *after* the conversion has been set up, and both apply to the whole file.

1. **The Save File As dialog** chooses the container and the file name. Two containers are available: **MP4 video(*.mp4)**, the default, and **AVI video(*.avi)**.

   The extension is forced to match your choice.

2. **The Video Compression Options dialog** chooses the codec, the audio handling and the quality. See §4.2 for what each option means and §4.3 for audio.

> **Which container?**
> Use
> **MP4**
> for anything a headset, phone or browser will play: it is the default, and the only container that can carry the stereo tagging that lets players recognise the file as 3D. Use
> **AVI**
> only for the lossless or MJPEG options.

2.5 Batch processing

`File → Batch Process Videos` applies the current filter list to a queue of files.

The dialog takes a list of videos, an output folder, a container and a filename suffix. The suffix defaults to the one belonging to the currently selected stereo format — `_LRF` for Parallel, `_TB` for Above/Below — and is inserted before the extension, so `dolphins.mp4` becomes `dolphins_LRF.mp4`. The field is editable, and `<none>` suppresses it.

The Video Compression Options dialog is then shown once and applies to every file in the batch. You are asked whether to overwrite existing files, with options to apply the answer to all of them.

Batch processing cannot use the manual 2D → 3D filter, because that filter needs interaction with each file: 3DCombine says so and refuses rather than failing partway through.

3 The manual four-step video workflow

The automatic method suits most work. The manual method exists for when the automatic one is not good enough — typically when you want to correct the depth by hand for a critical scene, or when the footage defeats the neural network entirely.

It is a guided four-step process, and each step consumes the previous step's output. It is slow, and it is not available in batch processing.

Step	What happens
1. Generate Keyframes	3DCombine identifies the scene changes in the film and writes out the frames at those points. It offers to generate depth maps for them automatically; if you decline, it still writes the frames, and you supply the depth maps yourself. The key frames are written as numbered BMP files into the output folder, along with a video holding the source frame beside its depth map.
2. Process Depthmaps	The depth maps are merged with the source movie into a single parallel video — source on the left, depth map on the right — with each frame's depth map interpolated between the surrounding key frames. This is the file you correct.
3. Post-process (optional)	Matches and smooths the depth maps to reduce blockiness. It can take a very long time and often changes little, so it is worth skipping unless the depth maps are visibly rough.
4. Convert to Stereo	Applies the depth maps to the source frames and produces the final stereo video in the format of your choice.

Adding keyframes by hand

Automatic keyframe detection misses many scene changes. To add more: reopen the movie with `File → Open Video → 2D Left`, enter Preview 3D mode, navigate to the frame you want, and press *Save Image* to write it out. The file type does not matter — these frames are only used to make depth maps and can be deleted afterwards.

3DCombine's own key frames carry a `b` suffix and mark the *start* of a scene; the frames without it mark the end. To use your own depth map, replace the corresponding `depthmap_nnn.bmp` file, keeping the name exactly — no leading zeros — so the correspondence with the frame is preserved. BMP is lossless and the safer choice.

> **Where the depth map comes from**
>
> When you decline automatic depth map generation, 3DCombine still writes `depthmap_nnn.bmp` files, but they contain the *original frame*, not a depth map. They are placeholders for you to replace. If you leave one in place, that scene will be converted with a nonsense depth map — the picture will look wrong in a way that is hard to diagnose later.

Editing depth maps by hand

A depth map is a greyscale image, so any layer-based image editor will do. Build a stack of transparent layers from nearest to farthest, painting each in a progressively darker grey, with a black layer at the bottom for infinity. Where a surface recedes gradually — a road, a sea, a floor running away from the camera — use a gradient rather than a flat grey: a flat grey says "this is all at one distance", which is one of the most common reasons a hand-made depth map looks wrong in 3D. GIMP is the usual choice, and there are free macro sets that paint a whole object with a single grey in one stroke.

4 Stereo output formats

4.1 Format reference

These are the formats offered in the output drop-down in Preview 3D mode. The format you select there is the format you save.

Format	Layout	MP4 tag	Suffix	Notes
Parallel	Side by side, full	Side by side	`_LRF`	Full-resolution stereo pair, left then right. The default in Preview 3D. The right choice for a VR headset, for two projectors, and as a master to edit from.
Above/Below Full	Over-under, full	Top-bottom	`_TBF`	Full-resolution stereo pair stacked. The traditional VR projection layout.
Above/Below	Over-under, half height	Top-bottom	`_TB`	Each eye at half height. For displays that expect the half-height packing.
Cross Eyed	Side by side, reversed	—	`_RLF`	Right eye first. Viewed by crossing your eyes rather than parallel-fusing.
SBS/Sensio	Side by side, half width	—	`_LR`	Both eyes squeezed into one frame. For 3D TVs and projectors at 1920×1080 or 1280×720.
Bigscreen VR	Side by side, full	—	`_LRF`	Side by side, full width. Choosing this format also switches the app into its big-screen display mode, where each eye is stretched to twice the height and padded to the screen aspect — the shape wanted by a large curved display rather than a headset lens. The file itself is an ordinary side-by-side pair.
Colour Anaglyph	Single, red/cyan	—	`_CANA`	Red/cyan glasses. Left eye greyscaled, right eye keeps its colour in the green and blue channels.
Grey Anaglyph	Single, red/cyan	—	`_GANA`	Both eyes greyscaled. The most robust anaglyph to watch and to encode.
Yellow Anaglyph	Single, yellow/blue	—	`_YANA`	Yellow/blue glasses. Right eye greyscaled.
Grey Yellow Anaglyph	Single, yellow/blue	—	`_GYANA`	Yellow/blue glasses, both eyes greyscaled.
Magenta Green Anaglyph	Single, magenta/green	—	`_MANA`	Magenta/green glasses, both eyes greyscaled.
Interlaced	Field-interlaced	—	`_INT`	For older shutter and projector hardware.
Vertical Interlace	Vertically interlaced	—	`_VINT`	As above, with a vertical split.
DLP 3D	Checkerboard	—	`_DLP`	For DLP projectors that use a checkerboard pattern.
iZ3D	Shutter, interpolated	—	`_IZ3D`	An interpolated shutter format for iZ3D 3D glasses. Each frame carries a blended image, so the depth depends on the glasses' timing rather than on a clean stereo pair.
Vuzix	Side by side, half width	—	`_LR`	A side-by-side pair with a Vuzix linecode strip along the bottom edge, which tells the headset which playback mode the file is in. Each eye is scaled to half the width and padded to the Vuzix native aspect.
2D Left / 2D Right	Single view	—	`_L` / `_R`	One eye only. Useful for checking the conversion, and for extracting a flat view.

The suffix is the filename marker used by batch processing — see §2.5. The *Dual*, *Vuze*, *Vuze Full*, *VR180 Panorama* and *VR180 Photo* entries exist in the format catalogue but are input-only, so they never appear in the output list.

> **What the MP4 stereo tag does**
> When the container is MP4 and the format is a plain stereo pair — Parallel, Above/Below or Above/Below Full — 3DCombine writes a stereo packing flag into the file. Stereo-aware players and video services then recognise the file as 3D and present it correctly without the user choosing. Anaglyph, interlaced and single-view outputs carry no such tag, because there is no standard way to describe them.

4.2 Choosing a container, codec and quality

The Video Compression Options dialog offers three choices plus an estimate of the result.

Video codec

Codec	Hardware?	Containers	Notes
H265 (hardware)	Yes	MP4 only	GPU hardware encoding — NVIDIA NVENC, Intel QuickSync or AMD AMF, whichever your machine provides. Small files at good quality, and fast enough that the encoder is rarely the bottleneck. Offered only if a hardware encoder is available.
AV1 (hardware)	Yes	MP4 only	NVENC hardware encoding. Smaller still than H265 for the same quality. Offered only if your GPU supports it.
MPEG4	No	MP4, AVI	Software encoding. The most compatible choice, and the fallback when no hardware encoder is available.
MJPEG	No	AVI only	Software. Each frame is a JPEG. Large files, but an easy target for editing.
FFVHUFF (lossless)	No	AVI only	Lossless and fast. The quality setting does not apply.

3DCombine defaults to the best available codec for the container: H265, then AV1, then MPEG4. It deliberately does not offer H.264, for licensing reasons.

Hardware encoding is offered whenever a working hardware encoder was detected at startup. It is not affected by whether the neural network is running on the GPU — the depth network and the video encoder share the GPU happily.

Quality

Quality	Bitrate multiplier	Effect
Low	0.6×	Smallest file. Visible compression artifacts, which the inpainting then has to work around.
Medium	1.0×	The default. A sensible balance for most material.
High	1.6×	Largest file. Worth it for a final render, particularly on a detailed 3D scene where the compression artifacts would otherwise read as inpainting artifacts.

The dialog shows an estimated bitrate, computed from the output frame size, frame rate, codec and quality. It is a guide rather than a promise, and for MPEG4 and MJPEG — which use a constant-quantiser mode — it is labelled *informational* because the actual rate depends on the content.

> **Use the estimate as a planning number**
> Estimated bitrate × duration gives you the expected file size, which is worth checking before a long render — especially since the stereo file is roughly twice the width of the source, and the converted result is often considerably larger than the original.

4.3 Audio

Audio options depend on the container you chose.

Option	Available in	Notes
Copy (original …)	Offered when the source's audio muxes cleanly into the chosen container	Passes the source's compressed audio through untouched — no decode, no re-encode, no quality loss. **Choose this whenever it is offered.** It is the default when available, and it preserves surround sound.
MP2	Both containers	The universal fallback. Modest quality, maximum compatibility.
FLAC	MP4 only	Lossless compression. Larger than MP2 but no quality loss.
Opus	MP4 only	Good quality at a low bitrate. Note that some older players do not handle Opus in MP4.
PCM (lossless)	AVI only	Uncompressed. Very large files.

If the source has a DTS or AC3 soundtrack, **Copy (original …)** is strongly preferable — re-encoding to MP2 throws away the surround mix. The option is labelled with the source codec, for example **Copy (original AC3)**.

One caveat to be aware of: the audio written from a re-encoded source can drift out of sync with the video over a long file. If that happens, re-saving with the original audio track is the reliable fix.

5 Image 2D → 3D conversion

Images follow the same logic as video, with one fewer problem: there is no temporal stability to worry about, so each image is converted independently.

Automatic

1. Choose `Guides → Image 2D to 3D → Automatic`.
2. Select the source image when prompted.
3. Enter Preview 3D mode, choose the output format, and save with *Save Image*.

The guide runs at the default depth strength and asks nothing else. To choose the strength yourself, use the **Image 2D->3D** button on the Depth toolbar instead — it asks for the 3D effect magnitude, which defaults to 1.00.

You can select several images at once and they are processed as a batch.

With Depthmap

Use this when you already have a depth map — one you made in an image editor, or one exported from a previous conversion. 3DCombine asks for the 2D image first and the depth map second, and then for the 3D effect magnitude.

To export a depth map for editing, use the *Create Depthmap* filter on the Depth toolbar, then `File → Save 2D Image(s) → Right` to write the depth map out as a flat image.

From Portrait Mode Photo

Phones store depth information in portrait-mode photos. This guide reads it and produces a 3D image directly, with no neural network involved.

For Looking Glass

Produces a tiled image for a Looking Glass or any other tiled display.

Image output formats

Images can be saved in a wider range of formats than video. The file format drop-down in the save dialog offers:

Format	Notes
JPEG (*.jpg)	An ordinary flat image.
PNG (*.png)	An ordinary flat image, with alpha.
Windows Bitmap (*.bmp)	Lossless and uncompressed.
TIFF (*.tif;*.tiff)	An ordinary flat image.
Photoshop (*.psd)	A single flattened image.
RAW YUV 444 (*.raw)	Uncompressed YUV 4:4:4.
JPEG Stereo (*.jps)	A stereo pair in a single JPEG file.
Bitmap Stereo (*.bms)	A stereo pair in a single bitmap file.
Fuji (*.mpo)	The multi-picture format used by Fuji 3D cameras.
Stereo JPG (*.mpo.jpg)	An MPO file with a JPEG extension.
Stereo PSD (*.dual.psd)	Two layers, one per eye.
H4V JPG (*._2x1.jpg)	A four-tile 2×2 layout, for H4V-compatible viewers.
Quilt (*.qs5x4.jpg)	A 5×4 quilt of 20 tiles, for a Looking Glass.
Large Quilt (*.qs8x6.jpg)	An 8×6 quilt of 48 tiles.

> **A note on still-image anaglyphs**
> The same anaglyph caveat applies to stills as to video: the left eye is greyscaled in a red/cyan anaglyph, so colour information is lost permanently. Produce it from a full-colour master rather than from an already-anaglyphed intermediate.

6 Working with existing stereo material

Not every job is a conversion. These are the guides for material that is already stereo, or that needs repair before use.

Combine left and right

Two separate videos — one left eye, one right eye — are combined into a single stereo video in the format of your choice. Load them with `File → Open Video → 2D Left` and `2D Right`, then run the guide.

This is also the way to make a full-colour stereo video from an anaglyph master, or to re-pack a stereo video into a different layout.

Repair Anaglyph

Converts a colour anaglyph back into a full-colour stereo pair. The colour that the anaglyph discarded cannot be recovered, so the result is never as colourful as the original recording — but the depth is preserved, and you get back a clean stereo pair you can re-pack as you like. The *De-Ghost* filter on the Automatic bar removes ghosting from anaglyph or interlaced sources.

Align

Corrects a stereo pair for vertical misalignment and exposure differences. Useful when combining two separately recorded sources.

Vuze

Stitches a Vuze video into a panorama for a Vuze headset. The VR toolbar also carries *Crop Vuze*, *Vuze Auto* and *Vuze Stitch* for finer control.

Panorama preparation for headsets

The VR toolbar converts a wide or fisheye image into the spherical projection a headset expects:

Filter	What it does
Cardboard	Converts a panoramic image for Google Cardboard. The default aspect ratio is 1.79, which suits most devices.
Oculus	Produces a 180° or 360° panorama for a headset or a panorama viewer, and asks how much of the view to fill. A 180° panorama is square; a 360° panorama is twice as wide as it is high. The 180° output is capped at 2880 pixels high, the Quest's display height.
Fisheye	Corrects fisheye lens distortion, and can convert between equirectangular and perspective projections.

7 Filters reference

Filters are the manual equivalent of the guides. You add them to a list, and they are applied to every frame in order. The *Filters* button shows the list and lets you edit or remove entries.

> **Filter order is fixed**
> Filters cannot be reordered, and each one works on the result of the one before it. A new filter is always added at the end of the list, so to change the order you must remove the offending filter and add it again — or clear the list and start again.

Filters are grouped onto the ribbon toolbars. The complete list is:

Toolbar	Filters
Simple	Double Vertically, Halve Vertically, Mirror, Flip, Swap, Convert to Grey, Get Saturation, Smooth, Crop to Same Size, Invert, Display Alpha, Remove Alpha, Add Alpha, Get HLS
Automatic	Automatic Align, Auto Converge, Match, Auto-Scale, Crop Borders, Automatic Rotate, Find Edges, Find Brightest, De-Ghost, Repair Colour Anaglyphs
Tweak	Align Horizontally, Align Vertically, Auto Crop, Colour Filter, Rotate Images, Resize Images, Scale Image, Crop, Pad, Smart Pad, Converge, Gamma
Depth	**Video 2D->3D**, **Image 2D->3D**, Create Depthmap, Make 3D, Make 3D Tiled, Invert, Process Depthmap, Rotate 3D Image, Generate Lightfield, Interlace
VR	Cardboard, Oculus, Fisheye, Crop Vuze, Vuze Auto, Vuze Stitch

The **Adjust** bar is different from the rest. It offers eight adjustments that act on the displayed image through a pair of sliders. The filters are not added to the list until you say you want to keep them: when you move on from the image you are asked whether to store the changes, and if you keep them the matching filter is added to the list:

Adjustment	Left slider	Right slider
Adjust Alignment	Shift left image horizontally	Shift left image vertically
Rotate	Rotate left and right images	No function
Horizontal Crop	Crop left side	Crop right side
Scale Right	Adjust horizontal scale	Adjust vertical scale
Non Linear	Move in and out of screen	No function
Vertical Crop	Crop top side	Crop bottom side
Brightness	Darkest colour	Brightest colour
Gamma	Gamma	No function

Non Linear moves the 3D image in and out of the screen more naturally than a plain horizontal shift, and is usually better for close-up subjects where a standard convergence feels harsh.

The **Video 2D->3D** and **Image 2D->3D** buttons on the Depth bar are mutually exclusive: which one is shown depends on whether a video or an image is open.

The depth filters in detail

Video 2D->3D

The automatic video conversion described in chapter 2. Asks for the method, the output and the depth strength.

Image 2D->3D

The automatic still-image conversion. Asks for the depth strength.

Create Depthmap

Generates a depth map from the source and places it in the right pane. Choose *Left Only* for a 2D source; the other options — Stereo Pair (default), Stereo Pair (Algorithmic), Stereo Pair (Eth3D) and Stereo Pair (Middlebury) — derive a depth map from a real stereo pair.

Make 3D

Applies the depth map in the right pane to the image in the left pane. Asks for the depth strength, and whether to warp both images. Warping both spreads the displacement across the two eyes, which looks more natural and shows fewer artifacts.

Process Depthmap

Nine operations on the depth map: Match (Fixed), Match (Invariant), Align with image, Sharpen falling edges, Auto-scale brightness, Median filter, Maximise filter, Flatten background, and Flatten and Normalise.

Rotate 3D Image

Rotates a 3D image without destroying the 3D effect, by rotating the depth map with it. Asks for the 3D effect magnitude.

8 Troubleshooting

The result has a wobbly or vibrating picture

Usually a low-bitrate source. The converter has little real information to work from when it has to fill the background, so the fill is coarse and changes from frame to frame. Use a higher-bitrate source. Some users also report that upscaling the source to 4K before converting removes most of the wobble, because there are more pixels to draw the fill from.

If the wobble is in the depth rather than the fill, try a different Intelligent mode — Dynamic for a scene where the blending misbehaves, Static for a static scene that is picking up noise.

Objects appear at the wrong distance

The depth network guessed. A head may appear detached from shoulders, or a flat object like a painting may appear to have depth. Lowering the depth strength reduces how obvious the error is. If a specific scene is important, use the manual workflow and correct its depth map.

Background around foreground objects looks smeared or torn

That is the fill. Raise `Settings → Background Fill Quality` to Balanced, and then to High Quality. Raising the depth strength makes it worse, because a larger warp exposes more area to be invented.

The output will not open in a video editor

Some editors expect a constant bitrate. Remux the file to a fixed bitrate with a converter of your choice.

Audio drifts out of sync

Most often the audio was re-encoded rather than copied. Re-save with **Copy (original …)**. If the film was assembled from several parts, check that the original audio was trimmed to match each part.

The conversion aborted partway through

Most often the output disk filled up — the stereo file is much larger than the source. 3DCombine stops when it runs out of space. If it happens repeatedly, convert the film in parts and join them in a video editor. There is no resume function, so a long conversion that fails near the end costs you the whole run.

Output files will not delete

Occasionally a handle to an output file is left open after a long session and Windows refuses to delete it. Rebooting releases it.

Conversion is much slower than expected

Processing time is dominated by the number of frames. The depth network runs at a fixed working resolution, so a 4K source is not proportionally slower than a 720p one, although the background fill still costs more at higher resolutions. Frame rate matters: a 50 fps source takes twice as long as a 25 fps one. The depth model matters far most: the Detailed model is substantially slower than the Quick one. Higher fill quality adds to it.

As a rough guide, expect the conversion to take a substantial multiple of the film's running time. A 90-minute film at 24 fps is nearly 130,000 frames, and each one costs a neural network inference and a warp.

The Detailed model or the High Quality fill is not being used

Both are separate model files that must be present in the program folder. If they are missing, 3DCombine says so once and falls back to the next best option. Check for `DA3MONO-LARGE.onnx` and `inpainting_lama_2025jan.onnx` in the installation directory.

No hardware video compression is offered

Your GPU does not expose a supported hardware encoder. MPEG4 is the software fallback and works everywhere.

3DCombine v6.231 — User's Guide, concise edition. Prepared from the application source. For the latest version, the online format reference and tutorial videos, see www.3dcombine.com.