DE:NOISE for OFX Hosts

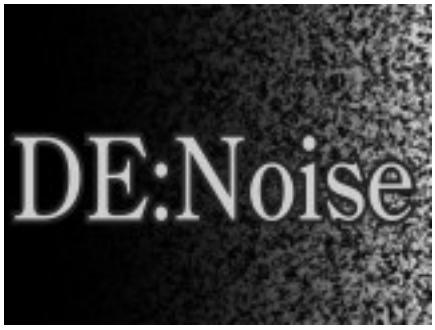

DE:Noise is a tool that removes excessive noise out of an image sequence. DE:Noise handles spurious frame-to-frame incoherence defects ranging from fine "digital/electronic" noise to blotchy spots (e.g. dirt on the film) in one simple easy-to-use tool. DE:Noise combines motion estimation techniques with feature-sensitive (edge preserving) spatial filtering methods to reduce the visual impact of various moving-picture problems such as: noisy video (as seen in many low-light capture contexts), excessive film grain, computer graphics renders affected by ray-tracing sampling artifacts, fingerprints and dust captured during film scan/transfer and printing, snow, drop-outs…

Internally, DE:Noise uses RE:Vision Effects' Academy Award ® winning optical flow technology to estimate the motion of each pixel between two frames. To create a noise-reduced result, DE:noise uses the estimated per pixel motion to warp the surrounding frames to match the current frame (i.e. to locally motion-compensate the surrounding frames), and then combines these intermediate frames with the current frame using one of 8 included modes in order to reduce noise. DE:Noise supports field-based material and allow you to mark cut points so that the temporal process does not attempt to work across cut boundaries

Combining motion-compensated images in this way tends to produce results that are sharper when compared to processes that rely only on spatial filtering and is usually as well resistant to inter-frame flickering. However, there are times when spatial filtering is necessary; for example, when there is little or no motion in the sequence, or the sequence is problematic to track. As such, we also provide feature-sensitive filtering options that attempt to smooth small artifacts while at the same time preserve object edges.

After all these feature-preserving noise-reduction methods are used, the final image may be too soft for your taste. As such, DE:Noise also integrates contrast and sharpening enhancement for a more complete "one-stop" final look control.

DE:Noise can be useful to generate better looking results, help other image processing tasks that are noise sensitive, and even help generate a cleaner still image from video for print.

NEW IN V2

Users of DE:Noise V1 should know:

- In V2 an additional input is added that allows you to apply DE:Noise to one clip but use an alternative clip as the tracking source.

- There has been a name change for the first two Temporal Processing modes. What used to be called Super is now called "Best - Forward Warping." And what used to be called "High" is not called "Best – Inverse Warping." They are in fact the same modes as before with just better descriptors for what they do.

OFX version of DE:Noise

DE:Noise supports 8 and 16 bits per pixel and floating point processing. The plugin has been tested on Mac, Windows (32 and 64b) and Linux (32 and 64b), and in the Foundry Nuke, Autodesk Toxik, and Eyeon Fusion. Other hosts probably work but are left to you to properly test for now.

For all applications:

- **DE:Noise** works most intuitively on progressive material and in projects, compositions or sequences that are specified to be progressive.
- If you have 3:2 pulldown in your source footage, you will want to **remove 3:2 pulldown** before processing with **DE:Noise**. If there is no motion between two frames then there will be no blur. This is useful to know when trying to figure out why **DE:Noise** doesn't apply proper processing to frames from a 3:2 pulldown sequence or animation done on "2"s, etc.
- DE:Noise usually works best on material with fields when the material has been deinterlaced with field blending. For example Toxik provides a deinterlacer operator which allows you to perform that operation prior to applying DE:Noise.

The plugins of **DE:Noise** are installed on Windows in C:\Program Files\Common Files\Ofx\Plugins , on Mac in "/Library/OFX/Plugins" and on Linux in "/usr/OFX/Plugins". The path defined by the environment variable OFX_PLUGIN_PATH is supposed to be searched as well by the host application. Consult your host user manual for alternative OFX path location.

Toxik 2008 and over (including Maya and Max Composite)

- Within Toxik the two plugins of DE:Noise show up in a grouping in the tools palette named OFX REVision Effects..

Known issue in Toxik 2008: occasionally while interacting you might end up with a red frame and that red frame will be kept in the cache. This is due to a problem in Toxik 2008 when interrupting (that is changing a slider during a render will cause an

interruption). This should not happen during a real sequence render. If you get a red frame, you will have to interact with a non animated parameter to force a redraw. This has been fixed in Toxik 2009.

Updating Plugins: For the future when we update a plugin and add and remove parameters it might break your comp/project/database when reopening a comp with the effect in it, either the tool will be said "not present" in which case you need to create a new instance and copy the values OR it might be missing some parameters entry in which case you will need to destroy the .xml file of the tool cached for example on Win32 here: C:\Program Files\Autodesk\Autodesk Toxik\resources\toolUi\ofx

Premult-Unmult: Toxik assumes unpremulted (straight) rgba inputs, and we provide an unpremultiplied image back to Toxik. There is an Unpremult button in the image import tool as well as a unpremult node tool which should be applied before our tools when images are in a premultiplied format.

Nuke

- Within Nuke the two plugins of DE:Noise show up in a RE:Vision Effects menu in the main menu bar

Premult-Unmult: Nuke assumes premultiplied rgba inputs, and we provide that back. If your input is straight/unpremultiplied, use the premultiplied tool before our tool.

Fusion

- Within Fusion the two plugins of DE:Noise show up in a REVision Effects menu in the tools menu. Note if you have native Fusion tools from us, it will show up in a different menu.

Premult-Unmult: Fusion assumes premultiplied rgba inputs, and we provide premultiplied images back. If your input is straight/unpremultiplied, use a tool with a post-multiply before our tool.

Others

- We don't know all the possible OFX host application in the universe, If you get unexpected behavior from another host then the listed above, please forward the problem to your host application and to us. Currently we only officially support the listed host applications in the Compatibility section of our website.

DE:Noise Controls

Pre Processing Controls:

Particularly when shooting indoors at night, it's often necessary to try to increase the dynamic range of the image in order for DE:Noise to "see" the noise or other artifacts. Note that the pre-processing controls are reversible (using the appropriate selection in the post-processing controls), so you can also use contrast enhancement to help the tracking/noise reduction process and then "undo" the preprocessing enhancements.

Pre Processing:

We provide the following methods to adjust contrast.
None: Performs no preprocessing. None also provides a quick way turn off preprocessing without having to set the Pre Contrast % to 0.0.
Contrast using global avg: Contracts or expands colors towards or away from the image average color.
Contrast using mid-grey: Contracts or expands colors towards 0.5 mid-grey. This is what traditional color contrast tools do.
Contrast using global avg is usually more useful, except in some particular cases such as a big white object appearing in a dark scene. In such shots, the traditional contrast (to mid-grey) can be used.

Of course, changing the contrast changes the noise characteristics

Note that turning Pre Processing to anything else other than None forces DE:Noise to work internally at 32 bit floating point per channel (even if the source is 8 or 16 bits per channel). This will increase the memory requirements by 48 bytes per pixel in a 16 bpc project and 64 bytes per pixel in an 8 bit project.

Pre Contrast %:

A value under 0 contracts the colors where a value of -100% produces a flat color. A value over 0% will expand the contrast range (moving values towards black or towards white).

Source Footage (left), Contrast Enhanced on the right.

Again remember that when you enhance the contrast, you also typically as well make the noise much more noticeable which might not be visually pleasant (and this tool is DE:Noise and not RE:Noise). Note DE:Noise has post processing controls (described later) whose purpose is to: a) undo the pre processing contrast enhancement applied to help find the noise and/or b) help give more punch back to the result image if the DE:Noise result is judged too blurry.

Spatial Noise Reduction

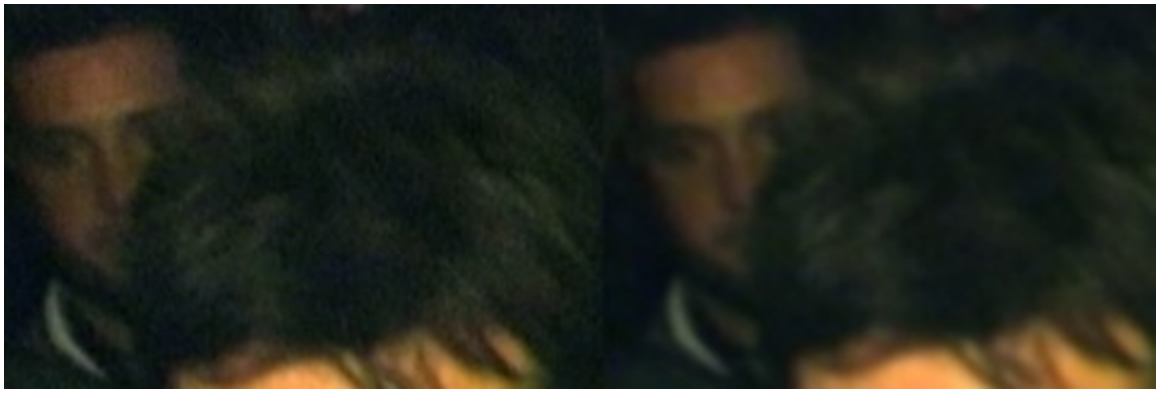

Example: On the left is without Spatial Noise Reduction, on the right is with the "Smarter Blur" mode. This is from a test sequence provided by Grant Davis (vjculture@yahoo.com)

Spatial Noise Reduction:

None: USE THIS TO TURN OFF Spatial Noise Reduction
This controls helps you spatially filter the residual noise that the temporal noise reduction component cannot deal with. Note the Spatial Threshold control greatly affects the result. Internally Spatial Noise reduction is applied then Temporal Noise Processing. If for some reason you wanted to change that order, just apply the effect twice in a row and simply turn OFF the appropriate control.
Diffuse: Averages values within the area defined by Spatial Radius. Pixels are averaged with the center pixel of the blur only if they vary from the center pixel by less than the Spatial Threshold % (a floating point setting described below). In this way, edges and

feature details are more likely to be preserved because pixels are only blurred together if they are "similar" (where similar is defined by the Spatial Threshhold %).

Smarter blur: Performs a blur constrained by the actual luminance. Similar to the Diffuse setting, however pixels are Gaussian blurred together instead of averaged. Pixels are Gaussian blurred and the resulting blurred pixel is only allowed to change by Spatial Threshold % or less. In this way, edges and feature details are more likely to be preserved.

Blur biased towards darks: Performs a blur biased towards the darks (that is, darker regions get blurred more than lighter regions). This mode is probably appropriate for footage that have large black areas (e.g. was shot on black).

Blur biased towards lights: Performs a blur biased towards the light areas. This mode is probably appropriate for footage that have large and noisy white areas. (more typical with widely manipulated HDR sources).

Spatial Radius:

This controls the kernel size for spatial denoising (and subsequently, how much spatial filtering can take place). A value of 0 effectively turns spatial denoising off. It's recommended to use spatial denoising when you have large patches of near constant values in your images (where it is the most perceptive).

As you know, the larger the radius of a blur is, the more pixels that are used to calculate a pixel. Notice the source on the left, the noise is very fine, so we probably don't need a large radius here.

Spatial Threshold %:

This control defines how much difference can take place between neighboring pixels when performing the spatial filtering (the neighborhood considered is set by the Spatial Radius setting). Pixels greater than the Spatial Threshhold % are either thrown out or weighted less when filtering. A value of 0 essentially turns off spatial filtering and a value toward 100% will include pixels in a greater range of pixel values.

Keeping the same large Radius as shown on the right in the previous figure, we now threshold it with a succession of small to large Spatial Threshhold values.

Here's an example of spatial noise reduction in action (sequence provided by Ami Sun). We have on the top left the "film" grainy look. The mode "Diffuse" usually works pretty well for this sort of fine white noise as shown right of it, with a smallish radius and a small threshold. However as

per the image on the lower-left side, as you increase the radius and raise the threshold (try to reduce more noise), the spatial filter tends to create overly smooth areas (look at the hair on the right of the frame for example). Then you probably want to switch to a blur method as on the bottom right. Note that if after this stage the result is a bit too soft for your taste, we provide post processing sharpening to help you restore details. The post-processing controls are described later.

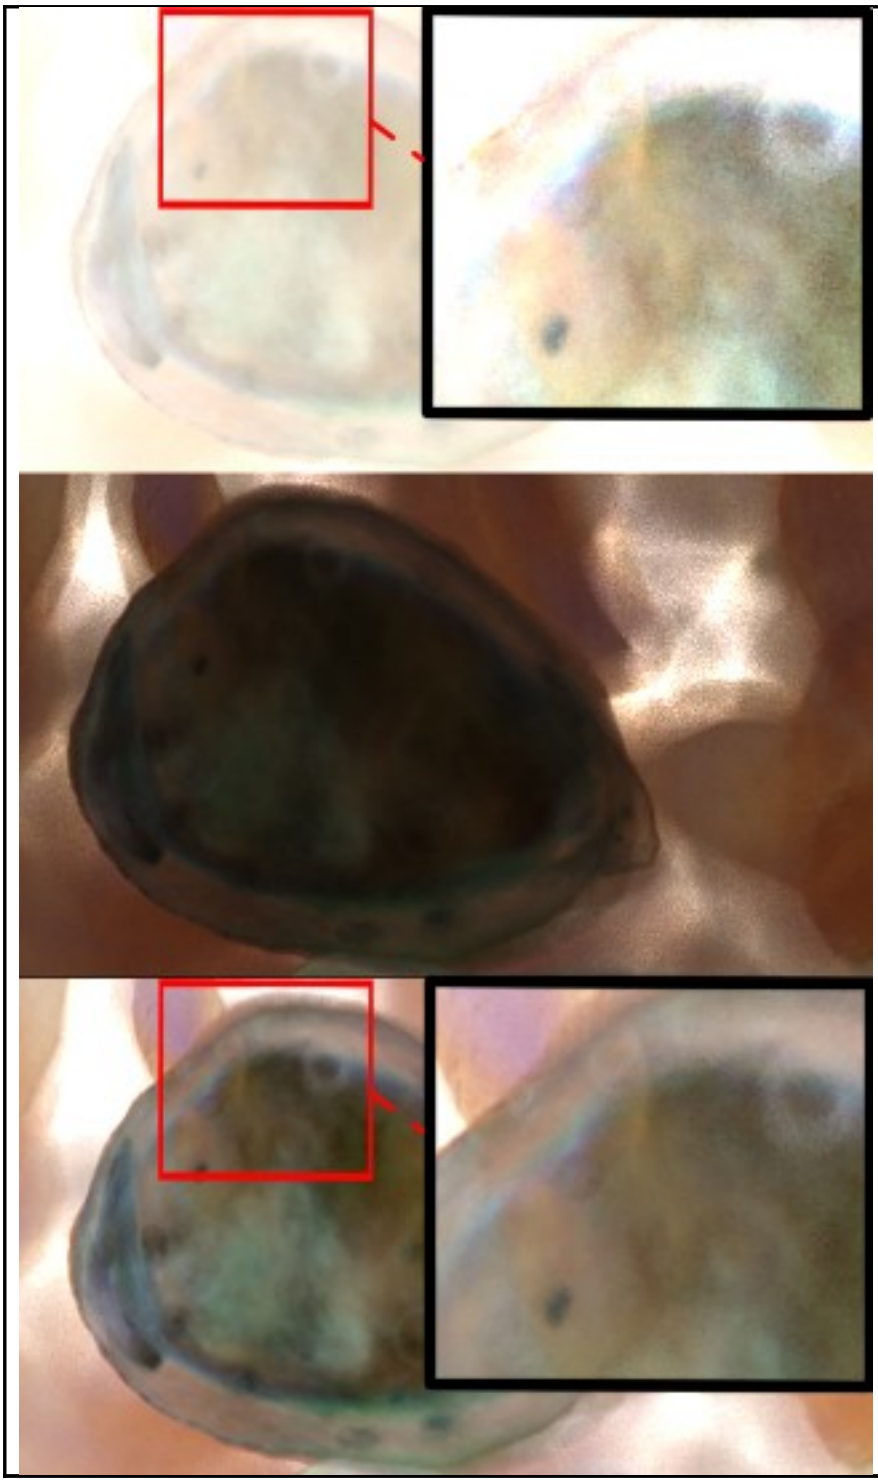

HDR and Noise

When you deal with high-dynamic range (HDR) images you benefit from a large color range, which is very helpful in color matching graphics to live-action and offload some of the rendering process to compositing.

However when dealing both with computer graphics and HDR photography,after merging the color spaces in a non linear manner can reveal some artifacts not visible (as shown in the top picture insert and then in the middle).

You can usually help out (as in the bottom picture) with DE:Noise, however you might prefer to denoise the result looking at color corrected results in the ballpark of the final result wanted.

Test provided by Chad Capeland, Anatomicaltravel.com

Temporal Processing

Temporal Process Mode:

The following modes are available. Note that each mode (except for None) uses optical flow to match up features between the current frame and the surrounding frames, and then performs per-pixel operations to reduce noise once the surrounding images are tracked and warped to match the current frame being processed.

None (Copy): This allows you to ignore temporal processing. USE THIS TO TURN OFF TEMPORAL PROCESSING.
Average: Warps the previous frame to the current AND the next frame to the current frame and mixes the 3 equally.
Median: Warps the previous frame to the current AND the next frame to the current frame and at each pixel keeps the median of the 3 values.
Average 2 most similar pixels: Warps the previous frame to the current AND the next frame to the current frame, discards the most different pixel and averages the remaining 2. This should work pretty well with drop-outs and similar single frame (non recurrent) noise.
Motion-weighted average: Warps the previous frame to the current AND the next frame to the current frame. This mode does not blindly blend the 3 frames, but attempts to preserve the sharpness in the image. It is longer to process as it does a lot more computations.
Average with prev: Warps the previous frame to the current and blends the two equally
Average with next: Warps the next frame to the current and blends the two equally. One reason to provide the two directions (to Previous and To Next) is for shots when you zoom in or out. Ideally you would like the other frame to be the wider field of view (zoom-in use previous, zoom-out use next).
Min: Warps the previous frame and next frame to the current and keeps the minimum value of the 3. This is usually good on things like little white dots, e.g. a shot with too much rain specularity, should do ok on shots with sporadic flash frames as well.
Max: Warps the previous and next frame to current and keeps the maximum value of the 3. Good for black dots.

Some practical examples are provided at the end. See the website gallery for more.

Temporal Quality

No MV: Turns off Motion Estimation completely
Best - Forward Warping: High Quality Motion Estimation and high quality warp filtering (slower but more precise)
Best - Inverse Warping: Same High Quality Motion Estimation but with a much faster warping technique. This setting will often prove to have enough quality.
Medium: Faster Motion Estimation
Fast: Sloppy result. It actually turns out that sometimes you want the motion estimation to be a bit sloppy (for example relatively static scene with falling rain would be such a

case) Basically you just want to make sure that the areas of the frame without motion and those with motion are properly differentiated.

Temporal Threshold %:

This controls how much a pixel is allowed to vary after it is processed (see Process Mode above). A value of 0% will in effect turn off the temporal denoising, a value of 100% is the absolute maximum deviation (inter-frame difference) from the source any particular pixel is allowed to change. To control this slider is completely crucial in obtaining the look wanted. Note such threshold-based cutoff can cause unnatural ghosting if you don't pay attention.

There are some contexts where you want this slider to be a high value, like 100%. At other times you may wish to limit the amount of change to a really tiny number like 3%, in part because the internal optical flow tracking does not work perfectly. If you see double-edges in the result you may want to lower the threshold amount.

Be careful about the setting of this value, if you see no difference before and after, make sure this value is not set to too low of a value. Low values are most useful for cases where you have "white noise", that is where the noise is at each pixel and varies +/- some small percent from the expected real value. High values are typically useful for defects that are not typical digital/electronic sensor noise.

In the following example, a laser animation appears "instantaneously" (some pixels don't really have a correspondence in next or previous frame), a larger threshold like 60% in this case could create ghosting as on the top right image. A lower value as below-right helps to reduce the ghosting, because pixels in the neighboring frames-in-time are discarded if they are too different. The proper value to put for this value will depend on the actual content.

ALT Track Source (Input source).

You now have the option to calculate the motion using an alternative clip.

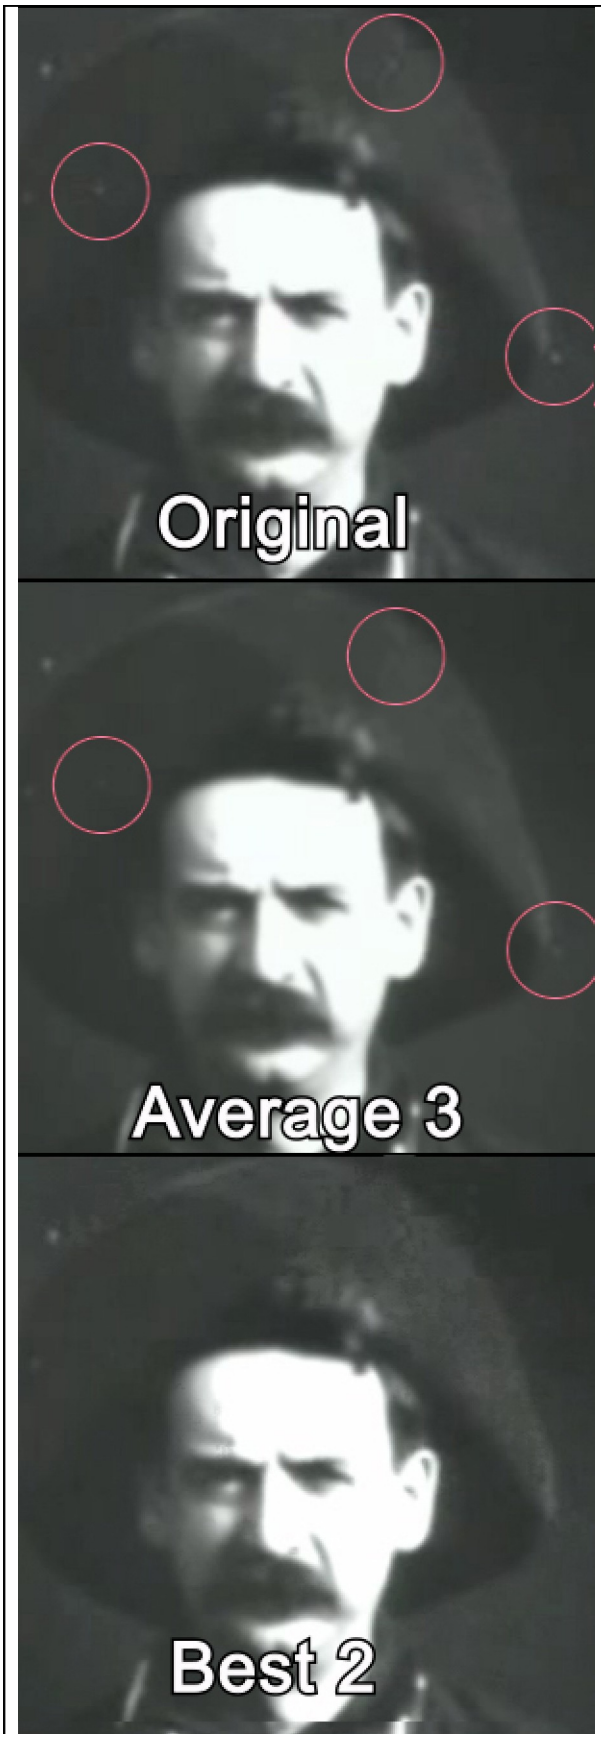

The different temporal process modes can produce dramatically different results. In this example we have another complication, which is that between certain frames in a sequence there is large luminance variance in certain areas. The variance comes from aging emulsion and is sometimes referred to as flickering.

Note the original image: in that particular frame there are 3 little blobs that we wish to remove. Note the circled blotch lower right, Average Mode (below) only attenuates it while the "2 Most Similar" mode (Best 2 at the bottom) completely eliminates it . However because the Previous frame has a sudden luminance shift, the "Best 2" mode creates some artifacts in the hat (see Best 2, right side of hat).

More discussion (all modes not pictured): Average, Motion Weighted Average, Average with Prev and Average with Next are "average-based" processes (mixes surrounding frames) while Median, Max, Min, Average 2 Most Similar have logical operators that might allow you to remove more noise but that may be more sensitive to large luminance shifts.

Sometimes, the best result will be to apply a mode with the appropriate logical condition (eg. "Min" to remove say the extra specularity of a rain machine gone wild) and then apply DE:Noise a second time to smooth the blending with the Average Mode.

This image is from Edwin S. Porter "The Great Train Robery".
http://www.archive.org/details/CEP146

Mark Segments:

If a menu is displayed for Mark Segments, it is because the host application supports the animation of menu choices. This setting allows you to specify edit (cut points) information you know about the source material. The purpose is to avoid artifacts at cut points, for example a frame with splicing tape is a large disruption of the inter-frame visual flow. Basically this setting allows you to identify the first frame of a sequence in a clip where the shots are not segmented so a frame from another shot is not mixed in with a frame from the current shot.

- **Cut A**
- **Cut B**
- **Cut C**
- **Spatial Only**
- **No MV**

DE:Noise will not blend frames between two segments marked with different Cut settings. For example, if one frame is marked as part of Cut A and the next frame is marked as part of Cut B, then DE:Noise for that frame will not interpolate between the two so will only use 2 frames. If your application does not support animated menus we might have turned this parameter into an integer with 3 states, note 0 and 2 are only these values if they are explicitely these values (consider the effect of default interpolation when setting such value).

Spatial Only and No MV are special modes meant to better match Twixtor 'Mark Segments' menu. When Mark Segments is set to Spatial Only then no temporal processing is done and only the spatial denoising is performed. . If Mark Segments is set to No MV the temporal processing is performed, but no motion estimation is calculated (optical flow) nor used to match up pixels. These special cases are to help streamline longer sequence where you might have dissolves or other effects that make optical flow tracking not very useful.

An integer setting is used if the host application does not properly support he animation of menu items. In this case the "Mark Segments" can range from 0 to 4, where each of the integer values match one of the settings described above:

- 0 is equivalent to "Cut A"
- 1 is equivalent to "Cut B"
- 2 is equivalent to "Cut C"
- 3 is equivalent to "Spatial Only"
- 4 is equivalent to "No MV"

Example:

The picture below is the source footage: on one frame there is a camera flash. The Mark Segments setting can be used to eliminate the spillover of the flash in to the surrounding frames by marking them as different cuts of the sequence.

This completely breaks the inter-frame comparison, look below how the bottom right of the image is pulled away from the frame

By animating the Mark Segments menu the problem is removed (see below)

Post Processing Controls

If the Denoising process produces results too smooth for your taste, we provide some post process enhancements to bring back details or remove haze.

PostProcess:

We provide two common methods to adjust contrast.

None (no post contrast, no sharpen): Turn off Post-processing altogether. Allows you quickly to visualize the result with or without post-processing.
Undo Pre-contrast: Inverse the contrast enhancement performed in pre-processing.
Contrast using global avg: Contracts or expands colors towards average color.
Contrast using mid-grey: Contracts or expands colors towards 0.5 mid-grey.

Post Contrast:

A value under 0 contracts the colors, a value of -100% is a flat color. A value over 0% will expand the contrast range.

Sharpen Amount:

A value of 0 turns the effect off. The sharpening used here is a form of unsharp blur. This control scales the amount of sharpening

Sharpen Radius:

Like unsharp masking, this is the radius that is used for the sharpening process. You typically want to use as large of a radius as you can get by while avoid restoring the small noise you work so hard removing.

Discussion

DE:Noise is a temporal noise reducer complemented by some spatial noise functionalities. Sometimes you might need to complement such process with other filters (including a standard blur or median filter).

DE:Noise does not address more "static" defects such as film scratches, hair in gate, sensor banding, wire or larger object removal and background filling, complete or partial damaged frame replacement, dust on lens, missing or partially defective (in a constant manner) pixels in a CCD array… Another product called RE:Fill might be useful for missing pixels problems. RE:Flex Motion Morphing tool can be used for frame replacement tasks.

DE:Noise also does not capture the signature of noise so you can use it later for example a re-graining process. Also, in some cases the temporal noise scheme of DE:Noise will not work very well (for example, when there is too much motion (the internal tracking fails) or no other motion than the noise itself). DE:Noise works best on a sequence (the temporal components won't work on a still frame). For the "no actual motion" video sequence case, we provided an assistant plugin "Frame Average" which is described above.

Because the optical flow process matches up features between frames DE:Noise's temporal noise reduction has the ability to remove noise without overblurring features. This is unlike a spatial noise that can oversmooth important features… The best practice may often entail using a bit of spatial noise reduction with a bit of temporal noise

reduction. Internally the filter first applies the spatial filtering and then temporally denoises. If for some reason you wish to reverse that order, you can always apply the effect/filter twice, turning on and off the spatial and temporal components appropriately.

DE:Noise does have some minimal spatial denoising controls. However sometimes other tools such as our SmoothKit filter set can have a great complementary role use in reducing spatial noise. For example if the object edges are very degraded by noise, SmoothKit Directional might help reconstruct better edges and so can be applied before DE:Noise. As well, one of the problems of SmoothKit Diffuse plugin is that in order to protect the object edge quality it will tend to leave "salt and pepper" noise over the image. By using DE:Noise temporal filtering after a filter that leaves some salt and pepper noise (speckles of darker and lighter points) such defects can be reduced or eliminated. Also sometimes the footage is so degraded that a Spatial Threshhold of 100% is all you can do and then all you can do to completely remove visible noise is to "over blur" the result. Then your option is to sharpen and play with contrast with the postprocessing controls or another tool after this effect. Also, look at the result in motion, often you want to reduce noise a bit but not completely eliminate it, and looking just at still images can be misleading as one naturally tends to over denoise the image when looking at stills.

Performance Optimization

Understand as well that if you use a multi-frame Temporal Process that any effects applied before this effect will generate multiple frame requests, so as you render the effects before might be requested 3 times to render the same frame, so in an application like After Effects it's best to do any preprocssing in a precomp, and then apply DE:Noise in a parent comp. This is because such application will then cache the intermediate renders of the precomp. For example it might be that one day you want to apply DE:Noise twice in a row (apply it to the first instance result). Since the first makes a request 3 times and the second one also 3 times, 3x3 = 9. You might want to organize your project to avoid longer using nesting (precomposition). Rendering a first pass to disk is always another way.

Useful Tricks

A. Removing Snow: Min and Max mode.

Here on the left we covered the image with "snow". On the right we used the "Min" temporal mode. This will remove the white speckles. If it was dark spots on a light background, you would then use "Max". Note you then have to raise the "Inter Frame Difference" (the Temporal Threshold) to 100% to completely remove the white speckles. Note as well that in certain cases you might want to reduce the excessive rain specularity etc… rather then try to eliminate so then less then 100% Difference can be appropriate. Also note the first and last frame will not have 3 frames to compare so might not do as well. Do not make setups using these frames.

This should work relatively well if the “motion pattern” is relatively fast moving or completely random. This will not work with static patterns like scratches or dust on the lens… (see our other tool RE:Fill for that). Also sometimes if the “motion pattern” is changing too slowly, you might want to reduce the quality of the Motion Estimation to Fast so the motion estimator does not catch these little overlays.

B. Handling Random Drop-Outs with 2 Most Similar Mode

On the left you see the frame before, the current frame and the frame after.
On the current frame we have two unwanted black bars.

In a case like this where the “drop-outs” are sporadically happening on a single frame the mode “Average 2 Most Similar Pixels” should work pretty well as it will discard the frame with the large difference compared to the two others. If the drop-outs appear here

and there in your sequence, you might want to animate off the effect on the frames without a problem to preserve the most overall quality. This technique should work as well for scanned film prints with dust on the film.

C. Isolated Problem: Dots in one area of the frame

With computer graphics ray tracing programs there are problems that can occur where a single object ends up producing little white dots here and there (whatever the reason: reflections, bad normals, not enough sampling, bad filtering). That noise is usually not of the same nature as video or film grain noise. It's more like "hard white" dots that come and go. In that case, it's often possible to handle this with temporal processing (eg set the mode to Min if it's white dots, remember to set the Temporal Threshold slider to 100%) and then to avoid affecting too much other areas of the image sequence, you might consider to make a rough track matte with a large feather and overlay that result over the original. By track matte here we mean a matte that is applied after the effect. In general when the defects are localized to an area of the frame, you can just run DE:Noise on the whole frame then composite the result back over the original through a very feather matte. The matte feather purpose is to mask any visible transition between the processed and unprocessed areas.

D. Dust-Busting:

Since film printing, inter-negatives, film to video transfer handling, often end up with small piece of dust on the celluloid that block the optical path (and then are scanned/recorded), and that sort of artifact is typically a single frame problem, you can usually deal with it with DE:Noise. Often the "Average 2 most similar pixels" is the mode that what you want to remove the dust artifacts. Larger debris (like hair), and scratches might not be removed by DE:Noise. In such case, then DE:Noise will be a first automatic pass that saves tons of time to your restoration process before a more manual clean up pass. Our own internal test used footage sometimes also had emulsion deterioration producing some luminance flicker over time. You probably have the option of applying the "Average Mode" instead OR if you have a render time budget, you can actually apply DE:Noise two times in a row, first with "Average 2 most similar pixels" to remove the dust (again probably with a Temporal Threshold – the Max Inter-Frame Difference -- of 100%) AND then with the Temporal Process Mode "Average" to eliminate flicker artifacts that might be locally amplified by "Average 2 most similar pixels".

E. Noise Localized in a Channel:

It can happen that the noise is localized in one color channel. In that case you either simply turn the image sequence to black and white using the channel as RGB of the to denoise source, apply DE:Noise then reset the source channel to that. Alternatively if the noise is just in the chroma, you could always use the transfer modes of your compositing applications to blend back the result over the original.

F. Additive White Noise:

Sometimes you can have very defective source with massive “white” noise added to the image.

Here (see image below) to restore the readability of the image (since it’s probably past the point of resulting into a perfectly nice image), we apply DE:Noise 2 times in a row.

First Pass (top right): a) We add a bit of contrast b) We blur using “Blur darker “
Second Pass: (bottom left) c) We now blur again using “Blur lighter” and then we turn the Temporal Denoising on to help reduce the blotchy look of oversmoothing.

DE:Noise Frame Average

DE:Noise comes with a companion tool that can be particularly useful for shots with no motion (or with very little motion). If there is no motion in the shot, it's probable you won't gain much from motion estimation). The frame averaging performed with this plugin is not your normal frame averaging because DE:Noise Frame Average contains a threshold parameter. Pixels are only averaged from other frames that are within the tolerance of that threshold value. You often want small threshold like 10-15 or 20 over a large number of frames (like 10 frames on each side) to completely cancel out the noise. This is a simple way to make a good quality still from video sequence in a low-light context.

Although you can apply the "Frame Average" plugin after "DE:Noise", be aware that in an application like After Effects you might want then to nest "DE:Noise" into a precomp so After Effects memory management helps you out, otherwise you might end up tripling the time to process.

Threshold:

Pixels are only averaged from other frames that are within the tolerance of that threshold value.

Frames Before and After

A value of 1 before and after becomes a 3 frames average (as it includes the current frame).

Auto Adjust In-Out

Let's say you set a "Frame After" of 5. What happens at the last frame? When this setting is off, the last frame is used 5 times in the averaging process (because there are no more frames after the last frame). When checked, this setting grabs more frames from the "before" side of the sequence so that all unique frames are used in the averaging process. Said another way: when checked, the plugin adjusts the number of "before" and 'after" frames at the beginning and end of the sequence so that all unique frames are used.