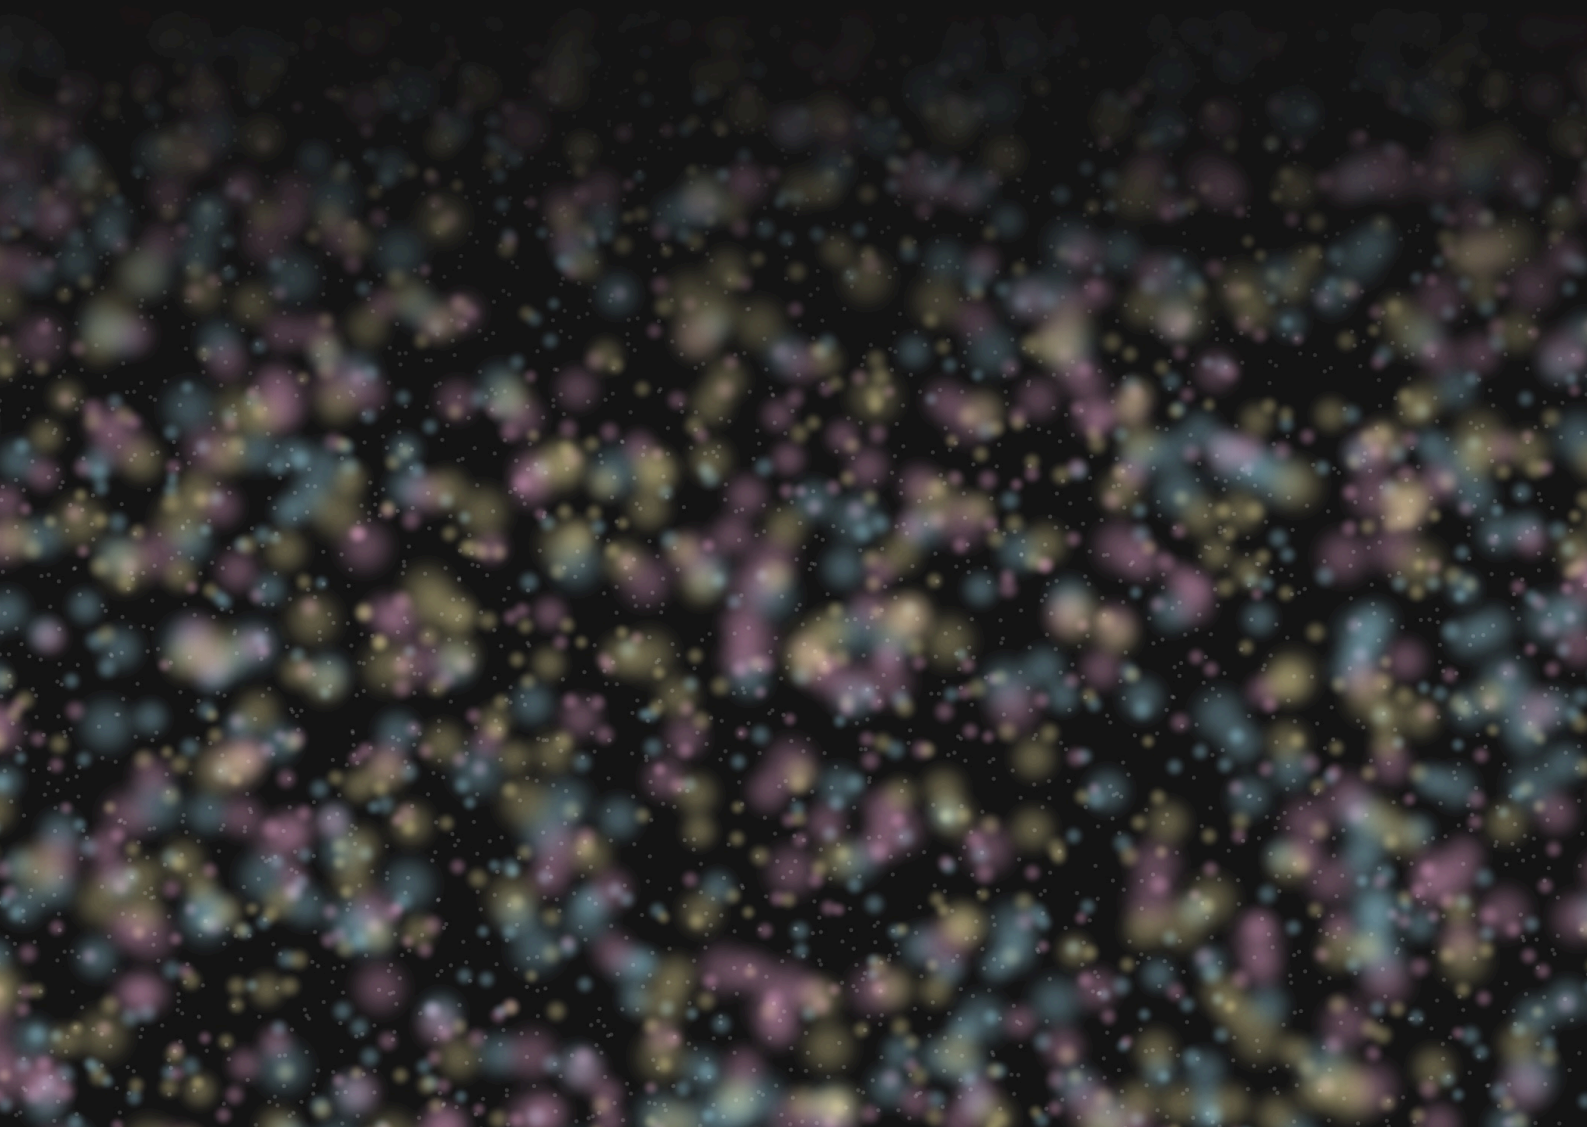
Filmic Texture

Technology Overview

November 2026



Contents

Summary **3**

How the Effect Works **4**

Film Size 5

Grain **6**

Crystals 6

Exposure and Development 6

Dye and the Scan 7

Adding Grain Without Changing the Picture 8

Grain in the Plug-in 8

Film Sharpness **9**

Matched to Published Curves 9

Softening at Small Pixel Sizes 10

Edge Outline in Density 10

Halation **11**

How the Glow Is Made 11

Clipped Highlights 12

Dust **13**

Colour and Signal **14**

Colour Space 14

What Is Preserved 14

Verification **15**

The Reference Against a Second Calculation 15

Across Resolutions 15

The Plug-in Against the Reference 15

Limits **16**

Technical Specifications **17**

References **18**

Summary

Filmic Texture is a plug-in for DaVinci Resolve. It adds the texture of film to a digital picture: grain; halation, a reddish glow around bright lights; the slight softening of fine detail; and specks of dust picked up when film is scanned. This paper describes how it calculates them, which published data it uses and what it leaves out.

Three ideas shape the design. First, everything is sized on film, not in pixels. You choose a film format, such as Super 35, and grain, glow and softening are measured in micrometres (µm, thousandths of a millimetre) on a frame of that size. A grain effect set in pixels gets finer as the resolution goes up. Here an HD and an 8K render of the same shot show the same texture, and the 8K render shows it in finer detail.

Second, the grain comes from a simulation of the light-sensitive crystals in film, which develop according to the light they receive. A Super 35 frame holds about 13 million of them, too many to develop for every frame of video, so the plug-in uses grain computed in advance from that simulation.

Third, the steps run in the order they happen in film: light spreads inside the film, development sharpens edges, and grain forms last. The glow around a light therefore has grain like the rest of the picture. Development works on how dark the negative becomes, as it does in film, so it draws no dark rings around bright lights.

The effect adds texture without regrading the picture. The grain averages to zero, flat areas come out unchanged, and very bright values, negative values and transparency (alpha) pass through.

Everything is calculated from equations that imitate film. Filmic Texture uses no AI or machine learning: nothing in it is trained on examples, and nothing is generated by a neural network.

The sharpness is fitted to Kodak's published curves for VISION3 250D [1]. The halation, the grain's strength and colour, and the dust were set by eye, with published measurements and real film scans as guides. Filmic Texture is a designed look, not a copy of a film stock, and the section on limits lists what it leaves out.

How the Effect Works

The effect takes a picture through the steps a film negative goes through, in the same order. The order matters because each step depends on a different property of the image:

- Light spreading inside the film, which causes halation and softening, depends on the amount of light. It is calculated in linear light, where twice the value means twice the light.
- Development, the chemical step, depends on how dark the negative becomes. That darkness is called density, and it is measured on a logarithmic scale, as stops are.
- Grain forms from the image that those two steps leave behind.
- Dust comes last, because it sits on the negative when it is scanned.

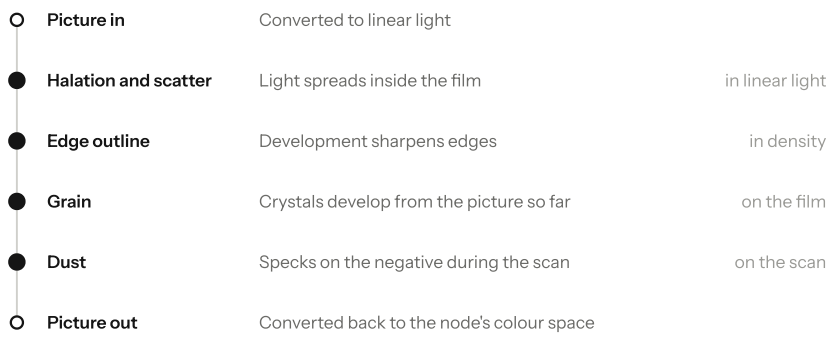


Figure 1. The steps of the effect, and the property of the image each one works on

Each step works on the result of the one before. The halation glow is part of the light the film receives, so it gains the edge outline and the grain like the rest of the picture. Grain forms after the softening, so it stays crisp even where the film has softened the detail.

Film Size

The first thing the effect needs is the size of the film frame. In a film camera, that size is set by the gate, the opening behind the lens, which is why the effect calls this setting Gate preset. You choose a format, and your picture is fitted onto it by cropping, never by stretching: a 16:9 picture on Super 35 uses the full 24.89 mm width and the central 14.0 mm of the height.

From the frame size in millimetres and in pixels, the effect knows how much film each pixel represents. Every size after that is in micrometres. The table gives the width of film that one pixel covers, for a 16:9 picture.

Format	Size (mm)	HD	4K UHD	8K UHD
Super 8	5.79 × 4.01	3.02 µm	1.51 µm	0.75 µm
Super 16	12.52 × 7.41	6.52 µm	3.26 µm	1.63 µm
Super 35 (4-perf)	24.89 × 18.66	12.96 µm	6.48 µm	3.24 µm
VistaVision (8-perf)	37.72 × 25.17	19.65 µm	9.82 µm	4.91 µm
65 mm (5-perf)	52.50 × 23.00	21.30 µm	10.65 µm	5.32 µm

On 65 mm a 16:9 picture is limited by the height of the frame, so it uses 40.89 mm of the width.

A speck of dye 8 µm across is larger than one pixel of a Super 35 frame at 4K and smaller than one at HD. In both cases each pixel shows the average of the patch of film it covers. A lower resolution therefore shows the same film with larger pixels, not a different texture.

The format also changes the look. A smaller format has to be enlarged more to fill the screen, and everything on it is enlarged with it. The grain that looks fine on VistaVision looks coarse on Super 16, and the glow around a lamp is three times wider, because Super 16 is a third of the size and is enlarged three times as much.

Grain

A colour negative holds its image as millions of tiny specks of dye. Each forms where a crystal of silver halide, the light-sensitive material in film, received enough light and was then developed. Filmic Texture makes its grain the same way, from a simulation of those crystals. The sections below describe the full simulation, in which every crystal in the frame develops from the picture. That is too slow for video, so the plug-in uses grain stored from it, as the last section, Grain in the Plug-in, describes.

Crystals

The simulated film has six groups of crystals: a fine, less sensitive group and a coarse, more sensitive group in each of its three colour layers. Their positions are chosen at random in three dimensions, through the thickness of each layer.

Crystals cannot overlap. A crystal placed too close to an earlier one is removed, which spaces them more evenly than purely random placement would. Within each group the crystals also vary: a quarter are smaller and less sensitive than average, and a quarter are larger and more sensitive.

At the default size the film holds about 38,000 crystals in every square millimetre, so a 16:9 Super 35 frame holds about 13 million: roughly 11 million fine and 2 million coarse. Which crystals a patch of film holds is decided by where the patch is on the film, the frame number and a seed, a number that picks one random pattern out of many. Pixels play no part in it, so a crop, a different resolution or a render done in pieces shows the same crystals in the same places.

Exposure and Development

Each crystal is given the brightness of the picture at its own position, after the halation and sharpness steps. Development then runs in three rounds. In each round, a crystal that has not yet developed either develops or does not, by chance, and the brighter its light, the better its chance. Once developed, it stays developed.

The coarse crystals are about four and a half stops more sensitive than the fine ones, so they develop first. In the shadows the image is carried by a few scattered coarse specks. Through the mid-tones the fine crystals fill in, and in bright areas nearly every crystal has developed.

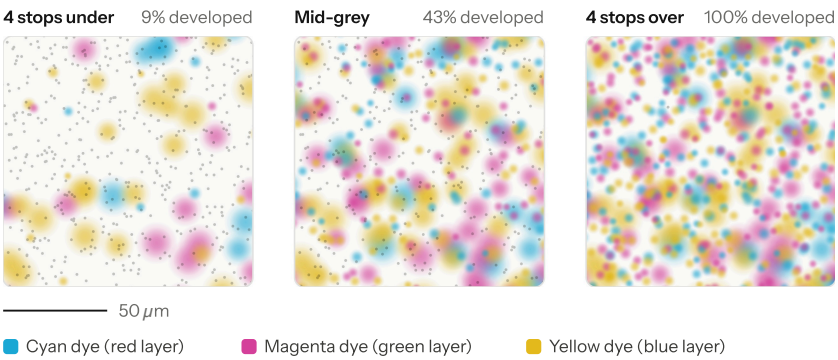


Figure 2. The same 0.12 mm square of simulated film at three brightness levels, drawn from the simulation's own output. Grey points are crystals that did not develop. Each disc is the dye speck of one that did.

Developed crystals also hold back their neighbours. A crystal surrounded by developed ones is less likely to develop in the next round, and its dye speck is slightly smaller. Real development does this too: a developing crystal uses up the developer around it, which on a larger scale draws the edge outline described under Film Sharpness. In the simulation it keeps dense areas from filling in evenly.

Dye and the Scan

Each developed crystal leaves a soft speck of dye: cyan in the red-sensitive layer, magenta in the green and yellow in the blue. At the default size the specks are 2.5 to 10 µm in radius. The colours of the three dyes follow the dye density curves Kodak publishes for VISION3 50D [2], a different stock from the 250D used for the sharpness.

Next the simulation works out how much light passes through the film. At every point it adds up the dye from all the specks above it and calculates the light that gets through at 31 wavelengths, from 400 to 700 nm. Each speck lets through a fraction of the light, and two overlapping specks let through a fraction of that fraction, so two specks darken the film by less than twice as much as one.

A simulated film scanner then reads that light. Its lens adds a blur of 0.45 µm, and its sensor has square pixels 2 µm wide. The scan is then averaged down to the pixels of your frame, each according to the exact area it covers. One pixel of a Super 35 frame at 4K averages about ten scanner pixels.

Adding Grain Without Changing the Picture

A scan of simulated film has a brightness and colour of its own, and if it replaced your picture it would change your grade. To avoid that, the simulation develops four versions of the same film, each with its own random arrangement of crystals, and combines them with plus and minus weights that add up to zero. The brightness and colour that the four versions share cancel out. What is left is grain alone, which averages to zero, and that is added to your picture.

How strongly the grain shows at each brightness is set by a table for each colour layer, and a final step adjusts its colour and contrast in the brightest and darkest areas. These were matched by eye to scans of processed negative viewed through a print film look. They are choices about the look, not measurements of a film stock.



Figure 3. Grain from the plug-in on five flat grey levels, at the default settings on Super 35 at 4K. Each tile is 150 pixels across, about 1 mm of film, enlarged 3×.

Three controls adjust the grain. Material size scales every length in the film together. Grain amount scales the grain that is added to the picture. Grain chroma sets how much of the dye colour the grain keeps, from neutral grey to full colour, without changing its brightness.

Grain in the Plug-in

Developing every crystal in the frame takes far too long to do on every frame of video. The crystal simulation is therefore run in advance, when the plug-in is built, on a square of film about 1 mm across at 33 brightness levels, half a stop apart from eight stops under to eight stops over mid-grey. The results are stored with the plug-in and make up most of its size.

When it renders a frame, the plug-in looks up the stored grain for the brightness at each point of the picture, blending between the two nearest levels. It covers the frame with overlapping patches of the stored film, each turned or flipped at random, and chooses a new arrangement for every frame. The stored grain was simulated on flat grey, so at a sharp edge the plug-in blends grain from two brightness levels, where the full simulation would develop crystals that see the edge itself.

Film Sharpness

Film does not record fine detail at full contrast. Light scatters inside the film, so a sharp line in the scene is recorded slightly wider and fainter. Film manufacturers publish this as a modulation transfer function (MTF), a curve that shows how much contrast survives at each level of detail. Detail is counted in cycles per millimetre, the number of pairs of dark and light lines that fit in a millimetre of film.

Matched to Published Curves

The effect's softening is fitted to the MTF curves Kodak publishes for VISION3 250D [1]. We traced 71 points from each colour's curve, from 2.5 to 80 cycles per millimetre. The fit follows them to within about 4 percentage points of contrast for red, 2 for green and 3 for blue.

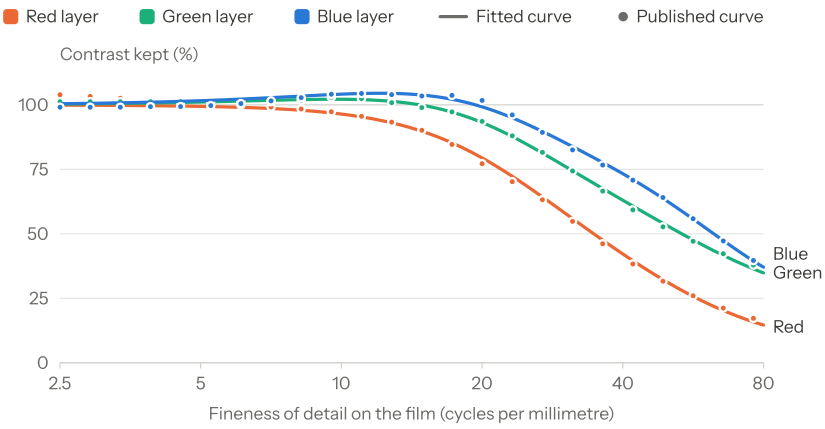


Figure 4. The fitted curve for each colour, against points traced from the published curves. Every third traced point is shown. Values above 100% at coarse detail come from the edge outline described below.

These are typical curves for the stock, as Kodak publishes them. We have not measured film ourselves. Film's sharpness comes from two things: scatter, which lowers contrast, and the edge outline from development, which raises it at some levels of detail. A published curve shows only the combined result, not how much of each is present. Where the data cannot decide, the fit uses the narrowest scatter that still matches, because wide scatter puts a haze around bright highlights.

The curve is defined on the film, so the format decides how much of it you see. At 4K, the finest detail a Super 35 frame can show is 77 cycles per millimetre on the film, where the green layer keeps about a third of

its contrast. On VistaVision the same 4K frame reaches only 51 cycles per millimetre, where film keeps about half. Larger formats look sharper at the same resolution, as they do on film.

Softening at Small Pixel Sizes

The scatter is a blur a few micrometres wide, with a sharp centre and soft sides. That is about the size of a pixel, and a blur that small, measured only at the centre of each pixel, produces faint ripples beside bright edges. The effect averages the blur over the whole area of each pixel instead. The result never produces negative light and never changes the overall brightness, at any pixel size.

Edge Outline in Density

During development, a heavily exposed area uses up the developer around it, and fresh developer drifts in from lightly exposed areas next to it. The result is a faint outline on both sides of an edge, which makes film look sharper than its scatter alone would.

The effect applies this outline to density, not to linear light. Density measures how dark the negative is on a logarithmic scale, where each increase of 0.3 halves the light that gets through. The effect converts light to density with the curve Kodak publishes for VISION3 250D [1]. The choice matters on bright edges. A 1000:1 edge outlined in linear light would push the dark side far below zero and draw a black ring around every lamp. In density, the outline stays within the range the film can record.

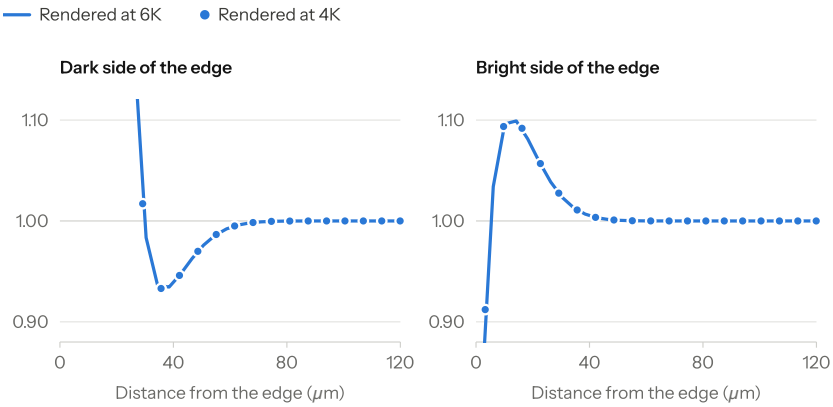


Figure 5. The two sides of a 1000:1 edge after the sharpness step, each relative to its own flat level. The dark side dips by about 7% and the bright side rises by about 10%. Right next to the edge, light scattered from the bright side fills in the dip. Nothing goes below zero. The green layer is shown. The line and the points are the same edge rendered at two resolutions.

Halation

Some of the light that reaches a negative passes through the light-sensitive layers, reflects off the back of the film, and exposes the film again a short distance away. The red-sensitive layer lies deepest and catches most of it. The result is a red-orange glow around bright lights.

How the Glow Is Made

The effect takes a share of the light, spreads it out, and adds it back. At the default settings that share is 20% of the red, 4% of the green and 1% of the blue, and the Halation radius is 8 µm. The spread is a soft blur about three times as wide as that setting. Beside a bright edge, the glow falls to half its strength about 16 µm into the dark side and has faded out by about 60 µm, as Figure 6 shows. The same share is taken away from the original picture, so no light is created: a flat area comes out exactly as it went in, and only edges change. The one exception is the extra glow given to clipped highlights, described below.

This happens to the light the film receives, before the edge outline and before grain. The glow therefore has grain in it like the rest of the picture, and it gains an outline where it meets a dark edge.

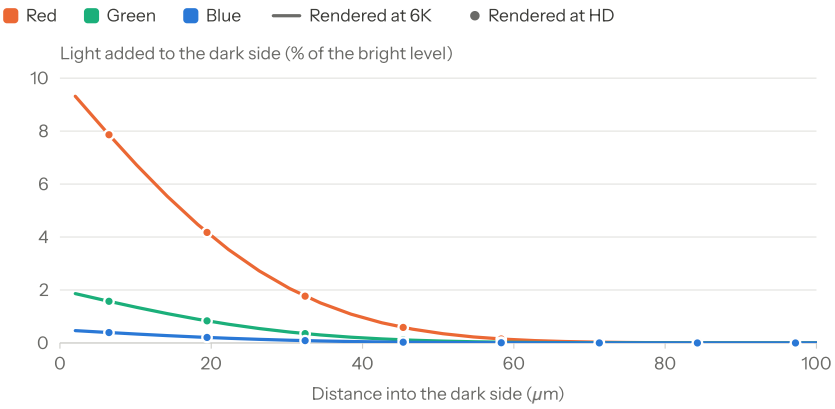


Figure 6. Halation alone on the dark side of the same 1000:1 edge. Because the spread is set on the film, a frame rendered at HD falls on the same curve as one rendered at 6K.

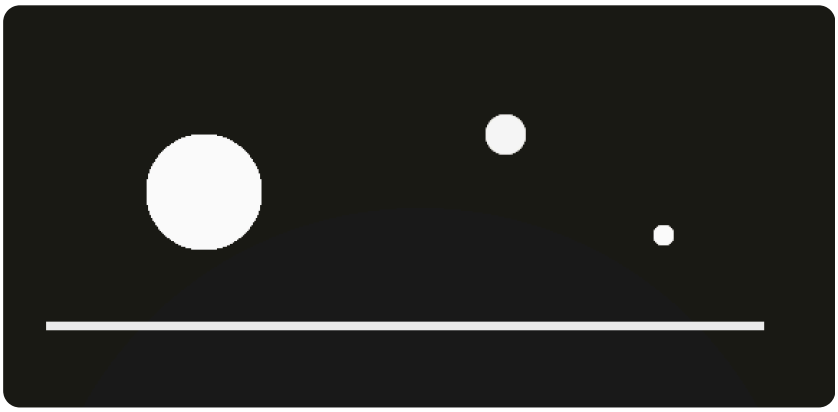
You can set the amount and radius. The default width and colour were chosen by eye, guided by a published measurement of how far halation spreads on film [3]. They are not measured from a particular stock, and most modern camera negatives have a backing layer that suppresses much of the effect.

Clipped Highlights

On film, a lamp five stops over the exposure produces far more halation than one two stops over. A digital camera records both at its maximum value, the clip level, so if the glow followed the recorded picture, both lamps would get the same faint glow.

With Extend clipped highlights on, as it is by default, the effect boosts the light used for halation. Above a threshold, by default 6.75 stops over mid-grey, the excess is multiplied by eight for the glow only. The colour of the light is kept. This glow is added light: the picture under it is not changed. This is an estimate, not a recovery of the real brightness, which the camera never recorded: every clipped light is boosted the same way.

Source 3.8 mm of film



Filmic Texture, default settings

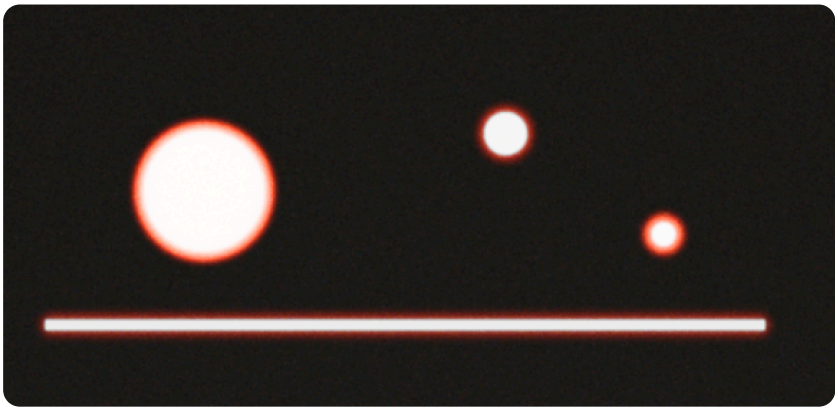


Figure 7. A computer-generated test picture on Super 35 at 4K, enlarged 2×. The largest and smallest lights are above the highlight threshold. The middle light and the strip are below it. Shown through the ACES 2.0 standard dynamic range display conversion.

Dust

Dust on a negative blocks light in the scanner. Once the scan is turned from negative to positive, the dust shows as a white speck. The effect places dust particles on the film and scans them with the picture.

Their sizes and number come from our own measurements of real film scans. A typical particle is 8 µm across, and nearly all are between 3 and 30 µm. At the default amount a Super 35 frame has about 30 on average, most of them close to one pixel wide at 4K. The number varies widely from frame to frame, as it does on a real scan, and each frame gets new dust.

Particles are long and ragged rather than round, and they are blurred by the scanner's lens before they reach the frame. Because they have a physical size, a smaller format shows fewer, larger specks. A speck is never brighter than about four stops over mid-grey, so dust reads as a bright fleck and not as a glint of light.

Colour and Signal

Filmic Texture is meant to be applied before the picture is converted for display, while it still holds the full range of brightness the camera recorded. Colourists call such a picture scene-referred: its values describe light in the scene, not the brightness of a screen.

Colour Space

The effect accepts ACEScct or DaVinci Wide Gamut / Intermediate. It converts the picture to linear light for its calculations and converts the result back the same way. Resolve does not tell an effect which colour space a node is in, so the effect's Colour space setting has to match the node.

What Is Preserved

Property	Behaviour
Flat areas	Softening and halation leave an even area unchanged, except above the clipped-highlight threshold, where the extra glow is added.
Tone and colour	Grain is built to average to zero, so it adds texture without changing the picture's overall brightness or colour.
Highlights	Values brighter than white are kept. Nothing is limited to what a display can show.
Negative values	Passed through, not cut off.
Alpha	Unchanged.
Bypass	With the effect off, or every amount at zero, the output is an exact copy of the input.
Repeatability	A frame rendered twice gives the same result, and a crop of a frame matches the same area of the full frame.

Verification

The calculations exist in two versions. The reference version is slow: it runs on the Mac's main processor and works with 64-bit numbers, which carry about 16 significant digits, so rounding errors are too small to matter. It is what this paper describes, apart from the section Grain in the Plug-in. The plug-in is the fast version, which runs on the Mac's graphics processor (GPU).

Filmic Texture is checked in three ways.

The Reference Against a Second Calculation

The sharpness step of the reference version is compared with a second program, written separately in a different programming language, that works out the same scatter by a more direct and slower method. The two agree to about seven decimal places.

Across Resolutions

Figures 5 and 6 show the same edge on the film rendered at different resolutions and landing on the same curve. A stricter check uses large frames: the central half of a 12K frame, rendered on its own as a 6K frame on a format half the size, matches the full 12K render to within 0.0002 in ACEScct. That is about a fifth of the smallest step a 10-bit video file can store.

The Plug-in Against the Reference

Every change to the plug-in is tested on a fixed set of test pictures. Its sharpness and halation are compared pixel by pixel with the reference version, after the standard display conversion. The limit is a quarter of one step of an 8-bit picture. The only pixels set aside, at most eight in a picture, are ones where the display conversion itself jumps on a change too small to see, and those must match the reference to one part in 10,000 before the conversion.

Grain is random, so it is compared by its statistics, not pixel by pixel: its strength in brightness and in colour, its balance of fine and coarse texture, and any shift in average brightness must stay within fixed limits of the release whose look was approved by eye.

Limits

Filmic Texture simplifies film, and its look is designed. In particular:

- **It does not reproduce a film stock.** The sharpness is fitted to curves published for VISION3 250D, and the dye colours to curves published for VISION3 50D, a different stock. Crystal sizes, sensitivities and development settings were chosen so that the result looks like film scans. They are not measurements of a real film.
- **Halation is not measured.** Its colour and width were chosen by eye, and both can be adjusted.
- **The grain look follows one set of scans.** Its strength and colour at each brightness were matched to scans of processed negative viewed through a print look. A different scanner or grade would have led to different choices.
- **The plug-in's grain is precomputed.** It is stored grain from the crystal simulation, looked up by brightness, rather than crystals developed afresh for each frame. The stored grain was simulated on flat grey, so across a sharp edge the plug-in blends the grain of the two sides, where the full simulation would develop crystals that see the edge itself.
- **Colour is not changed.** The effect adds texture. It does not reproduce a film's tone or colour, and it leaves out the ways the colour layers affect each other during development.
- **The frame is steady.** Gate weave, flicker, scratches and splices are not included.
- **Pixels are assumed square.** Footage from an anamorphic lens that was unsqueezed before the effect is treated like any other picture, so its grain is not stretched with it. Support for that is planned.

Technical Specifications

Effect	Filmic Texture, one OpenFX effect for DaVinci Resolve • Grain, halation, film sharpness (MTF) and scan dust, each with its own switch
Film formats	Super 8, Super 16, Super 35 (4-perf), VistaVision (8-perf), 65 mm (5-perf), or a custom width and height
Grain	Six crystal groups: fine and coarse in each colour layer • About 13 million crystals in a 16:9 Super 35 frame at the default size • Three development rounds, in which developed crystals hold back their neighbours • Light through the dye calculated at 31 wavelengths, 400 to 700 nm • Simulated scanner: 2 µm pixels, 0.45 µm lens blur • In the plug-in: precomputed at 33 brightness levels, half a stop apart • Controls: Grain amount, Material size (3 to 20 µm), Grain chroma
Sharpness	Fitted to the published VISION3 250D curves [1], 2.5 to 80 cycles per millimetre • Scatter averaged over the area of each pixel • Edge outline applied in film density, using the published density curve over 16 stops
Halation	By default, 20% of red, 4% of green and 1% of blue at a Halation radius of 8 µm • Controls: Halation amount, Halation radius • Stronger glow for clipped highlights, from 6.75 stops over mid-grey
Dust	Particles of 3 to 30 µm, typically 8 µm; about 30 per Super 35 frame at the default amount
Colour spaces	ACEScct; DaVinci Wide Gamut / Intermediate
Frame size	Up to 65,536 pixels on a side and 268 million pixels in total • Rendering tested to 17,520 × 8,040, outside DaVinci Resolve on a Mac with 18 GB of memory
Requirements	A Mac with Apple silicon, macOS 26 or later, DaVinci Resolve • Tested with DaVinci Resolve Studio 21. Earlier versions and the free edition are not yet tested

References

[1] Eastman Kodak Company, *KODAK VISION3 250D Color Negative Film 5207 / 7207*, product brochure: modulation transfer and sensitometric curves. https://www.kodak.com/content/products-brochures/motion-picture/KODAK-VISION3-250D-5207-7207-brochure.pdf. These are typical curves for the stock. The fit to them is ours: Kodak has not measured, reviewed or endorsed it.

[2] Eastman Kodak Company, *KODAK VISION3 50D Color Negative Film 5203 / 7203*, Technical Information H-1-5203: spectral dye density curves. https://www.kodak.com/content/products-brochures/motion-picture/KODAK-VISION3-50D-5203-7203-technical-information.pdf.

[3] C. Montmea, *Study of Film Texture: Grain, Halation and Rendition of Details*, master's thesis, ENS Louis-Lumière, carried out at ARRI, December 2023. https://www.ens-louis-lumiere.fr/wp-content/uploads/2024/02/ENSLL_PHOTO_MONTMEA_2023-English-Fev.pdf.

© 2026 Filmic Texture. Filmic Texture is not affiliated with or endorsed by Eastman Kodak Company, Blackmagic Design or any film manufacturer. Kodak and VISION3 are trademarks of Eastman Kodak Company. DaVinci Resolve is a trademark of Blackmagic Design Pty Ltd. Mac and macOS are trademarks of Apple Inc. Other names may be trademarks of their owners. Specifications may change. November 2026