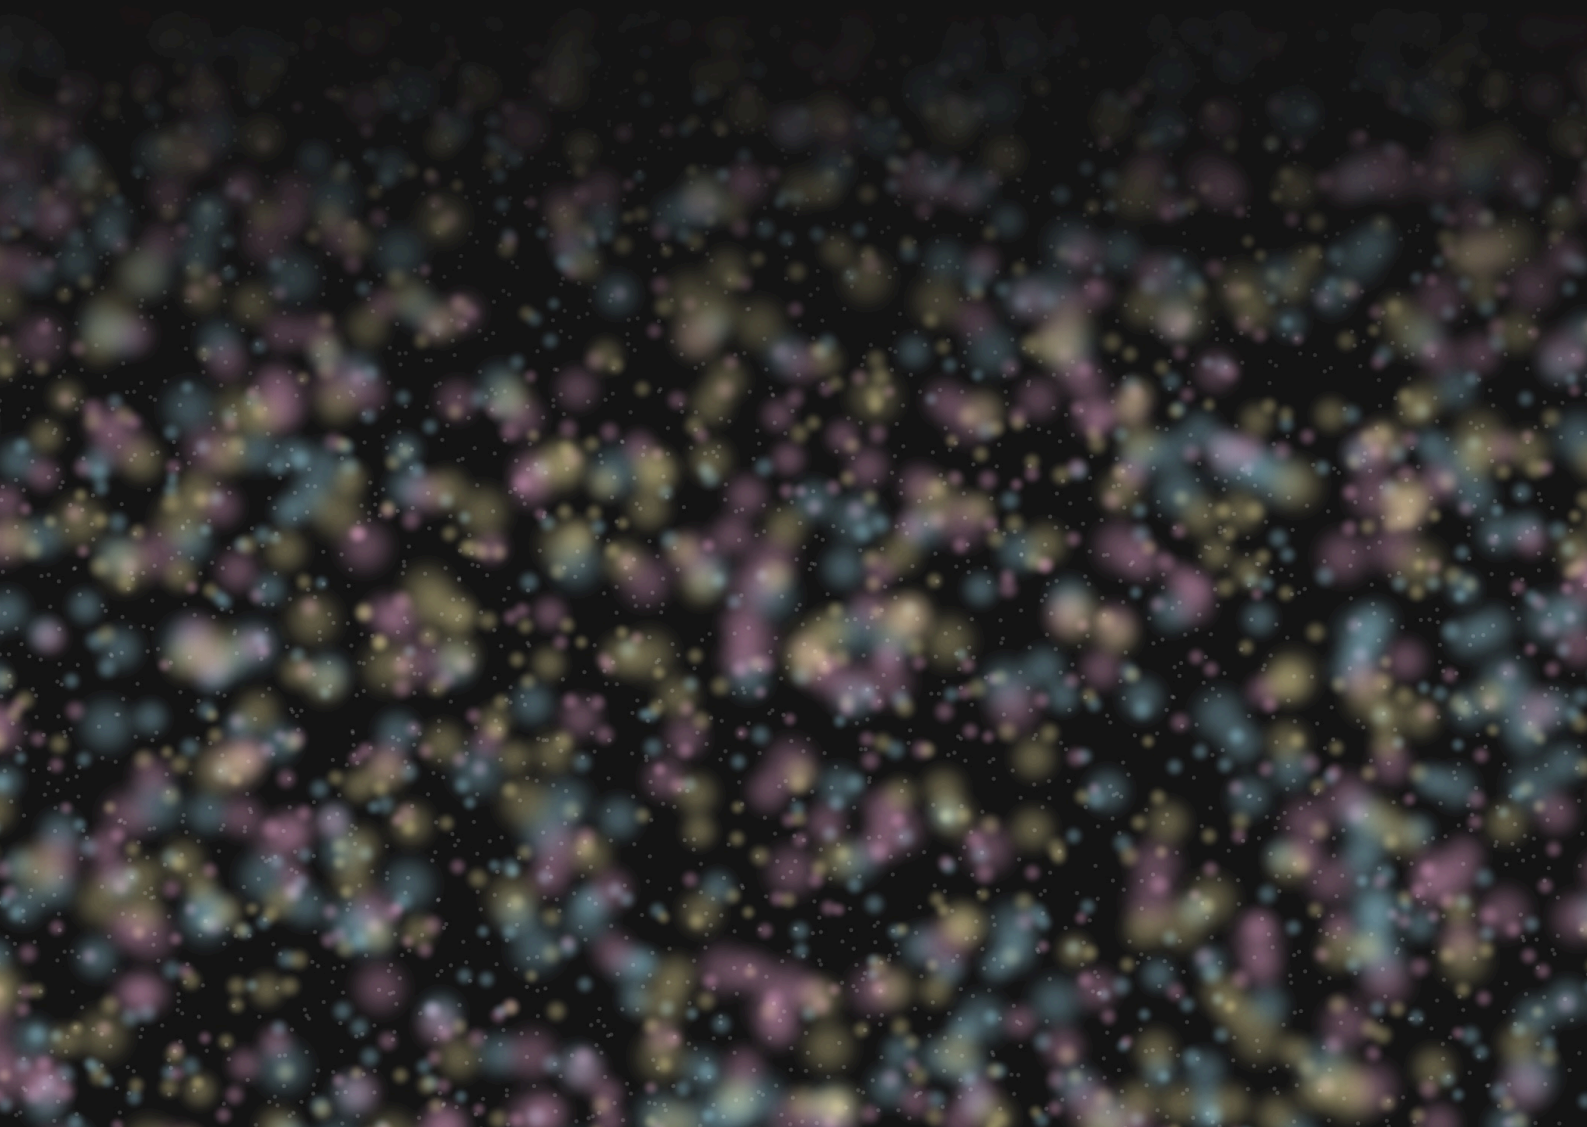
Filmic Texture

Technology Overview

November 2026



Contents

Introduction 3

Key Features 4

The Film Model 5

Physical Scale 6

Grain 7

Crystals 7

Exposure and Development 7

Dye and the Scan 8

Grain as a Difference 8

Film Sharpness 10

Fitted to Published Curves 10

Scatter Over the Pixel 11

Edge Effect in Density 11

Halation 12

A Return of Light 12

Clipped Highlights 13

Dust 14

Colour and Signal 15

Working Space 15

What Is Preserved 15

Verification 16

Against Independent Calculation 16

Across Resolutions 16

Against the Reference 16

Limits of the Model 17

Technical Specifications 18

Introduction

Filmic Texture is an effect for DaVinci Resolve that adds film grain, halation, film sharpness and scan dust to a scene-referred picture. It runs on Apple silicon.

The effect is defined on a virtual film gate, not on the pixel grid. The gate can be set to a standard format from Super 8 to 65 mm or to any width and height, and grain, halation and sharpness are specified in micrometres on it. The same settings therefore give the same texture, relative to the frame, at any delivery resolution.

This paper describes the model behind the effect: what it simulates, the data it is based on, and where it stops. The model is defined by a reference implementation in 64-bit arithmetic. The plugin's GPU renderer is tested against that reference and is not described here.

Filmic Texture is an authored look. It uses published film data and scans of real negative as references, but it is not a measured emulation of any one film stock. The sections below note which parts rest on data and which are judgement.

Key Features

Physical scale

- Grain, halation and sharpness are defined in micrometres on a film gate, not in pixels.
- Gate presets for Super 8, Super 16, Super 35, VistaVision and 65 mm, or any custom size.
- The texture keeps the same size relative to the frame at HD, 4K, 8K and above.

Grain from simulated crystals

- Grain is formed by simulating individual crystals: about 13 million in a Super 35 frame at the default size.
- Each crystal develops according to the exposure of the picture at its own position.
- The grain changes from frame to frame, and a given frame always renders the same grain.

Film sharpness

- A sharpness response fitted to the modulation transfer curves Kodak publishes for VISION3 250D.[1]
- The edge effect of development is computed in film density, which keeps very bright edges free of dark rings.

Halation

- A red-dominant glow, computed in exposure before the grain forms, at its physical size on the film.
- Clipped highlights can drive the glow as if they were brighter than the clip.

Scan dust

- Small white specks, sized and counted from measurements of real film scans, different on every frame.

Scene-referred signal

- Works on ACEScct and DaVinci Wide Gamut / Intermediate pictures.
- High dynamic range values, negative values and alpha pass through. With every stage off, the output is the input, bit for bit.

The Film Model

The model follows a picture through the stages a negative goes through, in the order they happen on film. The order matters because each stage works on a different quantity. Light spreads inside the emulsion in proportion to exposure. Development acts on density, which is logarithmic. Grain forms from the image that those two stages leave behind. Dust is added last, because it sits on the negative when it is scanned.

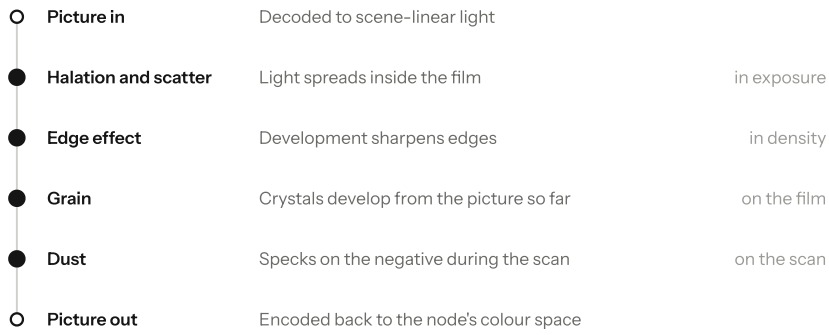


Figure 1. The stages of the model, and the quantity each one works in

Because the stages are kept separate, each acts on the result of the one before. Halation is part of the exposure, so it is sharpened by the edge effect and carries grain like the rest of the picture. Grain is formed from the softened image, so it stays crisp over detail that the film has blurred.

Physical Scale

The first thing the model needs is the size of the film. You choose a gate, and the frame is placed on it by a centre crop, so the picture is never stretched: a 16:9 timeline on a Super 35 gate uses the full 24.89 mm width and the central 14.0 mm of the height.

From the gate and the frame size the model knows how much film each pixel covers. Everything after that is expressed in micrometres.

Gate	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

The width of film that one pixel covers, for a 16:9 frame. The 65 mm gate is limited by its height at 16:9, so the frame uses 40.89 mm of its width.

A dye cloud 8 µm across is larger than a pixel of a Super 35 frame at 4K and smaller than one at HD. The model treats both cases the same way: every stage integrates over the area of film that each pixel covers. A coarser frame therefore shows the same film, averaged over larger pixels, and not a different texture.

The gate also changes the look. A smaller gate enlarges everything on it. The same grain that is fine on VistaVision is coarse on Super 16, and the halation around a lamp is three times wider, because the film is a third of the size and the picture is enlarged three times as much.

Grain

A colour negative holds its image as millions of small dye clouds, each formed where a silver halide crystal developed. The model builds its grain the same way, from simulated crystals.

Crystals

The emulsion is modelled as six populations of crystals: a fine, slow population and a coarse, fast one for each of the three colour records. Their positions are drawn in three dimensions through the depth of each layer.

Crystals cannot overlap. A candidate that falls too close to an earlier one is removed, which makes the field more even at small scales than purely random placement would be. Within each population, crystals also vary: a quarter are smaller and less sensitive than the norm, and a quarter are larger and more sensitive.

At the default size, a 16:9 Super 35 frame holds about 13 million crystals: roughly 11 million fine and 2 million coarse. Each one is identified by its position on the film, the frame number and a seed. It is not tied to the pixel grid, so a crop, a different resolution or a render split into tiles sees the same crystals in the same places.

Exposure and Development

Each crystal reads the picture at its own position, after sharpness and halation have been applied. Development then runs in three steps. In each step an undeveloped crystal either develops or does not, with a probability that rises with the exposure it received. 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 sparse set of coarse clouds. Through the mid-tones the fine population fills in, and in bright areas nearly every crystal has developed.

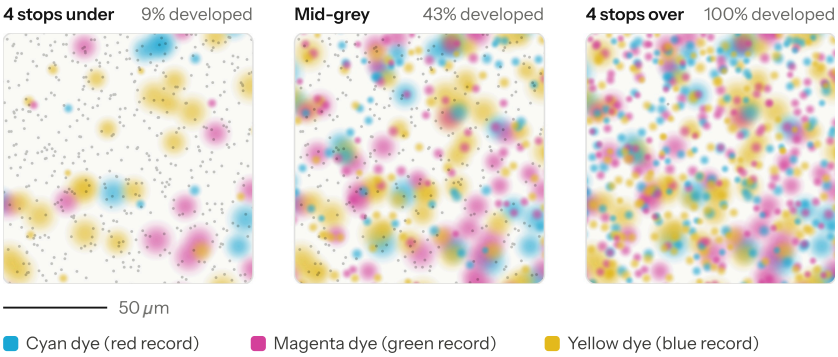


Figure 2. The same 0.12 mm square of the model's emulsion at three exposures, drawn from the model's own output. Grey points are crystals that did not develop. Each disc is the dye cloud of one that did.

Developed crystals also affect their neighbours. A crystal surrounded by developed ones is less likely to develop in the next step, and its dye cloud is slightly smaller. This is a simplified form of the inhibition that happens in real development, and it keeps dense areas from filling in evenly.

Dye and the Scan

Each developed crystal leaves a soft dye cloud: cyan for the red record, magenta for green, yellow for blue. At the default size the clouds are between 2.5 and 10 µm in radius. The shapes of the three dye absorption spectra are derived from the spectral dye density curves Kodak publishes for its colour negative film.

The model then looks through the film. At every point it adds up the dye from all the clouds above it and computes how much light passes, at 31 wavelengths from 400 to 700 nm. Light is absorbed, not added: two overlapping clouds darken the film less than twice as much as one.

A virtual scanner reads that transmitted light. Its lens has a blur of 0.45 µm, and its detector sites are 2 µm squares on a 2 µm pitch. The scan is then resampled to the output frame using the exact area of each pixel. One pixel of a Super 35 frame at 4K averages about ten scanner sites.

Grain as a Difference

A scan of the simulated emulsion has a tone and colour of its own. If it replaced your picture it would regrade it. The model avoids this by developing four related versions of the same emulsion and combining them with weights that sum to zero. Whatever the four have in common, which is the tone of the picture, cancels. What remains is the texture, and that is added to your picture.

How strongly the grain shows at each exposure level is set by a table for each colour record, and a final stage trims its colour and contrast in the brightest and darkest areas. These were fitted to scans of processed negative viewed through a print film look. They are look decisions, not measurements of a stock.



Figure 3. Grain from the plugin on five flat grey levels, at the default settings on a Super 35 gate at 4K. Each tile is 150 pixels across, about 1 mm of film, enlarged three times.

The three grain controls map onto the model directly. Material size scales every length in the emulsion together. Amount scales the difference that is added to the picture. Chroma sets how much of the dye colour the grain keeps, from neutral to full, without changing its brightness.

Film Sharpness

Film does not record fine detail at full contrast. Light scatters inside the emulsion, so a sharp line in the scene is recorded a little wider and a little weaker. Film manufacturers publish this as a modulation transfer function (MTF): the share of contrast that survives at each level of detail, measured in cycles per millimetre on the film.

Fitted to Published Curves

The model's response is fitted to the MTF curves Kodak publishes for VISION3 250D.[1] The curves were traced at 71 points for each colour record, from 2.5 to 80 cycles per millimetre. The model follows them to within 0.04 for the red record, 0.02 for green and 0.03 for blue, where 1.0 is full contrast.

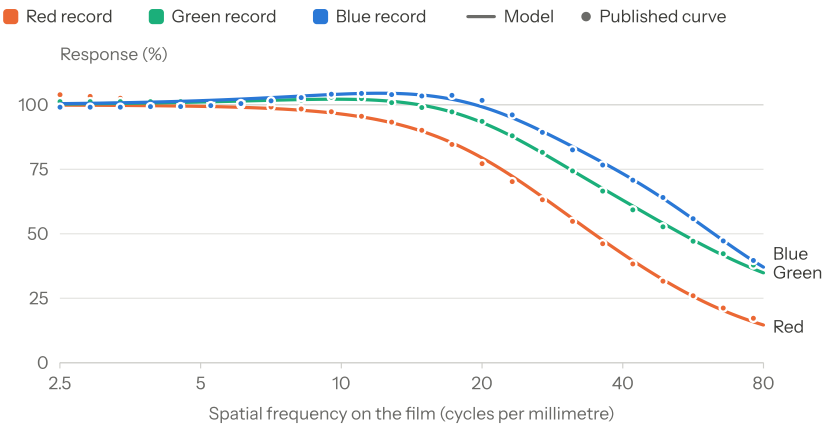


Figure 4. The model's response for each colour record, against points traced from the published curves. Every third traced point is shown. Values above 100% at low frequencies are the edge effect.

These are typical curves for the stock, as published. We have not measured film ourselves. The curves also leave one question open. Film's response is the product of two things: scatter, which lowers contrast, and the edge effect of development, which raises it at some frequencies. A curve cannot say how much of each is present. Where the data does not decide, the model uses the most compact scatter that still fits, because wide scatter puts a haze around bright highlights.

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

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

Scatter Over the Pixel

The scatter is a blur with a sharp core and a soft skirt, a few micrometres wide. That is close to the size of a pixel, and sampling such a blur only at pixel centres causes ringing beside bright edges. The model integrates the blur over the area of each pixel instead. The resulting filter is never negative and always sums to one, at any pixel size.

Edge Effect in Density

During development, a heavily exposed area uses up developer, and fresh developer drifts in from its lightly exposed neighbours. The result is a faint outline on both sides of an edge that makes film look sharper than its scatter alone would allow.

The model applies this effect to film density, using the characteristic curve Kodak publishes for the stock, and not to linear light. The difference shows on bright edges. Applied to linear light, an outline in proportion to a 1000:1 edge would push the dark side far below zero and draw a black ring around every lamp. Applied to density, the outline is limited by the film's own range.

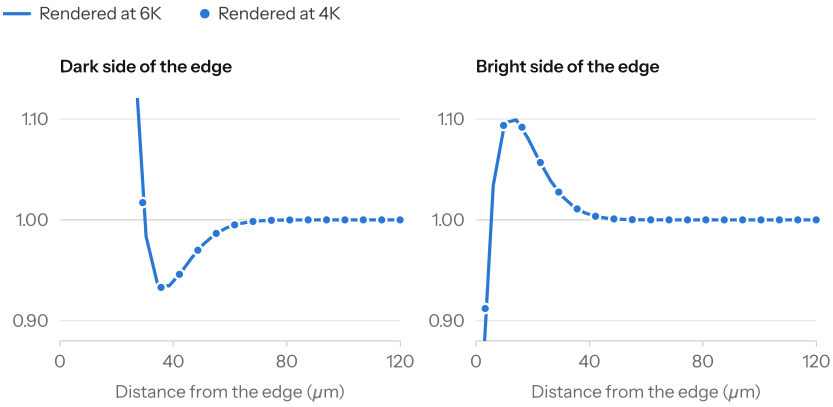


Figure 5. The two sides of a 1000:1 edge after the sharpness stage, each relative to its own flat level. The dark side dips by about 7% and the bright side rises by about 10%. Close to the edge, light scattered from the bright side covers the dip. Nothing goes below zero. The green record is shown. The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

Halation

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

A Return of Light

The model takes a share of the exposure, spreads it, and returns it to the film. At the default settings that share is 20% of the red, 4% of the green and 1% of the blue, spread with a radius of 8 µm. The glow reaches about three times that distance. The same share is taken away from the direct image, so no light is created: a flat area of any brightness comes out exactly as it went in, and only edges change.

This happens in exposure, before the edge effect and before grain. The glow is therefore part of the image the film develops. It has grain in it, and it sharpens where it meets a dark edge.

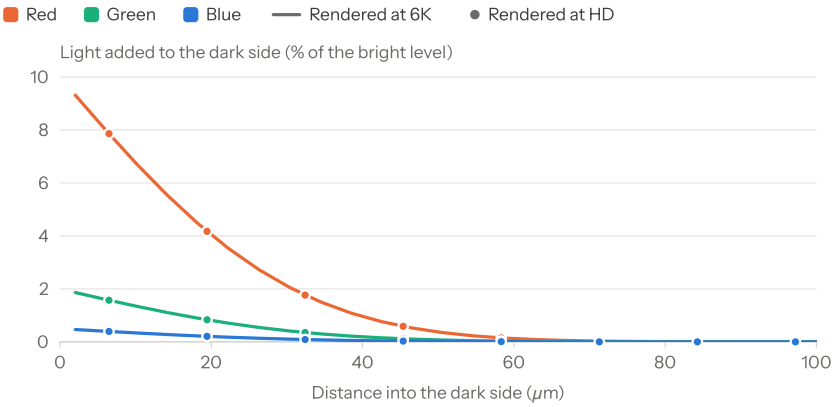


Figure 6. Halation alone on the dark side of the same 1000:1 edge. Because the spread is fixed on the film, a frame rendered at HD falls on the curve of one rendered at 6K.

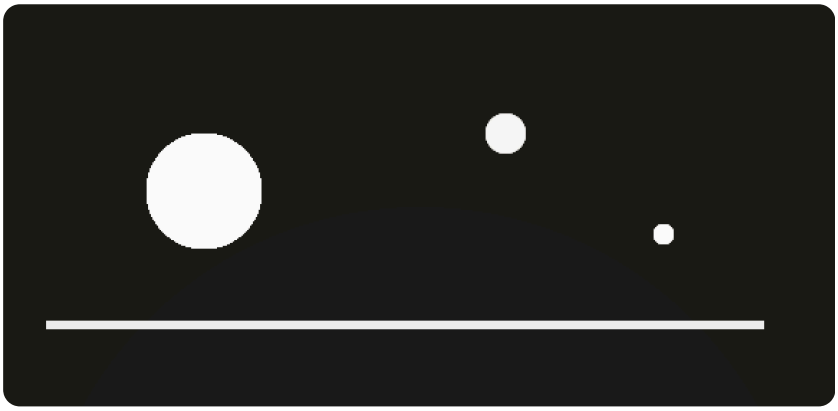
The amount and radius are yours to set. The width and colour of the default were authored, guided by a published measurement of how far halation spreads on film. They are not measured values for a particular stock, and modern stocks carry a backing layer that suppresses much of the effect.

Clipped Highlights

On film, a lamp five stops over the exposure produces far more halation than one that is two stops over. A digital camera that clips records both at the same level, so a model driven by the recorded picture gives both the same faint glow.

The model can extend the light that drives halation. Above a threshold, by default 6.75 stops over mid-grey, the excess is multiplied by eight for the purpose of the halation only. The hue is kept, and the direct image is not touched. This is a declared assumption and not a recovery of the scene: every clipped highlight is treated alike.

Source



3.8 mm of film

Filmic Texture, default settings

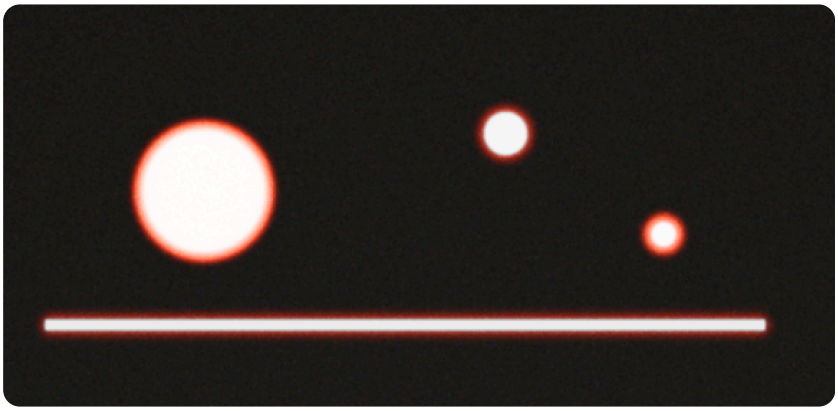


Figure 7. A synthetic test scene on a Super 35 gate at 4K, enlarged twice. The largest and smallest lights are above the highlight threshold. The middle light and the strip are below it. Shown through the ACES 2.0 standard dynamic range view.

Dust

Dust on a negative blocks light in the scanner, and after the scan is inverted it shows as a white speck. The model places particles on the film and scans them with the picture.

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

Particles are elongated and ragged, not round, and they are blurred by the scanner's optics before they reach the frame. Because they have a physical size, a smaller gate shows fewer and larger specks. The brightness of a speck is limited to about four stops over mid-grey, so dust reads as a bright fleck and not as a specular highlight.

Colour and Signal

Filmic Texture is a scene-referred effect. It is meant to sit before the output transform, on a picture that still holds its full range.

Working Space

The effect accepts ACEScct or DaVinci Wide Gamut / Intermediate. It decodes the picture to scene-linear light for the model and encodes the result the same way on the way out. Resolve does not tell an effect which colour space a node is in, so the setting has to match the node.

What Is Preserved

Property	Behaviour
Flat areas	Sharpness and halation leave a uniform area unchanged at any brightness.
Tone and colour	Grain is a zero-sum difference between realizations, built to add texture without regrading the picture.
Highlights	Values above diffuse white are kept. Nothing is clamped to a display range.
Negative values	Passed through, not clipped.
Alpha	Unchanged.
Bypass	With the effect off, or every amount at zero, the output is a bit-for-bit copy of the input.
Repeatability	A frame rendered twice, or rendered as a crop, matches the full frame.

Verification

Three kinds of check apply to Filmic Texture.

Against Independent Calculation

The sharpness stage of the reference implementation is compared with a second implementation, written separately in a different language, that computes the same blur by direct numerical integration. The two agree to about seven decimal places.

Across Resolutions

Figures 5 and 6 show the same physical edge rendered at different resolutions and landing on the same curve. A stricter check uses large frames: the centre half of a 12K frame, rendered on its own as a 6K frame on a gate of half the size, matches the full 12K render to within 0.0002 of the encoded signal.

Against the Reference

The GPU renderer in the plugin is tested against the reference implementation, pixel by pixel, on a fixed set of test pictures. Each test has a limit on how far the two may differ.

Limits of the Model

The model is a simplification, and the look it produces is authored. In particular:

- **It is not an emulation of a stock.** The sharpness response is fitted to curves published for VISION3 250D, and the dye spectra to published dye curves. Crystal sizes, sensitivities and development constants were chosen so that the result matches the appearance of film scans. They are not measurements of an emulsion.
- **Halation is not measured.** Its colour and width are authored, and both are adjustable.
- **The grain look has a reference.** Its strength and colour by tone were fitted to scans of processed negative viewed through a print look. A different scanner or grade would have led to different choices.
- **Colour is not changed.** The model adds texture. It does not reproduce a film's tone curve or colour response, and it leaves interactions between colour layers during development out.
- **The frame is steady.** Gate weave, flicker, scratches and splices are not modelled.
- **Pixels are assumed square.** Footage that was desqueezed from an anamorphic lens before the effect is treated like any other picture. Support for stretching the texture with the picture 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 gates	Super 8, Super 16, Super 35 (4-perf), VistaVision (8-perf), 65 mm (5-perf), or a custom width and height
Grain	Six crystal populations: fine and coarse for each colour record • About 13 million crystals in a 16:9 Super 35 frame at the default size • Three development steps with inhibition between neighbours • Dye absorption at 31 wavelengths, 400 to 700 nm • Virtual scanner: 2 µm pitch, 2 µm detector, 0.45 µm lens blur • Controls: amount, material size (3 to 20 µm), chroma
Sharpness	Fitted to the published VISION3 250D curves, 2.5 to 80 cycles per millimetre • Scatter integrated over the pixel area • Edge effect in film density, using the published characteristic curve over 16 stops
Halation	Return of 20% red, 4% green and 1% blue at an 8 µm radius by default • Controls: amount, radius • Highlight extension 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 (AP1); 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[2]
Requirements	A Mac with Apple silicon, macOS 26 or later, DaVinci Resolve

[1]Fitted to the typical curves in Kodak's published technical data for VISION3 250D Color Negative Film 5207/7207. The fit is ours. Kodak has not measured, reviewed or endorsed it. [2]Tested outside DaVinci Resolve on an 18 GB Mac. Very large frames need memory in proportion to their size and may be slow on machines with less.

© 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