DESIGN STUDY

Toward the world's best free video noise reducer

A nine-lens agent panel and a real-footage audit agree on an inversion: the frontier isn't removing more noise — it's reconstructing the image's optical character once the noise is gone.

Authored by Claude · Commissioned by Stephen Walter

2026-07-15 · from a 77-agent design panel · build postscript 2026-07-16 (v3.6)

68	22	33	13	9	77
IDEAS RAISED	KEPT	MAYBE	KILLED	LENSES	AGENTS

ABSTRACT

We asked one question — *how does a free plugin become the most effective video noise reducer on the planet?* — and answered it twice: by measuring the current denoiser on real 4K night footage, and by convening nine independent domain experts who read the actual source, generated 68 ideas, then adversarially stress-tested every one.

The central finding is a strategic inversion. Hush's **noise removal is already excellent** — good enough that nearly every idea whose only job was to remove *more* noise was killed or demoted for "merely lowering a number." The frontier is **texture reconstruction**: after the noise is gone, put back

what real optics and real film would have left — **optical edge acutance** (sharpness that reads like a lens, not a sharpen filter) and **emulsion-grade grain** (matched to the noise that was removed, not a white-noise dither). Nine lenses converged on this independently. That convergence, not any single idea, is the result.

> **Remove the noise, then reconstruct the optical image, and couple the whole thing to the color grade.**

---

01 What the panel did

The empirical baseline MEASURED, NOT ASSUMED

Before ideating, the exact CPU reference pipeline was run on real footage — a 4K DCI ProRes night portrait, with handheld drift and a focus pull. Native-resolution crops with shadows lifted 4× (to expose what a colorist sees after grading) gave a clear verdict.

REGION	RESULT
Flat shadow (sky)	`STRONG` ~⅔ grain removed, no blotching (effN ≈ 4.8 / 5)
Detail — hair, glasses, skin, edges	`STRONG` preserved, no smear, crisp edges
Bokeh	`STRONG` light shapes round, no ghosting on the drift
Deep shadow, 4× lift — **Weak 1**	`WEAK` residual **mid-freq chroma speckle** above the bands' ~23 px reach
Deep shadow — **Weak 2**	`WEAK` **plasticky / waxy**; grain off by default, spectrally white when on
Large shadow color patches	`CORRECT` mostly real lighting, rightly preserved

Those two weaknesses and the brief's north star — *remove noise while preserving or **enhancing** sharpness, with **film-like** grain, not video crunch* — became the fixed target for every agent.

The panel

Nine lenses ran in parallel — transform-domain, motion, deconvolution, perceptual/grain, grading-integration, signal-model, tiny-ML, systems, and a cross-domain wildcard. Each read the source, grounded ideas in real functions, avoided shipped features and a "graveyard" of measured-dead ideas, and attached a **control arm** to every idea: the cheap baseline it must beat, and by how much, on a prototype *before* any code. A plain post-blur once beat a ground-truth-offset super-resolution accumulator by 1.9 dB — so a win against a strawman is worthless. Every idea then faced an independent adversary judging thesis-fit, novelty, four-backend feasibility, real-time cost, and whether the control arm was real. Verdict: **22 keep, 33 maybe, 13 kill.**

---

02 The finding: texture reconstruction

The lenses weren't told to agree, yet they converged hard. Read across all nine, the same three-part architecture appears again and again, each expert arriving from a different direction:

1 · Optical acutance on the cleaned signal. Sharpen *after* denoising, gated by the plugin's own noise field so it can never amplify noise, and **overshoot-bounded** so edges gain slope without the bright halo that reads "video."

2 · Detail transfer, not blur inversion. Put back the high-frequency *structure* the denoiser removed by coring the denoise delta — amplify what survives a noise-scale threshold, discard the rest. Halo-free by construction.

3 · Reconstructed film grain. Replace the spectrally-white, mid-gray-peaked value-noise with grain whose **amplitude follows the measured shadow-loud noise curve**, whose **spectrum is shaped** near the eye's contrast-sensitivity peak, and which is **placed where the denoiser waxed the image** — refilling exactly what was removed.

Why this is the answer and not "denoise harder": **the kills prove it.** Almost every killed or demoted idea did so for the same reason — it only lowered a noise number, which the north star explicitly rules out. Hush stops competing on "cleanest" (a race against paid tools with far larger teams) and wins on *"most filmic"* — a race almost nobody is running, and one a small, transparent, free plugin can lead.

> **There is no stage anywhere that adds high-frequency energy. The whole opportunity is the un-built other half.**
>
> `↳ v3.6 built it — optical acutance is now that stage.`

---

03 The 22 survivors, as one architecture

The keepers aren't a list; they're four subsystems plus a substrate. Chips show the adversary's confidence.

✅ **Since this study, v3.6 shipped the marked items.** All of Phase 0 (acutance, the three grain fixes, coring, the confidence matte) and the Phase-1 chroma-speckle pass are now in the plugin, across all four backends. What the build taught — including three places it corrected the paper — is in the postscript below.

A Optical sharpening & detail transfer THE "ENHANCE SHARPNESS" PILLAR

`0.80 · RT✓` **Overshoot-bounded optical acutance** ✓ SHIPPED V3.6

Edginess-gated high-pass hard-clamped to the local 3×3 min/max — edge slope steepens with **zero ringing by construction**. Highest-confidence primitive; build first.

`0.72 · RT✓` **Detail-transfer by coring the denoise delta** ✓ SHIPPED V3.6

Discard `|delta|<k·noise`, re-inject the rest at gain ≥ 1. A strict upgrade over the old uncored texture path that re-noised the shadow.

`0.72 · WHT` **Subband acutance in a shift-invariant WHT**

Shrink noise octaves and **amplify signal octaves at once** — multi-scale acutance reads more optical than any single-radius sharpen. Walsh-Hadamard, never DCT.

`0.72 · Cinema` **SNR-regularized deconvolution**

Fixed-iteration Wiener / Richardson-Lucy on cleaned luma to **recover** lens / focus-pull softness — the only family that raises lost resolution.

B Film-grain reconstruction "FILM NOT VIDEO" · FIXES WEAK 2

0.80 · S **Brightness-amplitude-matched grain** ✓ SHIPPED V3.6

The old grain peaked at mid-gray; **real noise lives in shadows.** Drive amplitude off the measured 16-bin curve. Near-certain win.

0.80 · S **Contrast-masking placement** ✓ SHIPPED V3.6

Full grain into waxed flats, almost none on hair/glasses edges — reuses the NLM's own edginess signal. Near-free.

0.82 · top **Blue-noise / band-pass spectrum** ✓ SHIPPED V3.6

Shift grain energy toward the CSF peak; at matched RMS it masks plasticity and reads sharper — the eye reads HF texture as grain, LF wobble as dirt.

0.72 **effN-modulated octaves + RGB dye layers**

Most texture where the image was most cleaned; per-channel granularity (B>G>R) re-codes shadow chroma speckle as emulsion, not noise.

apex **Measured-NPS coloring FIR**

Measure the removed noise's actual spectrum and use it as the grain's coloring filter, with an in-viewer scope. The octave presets are its cheap control arm.

C Grade coupling THE "LEVERAGE GRADING" PILLAR

0.70 · RT✓ **Clean-confidence matte export** ✓ SHIPPED V3.6

Route per-pixel effN to the OFX alpha so a **downstream node** keys shadow-lift, grain, and contrast off where the denoiser actually averaged deep.

0.68 **Qualifier-mask input**

Accept an incoming matte that scales strength/bands/sharpen per region **and scopes the estimator** to keyed pixels — one seam-free node instead of a stack.

0.72 **Specular halation + negative sharpening**

Roll acutance off near clip, add a red-weighted highlight bloom (the film "glow"), full acutance in mids. Pure optical finishing.

D The chroma-speckle fix ✓ SHIPPED V3.6

Honest tension: the shadow chroma speckle **is** a genuine defect, but every fix is "denoising," which the thesis deprioritizes — so most landed in **maybe**. The credible route is a **small-scale collaborative chroma pass** (luma-guided shrinkage, or a spatially-variant chroma map that scores real low-frequency color as zero and leaves it alone). **Shipped in v3.6** as a luma-guided wide chroma pass — -70% shadow speckle on the clip, real color preserved (the postscript has the twist: the "block-mean subtraction" turned out to be the window size).

E

Substrate & efficiency WHAT BUYS THE BUDGET

`0.72` **Guided-filter detail transfer**

Cleaned luma as input, **noisy original as guide** — the linear model transfers edges while rejecting guide noise in flats.

`substrate` **Mip / à-trous pyramid + effN-adaptive NLM**

One parity-safe pyramid feeds the bands, the multi-scale acutance, and cheap large-radius chroma reach. Build once a filter that needs it passes its gate.

04 What was killed — the sharpest lessons

The kills encode discipline. The through-line: **an idea that only makes the image cleaner is, at this point, not good enough.**

~~**Temporally-correlated grain**~~

Killed on physics: real photochemical grain is **temporally independent** — that per-frame "grain boil" *is* the film look. Correlating it reads *more* video.

~~**PSF-from-noise**~~

Fatal physics error: sensor noise **never passes through the lens**, so you can't estimate optical blur from the removed noise's autocorrelation.

~~**Kalman-style IIR shadow accumulator**~~

Worsens its own target: more temporal averaging removes **more** micro-texture — making shadows *more* plasticky.

~~**Bilateral-grid chroma · grade-aware measurement**~~

Thesis-misses: they reduce a chroma-noise number but add no acutance, grain, or genuine grade-leverage.

05 A build sequence, with honest gates

Each item is prototyped on the real clip, gated against its stated control arm, then ported to four kernels only if it passes. Grain and acutance *add* high-frequency energy on purpose — they're gated on acutance / overshoot / matched-RMS / blind preference, **never on PSNR.**

PHASE 0 — CHEAP, HIGH-CONFIDENCE, SHIP-NOW ✓ SHIPPED V3.6

Fixes the plasticky read and adds the first real sharpness. **Shipped in full** across all four backends.

1. Grain overhaul: shadow-loud amplitude + contrast-masking + blue-noise spectrum `3×S, ~0.8`
2. Overshoot-bounded optical acutance on cleaned luma `the 0.80 primitive`
3. Detail-transfer by coring the denoise delta
4. Clean-confidence matte export

PHASE 1 — THE TRANSFORM & THE CHROMA FIX ◑ 2 OF 3 SHIPPED

Once Phase 0 is measured on the clip.

1. Shift-invariant WHT subband acutance — gated vs gradient-masked USM `next up`
2. Small-scale collaborative chroma pass for Weak 1 — real color untouched ✓ V3.6
3. effN-steered spatial `h` — shipped as **Adaptive Strength** ✓ V3.6

PHASE 2 — THE CINEMA TIER & THE FILM FINISH

The apex, opt-in.

1. Constrained deconvolution (Cinema flag) — recovers focus-pull / lens softness
2. Measured-NPS grain coloring + in-viewer spectrum scope
3. Specular halation + tone-keyed negative sharpening; qualifier-matte input

Code-level corrections the panel surfaced **(all since fixed in v3.6)**: the grain response peaked mid-gray (backwards); the luma-texture re-add was uncored (re-noising the shadow it was meant to fix); Detail Rescue was preserve-only (gain ≤ 1 — no sharpening existed). And `effN` plus the residual field were computed but under-consumed — they now ride the acutance gating and the confidence matte, exactly as predicted.

06 Splitting NR across Color-page nodes

Yes — and how you split it is itself a quality lever. Six situations, in rough order of everyday impact:

- **Isolation for caching (the biggest win).** Put NR on its own node, early, and let Resolve cache it — grade in real time without re-running the most expensive op on every tweak.
- **Before the grade amplifies the noise.** Lift/gamma/gain and saturation magnify grain — denoise upstream of the creative grade. Hush fetches ±3 neighbor frames, so the node before it *is* what it denoises; put it first so Render Boost's history stays coherent.
- **Region-split via qualifiers / Power Windows.** Shadows want heavy NR, skin wants gentle (keep pore texture), sky wants aggressive chroma, highlights almost none — exactly what the qualifier-matte-input keeper folds into one node.
- **Denoise early, sharpen late.** Cleanup is scene-referred; acutance and grain are display-referred finishing. Run NR on an early cached node and the texture mode on a final node, on graded contrast. *Clean early, reconstruct late.*
- **Temporal early, spatial late.** Temporal NR needs coherent frame-to-frame motion — run it before anything disturbs it; spatial NR after the balance.
- **Parallel blend & confidence downstream.** Mix NR back at reduced opacity for a gentler result; hand the **clean-confidence matte** forward so a later node keys shadow-lift and grain off where the denoiser actually succeeded.

The design implication: Hush should be **built to live in a node graph** — cacheable, matte-*in* for region control, confidence-*out* for the grade, and mode-selectable so the same plugin is both the early clean-up node and the final texture-reconstruction node.

---

07 Postscript: from study to shipped

ADDED 2026-07-16 · AFTER V3.6 BUILT PHASE 0 AND THE FIRST PHASE-1 FIX

The study was a design panel. The next day a second Claude instance **built** its Phase 0 and the start of Phase 1 — across the CPU reference and all three GPU kernels at parity, gated on the same night clip. Building it did what a panel cannot: it shipped the thesis, and it **corrected the paper in three places**. That correction is the entire point of the X3 "prototype before code" law — here is what it caught.

Shipped in v3.6. The full Phase-0 quartet — shadow-loud, contrast-masked, blue-noise grain; overshoot-bounded optical acutance; detail-transfer coring; and the clean-confidence matte — plus the Phase-1 shadow chroma-speckle fix. Everything is off by default, so existing projects render bit-for-bit identical until you reach for it.

> ### Three things the panel got wrong on paper only surfaced when the code met the real clip.

1 · The blue-noise grain formula was backwards for this codebase. The study proposed `gn − gHi·valueNoise(2·size)`. But here `valueNoise` at a coarser scale is an *independent* field — subtracting it *adds* low-frequency energy, the opposite of blue. On the real clip the high-frequency fraction moved the wrong way (0.41 → 0.29). The true high-pass subtracts a low-pass of the *same* grain field (its four-neighbour mean, a grain-cell away), which put the spectrum where it belongs: 0.41 → 0.69. A one-line error no panel could have caught without measuring.

2 · Auto-tuning correlated chroma noise needs *generalized* SURE — the build's most transferable result. The self-tuner's chroma pass read the real shadow speckle backwards and de-tuned a good profile by ~0.9 dB. The cause is textbook but easy to miss: Monte-Carlo SURE is an unbiased error estimate only for *white* noise, and the speckle is spatially correlated. Match the probe's own correlation to the measured noise — a block-constant Rademacher field at the estimated correlation length — and the estimator becomes generalized SURE (it measures `tr(AΣ)`,

not `σ²·tr(A)`); its argmin tracks true error again. The 0.9 dB loss became a **3.0 dB gain**. The lesson outlives this plugin: *if you tune against SURE, the probe must carry the noise's covariance, or it will confidently tune the wrong way on anything but white noise.*

3 · The chroma-speckle "safety" is just the window size. The panel asked for a block-mean-subtracted chroma pass so real low-frequency color scores zero. In practice no subtraction is needed: a luma-guided wide chroma mean at *moderate* reach averages the ~30 px speckle while the frame-scale lighting gradient — warm streetlight against cool ambient over dark fabric — is locally flat across the window and passes straight through. The protection falls out of the geometry: **-70% shadow speckle**, the real color untouched.

And two the build sharpened. Acutance's halo bound must reference the *pre-spatial* (post-temporal) 3×3 min/max, not the over-smoothed spatial output — exceeding the smoothed range is precisely how it *restores* the slope the smoother removed, while staying halo-free against the true signal. And coring belongs at the *input*-noise scale (`s.sy·gainY`), not the residual: the removed delta carries a single frame's full noise, so thresholding at the residual under-cores and re-noises the shadow — the correct scale cut shadow re-noising from 4.7× to 1.2×.

One honest limitation — for the graveyard. Scoring the whole temporal-plus-spatial pipeline against the *input* sigma over-credits the spatial stage once the temporal average has already crushed the residual, so on a near-lossless static shot the tuner can nudge toward over-smoothing (a 46 dB stack drifting to 45 dB — imperceptible, but the wrong direction). Hold-out cross-validation against an independent probe catches the *variance* form of this; the *systematic* bias wants scoring against the post-temporal residual, and is logged for a later pass. The panel's discipline was to publish its graveyard — this belongs in it.

The thesis grew a second tool

Section 06's conclusion — *clean early, reconstruct late, on a final node* — was a prediction about the shape of the node graph. Building it produced **both halves**. Hush denoises early and now exports the clean-confidence matte; the reconstruction half became **Speak**, a standalone film-emulation effect for the *end* of the tree — film-stock density, subtractive color, halation, bloom, grain — that reads Hush's confidence matte downstream to lay grain exactly where the image was cleaned most. Node one measures and cleans; the last node reconstructs the optical image; a matte crosses the graph between them. The study didn't just steer a denoiser — it specified an interface, and a second product grew into it.

> **The frontier the panel named — texture reconstruction — turned out to be big enough to be its own tool.**

ULTIMATE FINDING

The world's best free video noise reducer won't win by removing the most noise — that race is against paid tools with large teams, and Hush is already close. It wins by being the only free tool that treats denoising as **step one of texture management**: strip the noise, then reconstruct the optical image — lens-like acutance without a digital halo, emulsion-grade grain matched to exactly what was removed — and couple that reconstruction to the grade through confidence mattes, qualifier inputs, and film-optic finishing.

> **Clean early. Reconstruct late. Match the grain to the noise you took. Sharpen only what the noise field says is real.**

Do all of it transparently, portably across four backends, in real time — and give it away.

Authored by **Claude** (Anthropic), commissioned by Stephen Walter.

A design study for Hush · Open NR, an open-source tool of the Community AI Project in partnership with Brookline Interactive Group. MIT licensed — free for any use, forever.