-->
The FFT homogeneity report measures **spatial luminance ripple** inside the detected illuminated beam area. Rather than treating every brightness change as a defect, it separates variation by its physical spatial scale: broad beam fall-off, the ripple band of interest, and very fine detail/noise.

The implementation is in shared/eval_math.py, while the ROI sent to it is built in desktop/qt/dialogs.py.

1. Creating the evaluation ROI

The input is a rendered 2-D luminance map $L(y, x)$, in physical coordinates. The chosen evaluation resolution is 100–4096 pixels (default 2200), and the render supplies the pixel pitch d in mm/pixel.

The UI first identifies illuminated pixels with a relative threshold:

$$\mathrm{rawMask}(y, x) = L(y, x) > L_{\max} \frac{T}{100}$$

where T is the “Boundary Threshold (%)” setting, defaulting to 5%.

It then:

- morphologically closes the mask over roughly 0.15 mm, bridging small gaps;
- removes connected components below 4 mm²;
- falls back to the raw threshold mask if nothing survives.

This resulting `valid_mask` defines which samples contribute to the measurements. The report’s evaluation area is simply:

$$A = N_{\mathrm{valid}} d^2 / 100$$

in cm².

2. Normalizing luminance

It finds the mean luminance over valid pixels only:

$$\bar{L} = \frac{1}{N} \sum_{(x,y) \in ROI} L(x, y)$$

and converts every valid pixel into fractional deviation from that mean:

$$r(x,y) = \frac{L(x,y) - \bar{L}}{\bar{L}}$$

Thus a pixel at 110 cd/m² in a 100 cd/m² mean region becomes $r = 0.10$, i.e. +10%.

This makes the FFT metrics scale-independent: multiplying all luminance values by two does not change the percentage ripple. Invalid pixels are zeroed.

3. Windowing before the FFT

The ROI’s enclosing bounding rectangle is multiplied by a separable 2-D Hann window:

$$w(x,y) = w_x(x)w_y(y)$$

$$r_w(x,y) = r(x,y)\,M(x,y)\,w(x,y)$$

where M is the ROI mask.

An FFT assumes its input repeats periodically outside the image. Without a window, a mismatch between opposite edges creates artificial sharp discontinuities and spreads energy over many frequencies (“spectral leakage”). The Hann window tapers the outer rectangle toward zero, reducing that artifact.

A practical consequence: the FFT result is intentionally about the **interior spatial texture** of the ROI, not an unwindowed exact reconstruction of its edge values.

4. Transforming from position to spatial frequency

The code takes a 2-D FFT, then shifts zero frequency to the centre of the display:

$$F(f_x, f_y) = \mathcal{F}\{r_w(x,y)\}$$

The physical frequency axes are generated with the actual pixel size:

$$f_x = \texttt{fftfreq}(W, d), \qquad f_y = \texttt{fftfreq}(H, d)$$

in cycles/mm. Each FFT bin represents a 2-D sinusoidal pattern with a particular spatial frequency and direction.

Radial frequency is:

$$f_r = \sqrt{f_x^2 + f_y^2}$$

Using f_r makes the filter isotropic: a 5 mm-period horizontal stripe, vertical stripe, or diagonal texture is treated as the same spatial scale.

The default passband is:

$$0.05 \le f_r \le 0.50 \text{ cycles/mm}$$

which corresponds approximately to features with periods from:

$$1/0.50 = 2 \text{ mm} \quad \text{to} \quad 1/0.05 = 20 \text{ mm}$$

So, in plain language:

- below 0.05 cycles/mm: broad gradients and overall beam shape are de-emphasized;
- 0.05–0.50 cycles/mm: the ripple range the report evaluates;
- above 0.50 cycles/mm: very fine detail, pixel-scale variation, and numerical/sensor noise are de-emphasized.

5. The Butterworth band-pass filter

The report does not use a hard rectangular cutoff. It uses a fourth-order Butterworth high-pass multiplied by a fourth-order low-pass:

$$H_{\text{HP}}(f_r) = \frac{1}{1 + \left(\frac{f_{\text{low}}}{f_r}\right)^{2n}}$$

$$H_{\text{LP}}(f_r) = \frac{1}{1 + \left(\frac{f_r}{f_{\text{high}}}\right)^{2n}}$$

$$H(f_r) = H_{\text{HP}}(f_r) H_{\text{LP}}(f_r), \qquad n = 4$$

The DC bin is explicitly set to zero. This removes the mean component even if a small residual remained after normalization.

The filtered spectrum is:

$$F_{\text{band}}(f_x, f_y) = F(f_x, f_y) H(f_r)$$

Butterworth transitions avoid the ringing that a sharp frequency cutoff can introduce in the reconstructed image.

6. Returning to a ripple map

An inverse FFT reconstructs only the band-limited variation:

$$n(x,y) = \Re\left\{\mathcal{F}^{-1}(F_{\text{band}})\right\}$$

`n` is dimensionless, relative to the ROI mean. The displayed “FFT Ripple Map (% mean)” is:

$$100\,n(x,y)$$

So +3 means a locally positive band-passed ripple of roughly +3% of mean luminance; −3 means the corresponding negative ripple. It is not the original luminance image and should not be read as absolute luminance.

7. Reported metrics

The report evaluates the reconstructed ripple only at valid ROI samples.

Report field	Exact meaning
FFT RMS Ripple	$\sqrt{\text{mean}(n^2)} \times 100$. Overall band-passed ripple magnitude.
Ripple P95	95th percentile of (100
Ripple P99	99th percentile of (100
Ripple Peak	(\max(100
Dominant Frequency	Radial frequency of the highest-power bin after the filter.
Dominant Period	$1/f_{\text{dominant}}$, in mm.
Low/Mid/High Frequency Energy	Percentages of non-DC FFT power below 0.05, from 0.05–0.50, and above 0.50 cycles/mm.

The energy percentages are calculated from the original windowed spectrum—not the filtered ripple reconstruction. They are diagnostic proportions: for example, high “Low Freq Energy” suggests beam-scale variation dominates; high “Mid Freq Energy” suggests meaningful ripple is concentrated in the assessed band.

8. The homogeneity score

The score is a project-specific engineering score, not a photometric or regulatory standard. Its penalty is:

$$P = \max\left(RMS,\ 0.60P95,\ 0.45P99,\ 0.35Peak\right)$$

and:

$$\text{FFT Score} = \max(0, 100 - P)$$

The weights prevent any one statistic from being ignored:

- RMS detects widespread modest texture;
- P95 catches ripple affecting a substantial part of the beam;
- P99 catches severe near-local defects;
- Peak still penalizes the worst single anomaly, but less aggressively since a lone pixel may be noise.

For example, if RMS = 3%, P95 = 8%, P99 = 16%, Peak = 26%:

$$P = \max(3,\ 4.8,\ 7.2,\ 9.1) = 9.1$$

so the FFT homogeneity score is 90.9%.

9. Important interpretation details

The separately reported **RMS Contrast** is different: it is the ordinary unfiltered standard deviation divided by mean luminance across the ROI. It includes broad gradients, in-band ripple, and fine noise. FFT RMS Ripple isolates only the configured frequency band.

The dominant frequency is radial, not directional. It identifies the strongest pattern scale, but not whether it is horizontal, vertical, diagonal, or a more complex texture.

Resolution matters. The Nyquist limit is:

$$f_{\text{Nyquist}} = \frac{1}{2d}$$

If $d = 1$ mm/pixel, the 0.50 cycles/mm upper cutoff reaches Nyquist. If pixels are coarser than 1 mm, some of the nominal upper passband is unobservable. Frequency-bin spacing is also limited by physical ROI size:

$$\Delta f_x = \frac{1}{Wd}, \qquad \Delta f_y = \frac{1}{Hd}$$

Larger physical ROIs resolve lower frequencies more precisely; finer pixel pitch resolves finer ripple.

Finally, irregular or perforated ROIs are analysed as their bounding box with invalid areas set to zero before windowing. This is practical and stable for beam regions, but the ROI shape itself can contribute spectral structure. The result is most directly interpretable for a reasonably contiguous, well-thresholded illuminated region.