Film as a photon detector, read from Kodak's own data sheets

A preregistered estimate of detective quantum efficiency, grain noise and response shape for nine colour layers of three VISION3 motion-picture stocks

Zain Dana Harper

Version 1.0, 9 November 2026. Finding F147 in the findings ledger. Signal (tier B).

Abstract

Detective quantum efficiency asks how much of the information in arriving light a detector keeps, and it puts a film, a camera sensor and an eye on one scale. We derived it for the nine colour layers of three current Kodak motion-picture stocks, 5219 (500T), 7213 (200T) and 7203 (50D), from numbers Kodak prints on its data sheets, under a preregistered plan with planted checks. Peak detective quantum efficiency across the nine layers has a median of 2.1 percent and a range of 0.57 to 4.92 percent, an upper estimate because every photon is counted at the layer's most efficient wavelength. In the green layer, 500T and 200T reach 4.5 and 4.9 percent although 200T needs about twice the photons; 50D needs 15.7 times the photons of 500T and reaches 1.5 percent. Siedentopf's rule, that grain variance grows with density, lost on held-out points in 9 of 9 layers, and variance falls with density overall (median power slope -0.54); because our planted variance check failed, this is description only. A straight hinge in log exposure was the best held-out response form in 9 of 9 layers. A single Poisson hit-threshold model does not fit.

1 What this means for a reader

Modern colour negative film keeps a few percent of the information in the light that reaches it, by this estimate at most about 5 percent in its best layer. That puts a number on film that can be set beside a sensor or an eye. For anyone who emulates film grain, the sheets themselves give a measured noise curve per layer, and that curve does not grow with density the way the textbook rule says.

2 Introduction and motivation

A film is a photon detector with a long history of measurement. The detective quantum efficiency (DQE) of a detector compares the information in its output with the information in the light that arrived. Shaw introduced the photographic form of the idea in 1963 under the name equivalent quantum efficiency [2] and used it to compare the information capacity of photography with other detectors [3]. Dainty and Shaw's book on image science [4] treats DQE curves, granularity and the characteristic curve together. We found these works through searches run on 9 November 2026 and read the abstract of the second Shaw paper and the catalogue records of the others; the full texts were not read.

What a sheet gives is enough for a ceiling estimate. Kodak publishes, for each colour layer, the characteristic curve, the rms granularity against density, and the spectral sensitivity. From the first two come the noise-equivalent quanta (NEQ); from the third comes the number of photons. Their ratio is the DQE. A search summary we could not trace to a primary source put the best older films near 1 percent. We found no published DQE for the VISION3 stocks; that is what our searches returned, and it does not show that none exists.

The same sheets bear on two working rules. Siedentopf's rule says the variance of density grows in proportion to mean density, and it is the basis of signal-dependent grain models used in film emulation; Stephenson's grain simulation states it together with Selwyn's aperture law [5]. And a Poisson threshold model, in which a grain develops after it absorbs n photons, has been used for both grains and rod cells in the eye. The study asked

whether one such model describes the response and the noise of layered colour film together.

3 Preregistration and design

The plan was written and hash-locked in a preregistration on 31 October 2026 at 06:12:46 UTC, and the run followed at 06:12:48. It names the data, the measures, the planted checks, two controls, the forecasts and the kill conditions. The decision it informs is what the grain and response parts of our film renderer should use.

One disclosure comes from the file itself. Before the lock, the 5219 stored curves had been looked at and one hand calculation had been made for the 5219 green layer, giving a DQE of about 0.11 with assumed values of gamma and sigma. So the forecasts for 5219 were not blind. The stocks 7203 and 7213 had been looked at only for calibration.

The rules that decide each verdict were written into the run's plan before the run. Any failed planted check voids its block. The full list is in Appendix A.

4 Data and provenance

The data are the Kodak VISION3 technical data sheets for 5219 (500T), 7213 (200T) and 7203 (50D) [1]. The curves were digitised from the published vector sheets: for each of the blue, green and red layers, the characteristic curve, the diffuse rms granularity curve measured with a 48 micrometre aperture, and the spectral sensitivity curve at a net density of 0.2. The curves were resampled to a grid of 0.05 in log exposure. The sheets are Kodak's copyright. This paper reports only values derived from them and redraws none of their curves.

Kodak notes that its granularity figures use a modified measuring technique. We did not correct for it.

5 Methods

Photons

For each layer, let S be the sheet's spectral sensitivity at the layer's peak wavelength λ_p, in reciprocal erg per square centimetre, at net density 0.2. The exposure that reaches net density 0.2 is $1/S$, and dividing by the energy of one photon at λ_p gives the photon count per unit area there:

$$q_{02} = \frac{\lambda_p}{Shc} \tag{1}$$

with h Planck's constant and c the speed of light, converted to photons per square micrometre. Counting every photon at the most efficient wavelength makes q_{02} the fewest photons that could give net density 0.2, and so makes the DQE below an upper estimate.

At log exposure x, the photons inside the 48 micrometre aperture, of area A, are

$$Q(x) = A q_{02} 10^{x - x_{02}} \tag{2}$$

where x_{02} is the log exposure at which the curve reaches net density 0.2.

Noise-equivalent quanta and DQE

Let γ be the gamma of the curve at x, taken from a smoothed derivative, and σ the rms granularity at the same point. Then

$$N_{eq} = \left(\frac{\gamma}{\sigma \ln 10}\right)^2 \tag{3}$$

$$DQE = \frac{N_{eq}}{Q} \tag{4}$$

Each layer's peak DQE is the largest value of eq. 4 over points with net density and gamma both above 0.05.

Grain variance

For points with net density D above 0.1, four laws for the variance σ^2 were compared: Siedentopf's $\sigma^2 = aD$, a constant, an affine law $c + aD$ and a quadratic. Each was scored by BIC on all points and by its hold-out score in log variance on the top 30 percent of density after fitting on the rest. A power slope b was fitted as

$$\ln \sigma^2 = k + b \ln D \tag{5}$$

so that Siedentopf's rule corresponds to $b = 1$.

Response forms

Eight forms for net density against log exposure were fitted on the lowest 70 percent of the exposure axis and scored on the top 30 percent: Hill, Naka-Rushton with exponent 1, single hit, Erlang with real n, lognormal, Weibull, hinge and soft hinge. The hinge has two parameters, a threshold and a slope. BIC was reported beside the hold-out score, though the forms have different numbers of parameters.

Poisson threshold

The threshold model fits net density and log variance together, with hit number n in 1, 2, 3, 4 and 6 and a lognormal spread of grain area with free width, and picks n by BIC.

6 Controls and planted checks

Each block had a planted check with a known answer, run through the same code.

- DQE: synthetic single-size grains with a closed-form DQE. The pipeline came back within 11.3 percent, against a limit of 20. The DQE block stands.
- Response forms: five planted forms at the real data's scale of noise, 12 each, were recovered 80 percent of the time on hold-out, against a limit of 50. The forms block stands.
- Variance law: a planted affine law was chosen 70 percent of the time, against a locked requirement of 80. By the locked rule the variance verdict is void, and what we say about grain variance is description only.
- Poisson threshold: planted hit numbers 1 and 4 were recovered every time at both grain widths, 0.5 and 2.0.

Two controls built a null distribution for a statistic. Shuffling density against exposure 25 times gave a top-form share of 0.36 on average with a 95th percentile of 0.44; the real share of 1.0 is above all of them, a p value from draws of 0.038, which is the smallest value 25 draws can give. Pairing one layer's slope with another layer's noise 200 times gave a mean correlation of 0.66 with a 95th percentile of 0.76; the real 0.61 sits inside it, p 0.74, so that test could not separate the layers.

The preregistration called for a Benjamini-Hochberg correction at 0.05 over the control p values. The result file does not report it. With two p values, 0.038 and 0.74, the smaller would have to be at or below 0.025 to pass, and it is not, so after the correction neither control rejects. The hinge result rests on its hold-out errors, which Figure 3 shows, more than on the shuffle.

7 Results

Peak detective quantum efficiency across the nine layers has a median of 2.1 percent and a range of 0.57 to 4.92 percent, an upper estimate because every photon is counted at the layer's most efficient wavelength. Table 1 gives each layer and Figure 1 draws them. The peak falls low on the curve, at net density 0.06 to 0.31, in the toe.

Table 1. Derived values per layer. Photons are per square micrometre at net density 0.2, counted at the layer's peak wavelength. Hinge error is its hold-out error in net density.

Stock	Layer	Photons	Peak DQE, percent	Net density at peak	Gamma at peak	Variance slope b	Hinge error
5219	blue	1.2	2.12	0.26	0.55	-0.66	0.096
5219	green	4.2	4.51	0.31	0.64	-0.54	0.112
5219	red	10.8	2.51	0.12	0.44	-0.51	0.117
7213	blue	2.9	2.35	0.29	0.60	-0.32	0.047
7213	green	8.7	4.92	0.24	0.65	-0.49	0.071
7213	red	20.2	1.99	0.16	0.55	-0.56	0.054
7203	blue	33.6	0.57	0.14	0.46	-0.59	0.137
7203	green	65.8	1.52	0.06	0.37	-0.36	0.137
7203	red	87.6	0.90	0.17	0.54	-0.71	0.132

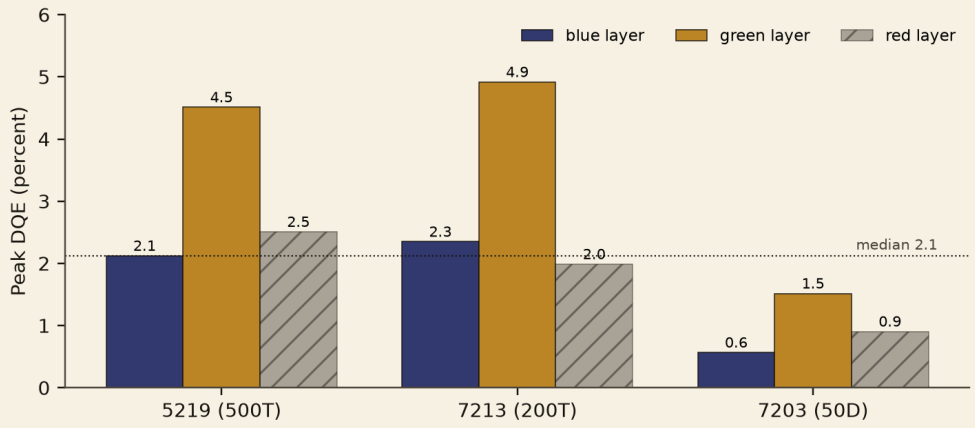


Figure 1. Peak DQE of each colour layer, an upper estimate. The dotted line is the median over the nine layers. Drawn from the run's result file.

In the green layer, 500T and 200T reach 4.5 and 4.9 percent although 200T needs about twice the photons; 50D needs 15.7 times the photons of 500T and reaches 1.5 percent. The preregistered rule for "same efficiency at every speed" asked for a DQE spread of at most 3 across the three green layers while photon counts spread by 5 or more. The photon spread was 15.7 and the DQE spread 3.24, so the rule failed narrowly. Figure 2 shows the three points.

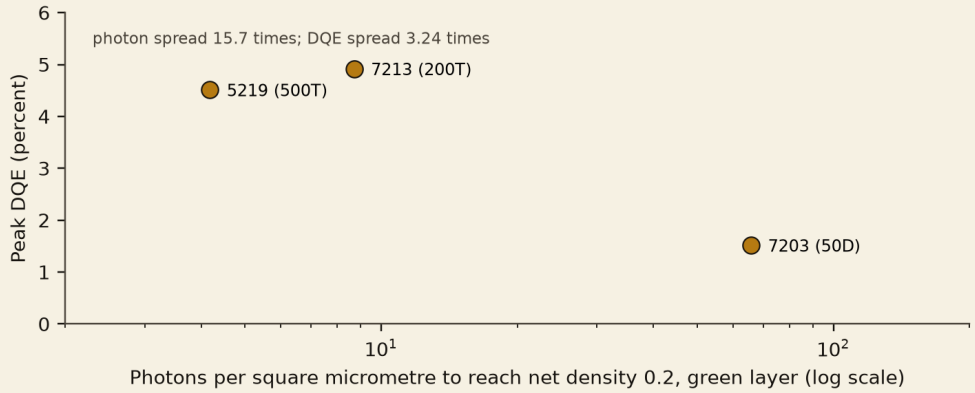


Figure 2. Green layer: photons per square micrometre needed for net density 0.2 against peak DQE.

Siedentopf's rule, that grain variance grows with density, lost on held-out points in 9 of 9 layers, and variance falls with density overall (median power slope -0.54); because our planted variance check failed, this is description only. The forecast slope was +0.15. The variance curves also show a hump in the toe. Our guess at a cause, untested, is the hand-off between fast and slow emulsions in a layered coating together with the colour masking couplers.

A straight hinge in log exposure was the best held-out response form in 9 of 9 layers. Its median hold-out error was 0.112 in net density, and every form that saturates did worse, as Figure 3 shows. Negative stock shows no shoulder within the range a sheet covers, so a form that bends over at high exposure loses when it extrapolates. Hill's exponent, fitted anyway, had a median of 0.615, which here is the gamma in another guise.

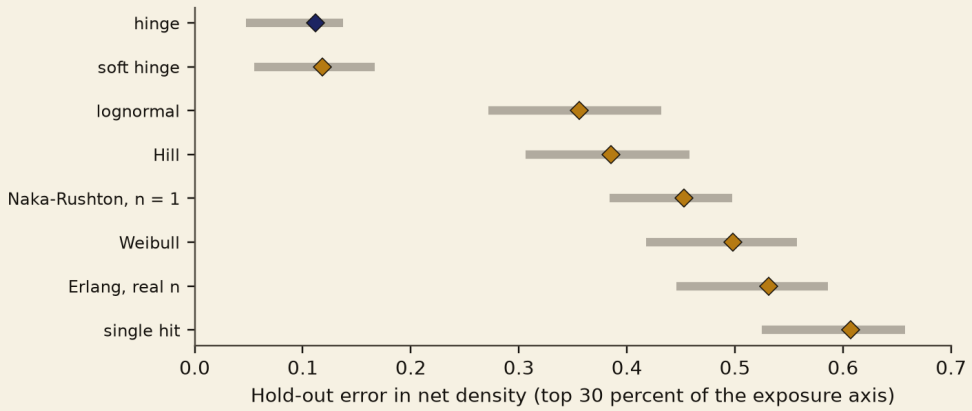


Figure 3. Hold-out error of each response form over the nine layers: median (diamond) and range (bar).

The Poisson threshold model does not fit. On the channel the run used, the median error was 0.134 in log variance and 0.043 in density, against limits of 0.05 and 0.04. A later study of ours found that this fit passed the logarithm of the standard deviation where the model expected the logarithm of the variance. On the corrected channel the errors are 0.186 and 0.068, and the best hit number per layer changes from 1, 1, 1, 6, 6, 6, 6, 6, 6 to 1, 1, 1, 1, 6, 1, 6, 1, 1. The verdict stands either way, and no hit number is read as a measurement.

8 Discussion

The DQE figures sit above the roughly 1 percent we saw quoted for older films, and the planted grains show the pipeline itself is good to about 11 percent. Counting photons at the peak wavelength pushes every value up. Two other choices move them by an amount and in a direction we did not measure: the single aperture with Kodak's modified technique, and the toe, where gamma is small, so that a small error in the digitised slope is a large share of it, and NEQ goes as the square of gamma. The order of magnitude, a few percent, is the safer reading than any one figure.

The green layers carry the clearest pattern. 500T and 200T keep the same share of the information while one needs half the light of the other. A faster film here buys speed by needing fewer photons for the same density, and it does not lose efficiency doing so. 50D breaks the pattern at 1.5 percent. One stock is not enough to say why.

For a film renderer the decision is plain. Grain should come from each layer's measured granularity against exposure, with no square-root-of-density rule, and the response of negative stock within the sheet's range needs no saturating shoulder.

The threshold result connects to a later study of ours on rod cells, which found that rods and film grain do not share one threshold model on the data available. Nothing in this paper bears on rods.

9 Limitations

- One maker, one aperture, and vendor curves read from printed sheets. Digitising adds its own error, which was not measured.
- The photon count assumes all light at the peak wavelength, so every DQE here is an upper estimate. No interval is reported for the DQE values; the planted error of 11.3 percent is the only measure of the pipeline's accuracy.
- The variance verdict is void by our own rule. The variance results are description.
- The Benjamini-Hochberg correction the plan called for was not run; done by hand above, neither control survives it.
- Forecasts for 5219 were not blind, as the preregistration states.
- Shaw's 1963 papers and Dainty and Shaw's book were not read beyond abstracts and catalogue records, and may already set these numbers in context.
- Scoring was done by our own code and was not blind.

10 Verified

Verified: on 9 November 2026 the nine peak DQE values (0.0057 to 0.0492, median 0.0212), the per-layer values in Table 1, the green-layer DQE and photon counts, the 3.24 and 15.7 spreads, the median slope -0.54, the planted results (0.113, 0.70, 0.80, 1.0), both controls and every verdict in Appendix A were re-read from the run's result file. They match the findings ledger. The preregistration's SHA-256 was recomputed from the file and matches the hash recorded in the lock and the run manifest. The Poisson threshold values on the corrected channel come from the later ledger entry and were not recomputed here. All three figures were drawn by a script that reads only the result file. The study was not re-run for this paper. The cited works were found by searches run on 9 November 2026 and read at the level of abstracts and catalogue records; their full texts were not read.

11 Corrections

- 31 October 2026: a later study of ours found that the Poisson threshold fit used log standard deviation where the model expected log variance. Corrected errors 0.186 and 0.068; the failed verdict stands and the pattern of hit numbers changes, as given in the results.
- 9 November 2026: the Benjamini-Hochberg step named in the preregistration was not run by the study. This paper states the outcome it would have had.
- 9 November 2026: first published version.

12 Reproduction

- Data: the Kodak VISION3 technical data sheets for 5219, 7213 and 7203 [1], from Kodak's motion-picture product pages.
- Digitise, for each layer, the characteristic curve, the diffuse rms granularity curve at the 48 micrometre aperture and the spectral sensitivity at net density 0.2. Resample to 0.05 in log exposure.
- Compute eqs. 1, 2, 3 and 4 on that grid and take each layer's peak.
- Fit the variance laws, eq. 5, the eight response forms and the threshold model as in the methods, using log variance and not log standard deviation in the threshold fit.
- Preregistration SHA-256: c0dbcacf4caf6bd35564dde5fd02ea7b651165162206b5b1c4565fd9d326329a. Plan SHA-256: 67bcd757fe62d97c70fef6dc556bc91e443a0842712cff5345841a4726373faa. Run 20261031T061248Z. Result file SHA-256: 108b6952f1807a99580fdddd641ba3ca9dd38a3319f747be

cab8a60f168ad03a. Digitised data SHA-256: f3ebb6174d4a9e0f36f01837a6860c934c7b8c6783826b6661080a8b65c81427.

- Cost: about two and a half minutes on one desktop CPU (56 s analysis, 96 s shuffle control, under 1 s pairing control).

13 Appendix A. The preregistered verdicts

Each verdict is a set of rules written before the run. A verdict passes only when every rule holds. Values are read from the run's result file.

Table 2. Preregistered verdicts and the rule that decided each one.

Verdict	Deciding rule	Value	Holds
Planted DQE	planted error at most 0.2	0.113	yes
Planted variance law	planted law chosen at rate 0.8 or more	0.70	no
Planted forms	hold-out recovery 0.5 or more	0.80	yes
Siedentopf fails	beaten share 0.7 or more, and slope b below 0.7	1.0; -0.54	yes, void
DQE in range	median peak DQE above 0.01 and below 0.2	0.021	yes
Same efficiency at every speed	green DQE spread at most 3, and photon spread 5 or more	3.24; 15.7	no
No one form wins	top-form share below 0.6	1.0	no
Form beats shuffled density	top-form share above the shuffle 95th percentile (0.44)	1.0	yes
Hit number unidentified, wide grains	planted recovery below 0.5 at width 2.0	1.0	no
Hit number recovered, narrow grains	planted recovery 0.7 or more at width 0.5	1.0	yes
Threshold form fits	density error at most 0.04, and log variance error at most 0.05	0.043; 0.134	no
Slope tracks noise	median correlation above the pairing 95th percentile (0.76)	0.61	no

14 Glossary

Each term is defined where it is first used in the text. The same definitions, with the field each comes from and what each excludes, are collected here.

- **detective quantum efficiency (DQE)**. The share of the information carried by the arriving photons that a detector's output keeps, measured as the square of its output signal-to-noise ratio divided by the square of the input's; 1 is a perfect photon counter. Field: imaging science and detector physics. Excludes the plain share of photons absorbed, and the speed of a film.
- **characteristic curve**. The developed optical density of a film plotted against the logarithm of the exposure it received. Field: photographic sensitometry; also called the H and D curve. Excludes the film's spectral sensitivity and its grain.
- **rms granularity**. The standard deviation of density read through a small round aperture scanned across a uniformly exposed patch; Kodak's sheets use a 48 micrometre aperture. Field: photographic sensitometry. Excludes how grainy the image looks to a viewer, and noise measured at other apertures.
- **spectral sensitivity**. For each wavelength, the reciprocal of the exposure energy per unit area needed to reach a stated net density, here 0.2. Field: photographic sensitometry. Excludes sensitivity to white light or to a named illuminant.
- **noise-equivalent quanta (NEQ)**. The number of photons a perfect detector would need to give the same signal-to-noise ratio as the real output. Field: imaging science. Excludes the number of photons that arrived.
- **Siedentopf's rule**. The classic statement that the variance of density grows in proportion to the mean density, so that rms granularity grows as its square root. Field: photographic science. Excludes Selwyn's law, which concerns the size of the measuring aperture.

- **Poisson threshold model**. A model in which each grain, or each rod cell, responds once it has absorbed at least n photons, with photon counts following Poisson statistics; n is the hit number. Field: photographic science and visual physiology. Excludes models in which response grows smoothly with every photon.
- **preregistration**. A dated written plan for a study, fixed and hash-locked before any result is seen: the question, the data, the statistic, the controls, the forecasts and the kill conditions. Field: research method (metascience). Excludes a plan revised after looking at the outcome, and a protocol written only for the record after the run.
- **planted check**. Data built so that the right answer is known, run through the same analysis before the real data, to show that the analysis can see the effect when it is there. Field: research method; called a positive control in experimental science. Excludes any check made on the real data, which cannot show what the analysis misses.
- **control**. A version of the data with the structure under test removed (for example shuffled or pooled), scored the same way as the real data, so the real score has something to beat. Field: statistics and experimental design. Excludes a comparison with a different dataset or with a published number.
- **kill condition**. A rule written in the preregistration that, when met, decides against the claim with no room for judgement after the run. Field: research method (our usage, after Popper's falsification). Excludes reasons found after the run for setting a result aside.
- **verdict**. The pass or fail of one preregistered rule set against the run's numbers, read off by code. Field: research method (our usage). Excludes an interpretation written after reading the numbers.
- **net density**. Optical density above the film's base and fog, so that unexposed film reads zero. Field: photographic sensitometry. Excludes the density of the base and the fog themselves.
- **gamma**. The local slope of the characteristic curve: the change in density per unit change in log exposure. Field: photographic sensitometry. Excludes the display gamma of video and computer screens.
- **BIC**. A model score that adds a penalty for each fitted parameter to the misfit, lower being better, so that a more flexible model must earn its extra parameters. Field: statistics; the Bayesian information criterion. Excludes a test on held-out data.
- **hold-out score**. The error of a model fitted on one part of the data and scored on a part it never saw; here fitted on the lowest 70 percent of the axis and scored on the top 30 percent. Field: statistics and machine learning. Excludes an error computed on the points used to fit.
- **hinge**. A response that is zero below a threshold exposure and rises in a straight line above it, here in log exposure. Field: curve fitting. Excludes any form that saturates at high exposure.
- **void**. The verdict of a run whose planted check failed or whose sample fell below its stated minimum; it says nothing for or against the claim. Field: research method (our usage). Excludes a null result, which comes from a working instrument.
- **null distribution**. The spread of scores a control gives over many random draws; its 95th percentile is the bar a real score must clear. Field: statistics. Excludes a theoretical distribution assumed rather than drawn from the data.
- **p value (from draws)**. The share of control draws scoring at or above the real score, with one added to the count and to the number of draws, so it is never zero. Field: statistics. Excludes a p value from a textbook formula that assumes a distribution.
- **Benjamini-Hochberg correction**. A rule for many tests at once: rank the p values from smallest, and reject the k smallest when the kth is no larger than k over the number of tests times the chosen false discovery rate, here 0.05. Field: statistics. Excludes a correction for one test, and the stricter Bonferroni rule.
- **findings ledger**. The append-only, dated list of every result we have recorded, numbered F1, F2 and on; corrections are added as new dated lines and old entries are never rewritten. Field: our records. Excludes a curated list of successes.

15 References

[1] Eastman Kodak Company. KODAK VISION3 500T 5219, 200T 7213 and 50D 7203 colour negative films, technical data sheets. Example: https://www.kodak.com/content/products-brochures/motion-picture/KODAK-VISION3-500T-5219-7219-brochure.pdf

[2] Shaw, R. (1963). The equivalent quantum efficiency of the photographic process. The Journal of Photographic Science 11, 199-204. doi:10.1080/00223638.1963.11736919

[3] Shaw, R. (1963). Photon fluctuations, equivalent quantum efficiency, and the information capacity of photographic images. The Journal of Photographic Science 11(6), 313-320. doi:10.1080/00223638.1963.11736937

[4] Dainty, J. C. and Shaw, R. (1974). Image Science: Principles, Analysis and Evaluation of Photographic-Type Imaging Processes. London: Academic Press. ISBN 0122008502.

[5] Stephenson, I. (2007). Simulating film grain using the noise power spectrum. Theory and Practice of Computer Graphics, Eurographics UK.