Universität Hamburg

Department Informatik
Arbeitsbereich Kognitive Systeme (KOGS)

64-350 Multidimensionale und Multimodale Signale

Gliederung:

Exkurs 1: Radon-Transformation

Radon-Transformation

Einleitung

CT- oder MRT-Aufnahmen, wie sie in der medizinischen Bildverarbeitung Verwendung finden, stellen nicht direkte Messungen von Sensoren dar, sondern müssen erst aus den Messdaten berechnet werden.

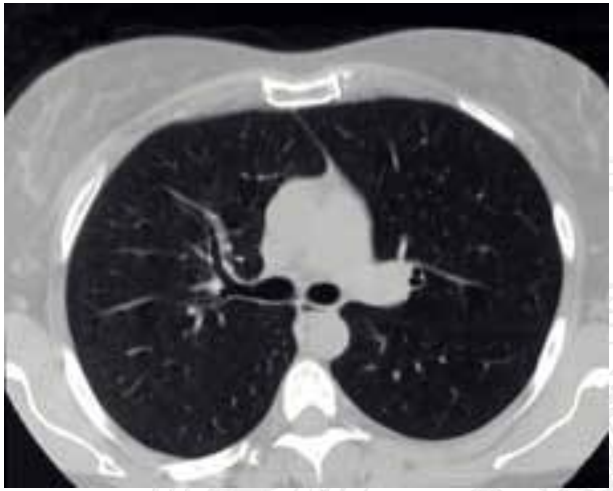

Radon-Transformation

Einleitung

CT- oder MRT-Aufnahmen, wie sie in der medizinischen Bildverarbeitung Verwendung finden, stellen nicht direkte Messungen von Sensoren dar, sondern müssen erst aus den Messdaten berechnet werden.

Jeder Punkt in einem kassischen Röntgenbild stellt das Ergebnis einer Integration entlang des jeweiligen „Sehstrahls“ dar (Transmissionsmessung).

Klassisches Röntgenbild (Projektionsradiographie):

-Modulierte Verteilung der durch Gewebe transmittierten Röntgenquanten

-2D-Projektionen der Schwächungseigenschaften des Gewebes

-*Überlagerungsbild:*

- **alle** durchstrahlten Volumenelemente tragen zur Schwächung bei

- Strukturen mit starker Absorption (z.B. Knochen) gut erkennbar. Weichteilgewebe kaum darstellbar.

Radon-Transformation

Idee: Erstelle Projektionsbilder aus verschiedenen Richtungen und berechne hieraus das gewünschte Schichtbild.

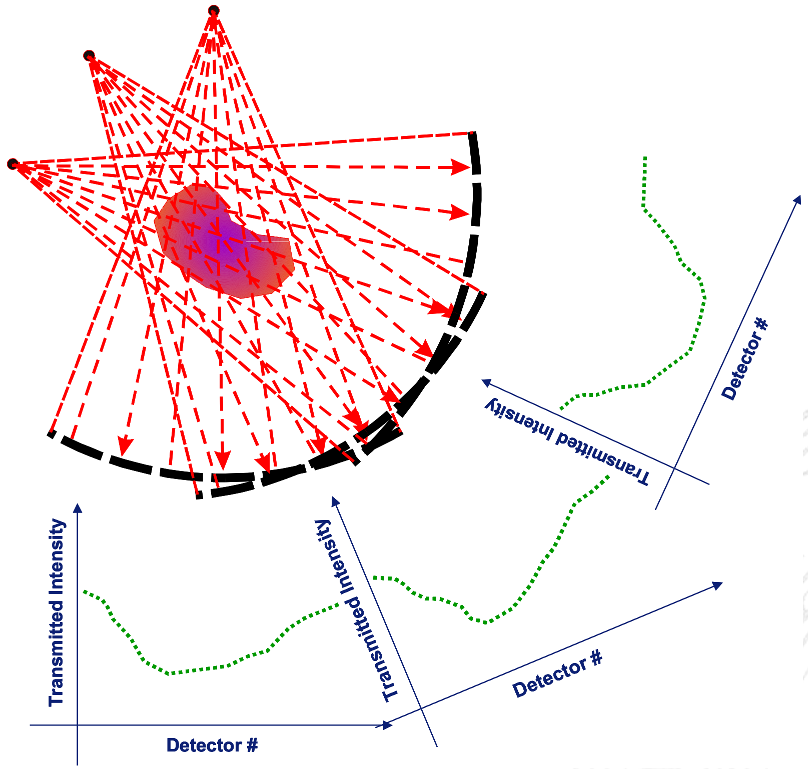

Radon-Transformation

Einfachstes Meßprinzip der CT

Mit Radon Vorgabe:
erst Translation
(ergibt Projektion)
dann
Rotation (min. 180°)

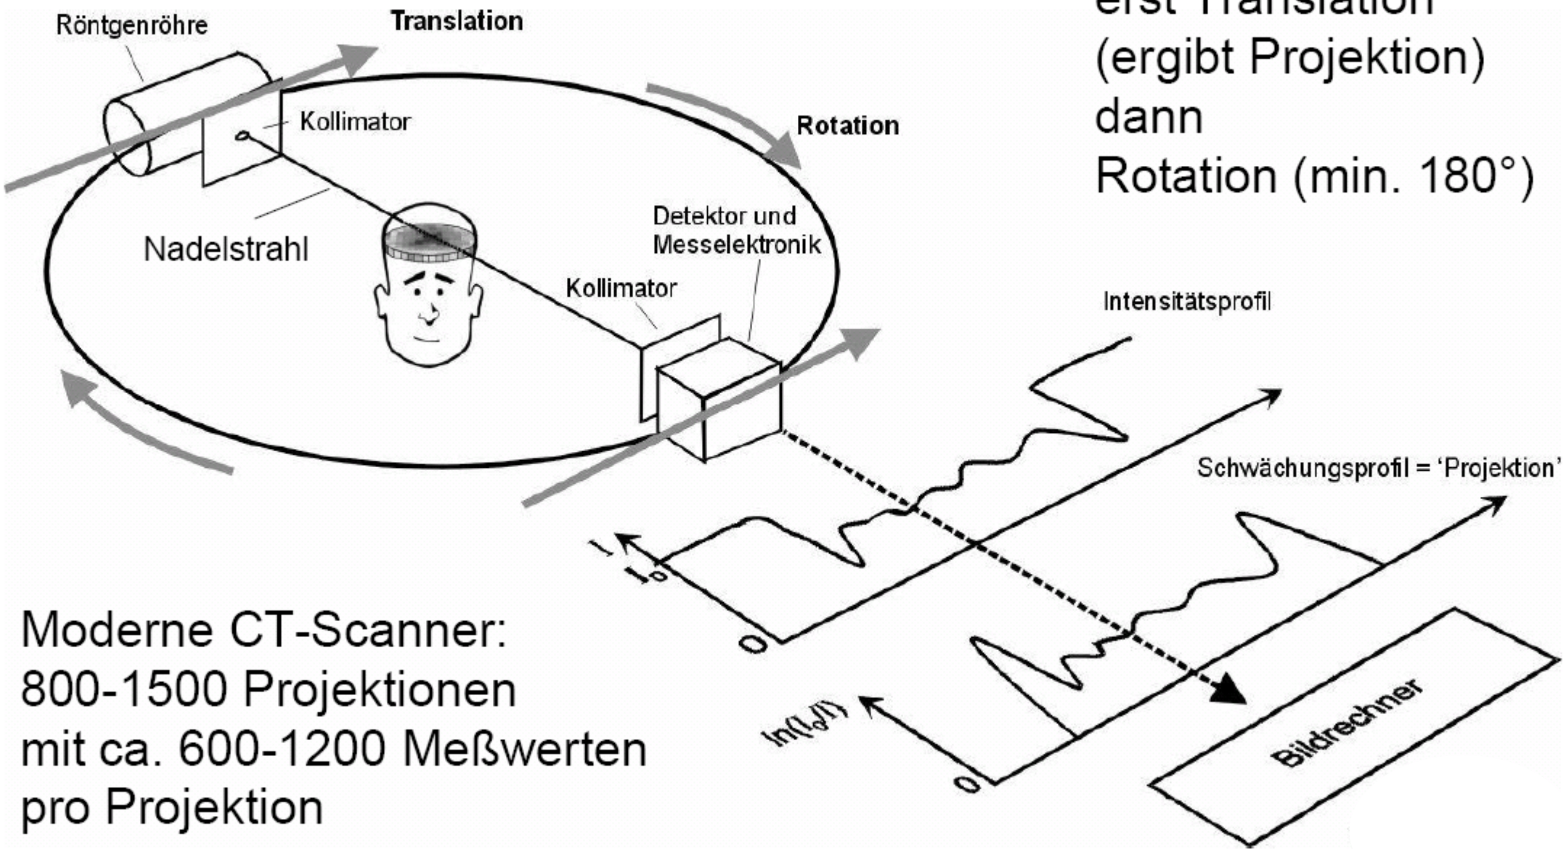

Moderne CT-Scanner:
800-1500 Projektionen
mit ca. 600-1200 Meßwerten
pro Projektion

Radon-Transformation

J. Radon (1917): Die 2D-Verteilung einer Objekteigenschaft kann exakt beschrieben werden, wenn eine **unendliche Anzahl von Linienintegralen** vorliegt.

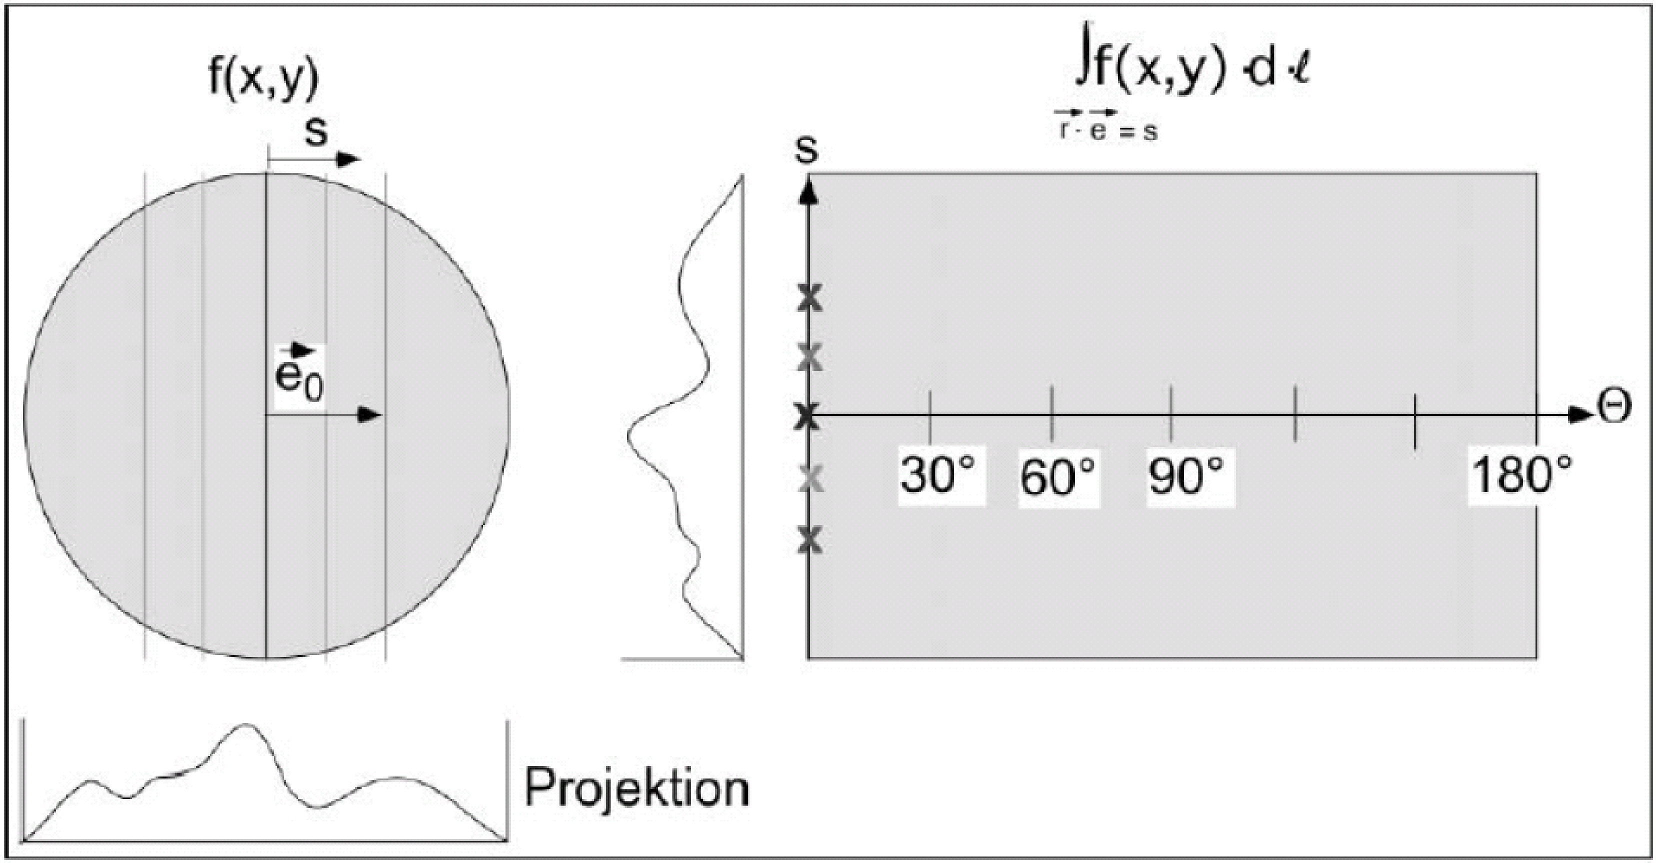

Radon-Transformation

J. Radon (1917): Die 2D-Verteilung einer Objekteigenschaft kann exakt beschrieben werden, wenn eine **unendliche Anzahl von Linienintegralen** vorliegt.

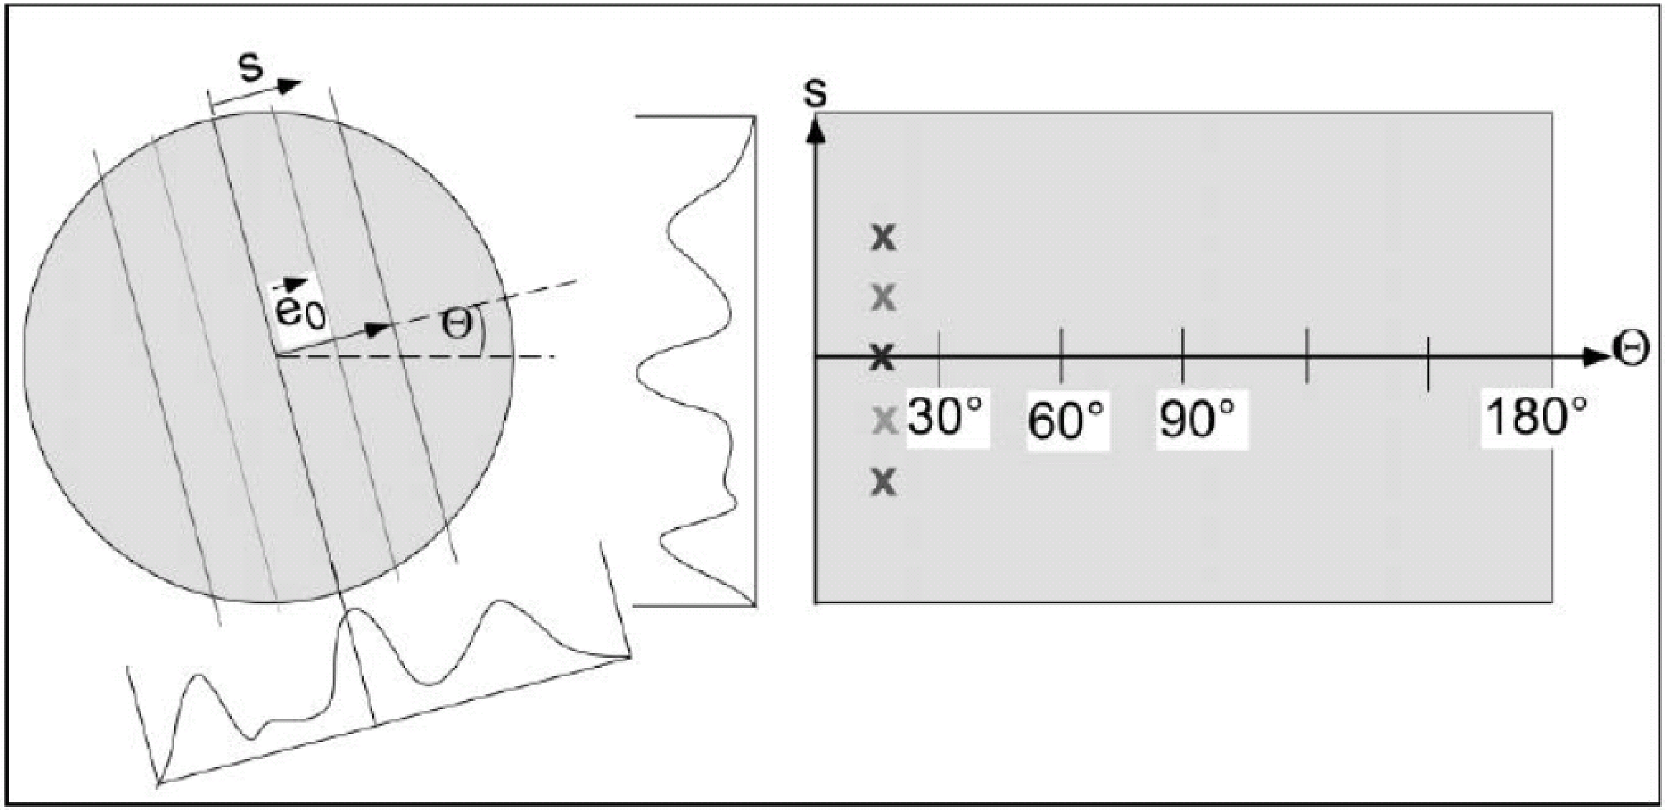

Radon-Transformation

-Geraden werden als lokale Maxima erkennbar

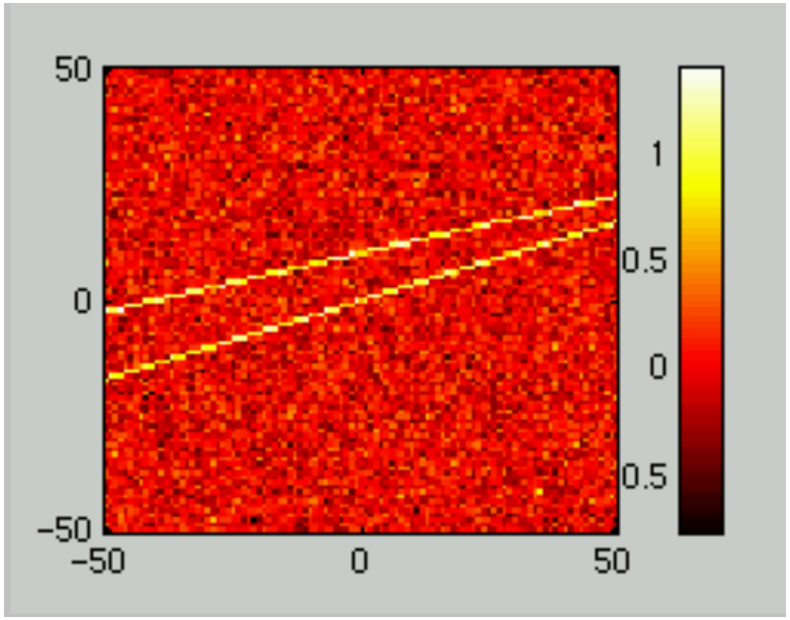

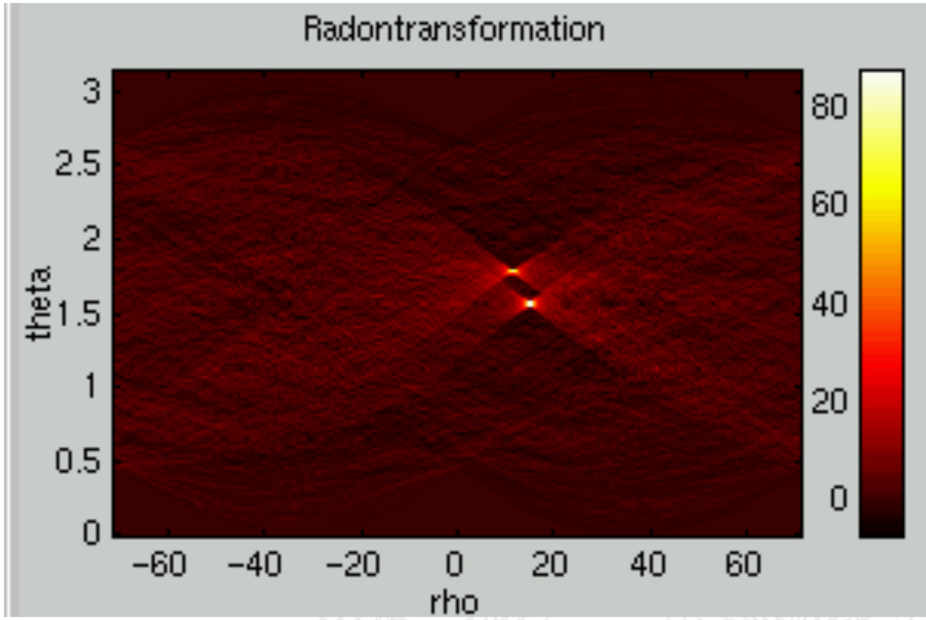

Radon-Transformation

-Geraden werden als lokale Maxima erkennbar

-hohe Rauschunempfindlichkeit

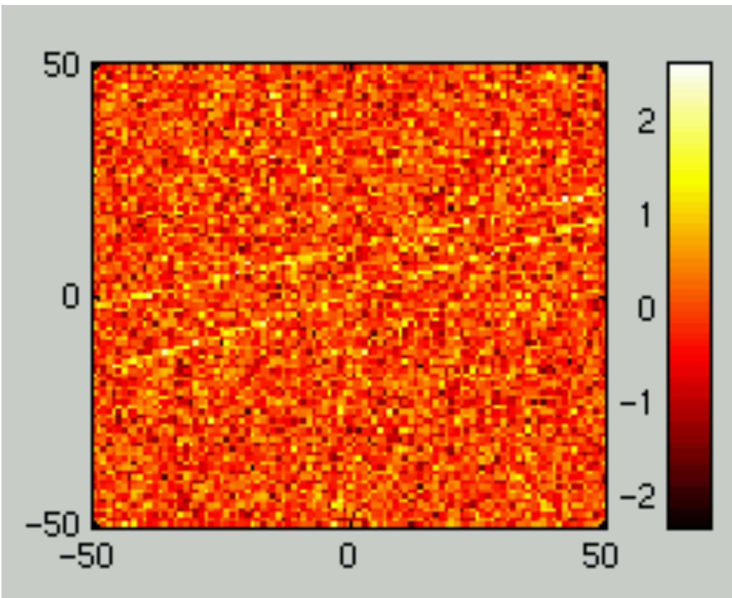

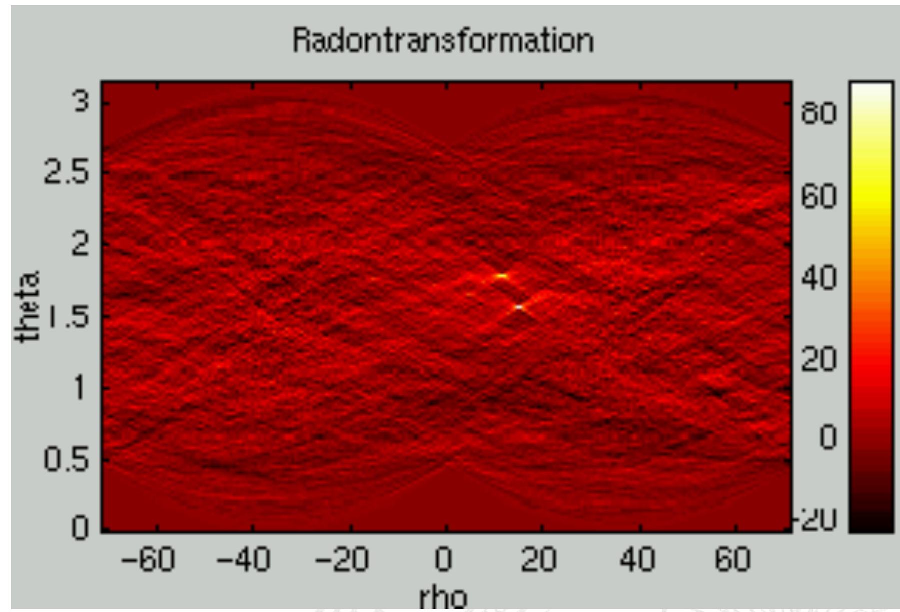

Radon-Transformation

-Geraden werden als lokale Maxima erkennbar

-hohe Rauschunempfindlichkeit

-Sinusoide Strukturen → „Sinogramm“

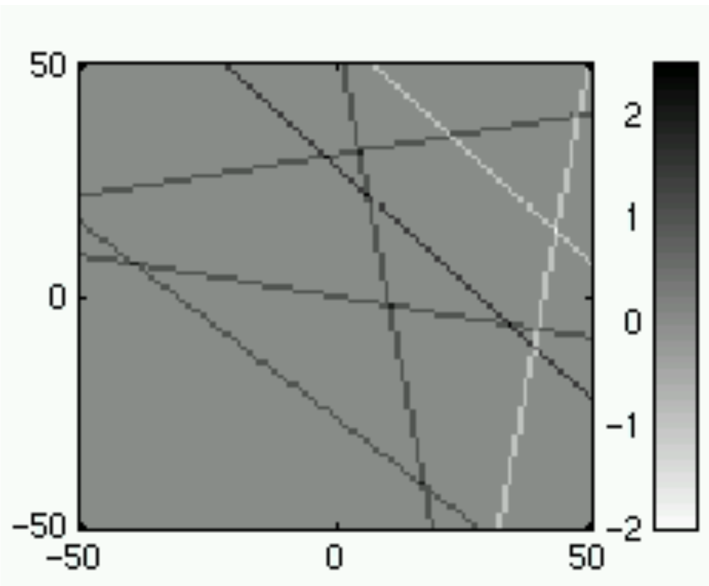

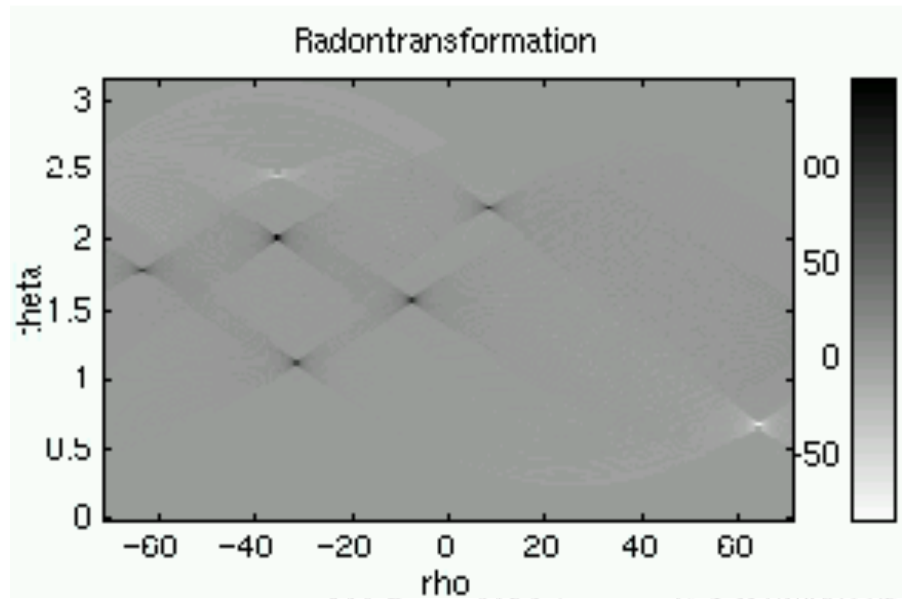

Radon-Transformation

-Geraden werden als lokale Maxima erkennbar

-hohe Rauschunempfindlichkeit

-Sinusoide Strukturen → „Sinogramm“

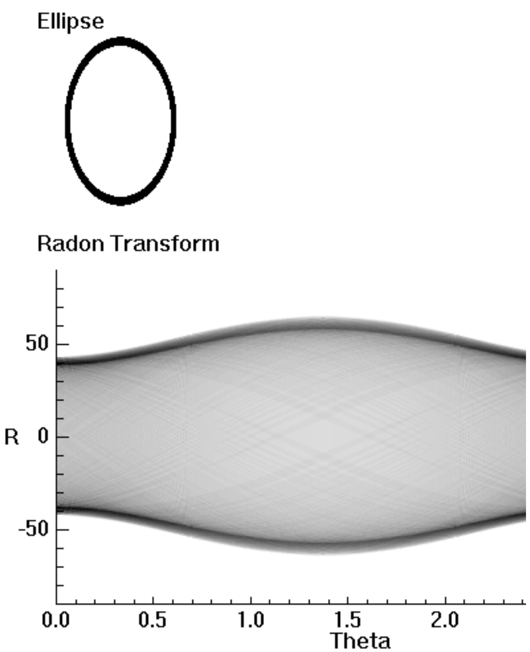

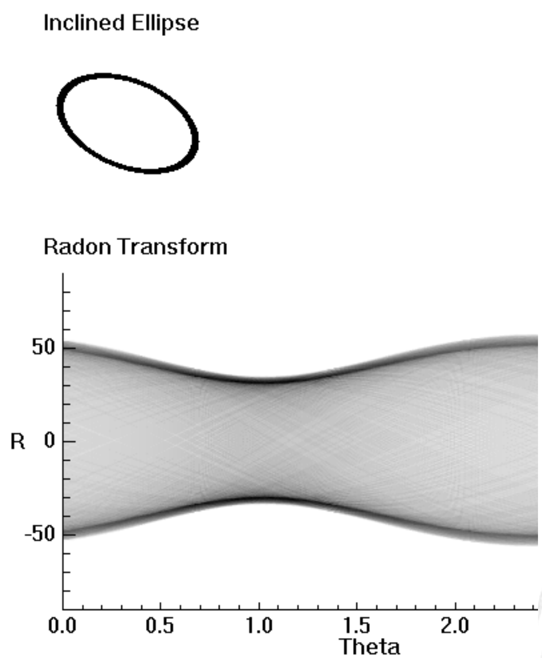

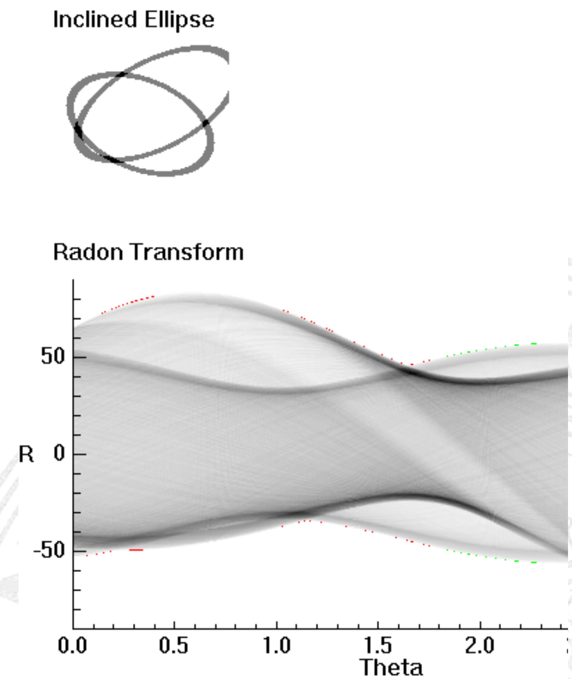

Inverse Radon-Transformation

Rücktransformation mit Hilfe des Fourier-Scheiben-Theorems:

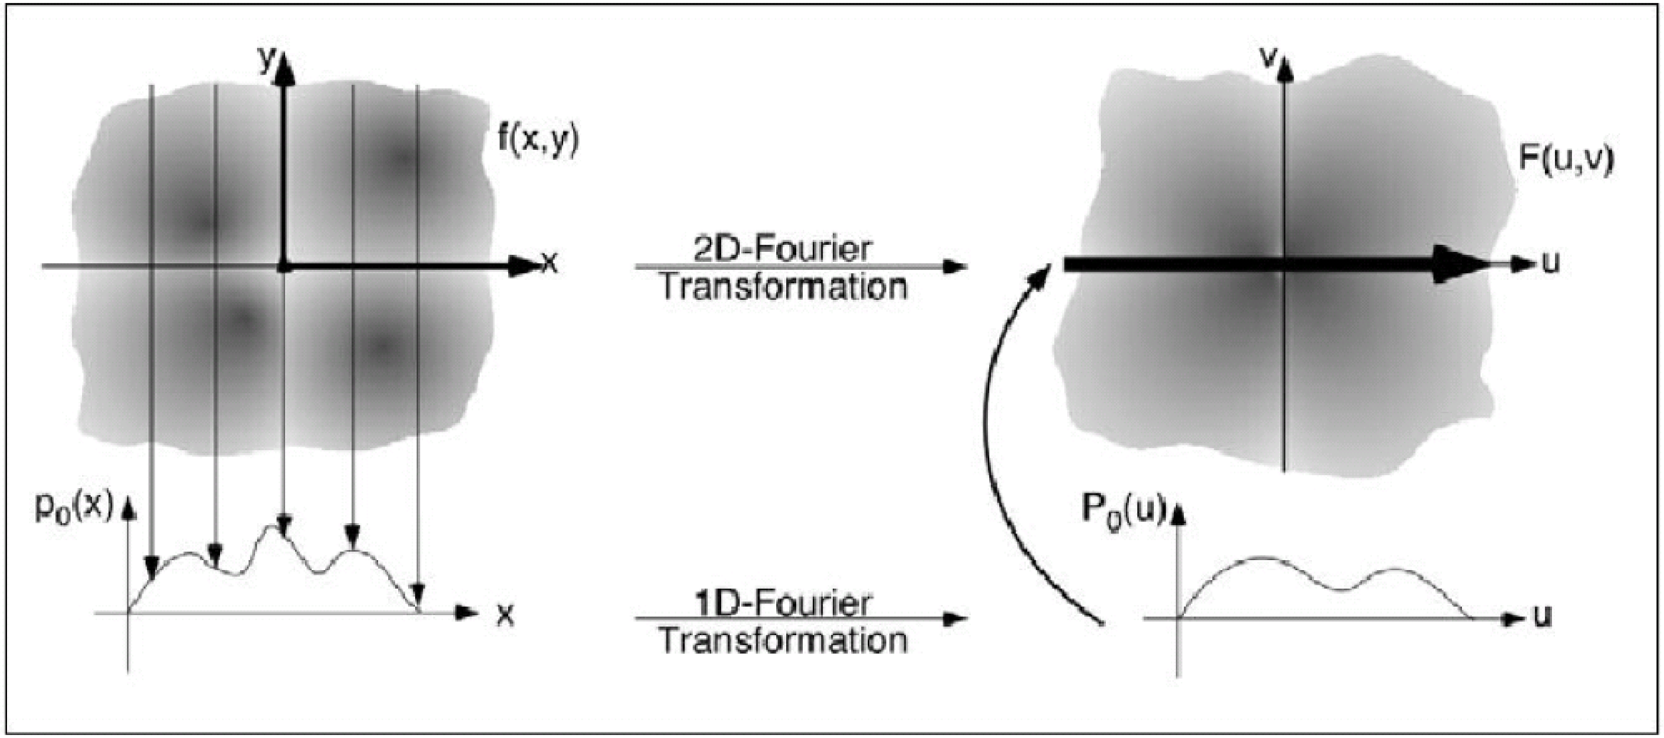

Inverse Radon-Transformation

Rücktransformation mit Hilfe des Fourier-Scheiben-Theorems:

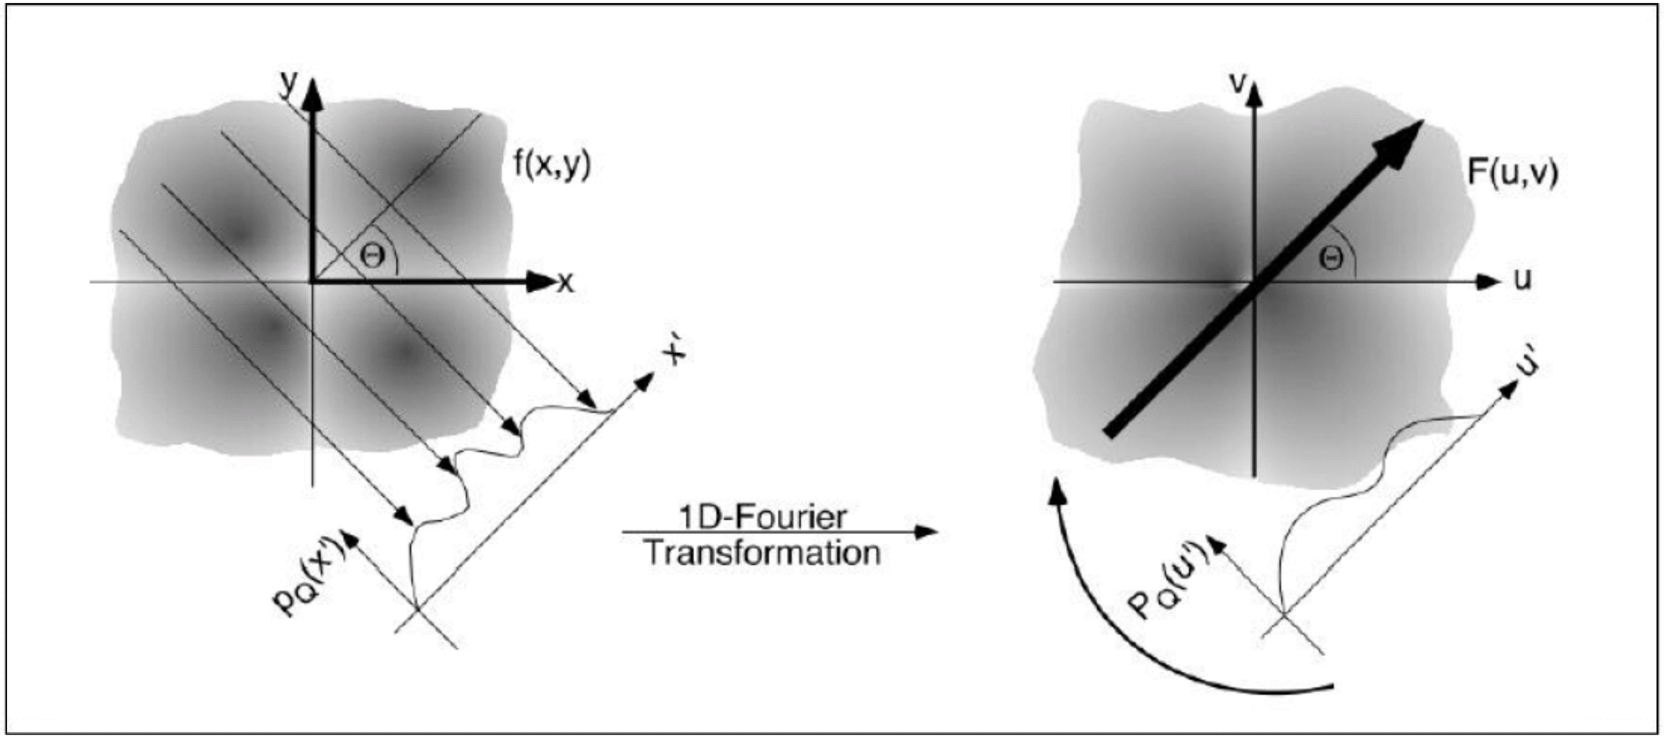

Inverse Radon-Transformation

Rücktransformation:

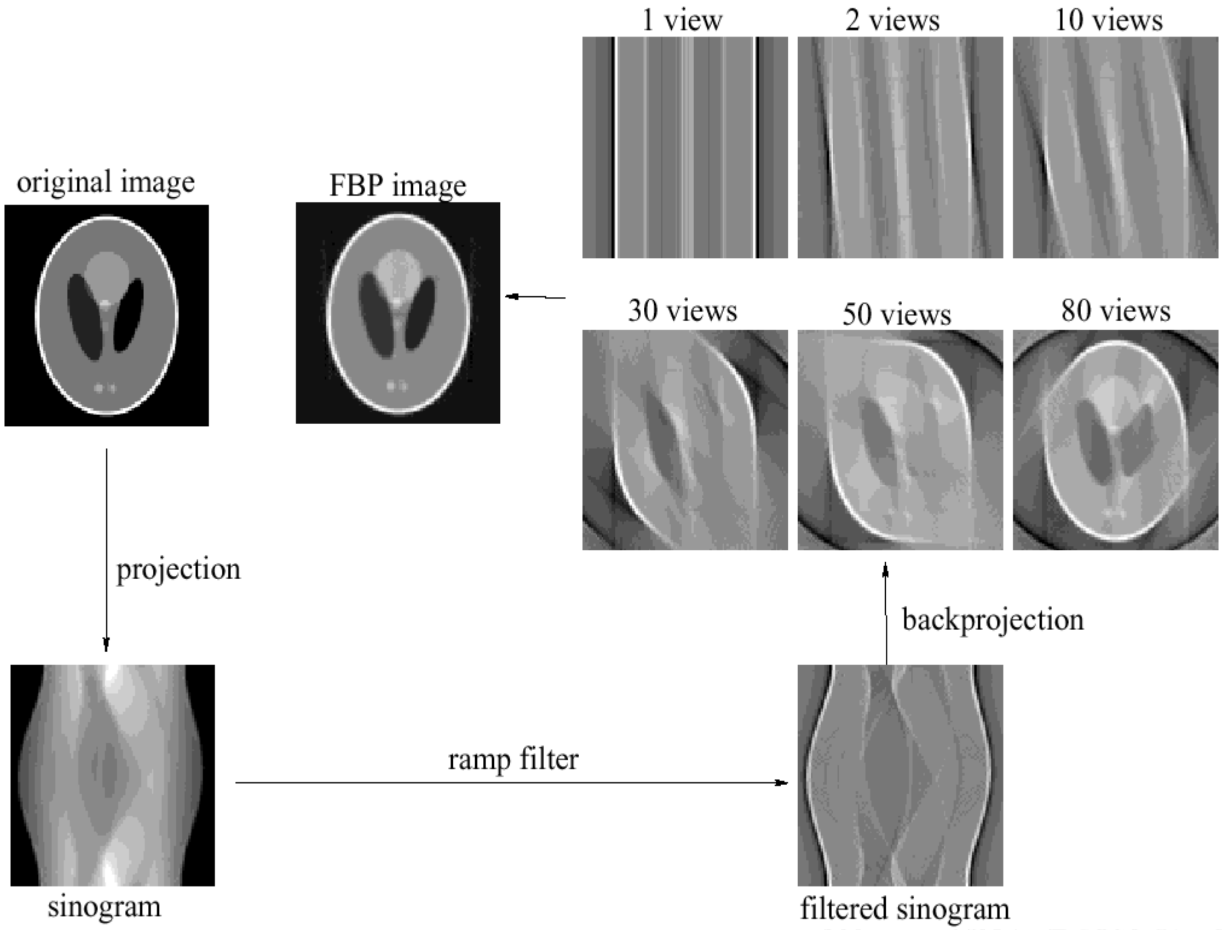

Exkurs 2: Blur-Invariance

Blur-Invariance

Wie in der Übung gezeigt, verändert sich das quadrierte Phasenspektrum eines Signals nicht, wenn das Signal mit einer rellwertigen geraden Funktion (z.B. Unschärfe oder Bewegungsunschärfe) gefaltet wird.

Der Grund hierfür liegt darin, dass die Faltung mit einer Filterfunktion im Zeitbereich einer Multiplikation im Frequenzbereich bedeutet. Die Multiplikation mit einer geraden reellwertigen Funktion ändert jedoch nur die Beträge einer Funktion, aber nicht den Phasenwinkel (ausser einer eventuellen 180°-Drehung durch negative Werte der Filterfunktion). Durch Quadrieren des Signals heben sich die 180°-Drehungen auf, so dass folgt, dass das quadrierte Phasenspektrum nicht durch die Faltung beeinflusst wird (ausser an Punkten, an denen der Betrag 0 wird und somit die Phaseninformation verloren geht).

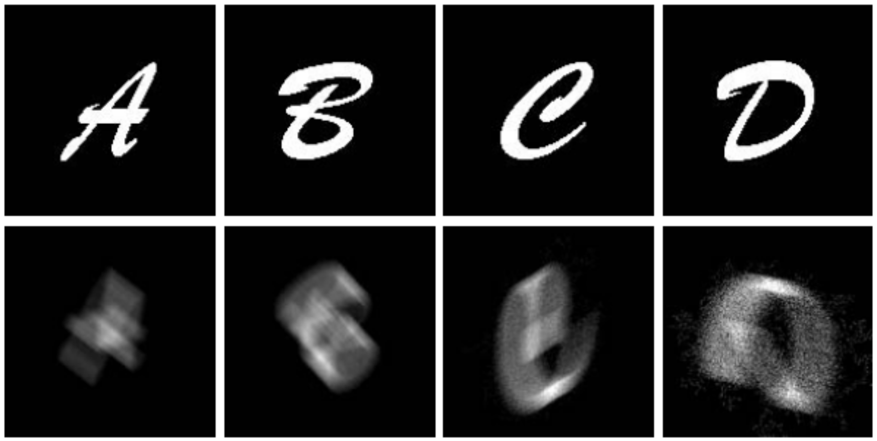

Figure 1. Top row: images of the first four letters in the alphabet. Bottom row: motion blurred, rotated, scaled, translated and noisy versions of the same images. PSNRs are from left to right: 50, 40, 30 and 20 dB. Noise is removed from the background using thresholding and connectivity analysis.

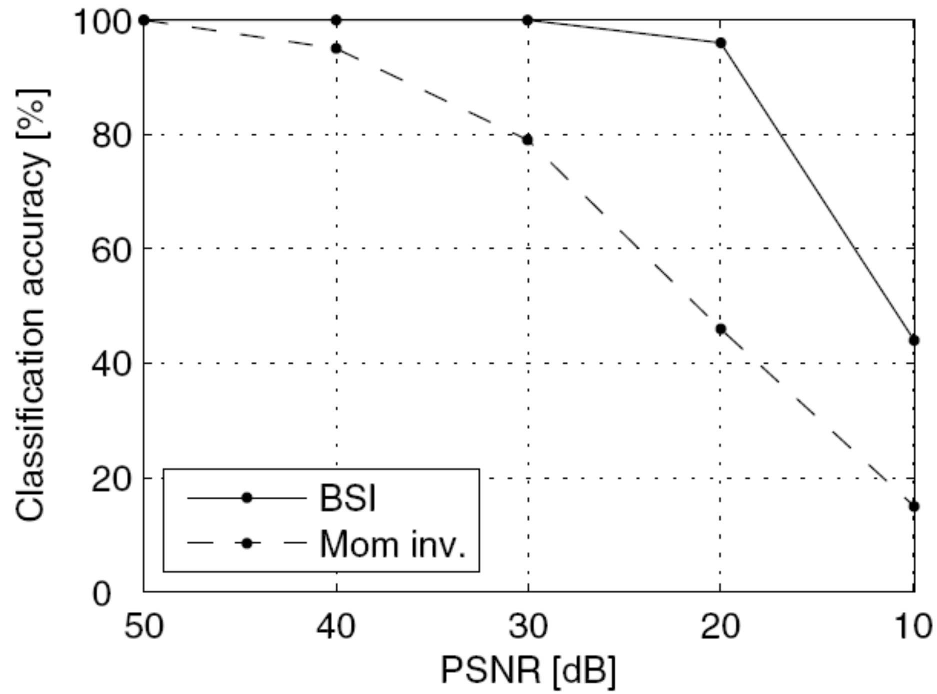

Figure 2. Classification accuracy of the nearest neighbor classification of distorted images of the letters of alphabet.

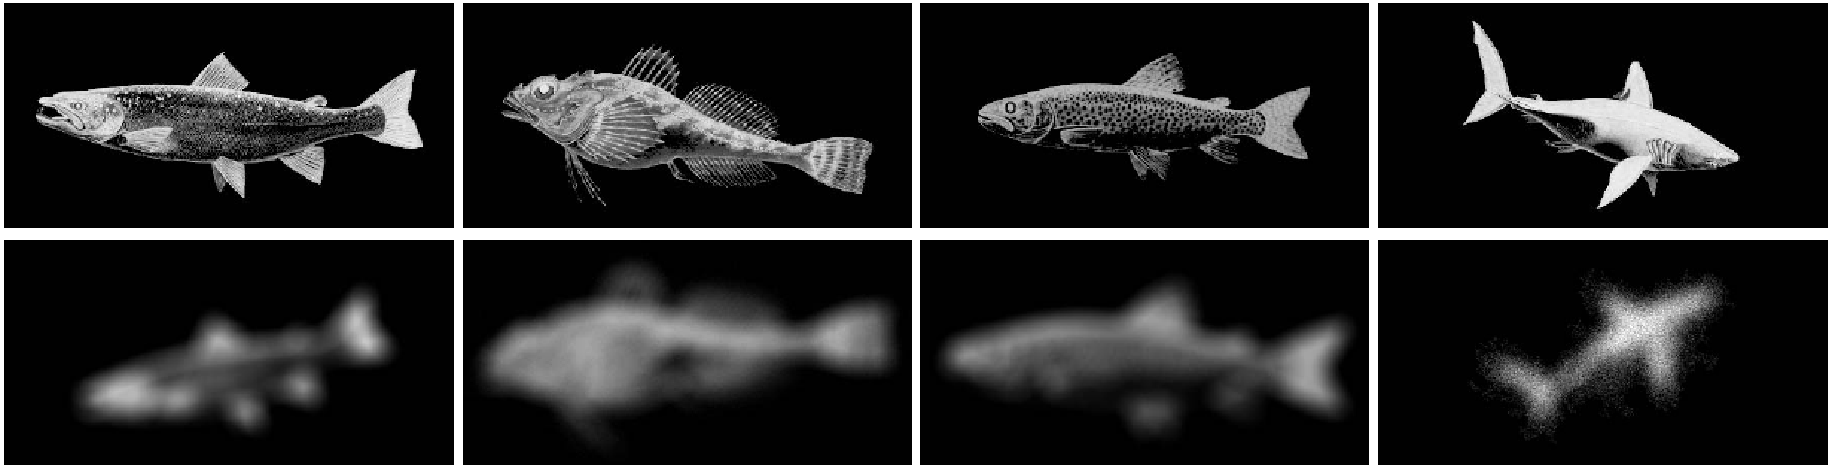

Figure 3. Top row: four examples of the 90 fish images. Bottom row: circularly blurred, rotated, scaled, translated and noisy versions of the same images. PSNRs are from left to right: 50, 40, 30 and 20 dB. Noise is removed from the background using thresholding and connectivity analysis. (200×400 images are cropped from the 500×500 images that were used.)

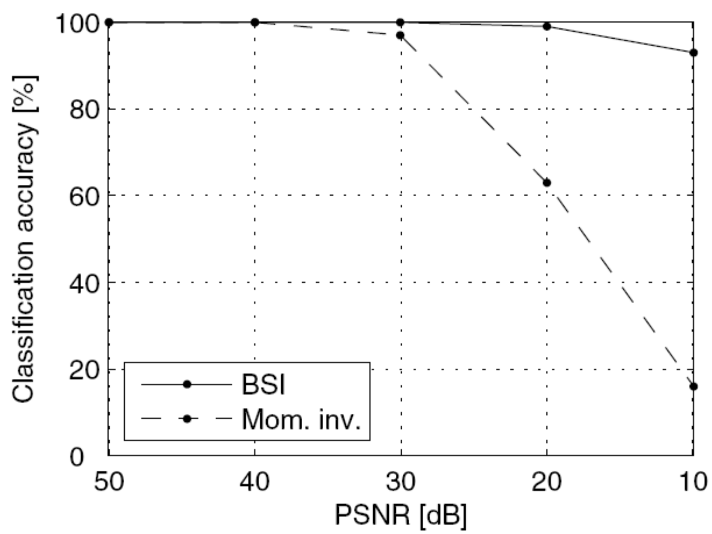

Figure 4. Classification accuracy of the nearest neighbor classification of distorted images of 90 fish.

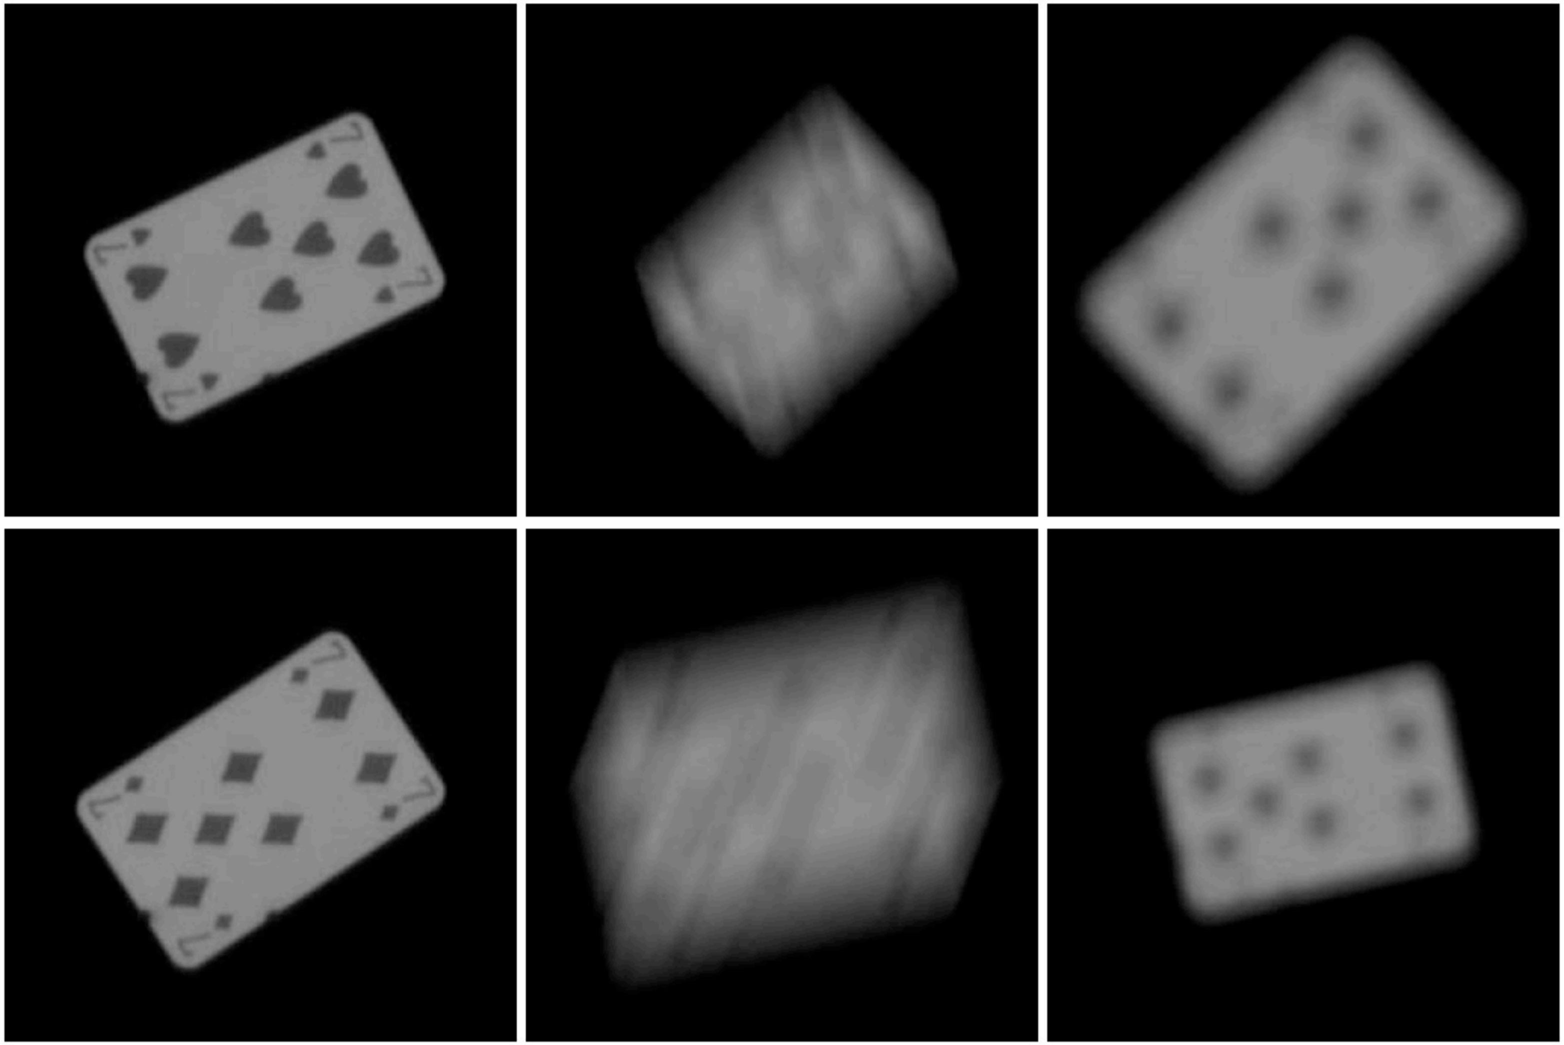

Figure 5. The rows contain one sharp and two blurred and transformed versions of the cards Seven of Hearts and Seven of Diamonds, respectively. Cards are photographed on a black background. Background clutter is removed using thresholding and connectivity analysis.

Teil VIII: Sensoren, Rauschen, Rauschreduktion

Quellen: P. Stelldinger, H.S. Stiehl: VL Systemtheorie, WiSe 07/08, Univ. Hamburg
H. Neumann, VL Computer Vision I, SoSe 07, Univ. Ulm

p=0.01

p=0.05

p=0.1

Impulse Noise: the effect of probability p

Median filtered image ($p = 0.1$)

$\sigma = 10$ $\sigma = 20$ $\sigma = 50$

Gaussian Noise: the effect of standard deviation (σ)

Noise in Images

All images created through optical projection onto a sensor array are noisy.

Correlated noise

- Due to electrical interference
- Due to source / sensor interference
- Halftone distortion / moiré patterns

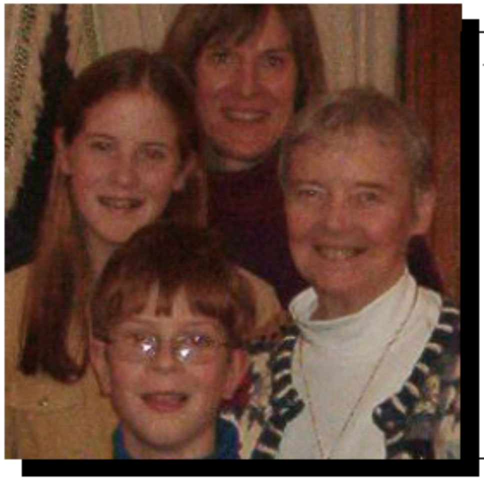

Uncorrelated noise

- Quantum noise in CCD arrays
- Silver halide grains in film photography
- Neuronal noise in a retina
- Quantization noise in digital photographs

Image with Additive Noise

noisy image — additive noise — undegraded image

$$J(r,c) = I(r,c) + N(r,c).$$

spatial domain

noisy image — additive noise — undegraded image

$$\mathfrak{J}(v,u) = \mathfrak{I}(v,u) + \mathfrak{N}(v,u).$$

frequency domain

Uncorrelated Noise

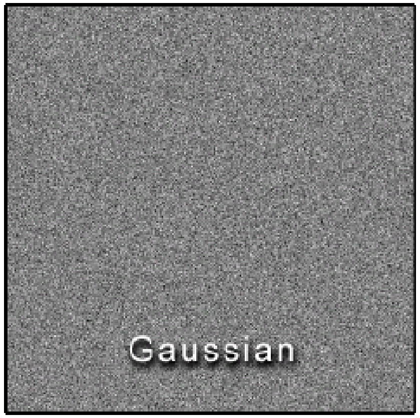

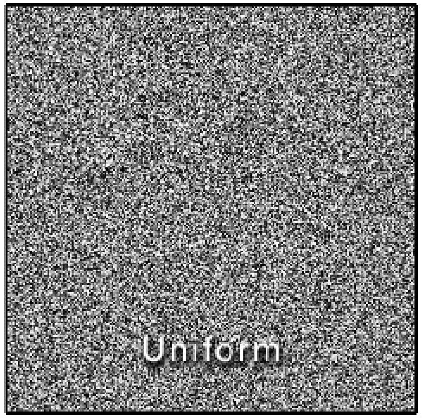

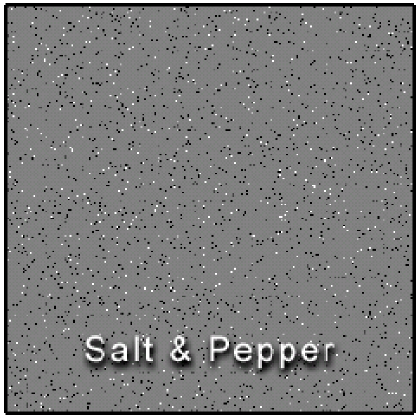

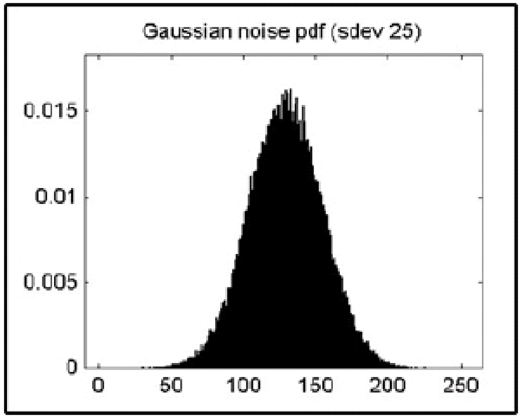

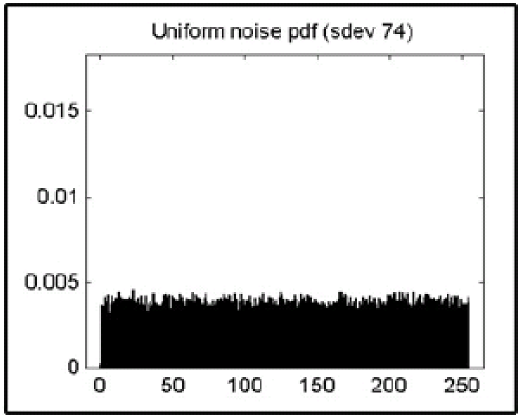

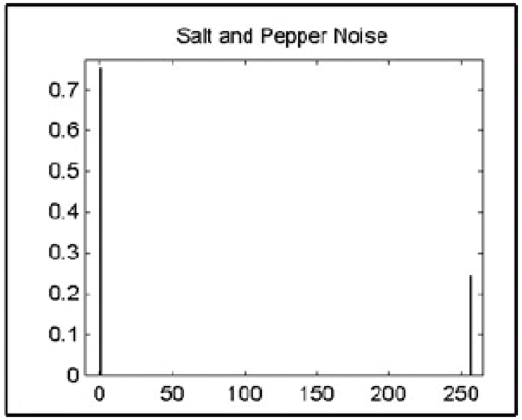

Each pixel's value has probability of occurrence given by the associated distribution.

Intensity distributions – normalized histograms

Uncorrelated Noise

The most likely value is 128 with an average difference of 25 from 128 (std. dev.).

This is sparse noise: Only 12.5% of the pixels contain noise. Of those 12.5% ...

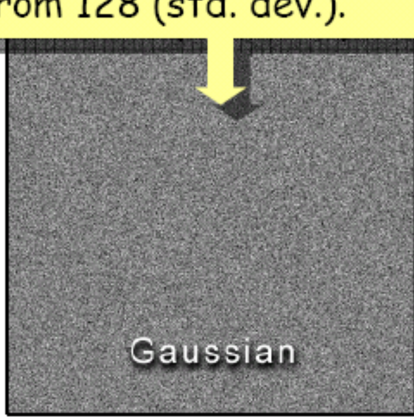

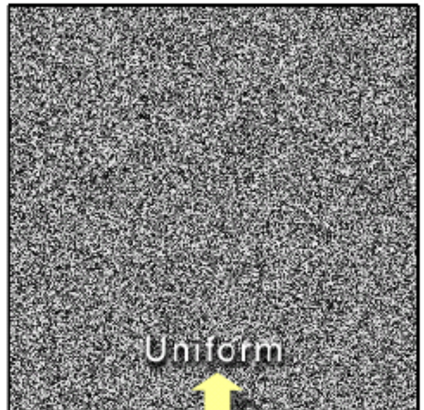

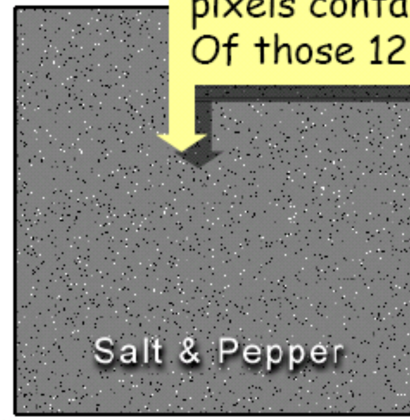

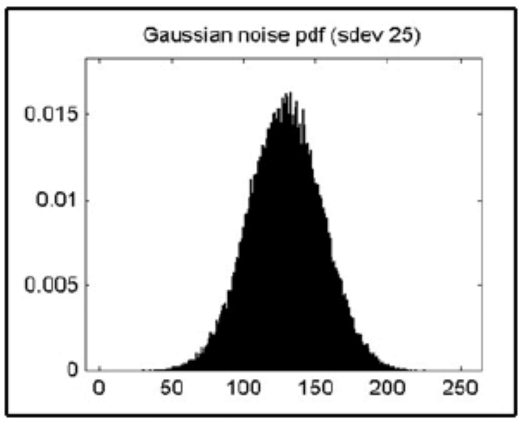

All values occur with equal probability.

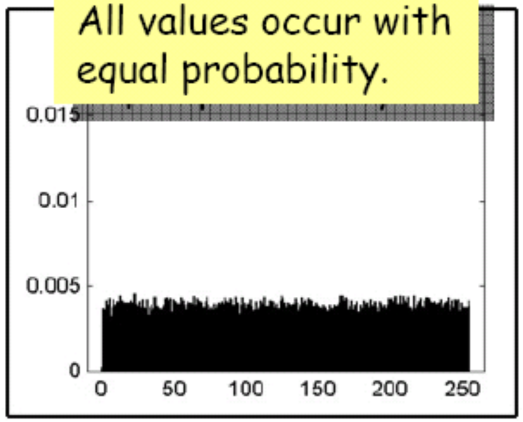

... black pixels occur 75% of the time and

white pixels occur 25% of the time.

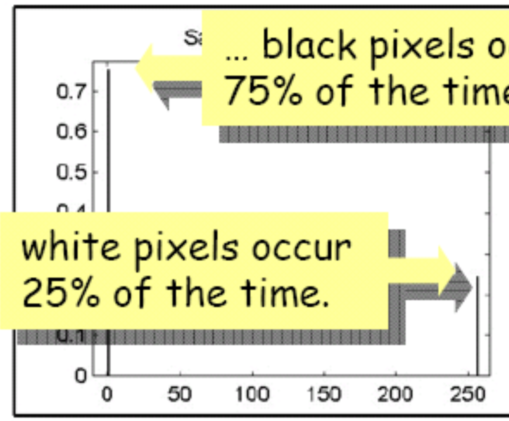

Each pixel's value has probability of occurrence given by the associated distribution.

Uncorrelated Color Noise: Gaussian

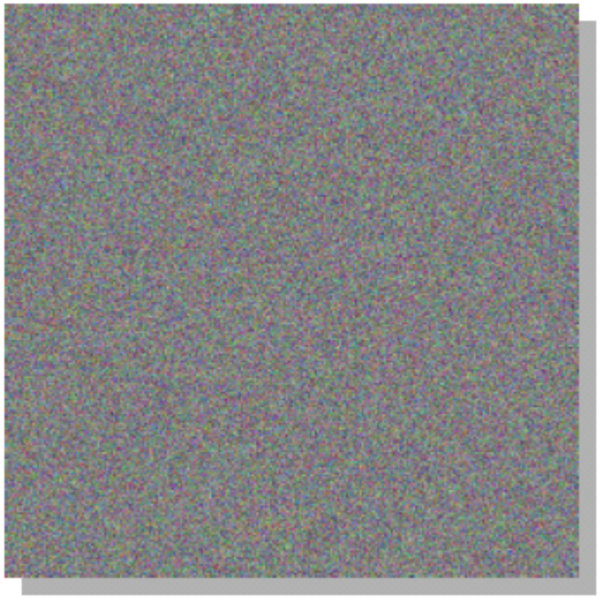

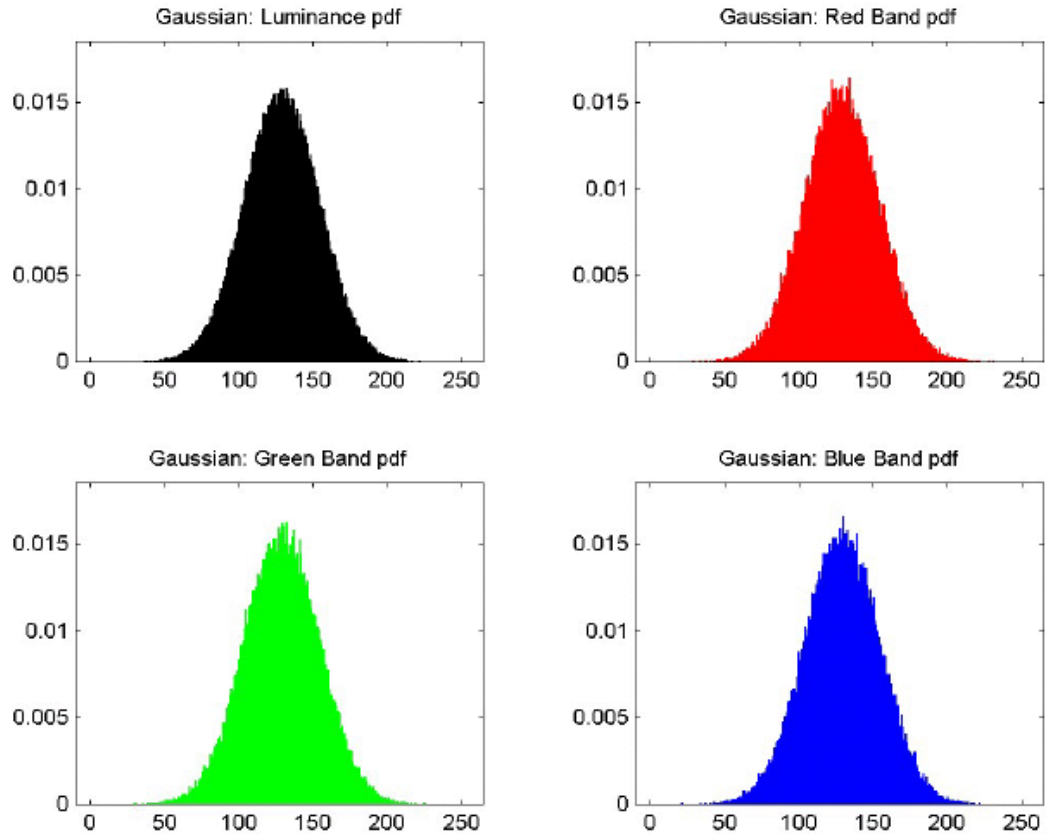

Uncorrelated Color Noise: Uniform

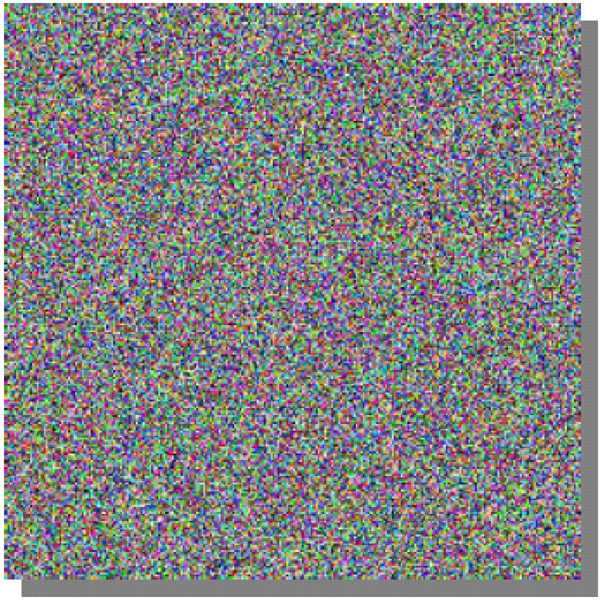

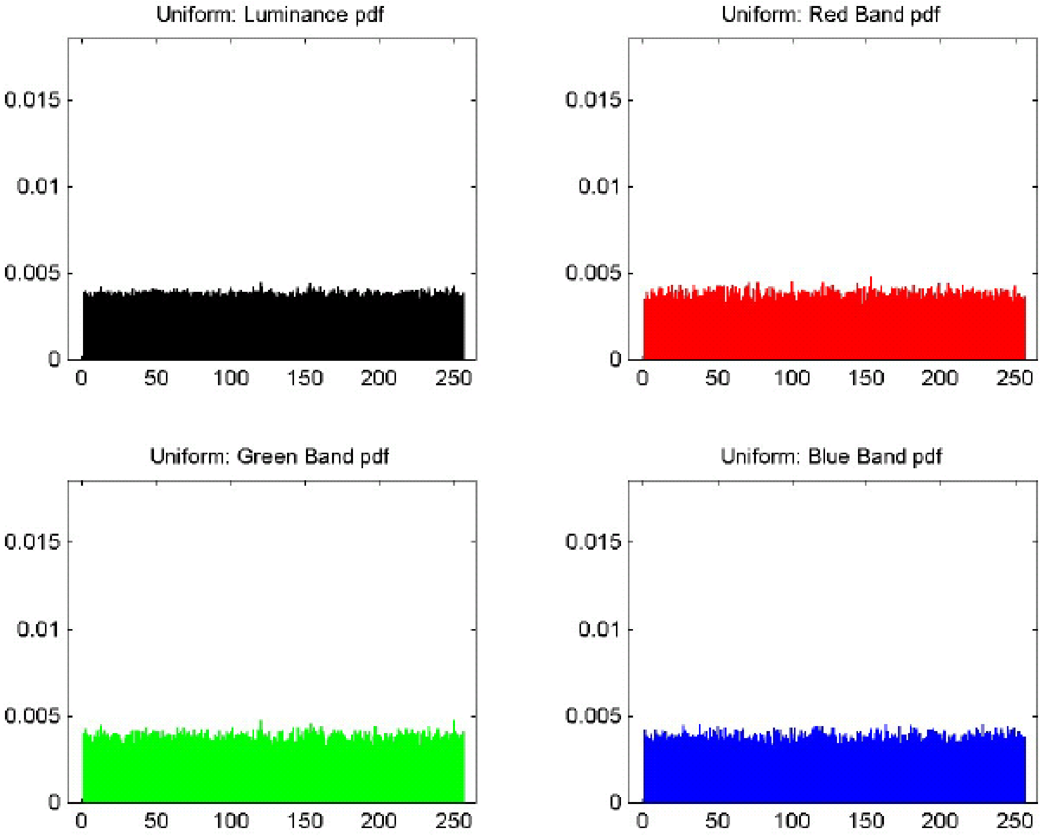

Gaussian IID Noise Field

IID ⇒ no spatial correlation

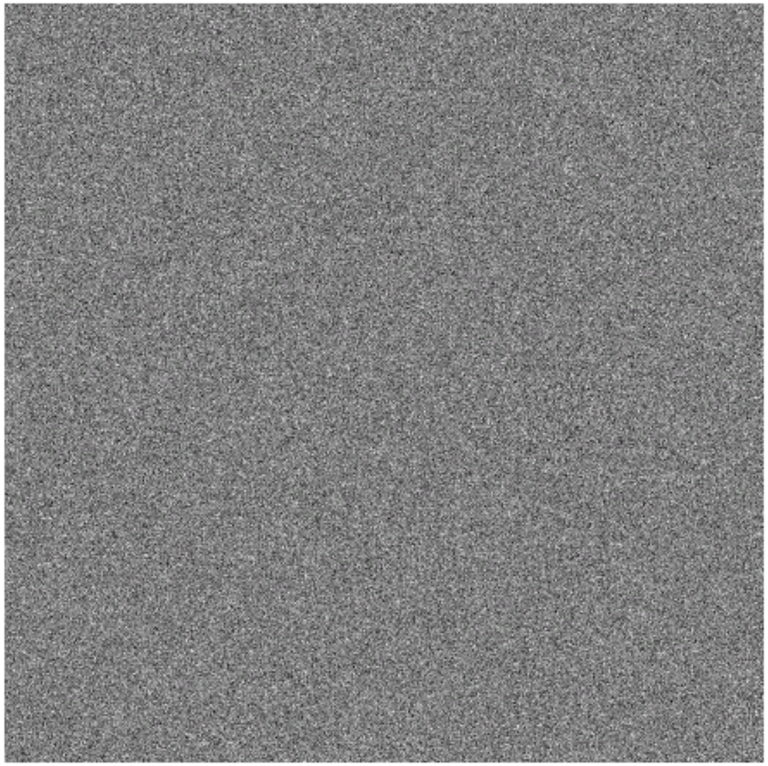

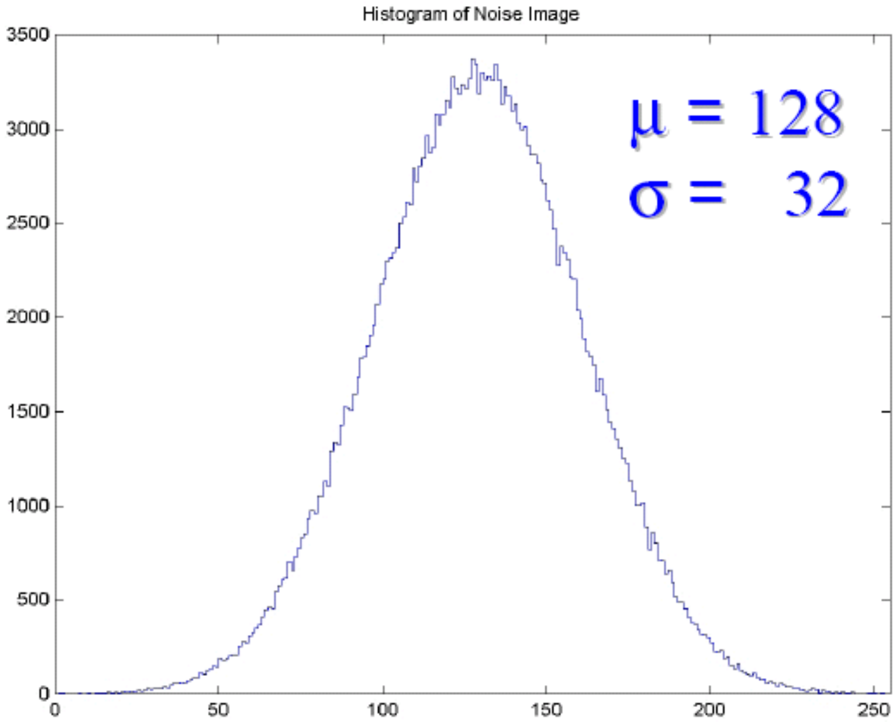

IID: Independent, Identically Distributed

Gaussian IID Noise Field

IID ⇒ no spatial correlation

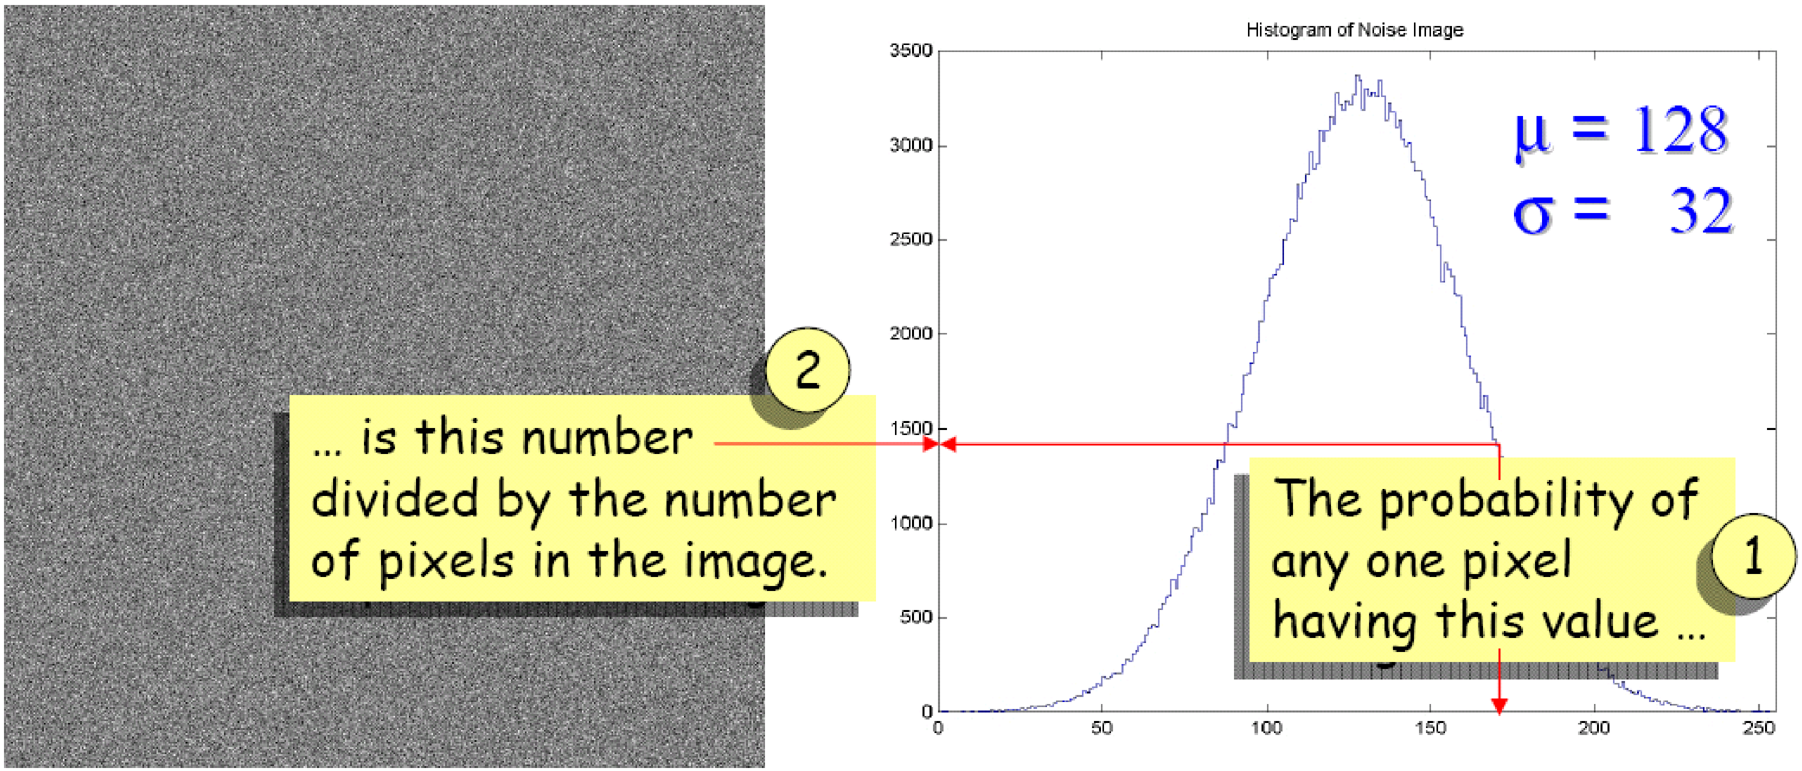

IID: Independent, Identically Distributed

Power Spectrum & Autocorrelation of IID Noise

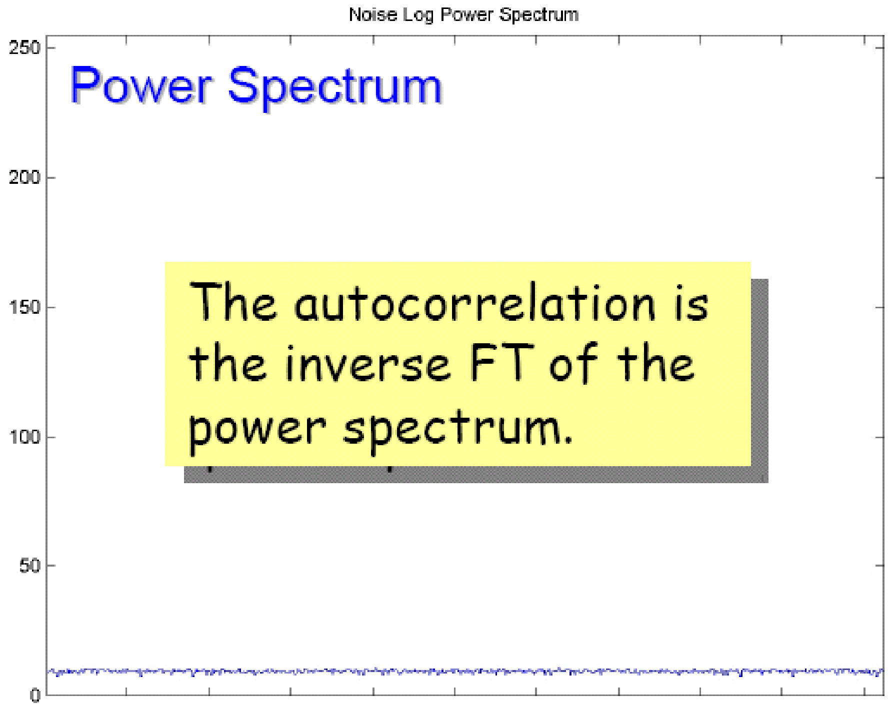

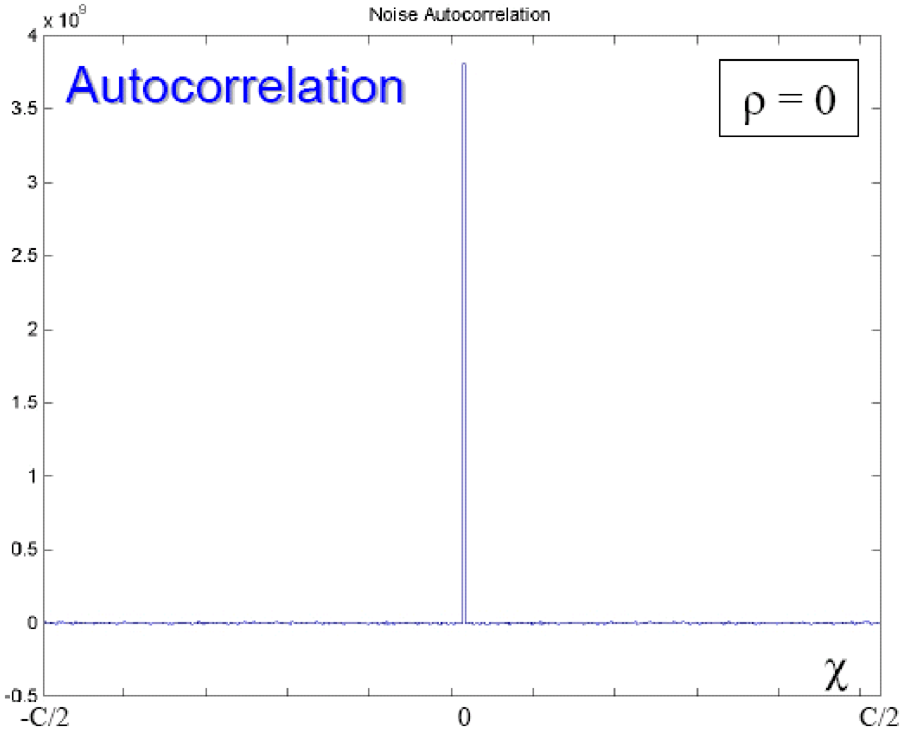

$$\mathrm{PS}(I) = \left| \mathfrak{F}(I) \right|^2$$

$$A_I(\rho, \chi) = \mathrm{Re}\left[\mathfrak{F}^{-1}\left\{ \left| \mathfrak{F}(I) \right|^2 \right\} \right]$$

Power Spectrum & Autocorrelation of IID Noise

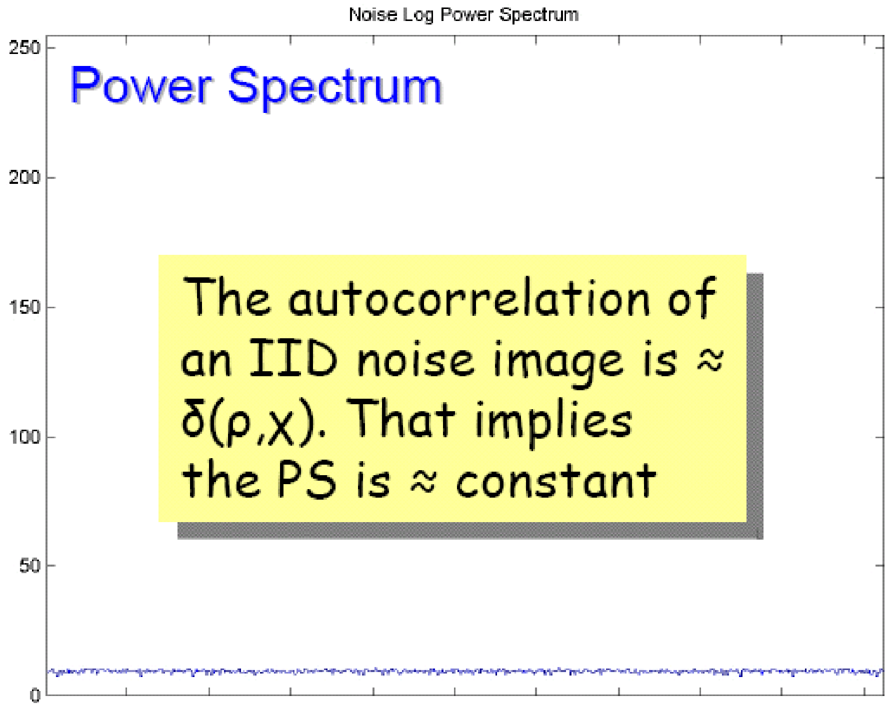

The autocorrelation of an IID noise image is ≈ $\delta(\rho,\chi)$. That implies the PS is ≈ constant

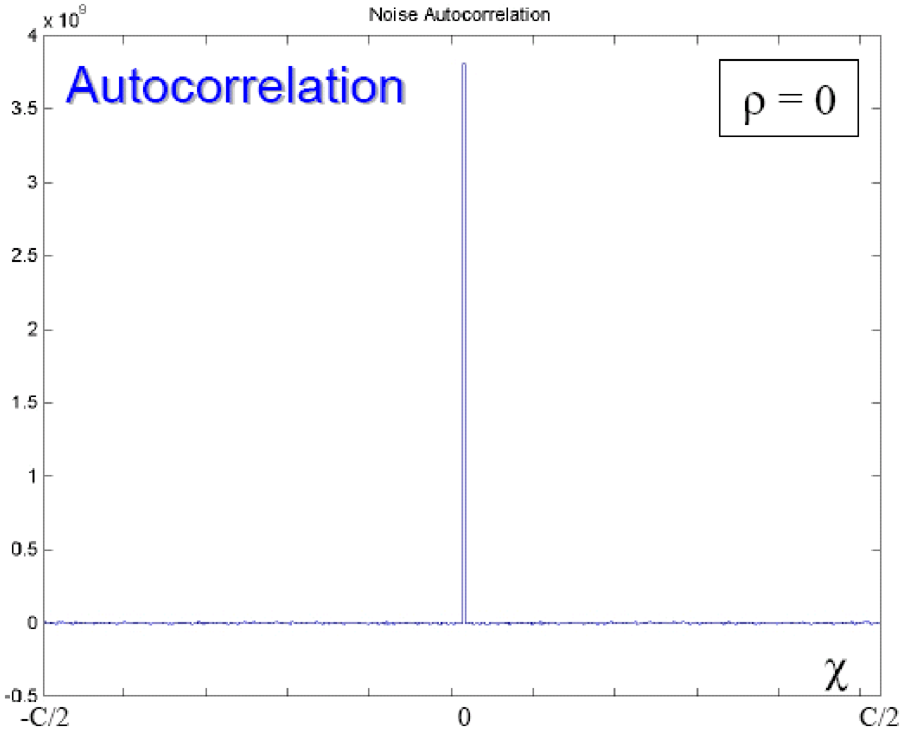

$$\mathrm{PS}(I) = \left|\mathfrak{F}(I)\right|^2$$

$$A_I(\rho,\chi) = \mathrm{Re}\left[\mathfrak{F}^{-1}\left\{\left|\mathfrak{F}(I)\right|^2\right\}\right]$$

Noise-Free Image and Uncorrelated Noise Field

image

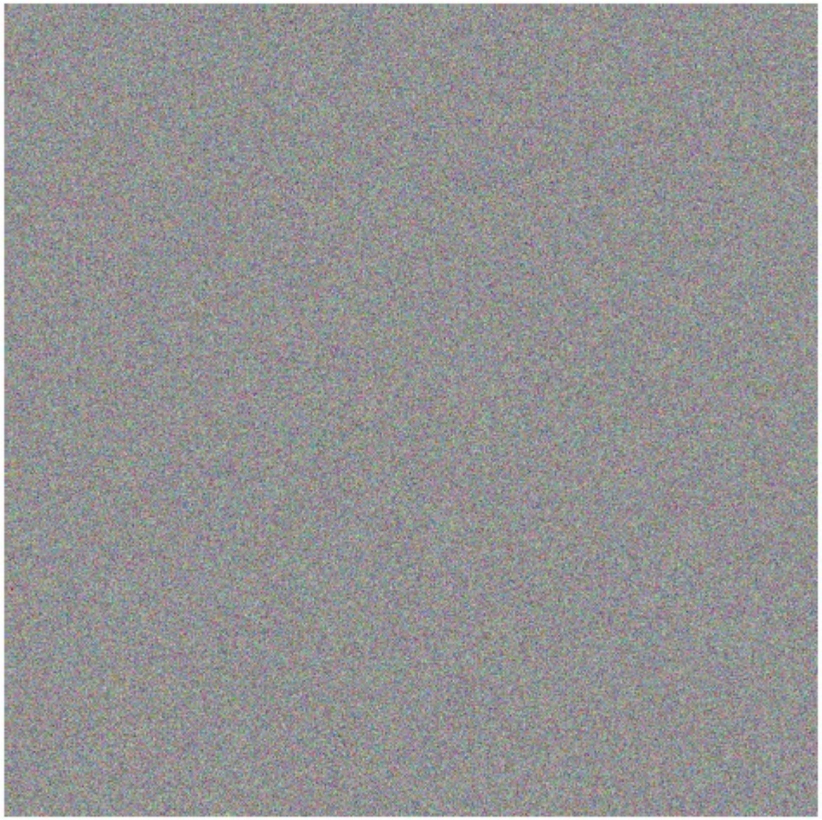

Gaussian noise field

Spectra of Noise-Free Image and Uncorr. Noise Field

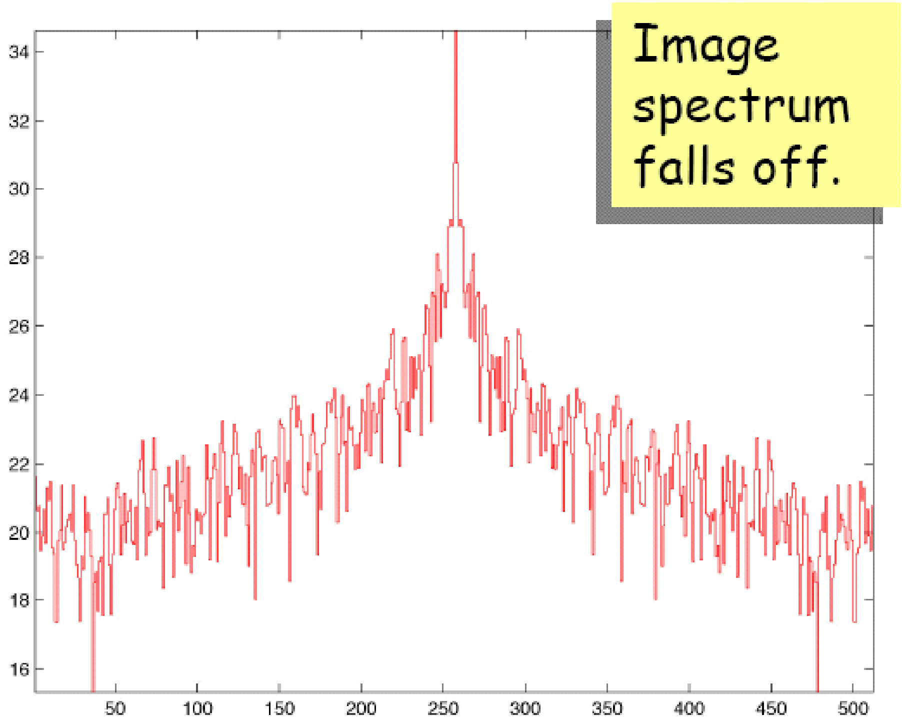

image center row log power spectrum

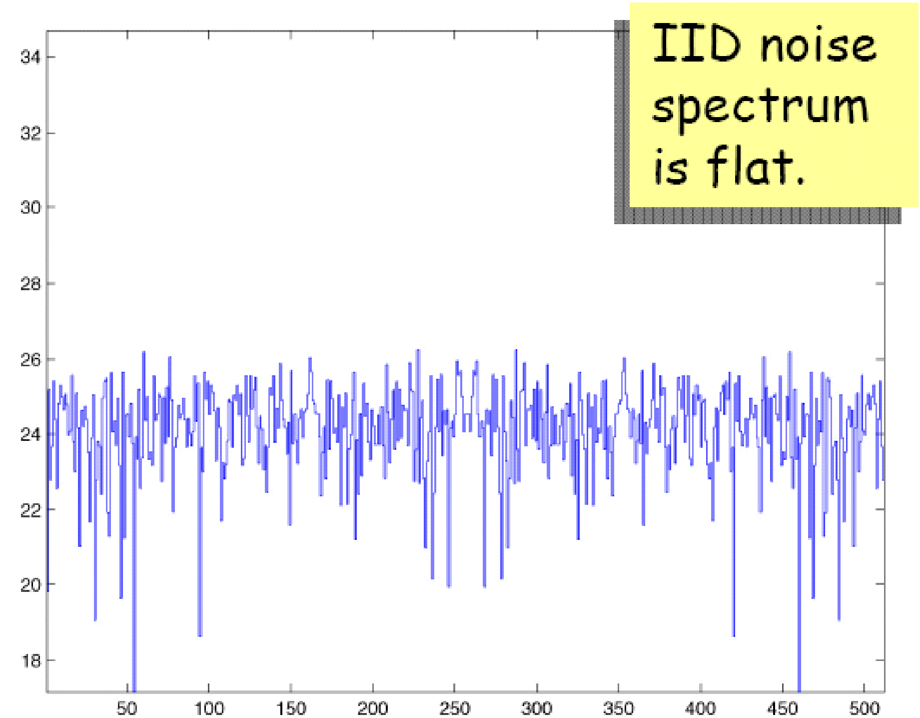

noise field center row log power spectrum

Sum of Noise-Free Image and Uncorrelated Noise Field

image + noise field

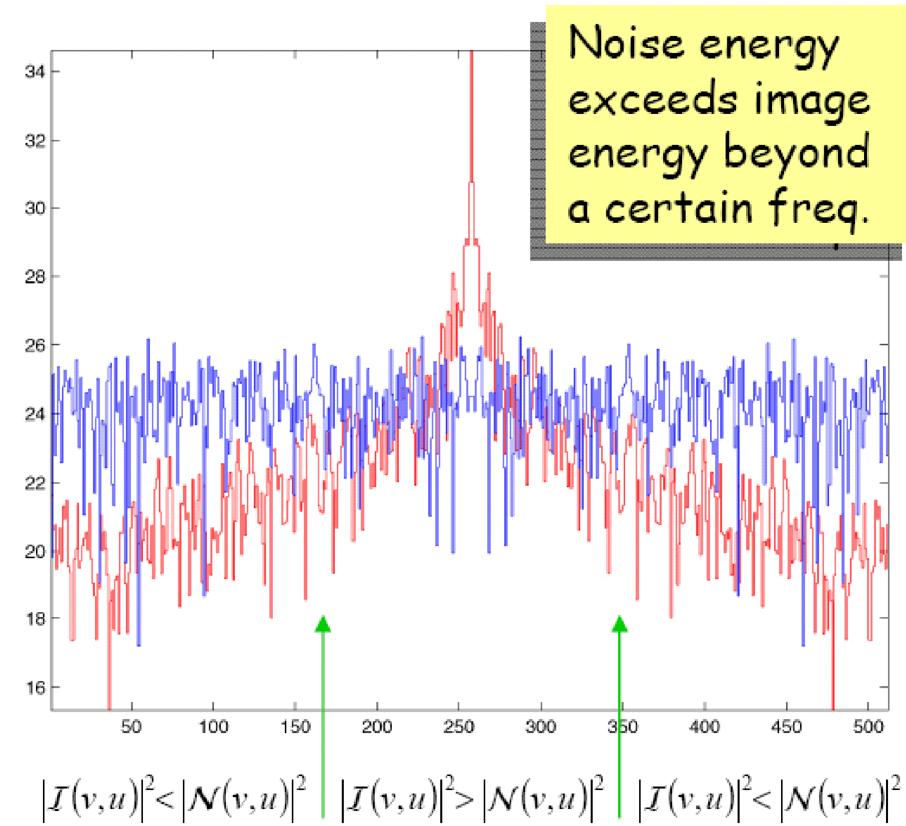

image + noise field center row log PS

Power Spectra of Noise-Free Image and Noise Field

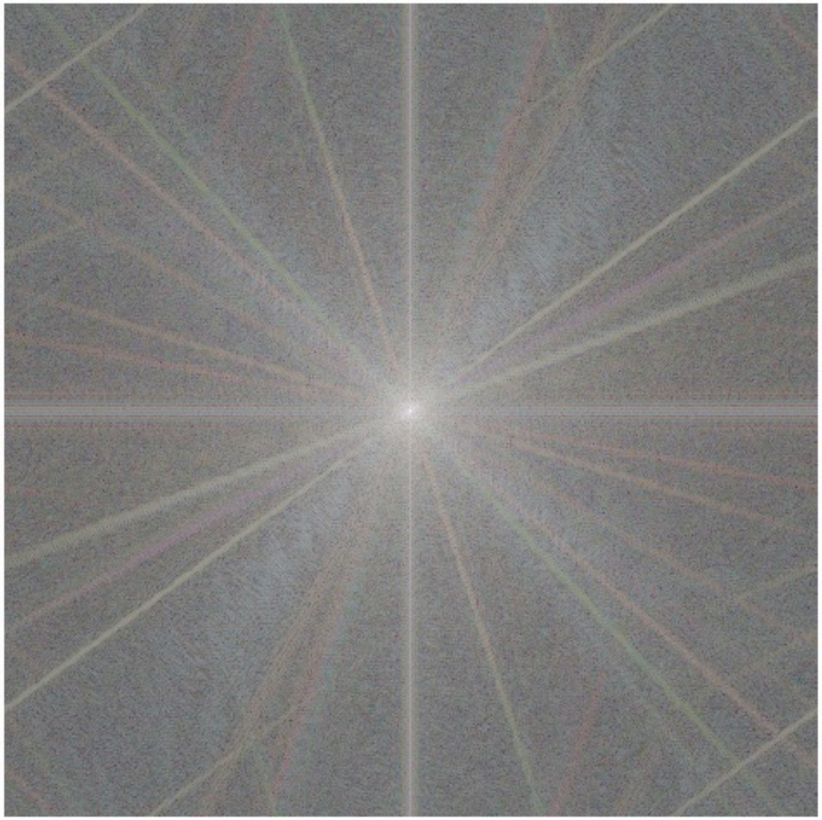

original image

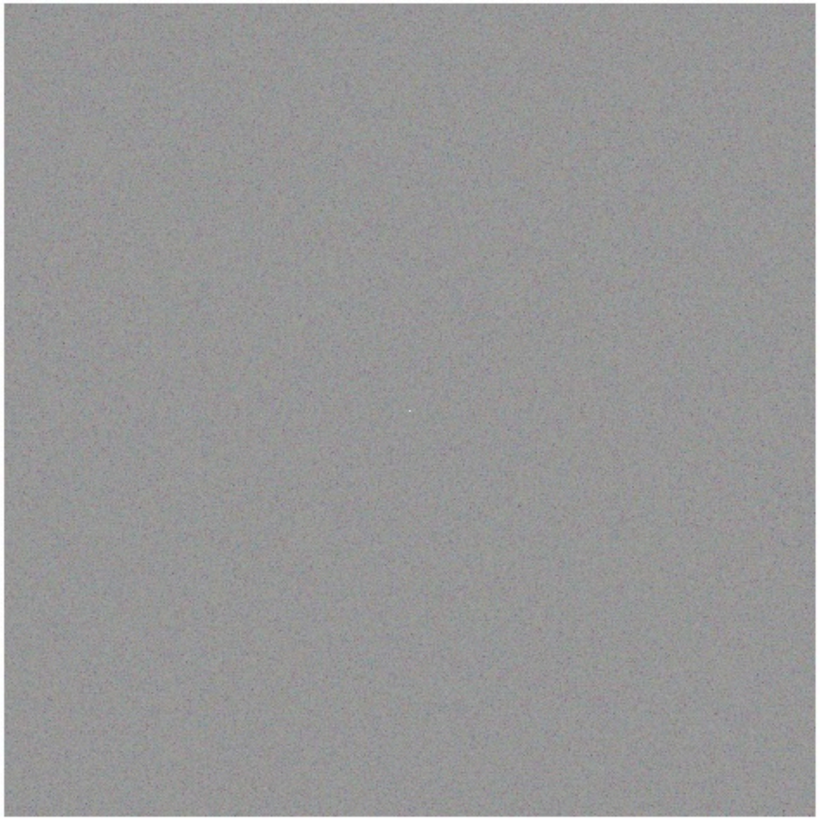

noise image

Power Spectra of Sum of Image and Noise Field

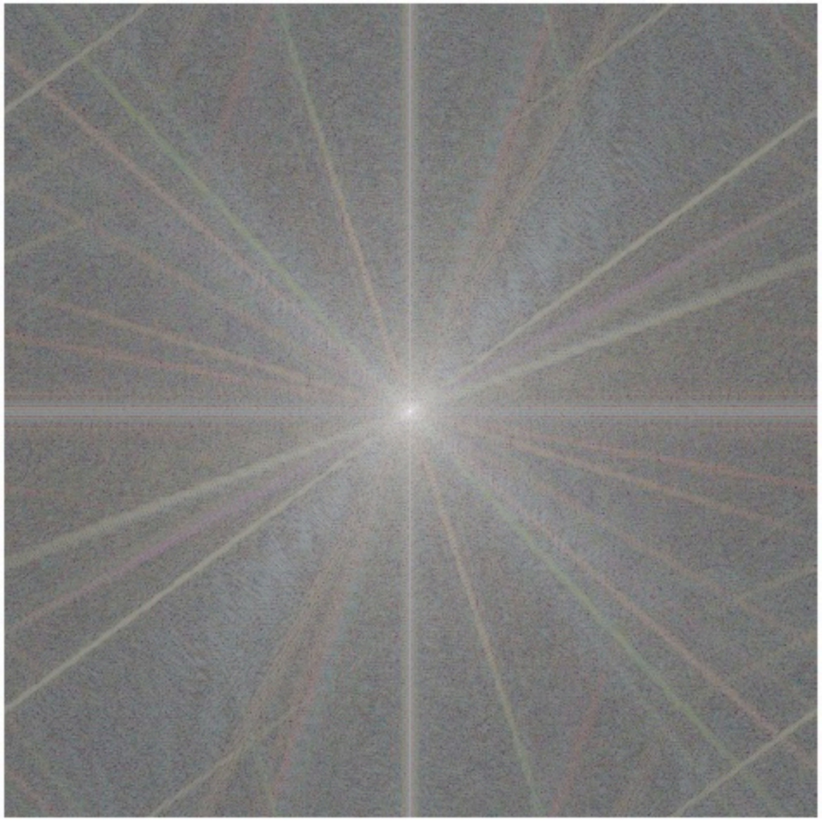

original image

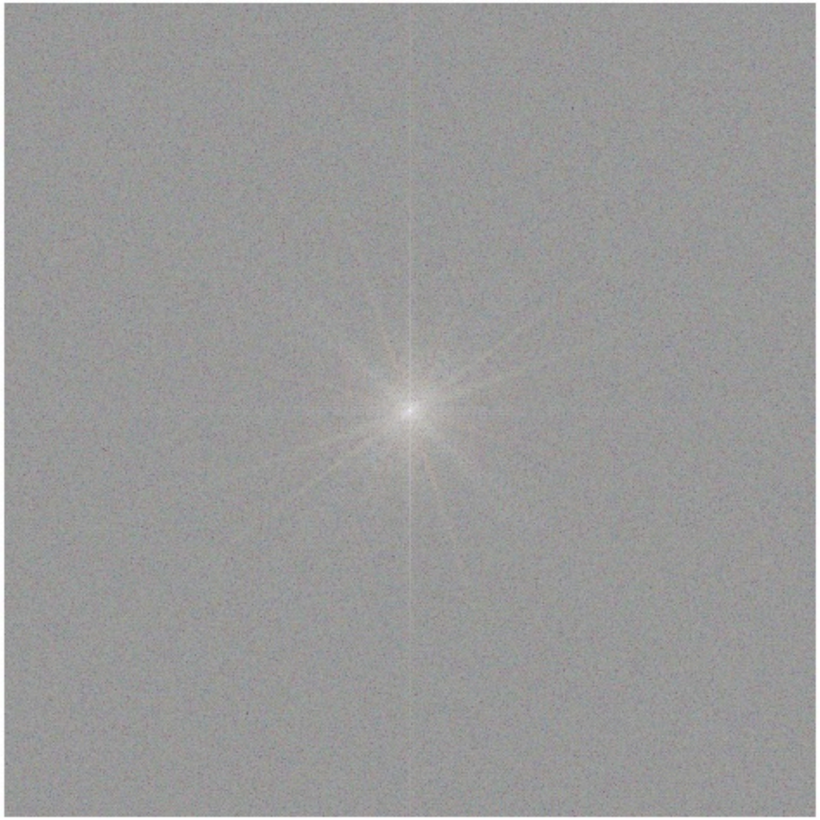

noisy image

Power Spectra of Sum of Image and Noise Field

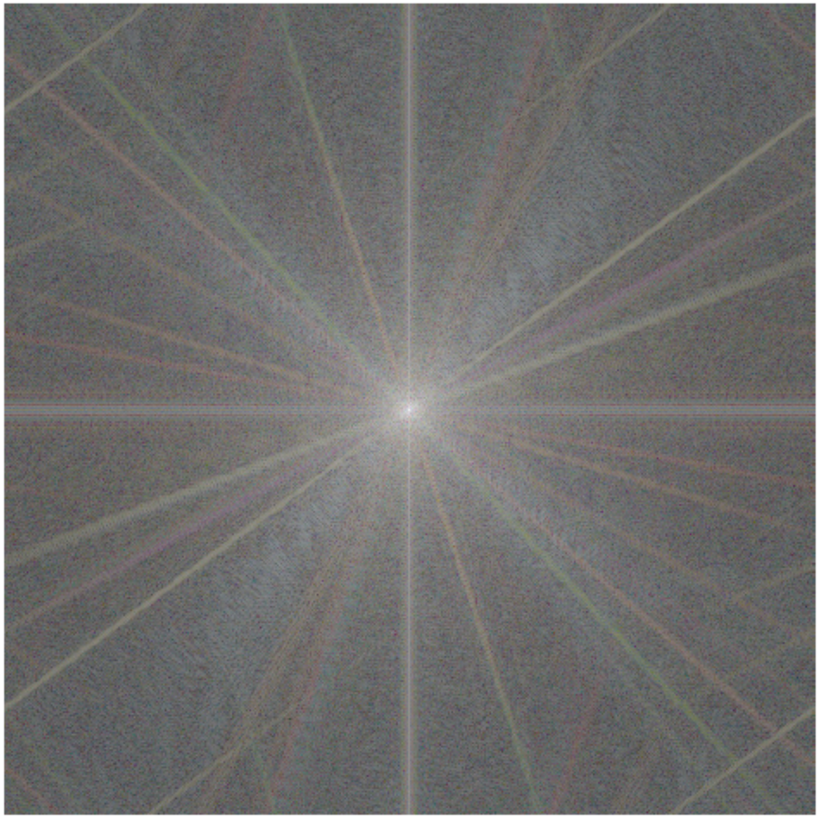

original image

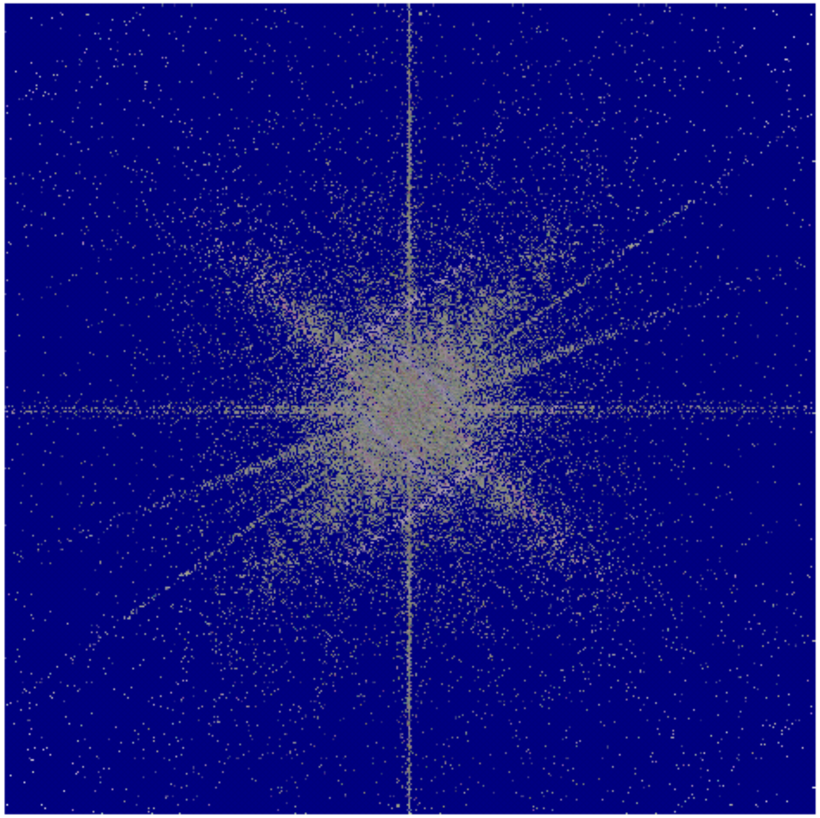

blue indicates noise > image

Power Spectra of Sum of Image and Noise Field

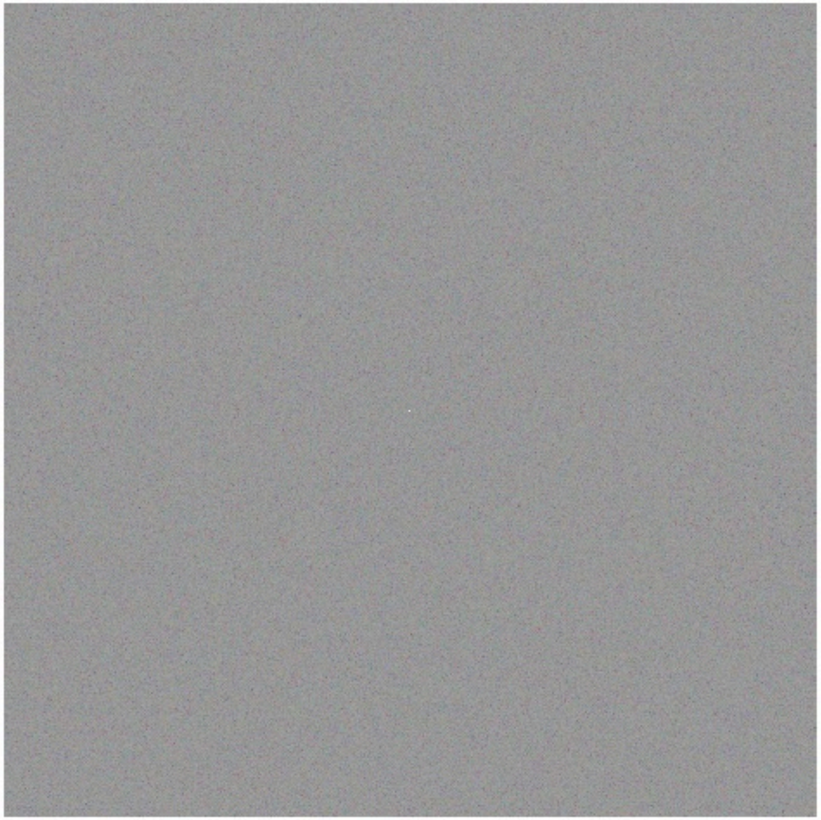

noise image

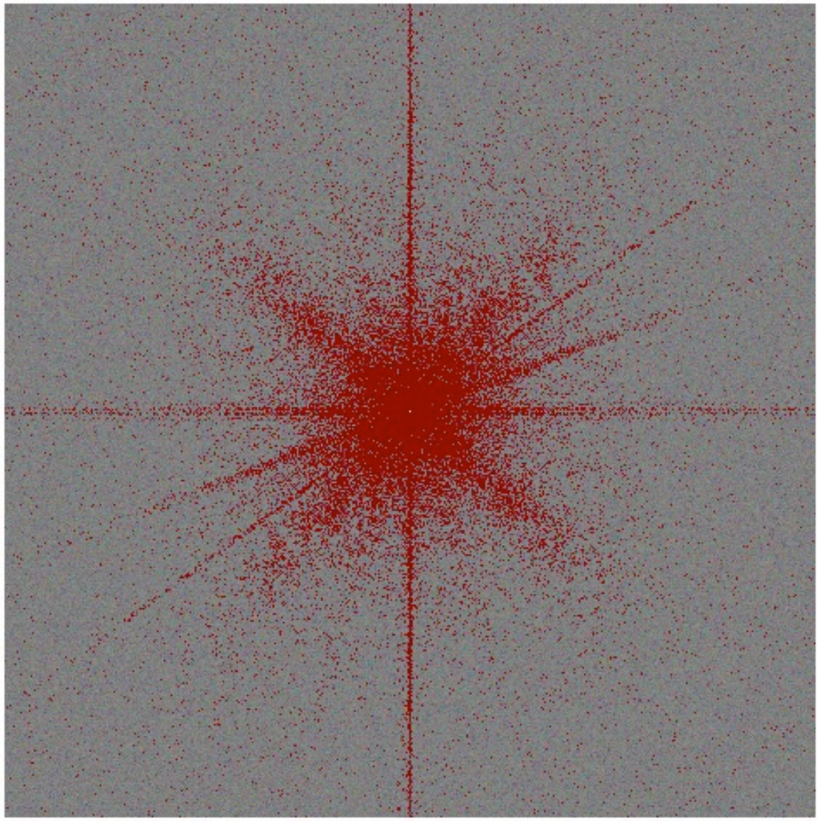

red indicates image > noise

Power Spectra of Sum of Image and Noise Field

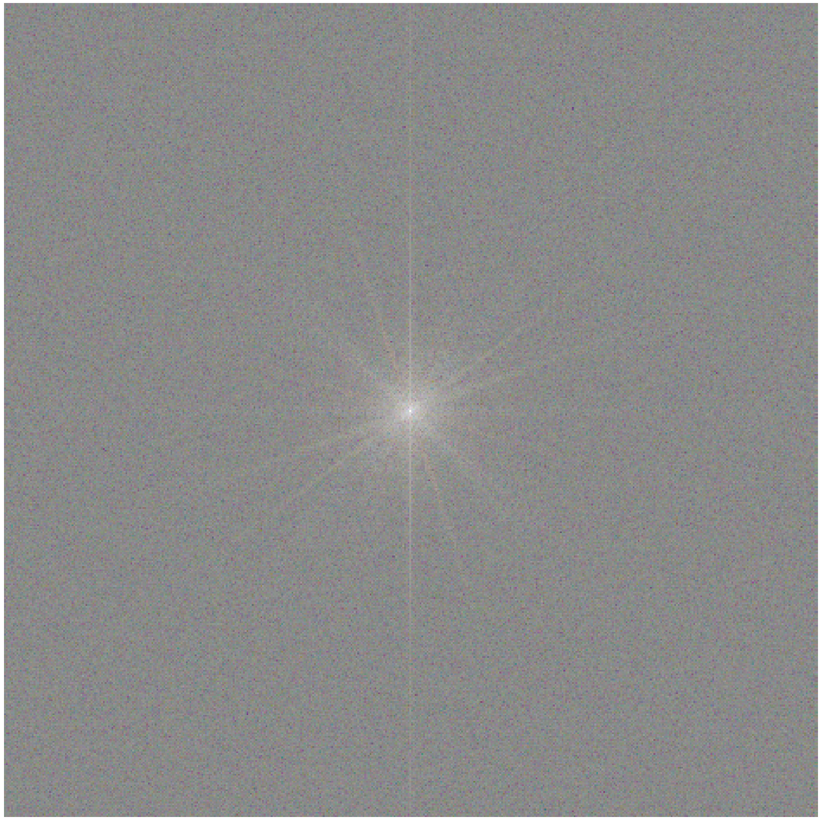

noisy image

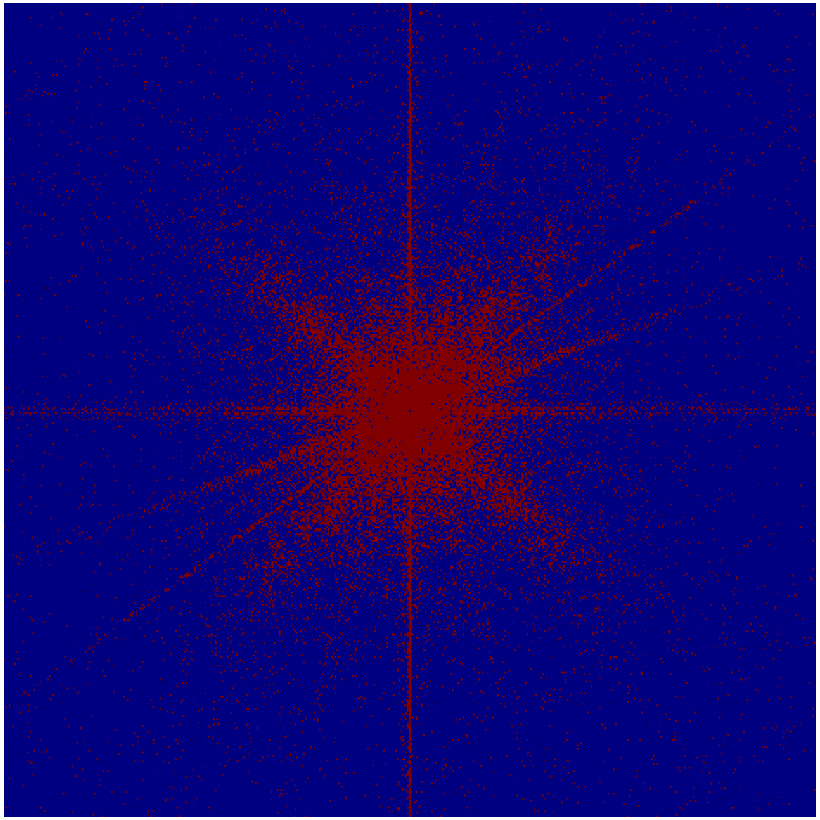

image & noise

Additive Noise: Another Example

original image

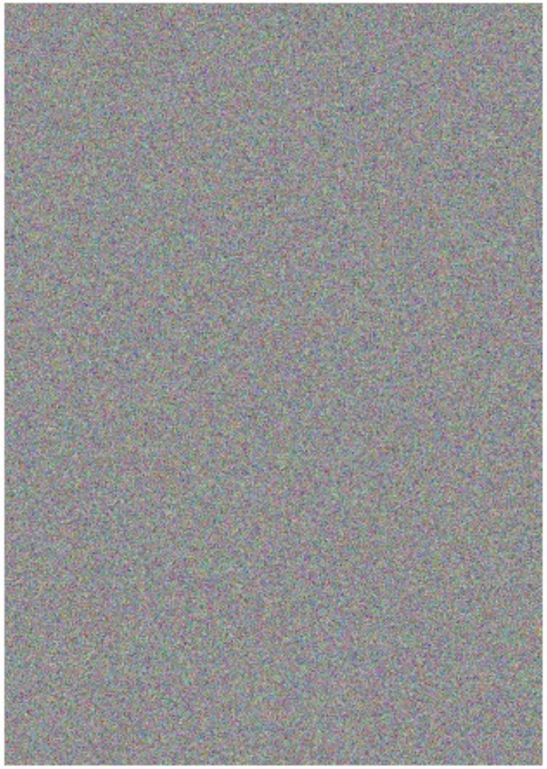

noise image

image+noise

Additive Noise: Another Example

displayed:
$\log\{|\mathfrak{F}(I)|^2+1\}$

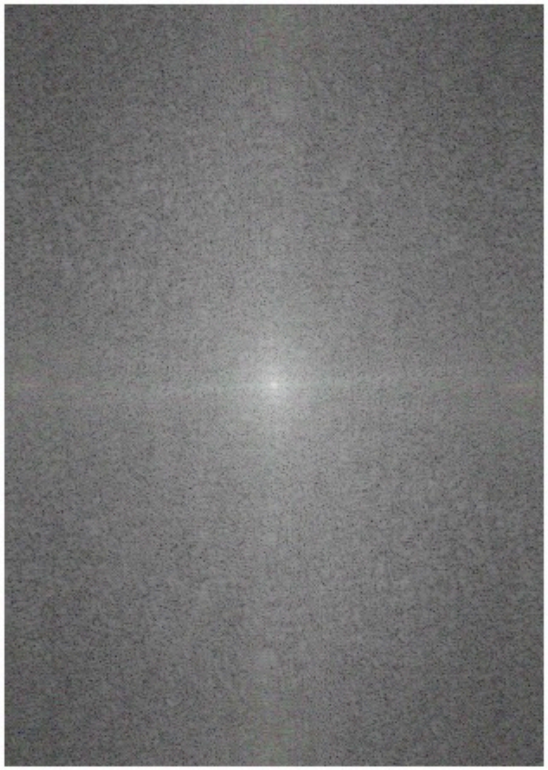

image PS

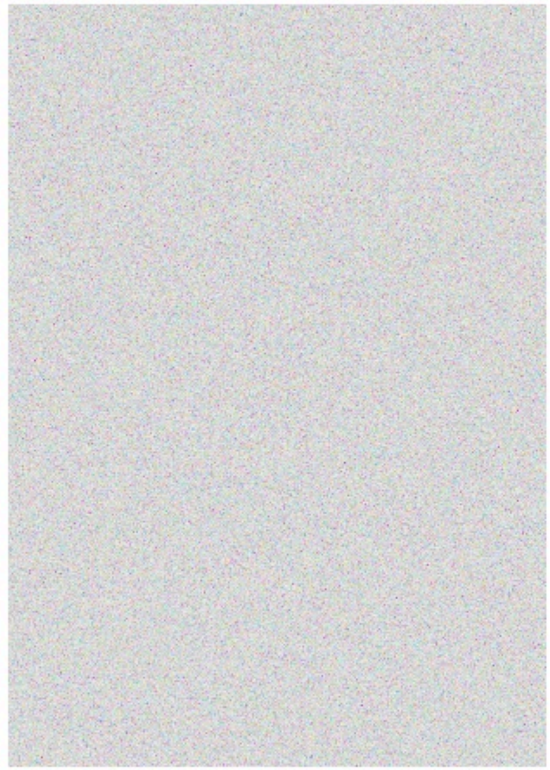

noise PS

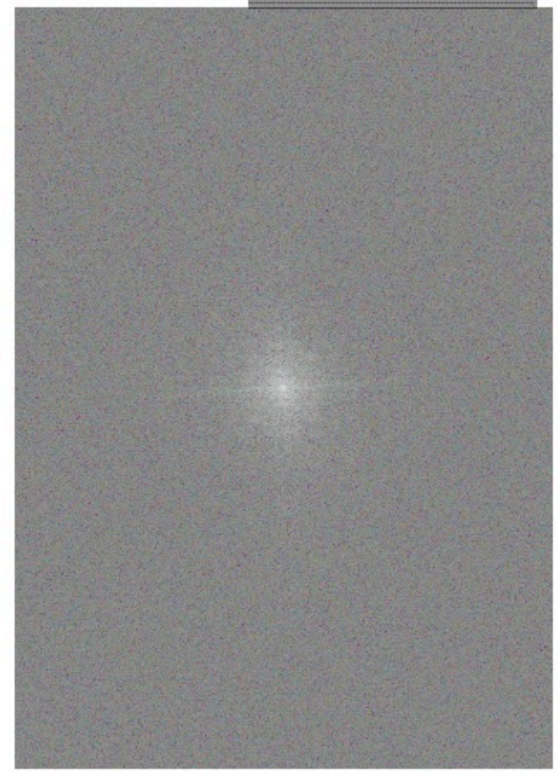

image+noise PS

Additive Noise: Another Example

displayed:
$\log\{|\mathfrak{F}(I)|^2+1\}$

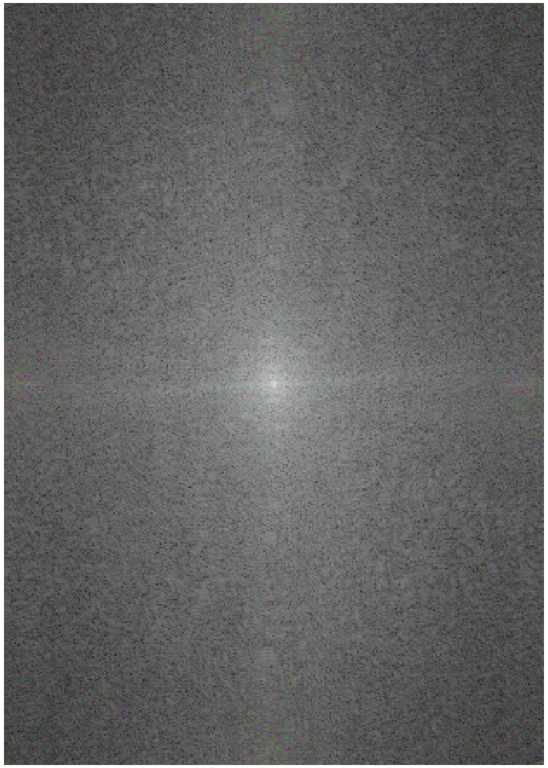

image PS

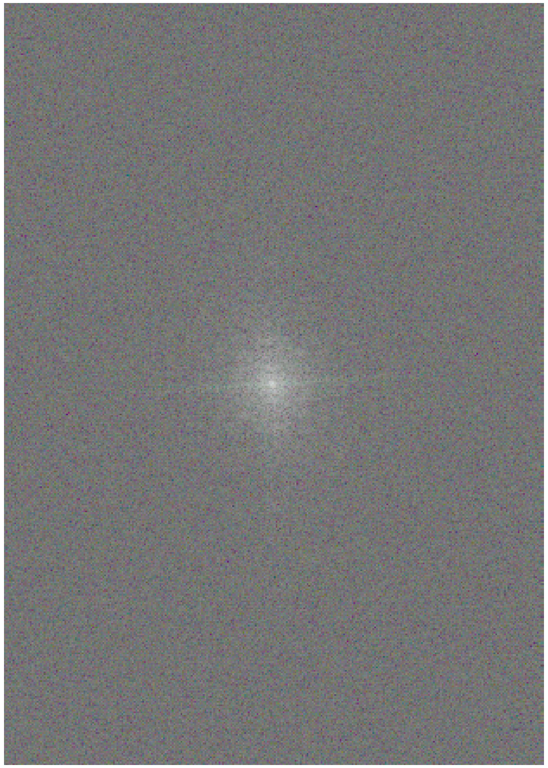

image+noise PS

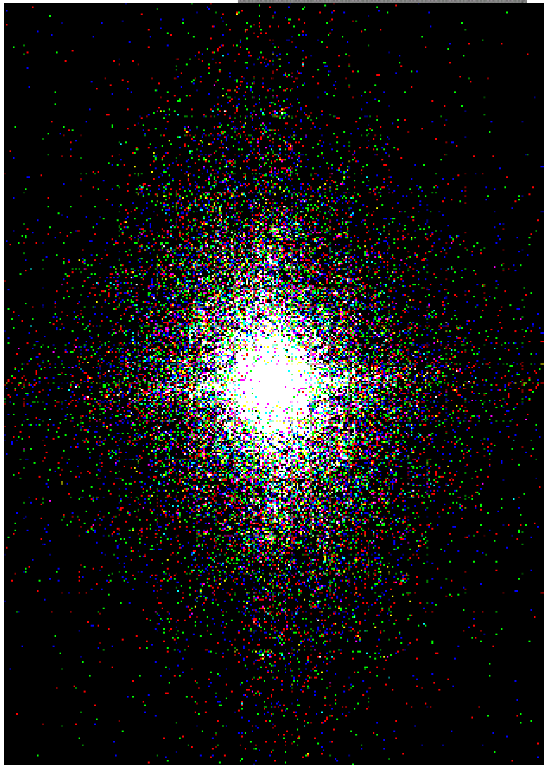

image PS > noise PS

Additive Noise: Reduce Through Blurring?

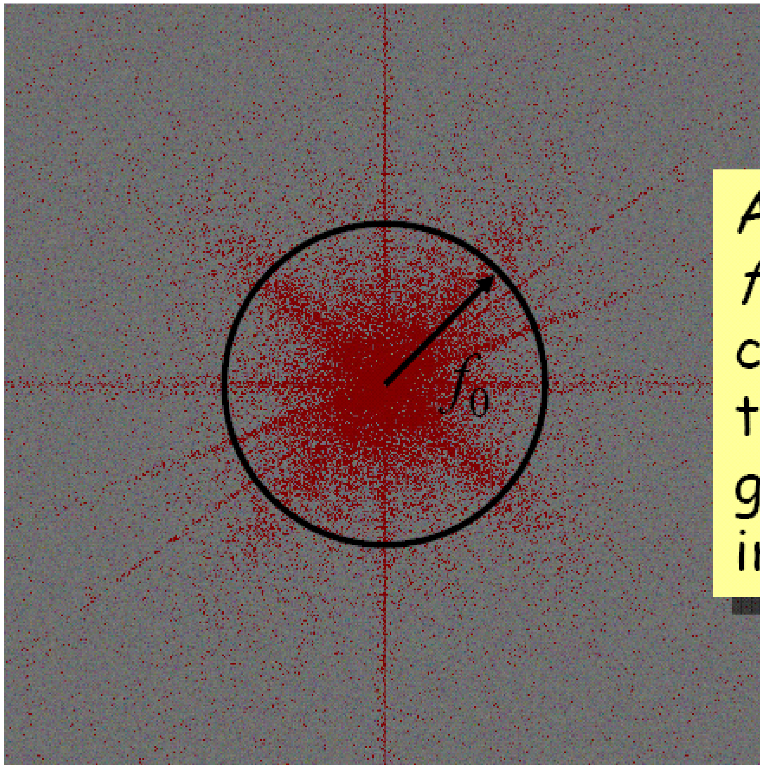

At some frequency, f_0, there are more components where the noise power is greater than the image power.

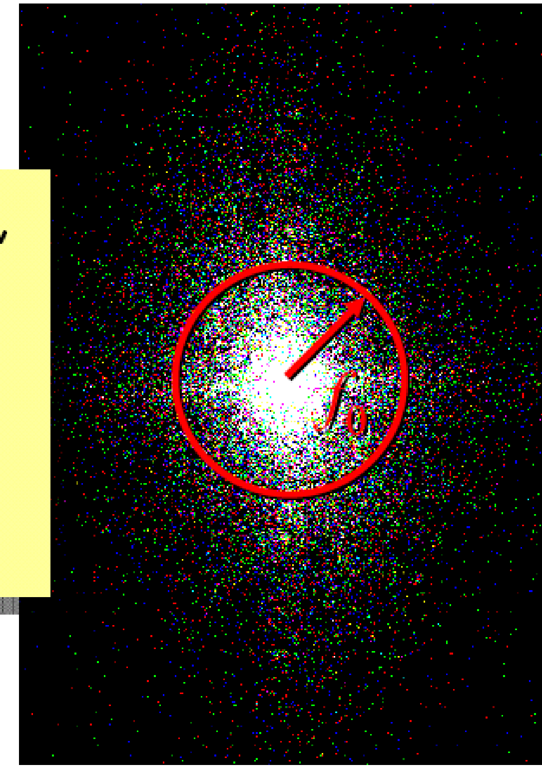

red indicates image > noise

image PS > noise PS

Additive Noise: Reduce Through Blurring?

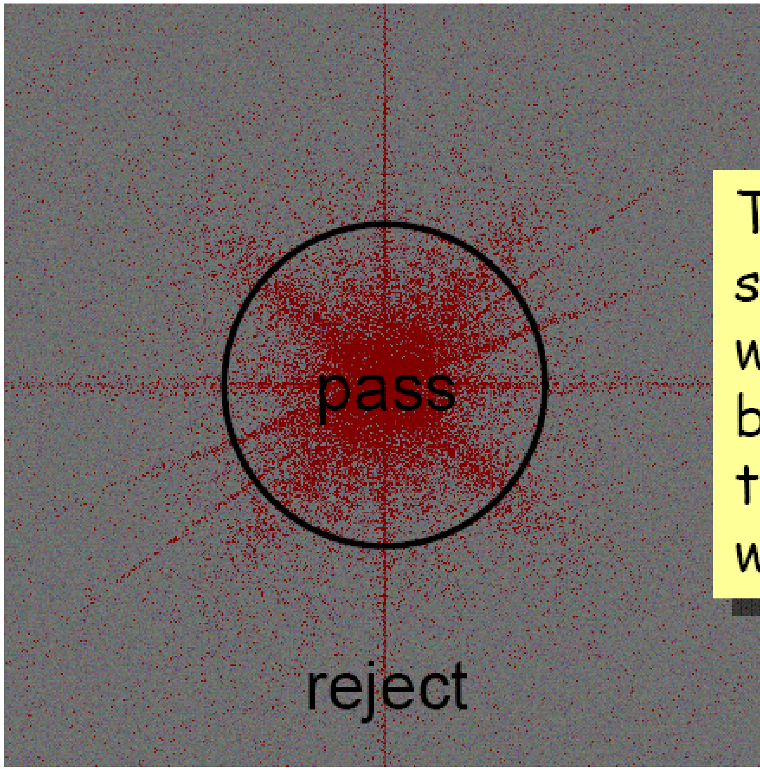

Thus, it makes sense to apply a LPF with cutoff f_0, (a blurring filter) to the images and see what happens.

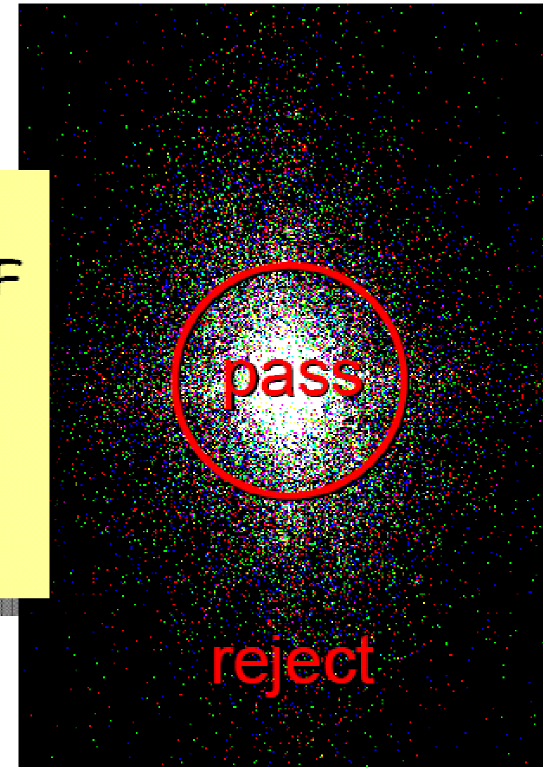

red indicates image > noise

image PS > noise PS

Additive Noise: Reduction Through Blurring.

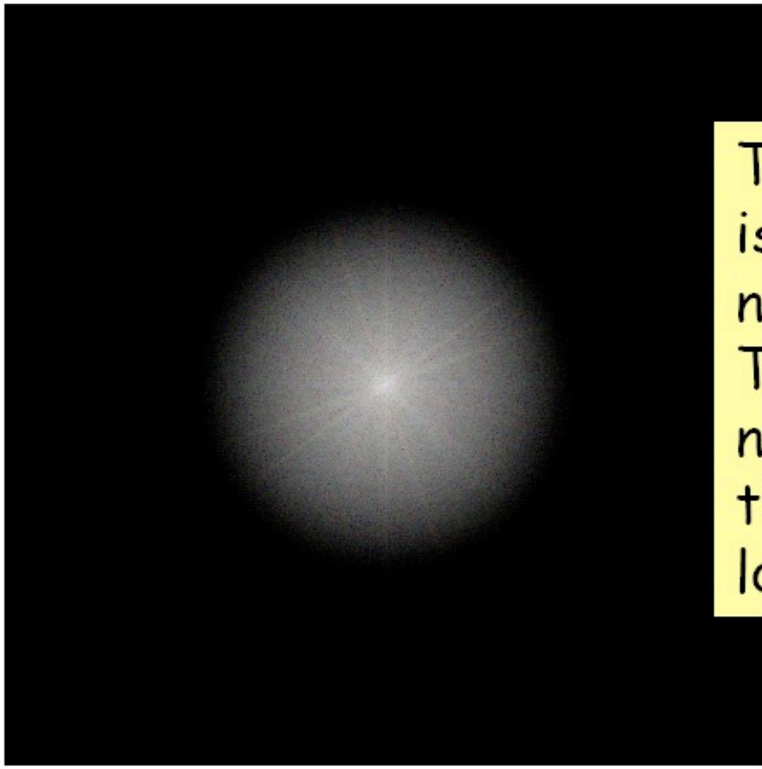

The result is actually no better. There's less noise but the blurring looks worse.

PS of Gaussian blurred image

Gaussian Blurred Image

Additive Noise: Reduction Through Blurring.

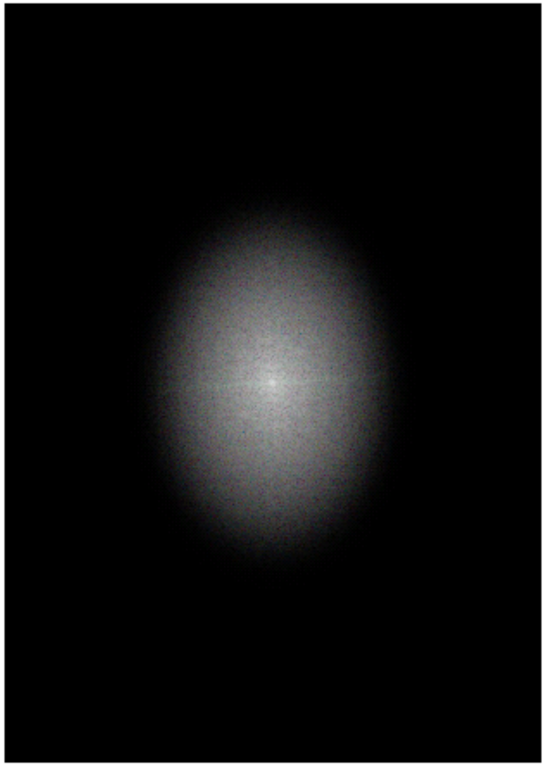

PS of Gaussian blurred image

The result is actually no better. There's less noise but the blurring looks worse.

Gaussian Blurred Image

Noise Masking

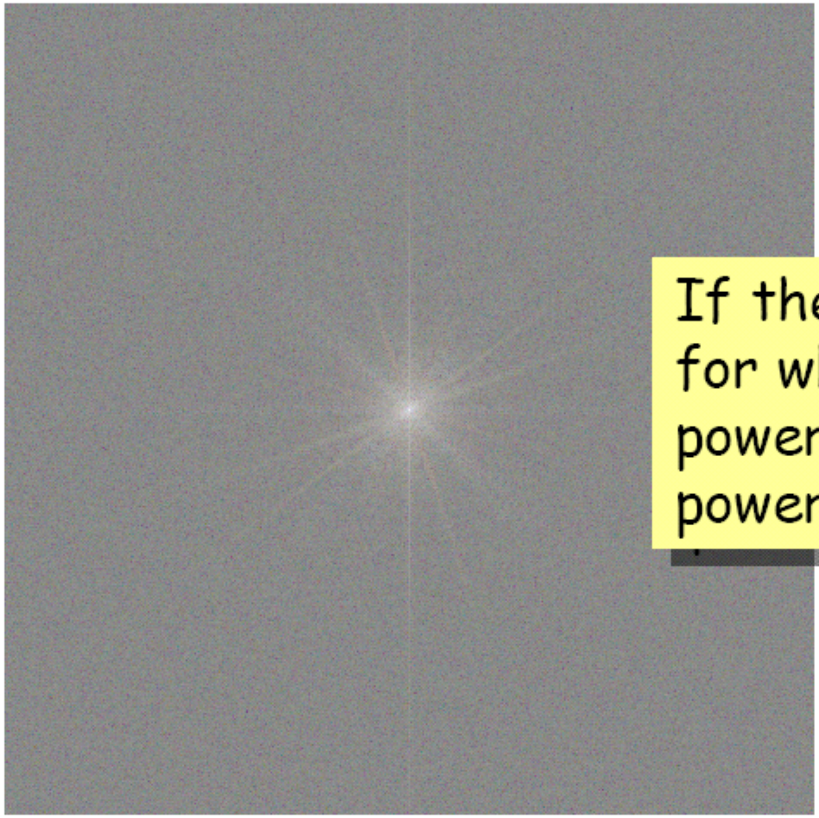

power spec. of noisy image

If the freq. comps. for which noise-power > image-power are known[1]...

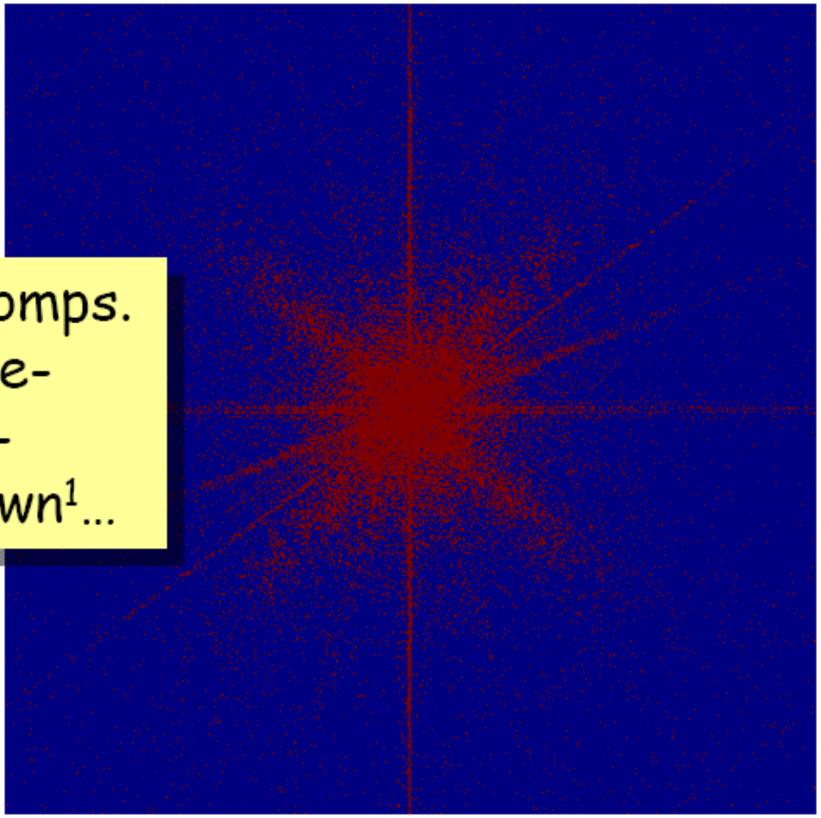

red: image > noise

blue: image < noise

[1]of course they almost never are.

Noise Masking

... mask those out of the FT and invert the result to get...

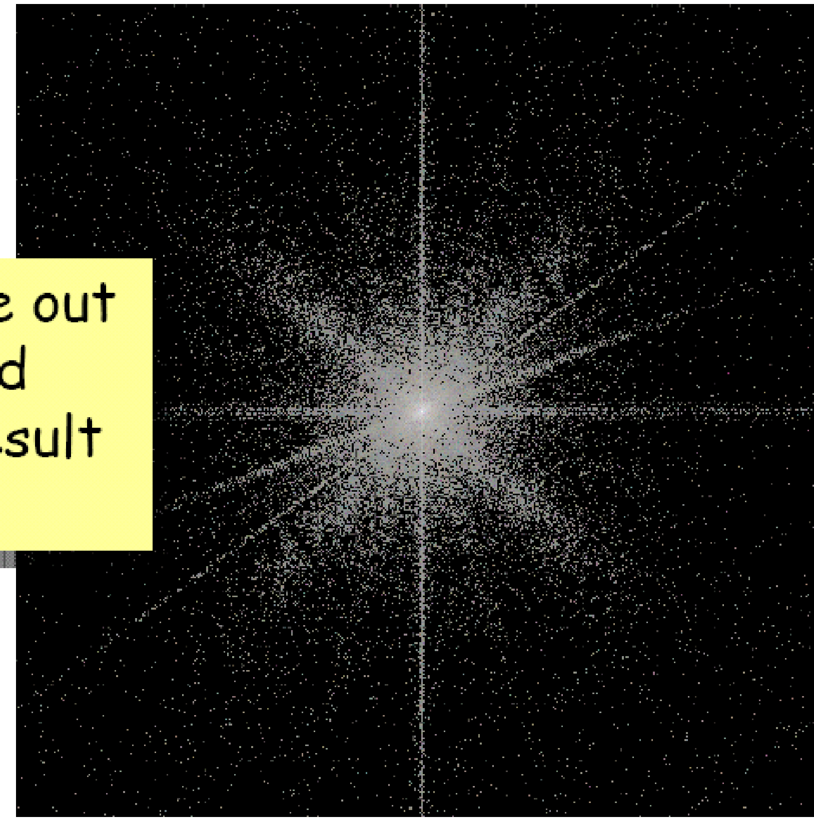

power spec. of noisy image

image < noise masked out

Noise Masking

... this:

noisy image

noise-masked mage

Noise Masking

Although the noise-masked image looks better than the blurred one, it is still noisy. Moreover, this example is unrealistic because we know the exact noise power spectrum. In any real case we will at most know its statistics.

blurred noisy image

noise-masked mage

Image Degradation Model

So far, we have considered only additive noise. Before going further it will be useful to consider a more general model of image degradation, one that includes convolution with a pointspread[1] function, H, as well as additive noise.

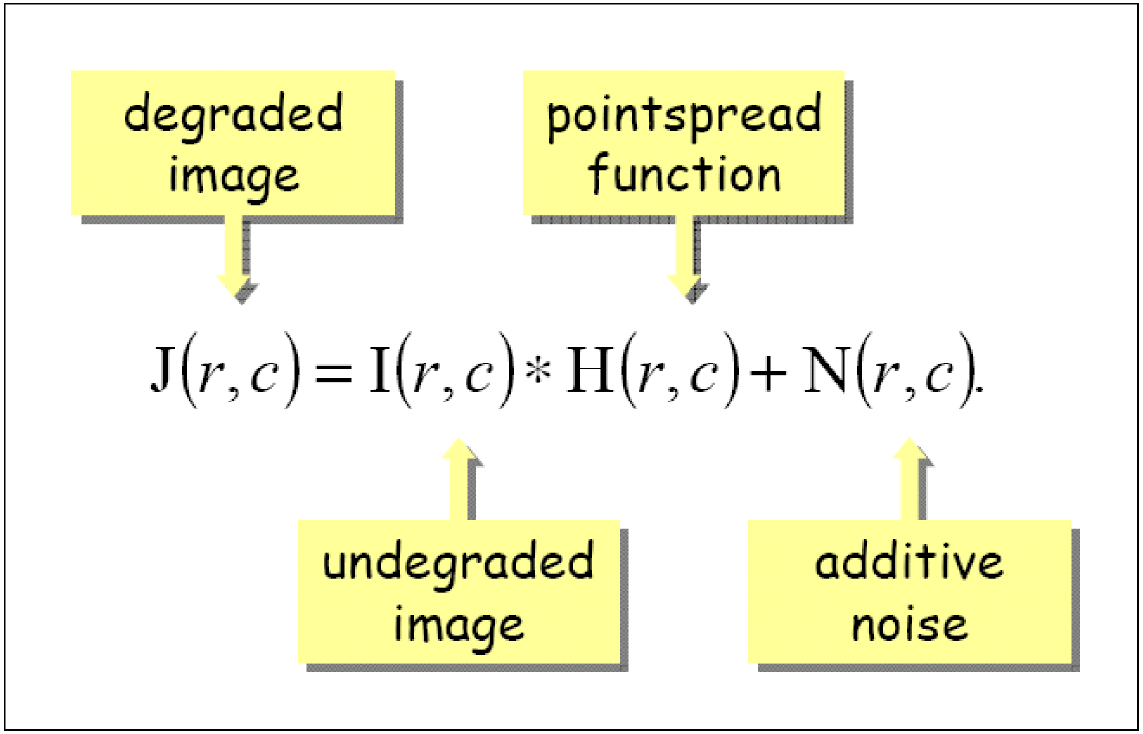

$$J(r,c) = I(r,c) * H(r,c) + N(r,c).$$

---

[1]H is also referred to as the optical transfer function.

Pointspread Operators

A pointspread operator is a linear model of the distortion acquired during the imaging process. Since it is a linear model, it is a convolution operator. One example of this is aperture distortion, an unavoidable consequence of making an image with a camera that has an opening larger than a point.

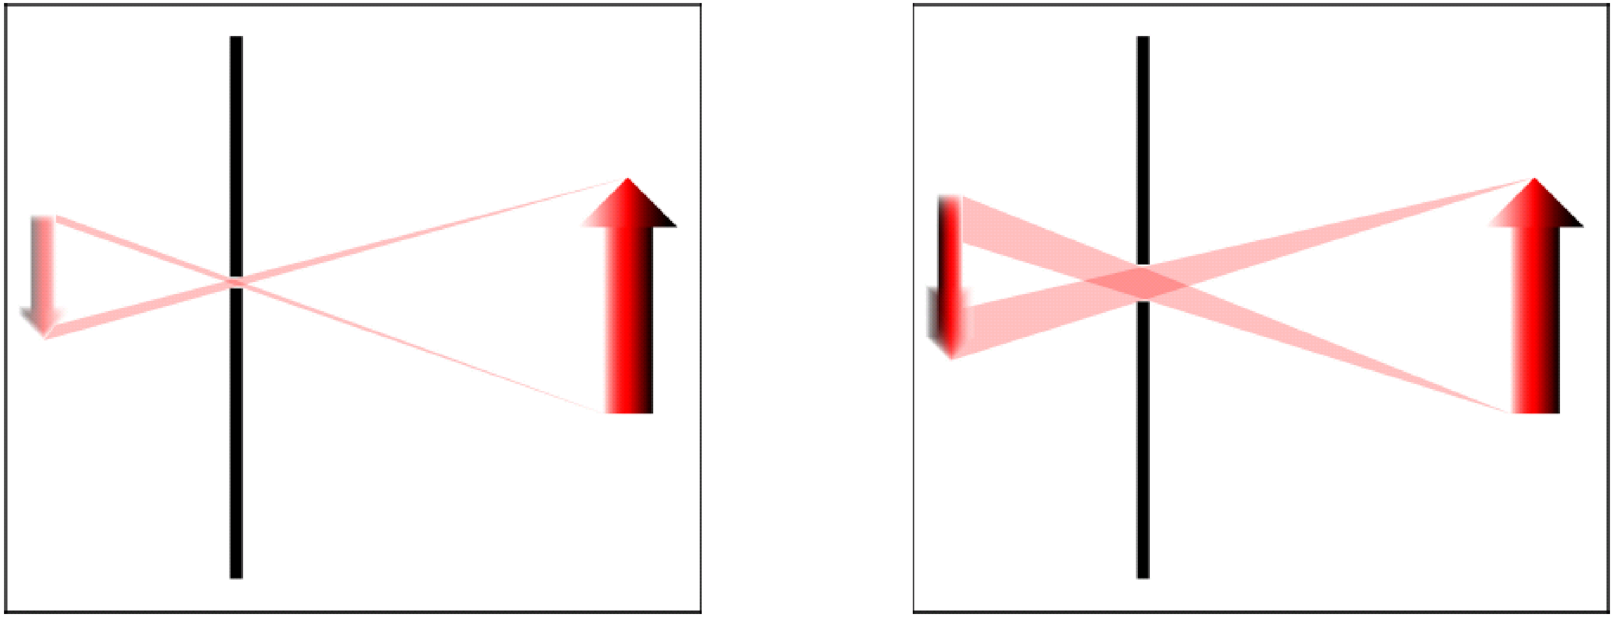

Pointspread Operators

pinhole camera

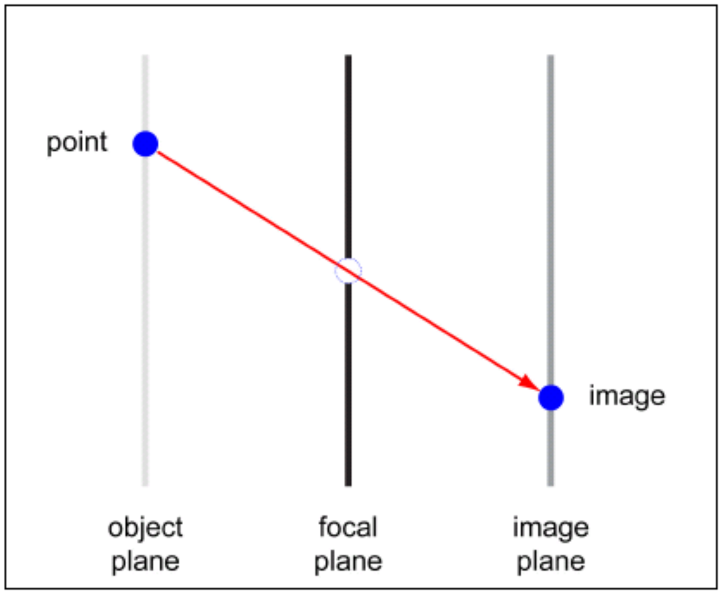

A pinhole camera maps one object point to one image point; it is one-to-one.

aperture camera

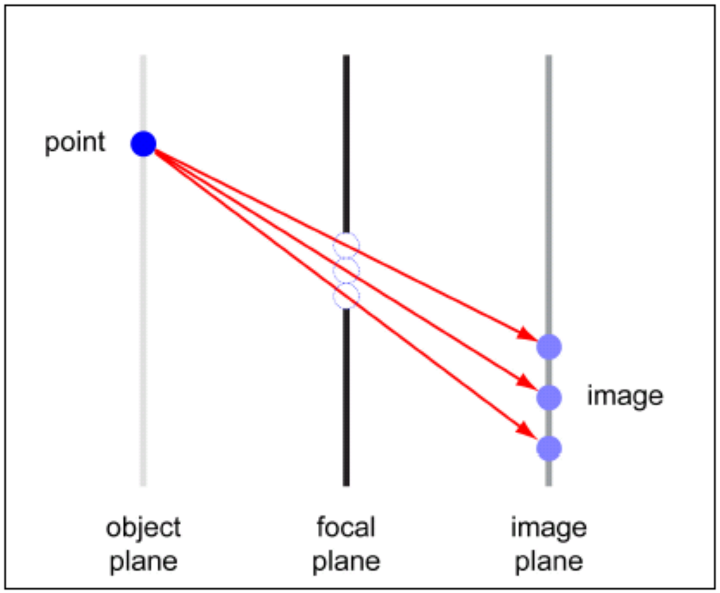

An aperture camera maps one object point to many image points; it spreads the points.

Pointspread Operators and Convolution

$I(r,c)$

$J(r,c) = I(r,c) * H(r,c)$

$H(r,c)$

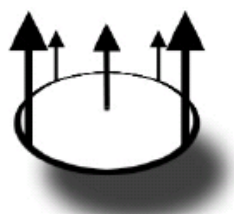

Recall how a convolution works through multiply, shift, and add (See Lect. 7 p. 25ff). That is precisely the effect of imaging through an aperture. It results in a blurry image.

Lenses

A properly designed lens will focus the light emanating from a point and thereby reduce the blurring. But no lens can do this perfectly. In fact, the lens adds its own distortion. The result is an optical transfer function, H(r,c), that is convolved with the image.

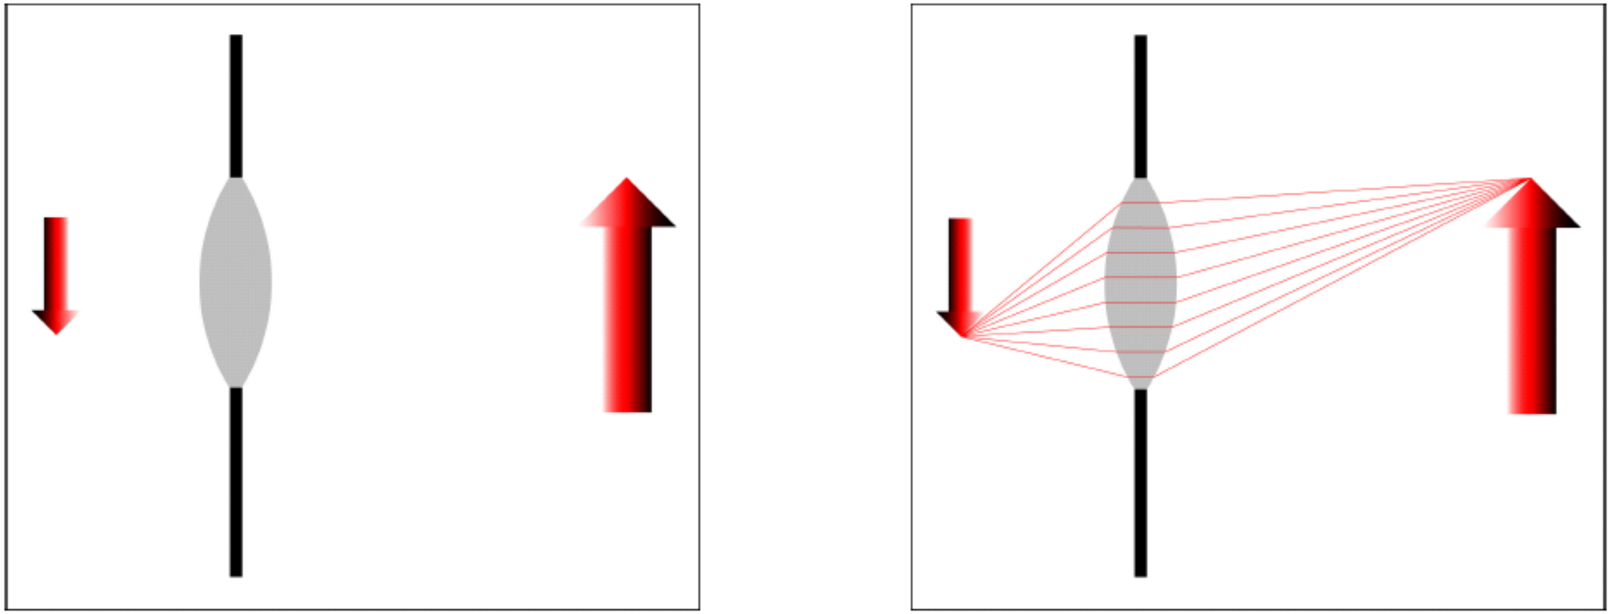

Image Degradation Model

degraded image ↓ pointspread function ↓

$$J(r,c) = I(r,c) * H(r,c) + N(r,c).$$

undegraded image ↑ additive noise ↑

Note: The term *pointspread operator* refers to convolution by the pointspread function.

Image Degradation Model

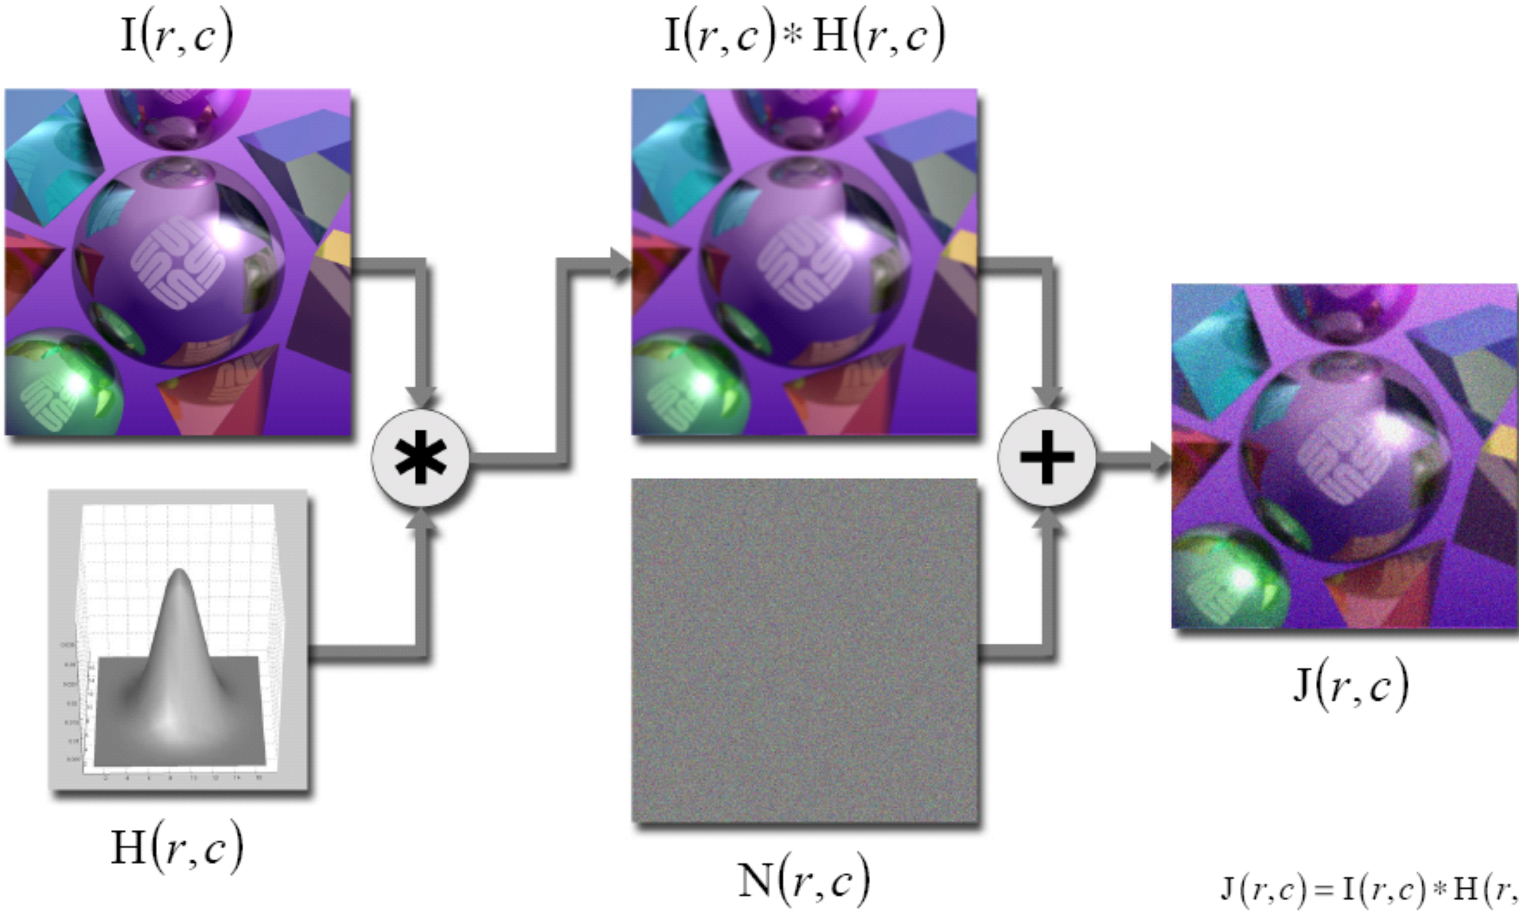

$$J(r,c) = I(r,c) * H(r,c) + N(r,c)$$

Image Degradation Model (Frequency Domain)

degraded image → $\mathcal{J}(v,u)$

pointspread operator → $\mathcal{H}(v,u)$

$$\mathcal{J}(v,u) = \mathcal{I}(v,u) \cdot \mathcal{H}(v,u) + \mathcal{N}(v,u).$$

undegraded image → $\mathcal{I}(v,u)$

additive noise → $\mathcal{N}(v,u)$

Image Degradation Model (Frequency Domain)

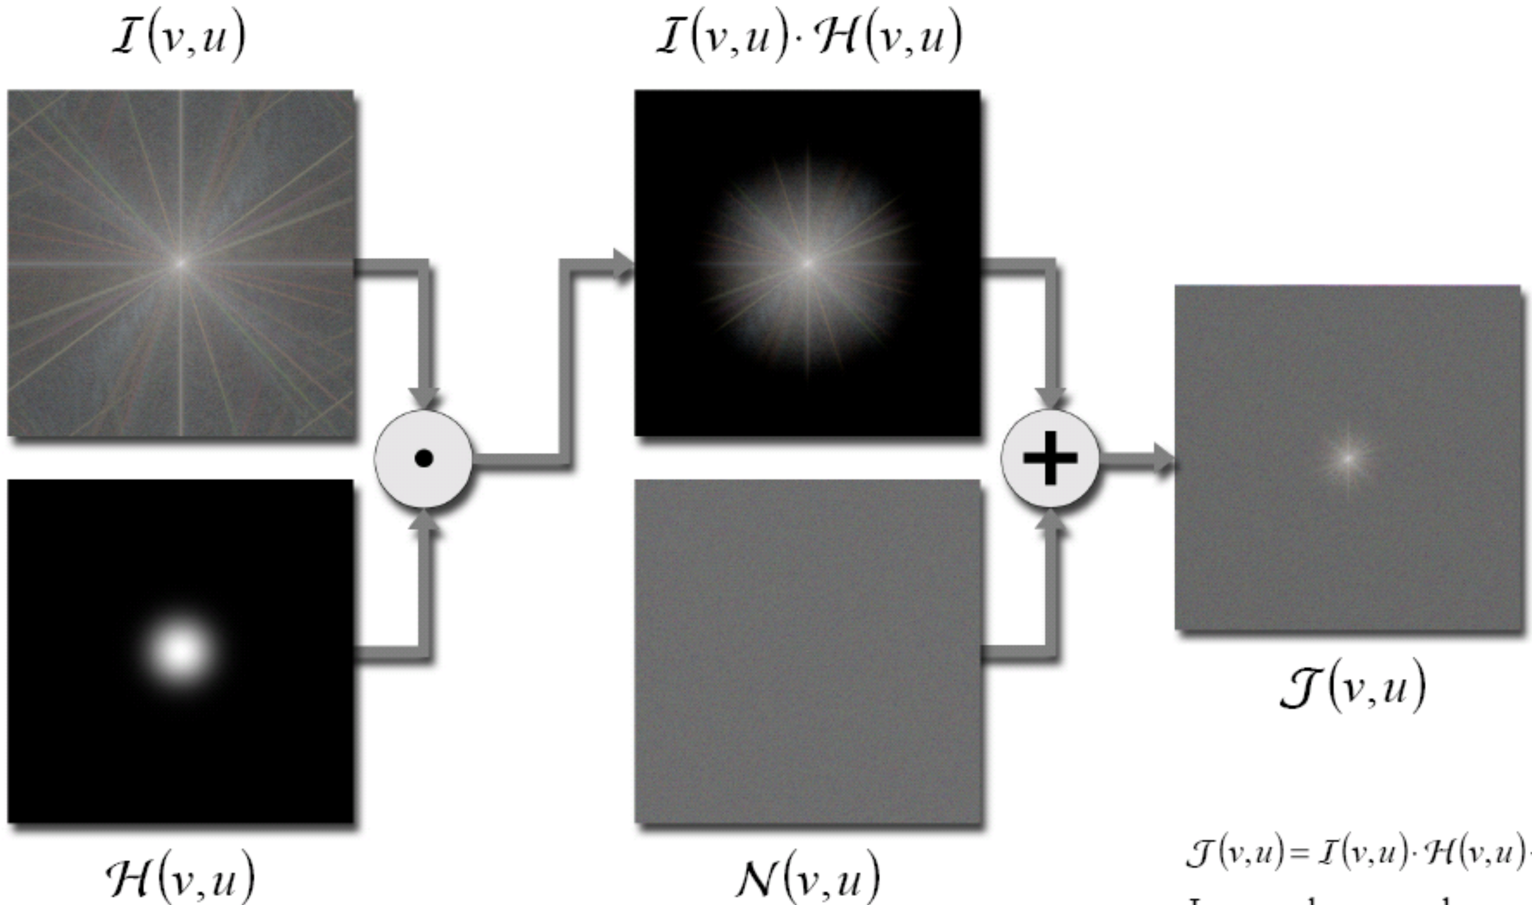

$\mathcal{J}(v,u)=\mathcal{I}(v,u)\cdot\mathcal{H}(v,u)+\mathcal{N}(v,u).$

Images shown are log magnitude.

Image Restoration

Let I be a perfect image and let K be the image convolved with a pointspread function, H. Then in the frequency domain:

$$\mathcal{K}(u,v) = \mathcal{I}(u,v)\mathcal{H}(u,v).$$

If the process of imaging adds noise then we get $J = K + N$, or in freq.:

$$\mathcal{J}(u,v) = \mathcal{K}(u,v) + \mathcal{N}(u,v).$$

We want a filter, W, to remove as much of the noise from J as possible:

$$\widetilde{\mathcal{K}}(v,u) = \mathcal{W}(v,u)\mathcal{J}(v,u).$$

Then an estimate of I would be the inverse Fourier transform of

$$\widetilde{\mathcal{I}}(u,v) = \frac{\widetilde{\mathcal{K}}(u,v)}{\mathcal{H}(u,v)} = \frac{\mathcal{W}(u,v)\mathcal{J}(u,v)}{\mathcal{H}(u,v)}.$$

We want to find the filter, W, that results in the closest possible estimate of I – the W that minimizes the energy of the difference between the estimate and I. That is we want to find W such that

$$\varepsilon^2 = \iint \left| \mathcal{I} - \widetilde{\mathcal{I}} \right|^2 du dv$$

is as small as possible. This is called least mean squared (LMS) minimization.

Image Restoration

There are a number of ways to solve for the minimum squared error. All make use of the assumption that the image and the noise are uncorrelated. Depending on how that fact is used, slightly different solutions are found. The most common one used in image processing has

$$\mathcal{W} = \frac{\mathcal{H}^* |\mathcal{I}|^2}{|\mathcal{H}|^2 |\mathcal{I}|^2 + |\mathcal{N}|^2}.$$

Then, with a little bit of algebra, we get

$$\mathcal{W}\mathcal{J} = \frac{|\mathcal{H}|^2 \mathcal{I} + \mathcal{H}^* \mathcal{N}}{|\mathcal{H}|^2 + \frac{|\mathcal{N}|^2}{|\mathcal{I}|^2}}.$$

For frequencies (u,v) where noise power is smaller than the image power $\mathcal{W}$ acts like an inverse filter since $\mathcal{N}(u,v)/\mathcal{I}(u,v) < 1$ and

$$\mathcal{W}\mathcal{J}(u,v) \approx \frac{|\mathcal{H}|^2}{|\mathcal{H}|^2} \mathcal{I}(u,v) = \mathcal{I}(u,v),$$

and at frequencies where the noise power dominates, $\mathcal{N}(u,v)/\mathcal{I}(u,v) > 1$ and

$$\mathcal{W}\mathcal{J}(u,v) = \frac{|\mathcal{I}|^2 \mathcal{H}^*}{|\mathcal{I}|^2 |\mathcal{H}|^2 + |\mathcal{N}|^2} \mathcal{N}(u,v),$$

the fraction is small so the noise power is diminished.

Image Restoration

$$\varepsilon^2 = \iint \left| \mathcal{I} - \tilde{\mathcal{I}} \right|^2 du dv$$

$$= \iint \left| \frac{\mathcal{K}}{\mathcal{H}} - \frac{\mathcal{W}\mathcal{J}}{\mathcal{H}} \right|^2 du dv$$

$$= \iint \left| \mathcal{H} \right|^{-2} \left| \mathcal{K} - \mathcal{W}(\mathcal{K} + \mathcal{N}) \right|^2 du dv$$

$$= \iint \left| \mathcal{H} \right|^{-2} \left| \mathcal{K}(1 - \mathcal{W}) + \mathcal{W}\mathcal{N} \right|^2 du dv$$

$$= \iint \left| \mathcal{H} \right|^{-2} \left[\mathcal{K}(1 - \mathcal{W}) + \mathcal{W}\mathcal{N} \right] \left[\overline{\mathcal{K}(1 - \mathcal{W}) + \mathcal{W}\mathcal{N}} \right] du dv$$

$$= \iint \left| \mathcal{H} \right|^{-2} \left[\left| \mathcal{K}(1 - \mathcal{W}) \right|^2 + \mathcal{K}(1 - \mathcal{W}) \overline{\mathcal{W}\mathcal{N}} + \mathcal{W}\mathcal{N} \overline{\mathcal{K}(1 - \mathcal{W})} + \left| \mathcal{W}\mathcal{N} \right|^2 \right] du dv$$

$$= \iint \left| \mathcal{H} \right|^{-2} \left\{ \left| \mathcal{K}(1 - \mathcal{W}) \right|^2 + 2 \operatorname{Re} \left[\mathcal{K}(1 - \mathcal{W}) \overline{\mathcal{W}\mathcal{N}} \right] + \left| \mathcal{W}\mathcal{N} \right|^2 \right\} du dv$$

$$= \iint \left| \mathcal{H} \right|^{-2} \left[\left| \mathcal{K}(1 - \mathcal{W}) \right|^2 + \left| \mathcal{W}\mathcal{N} \right|^2 \right] du dv + 2 \operatorname{Re} \iint \left| \mathcal{H} \right|^{-2} (1 - \mathcal{W}) \overline{\mathcal{W}} \mathcal{K} \overline{\mathcal{N}} \, du dv$$

This is one of the possible derivations of the Weiner filter

Image Restoration

From the previous page, the squared error is

$$\varepsilon^2 = \iint |\mathcal{H}|^{-2} \left[|\mathcal{K}(1-\mathcal{W})|^2 + |\mathcal{W}\mathcal{N}|^2 \right] du dv + 2\,\mathrm{Re} \iint |\mathcal{H}|^{-2} (1-\mathcal{W}) \overline{\mathcal{W}}\, \mathcal{K} \overline{\mathcal{N}}\, du dv.$$

The second term should be small compared to the first since it can be written

$$2\,\mathrm{Re} \iint |\mathcal{H}|^{-1} (1-\mathcal{W}) \overline{\mathcal{W}}\, \mathcal{I} \overline{\mathcal{N}}\, du dv$$

and the image and the noise are assumed to be uncorrelated. Thus the error can be approximated by

$$\varepsilon^2 = \iint |\mathcal{H}|^{-2} \left[|\mathcal{K}(1-\mathcal{W})|^2 + |\mathcal{W}\mathcal{N}|^2 \right] du dv.$$

The mean squared error, ε^2, is minimized when W is given by,

$$\mathcal{W} = \frac{\mathcal{H}^* |\mathcal{I}|^2}{|\mathcal{H}|^2 |\mathcal{I}|^2 + |\mathcal{N}|^2}.$$

Noise Reduction Through LMS Filtering[1]

image

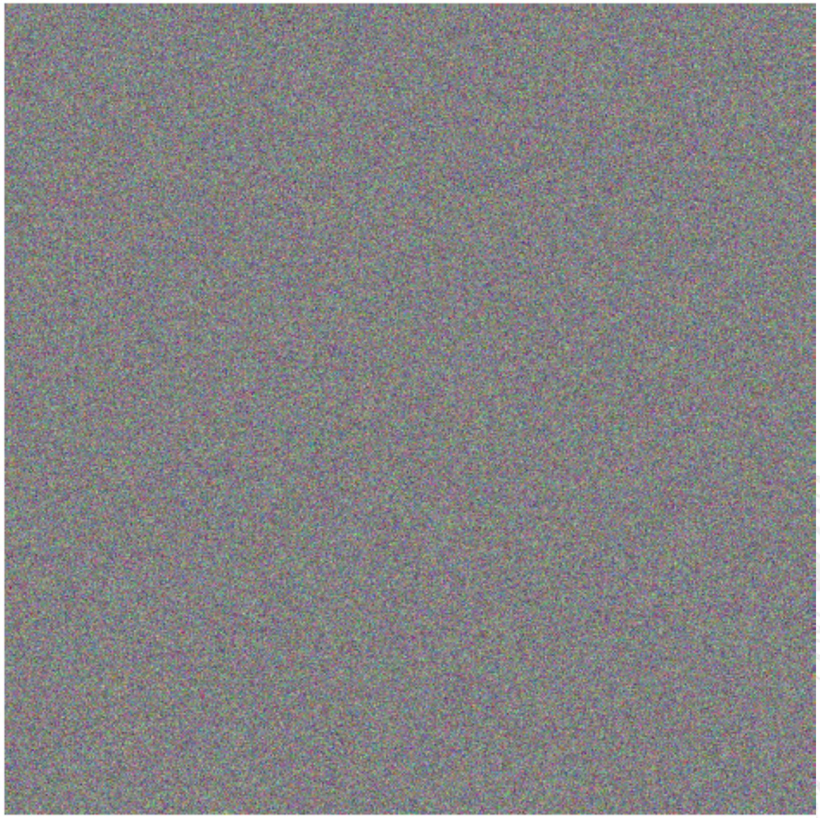

Gaussian noise field

[1]Here the PSF is the identity.

Noise Reduction Through LMS Filtering[1]

image

noisy image

[1]Here the PSF is the identity.

Additive Noise (Power Spectra)

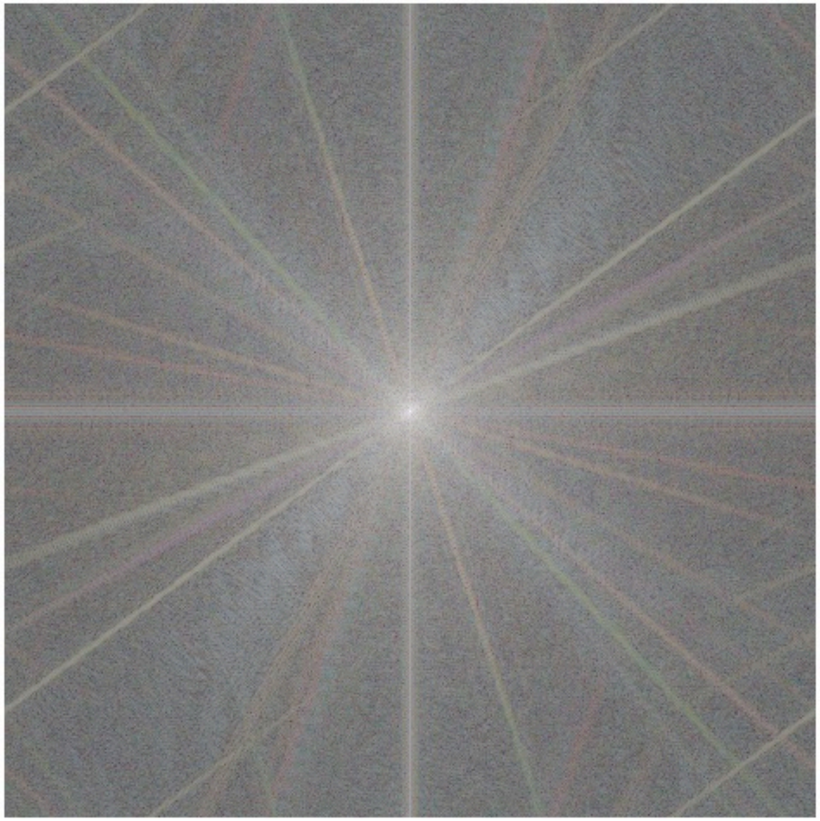

original image

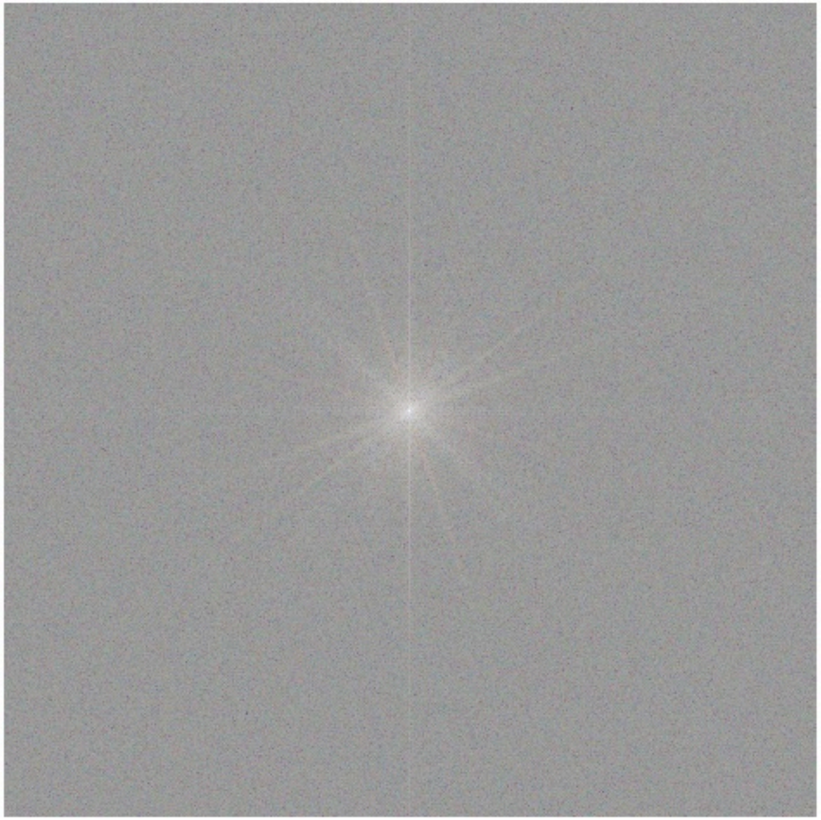

noisy image

Additive Noise (Power Spectra)

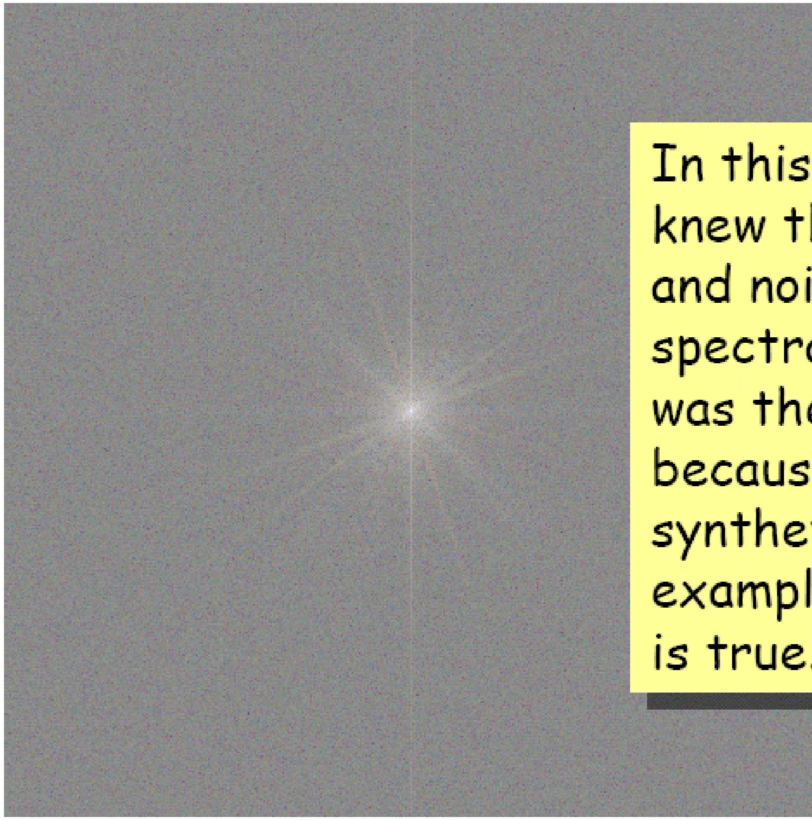

In this example we knew the exact image and noise power spectra and the PSF was the identity because the image is synthetic. In a real example, none of that is true.

Wiener filtered image

Wiener filter

Additive Noise (Power Spectra)

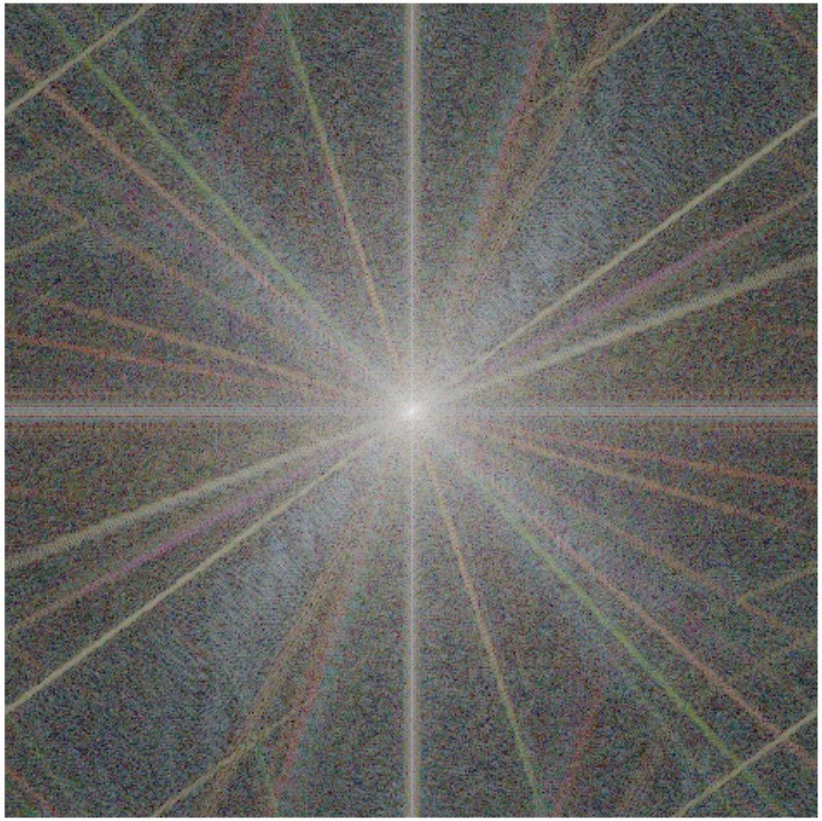

Wiener filtered image

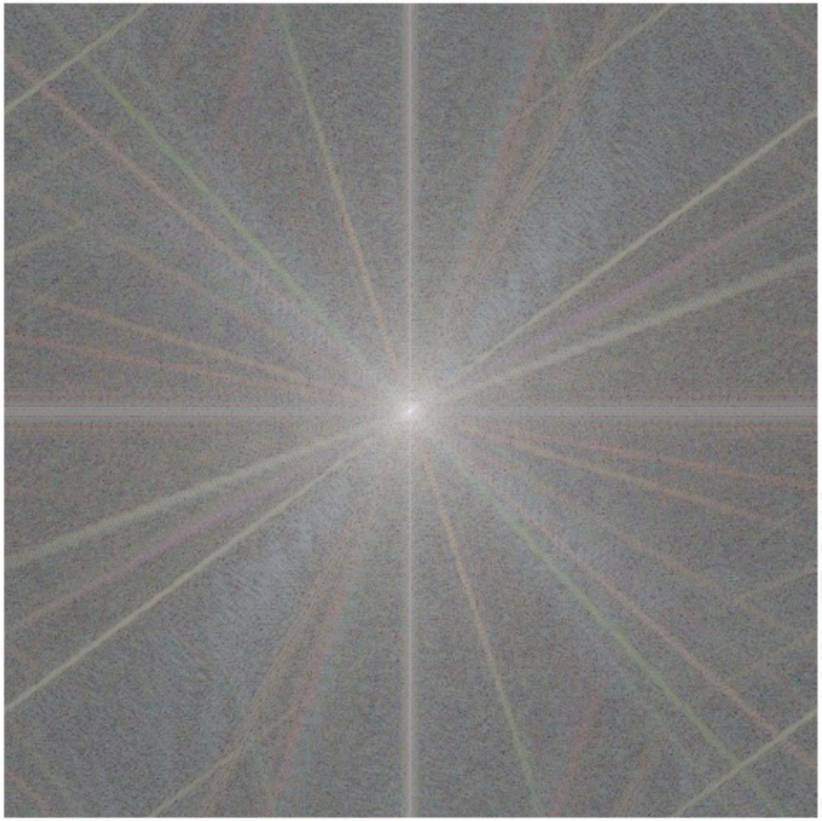

original image

Additive Noise

noisy image

Wiener filtered image

Additive Noise

Wiener filtered image

original image

Noise Reduction Through LMS Filtering[1]

image

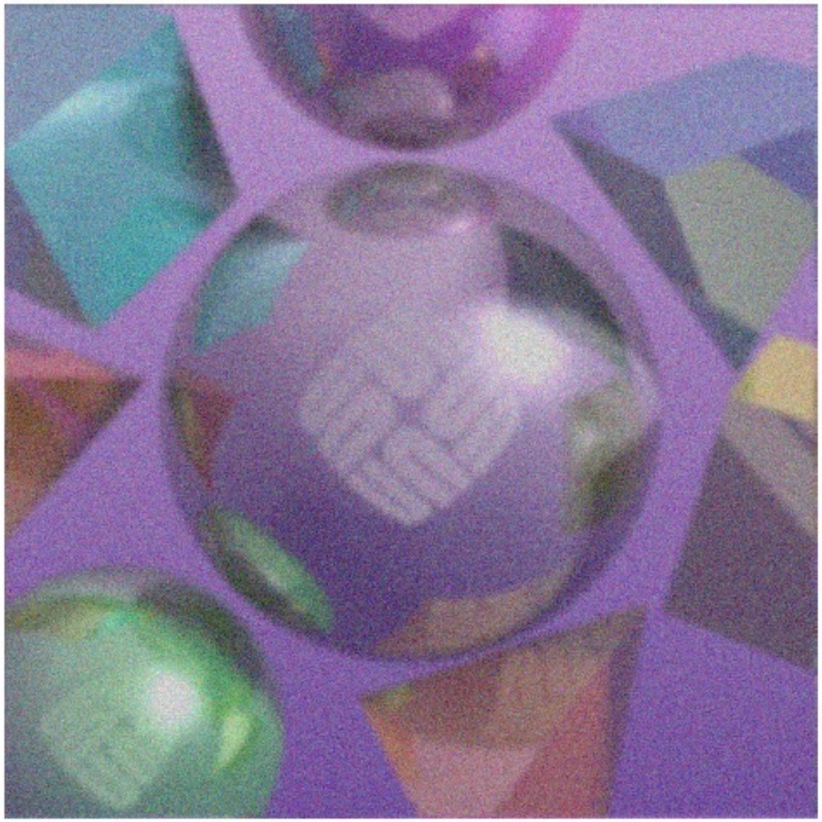

noisy image $J = I*h + N$

[1]PSF, h, is Gaussian $\mu=0$, $\sigma=2$.

Image*PSF + Noise (Power Spectra)

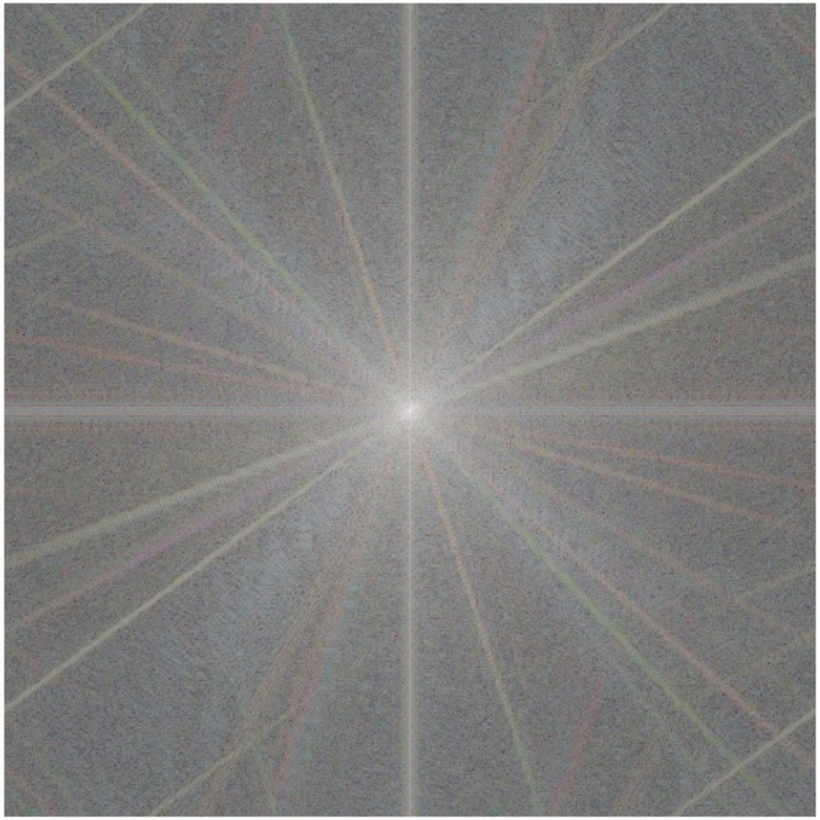

original image

noisy image J = I*h + N

Image*PSF + Noise (Power Spectra)

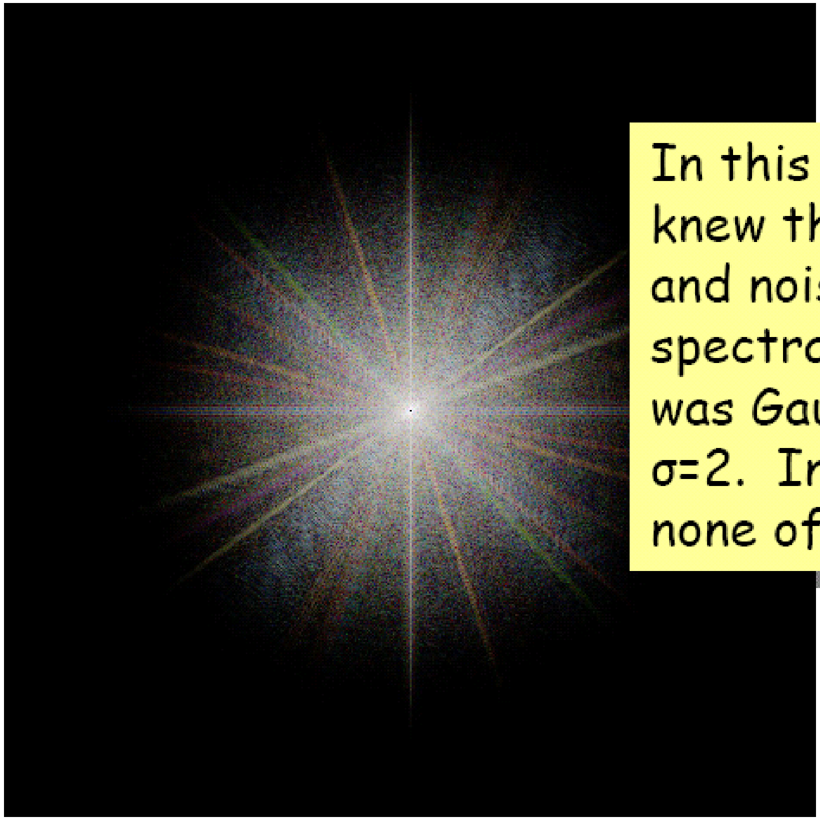

In this example we knew the exact image and noise power spectra and the PSF was Gaussian w/ μ=0, σ=2. In a real example, none of that is true.

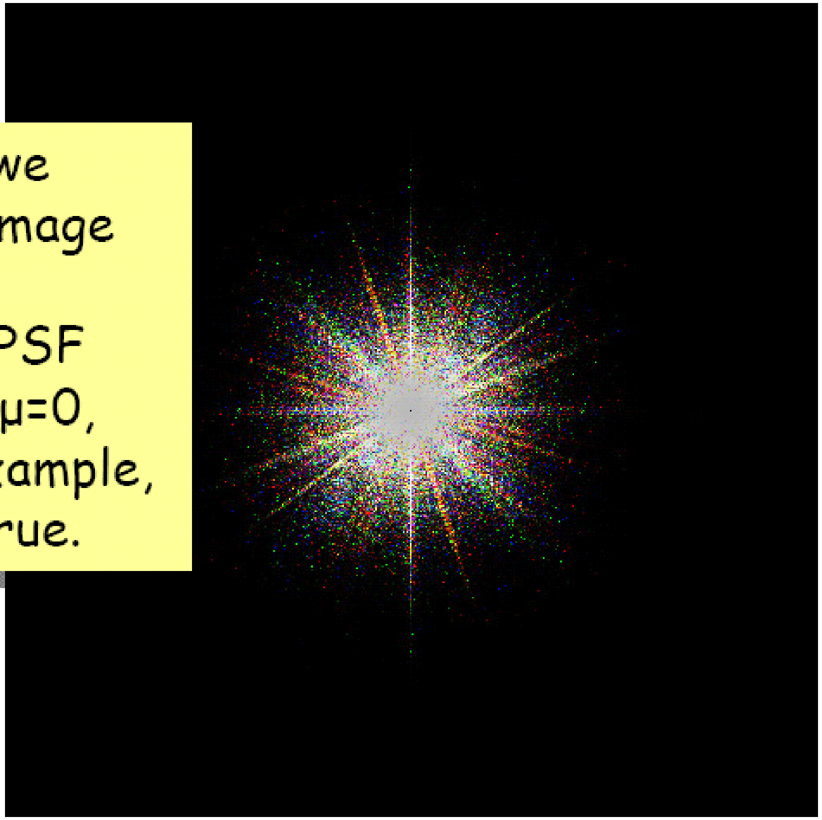

Wiener filtered image

Wiener filter

Image*PSF + Noise (Power Spectra)

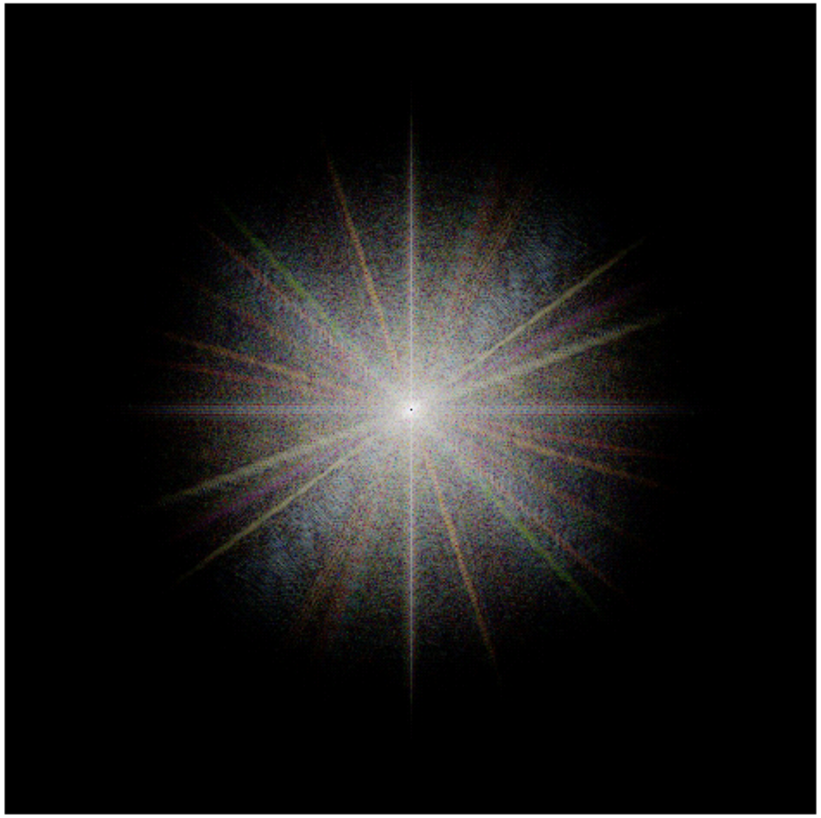

Wiener filtered image

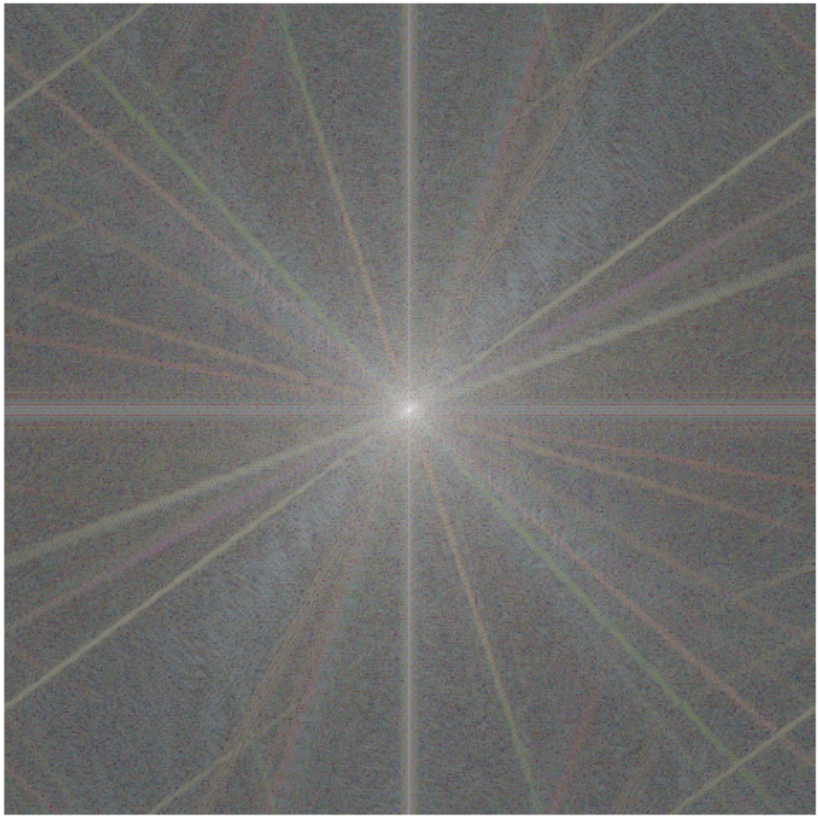

original image

Image*PSF + Noise

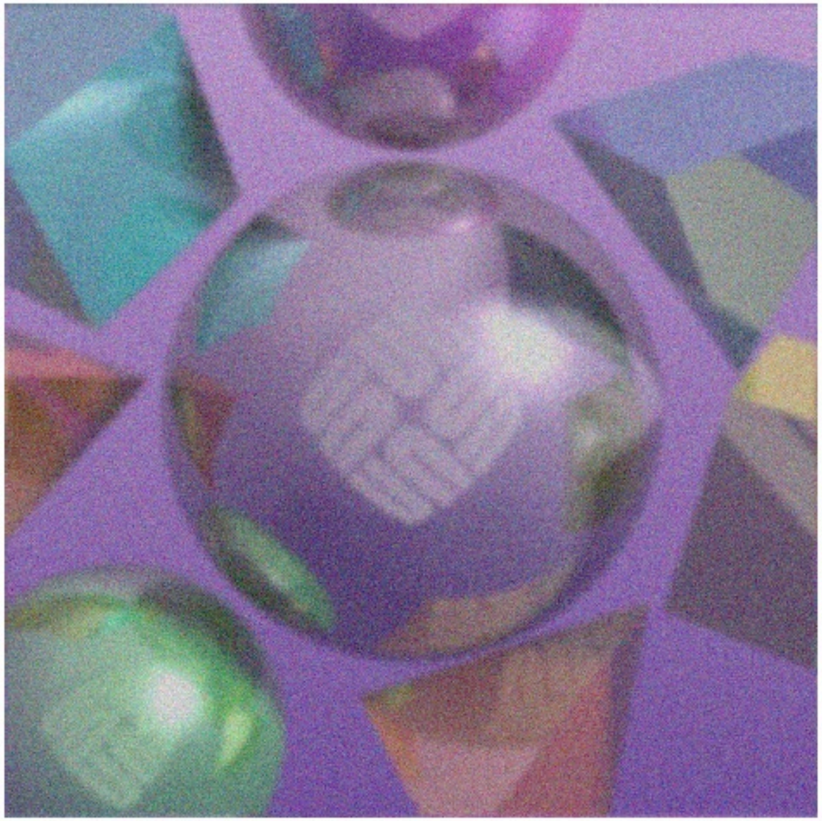

noisy image $J = I*h + N$

Wiener filtered image

Image*PSF + Noise

Wiener filtered image

original image

LMS Image Restoration (Real Example)

For this real example we need to estimate the image power spectrum, the pointspread function and the noise power spectrum.

LMS Image Restoration (Real Example)

To estimate the noise power spectrum, analyze a constant area from the image.

Noise Estimation

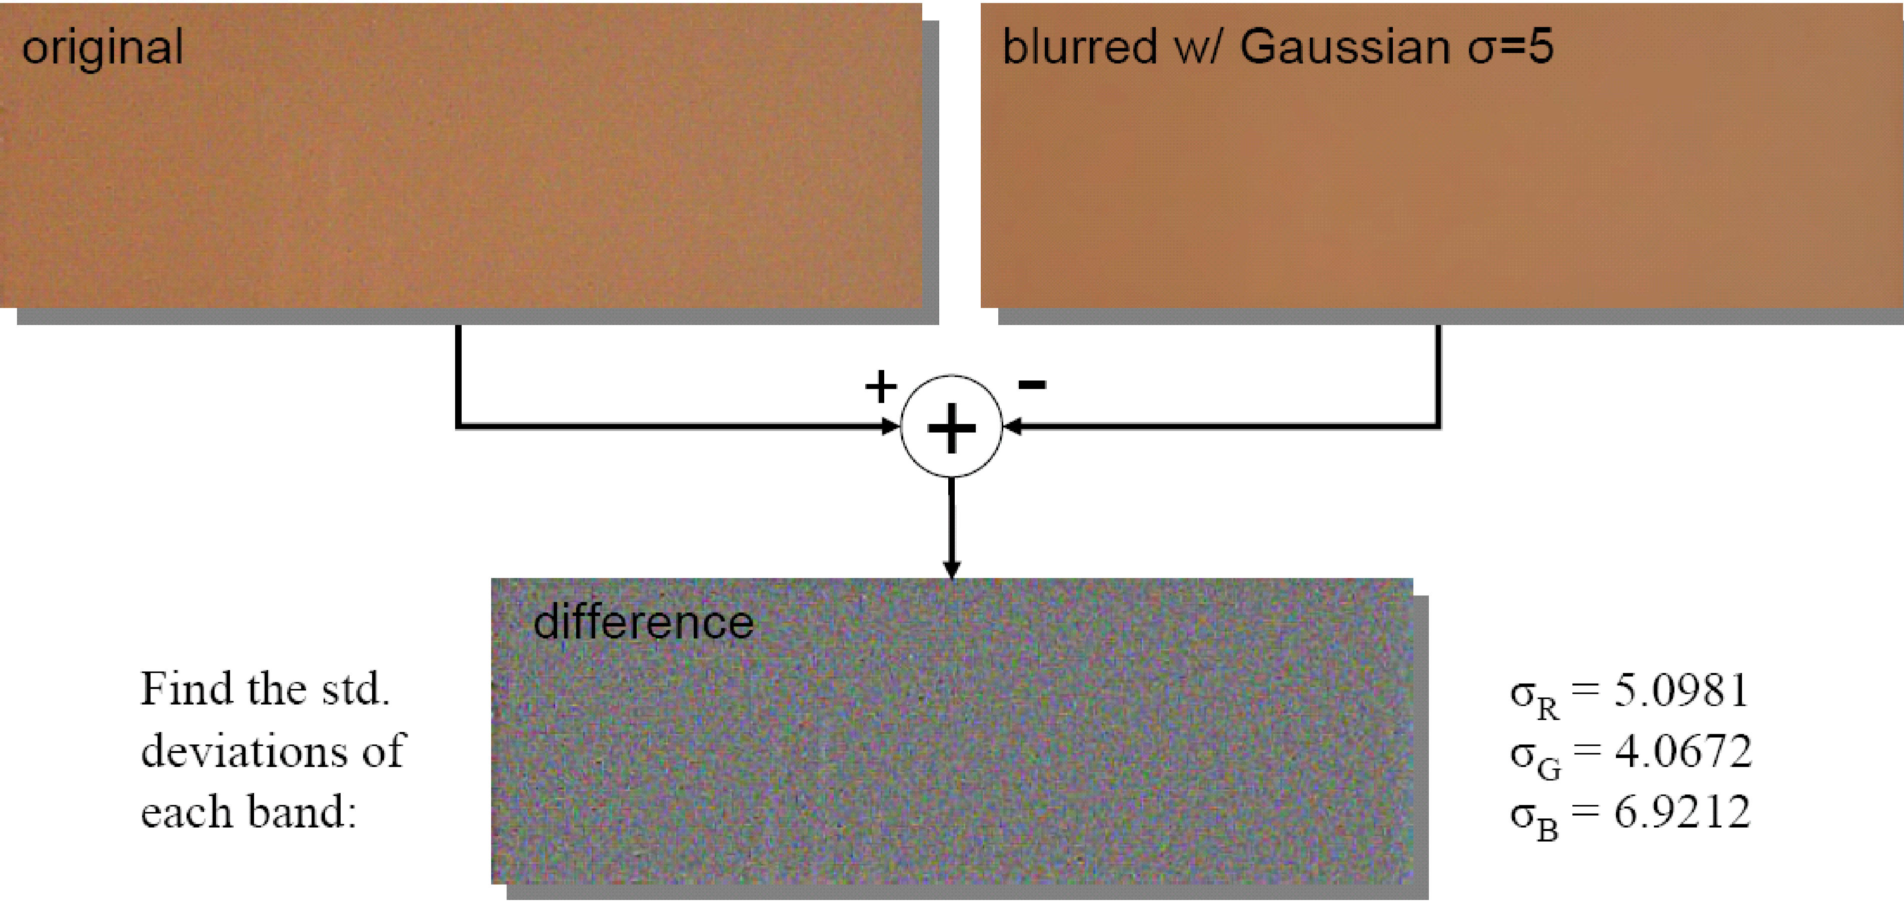

Find the std. deviations of each band:

$\sigma_R = 5.0981$
$\sigma_G = 4.0672$
$\sigma_B = 6.9212$

Pointspread Function Estimation

$$h = \begin{bmatrix} 0.0625 & 0.1250 & 0.0625 \\ 0.1250 & 0.2500 & 0.1250 \\ 0.0625 & 0.1250 & 0.0625 \end{bmatrix}$$

To estimate the PSF, find the image of a point and construct a convolution mask from it.

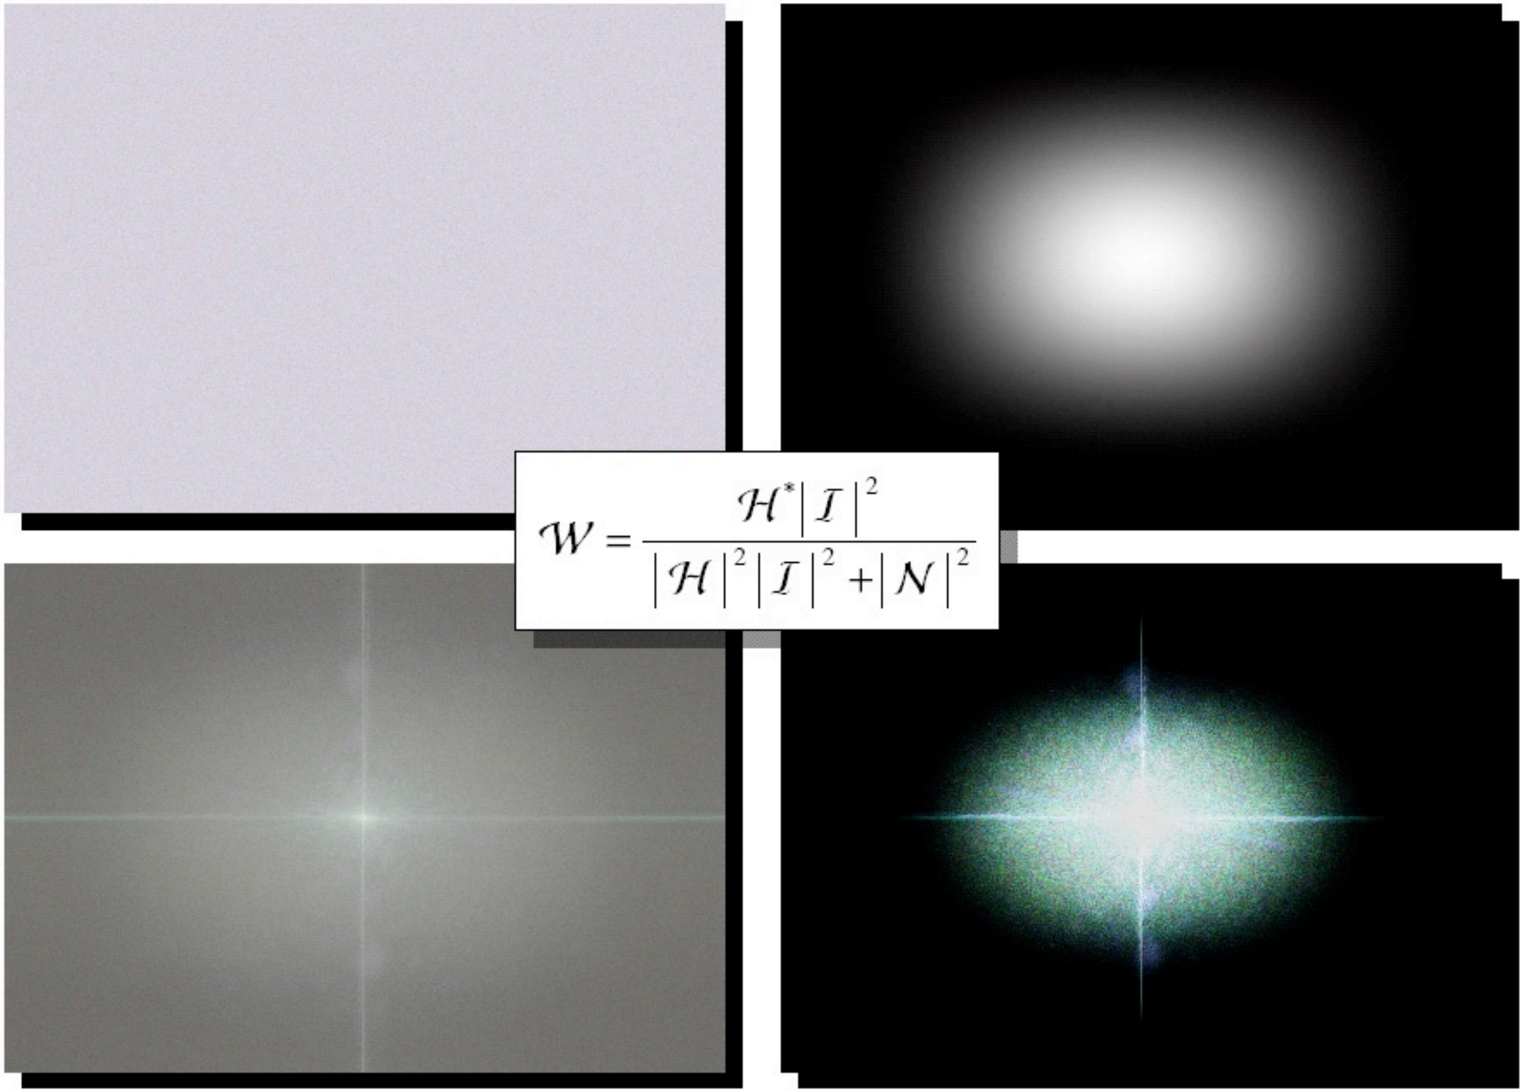

$$\mathcal{W} = \frac{\mathcal{H}^* |\mathcal{I}|^2}{|\mathcal{H}|^2 |\mathcal{I}|^2 + |\mathcal{N}|^2}$$

LMS Image Restoration (original)

LMS Image Restoration (filtered)

Blur+Noise

Motion Blur

Original

Restauration

Restauration

Periodic Noise

original image

image + noise

Noise Reduction through Directional Blurring

image + noise

*

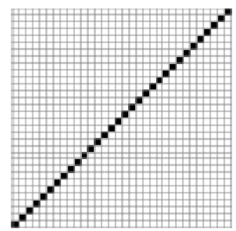

diagonal convolution mask

=

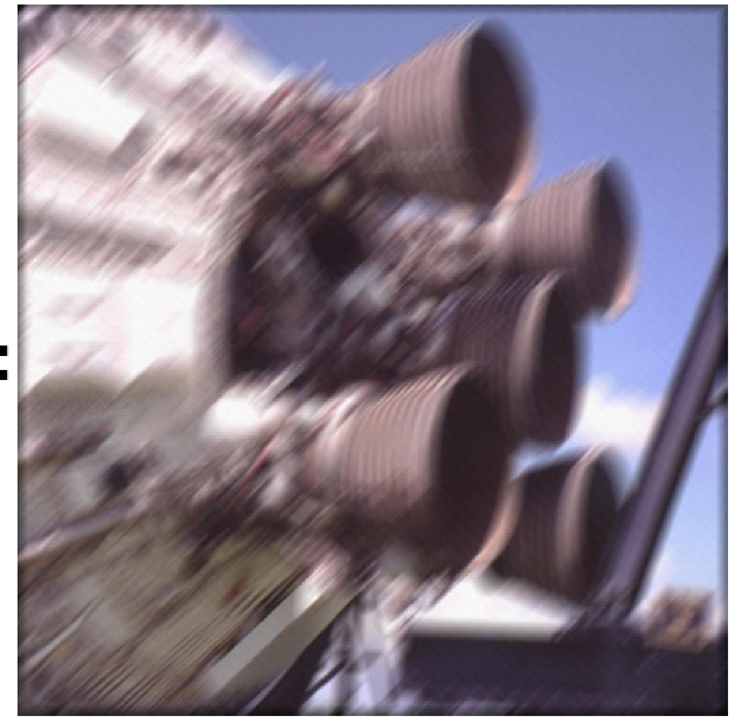

blurred image

Power Spectrum of Image with Periodic Noise

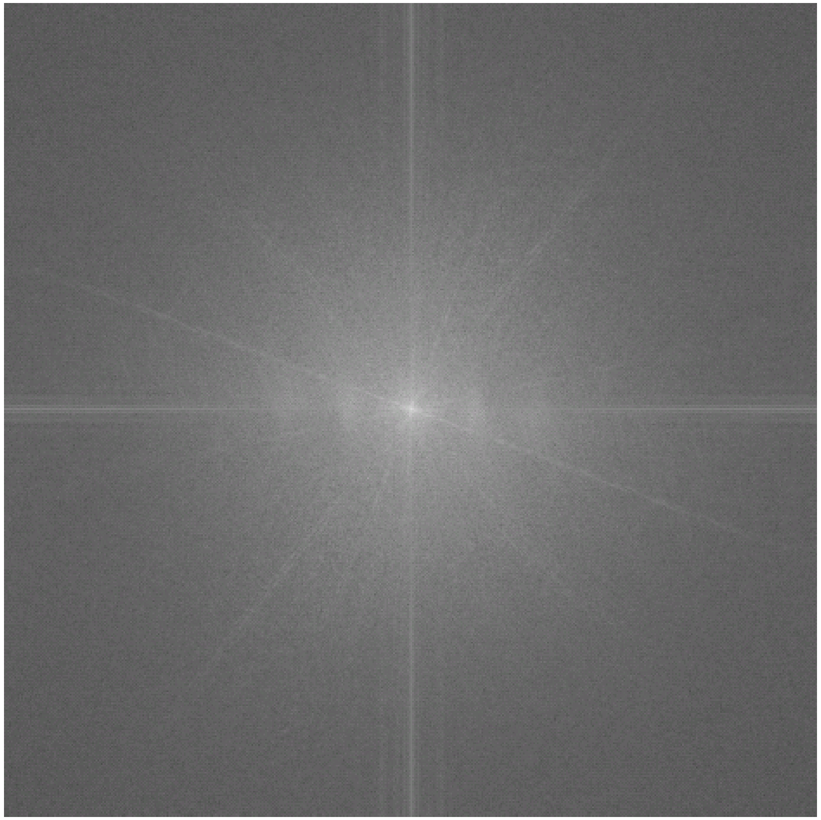

original image

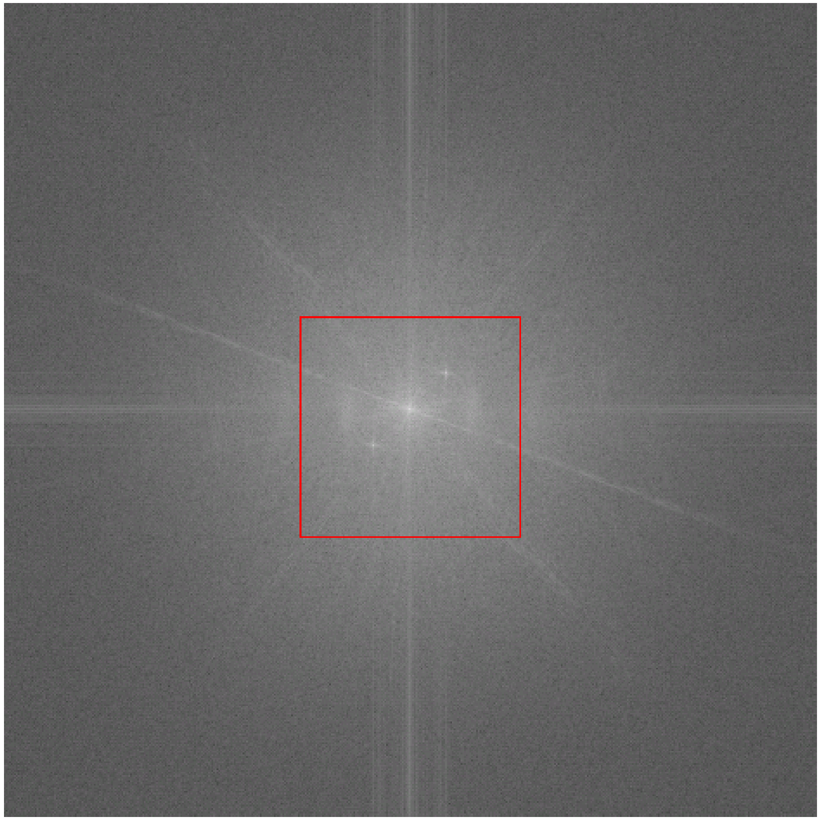

image + noise

Low Frequency Region

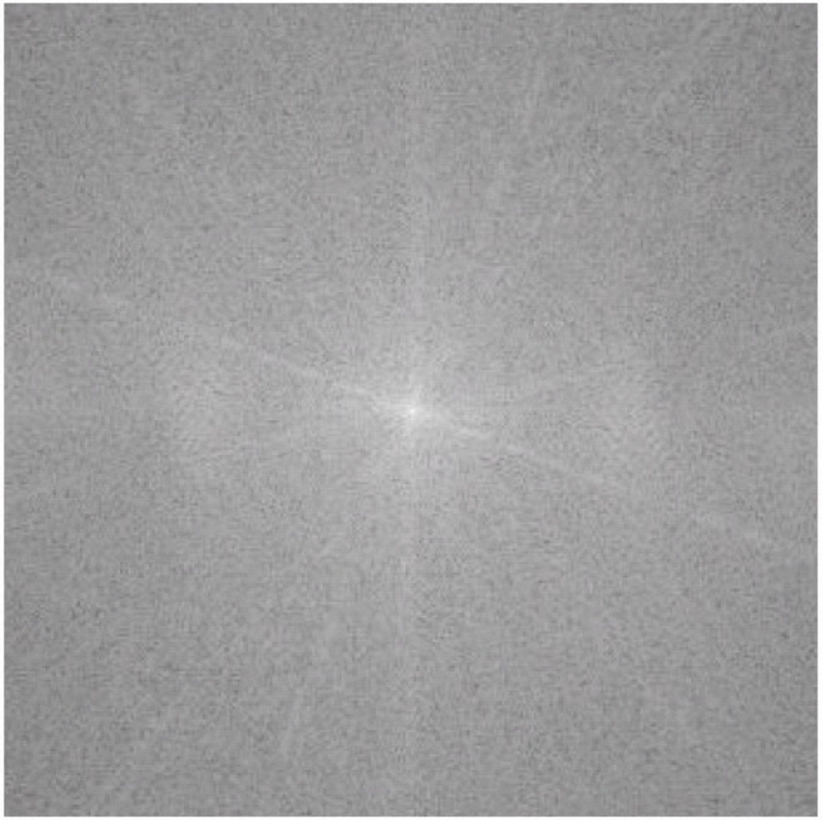

original image

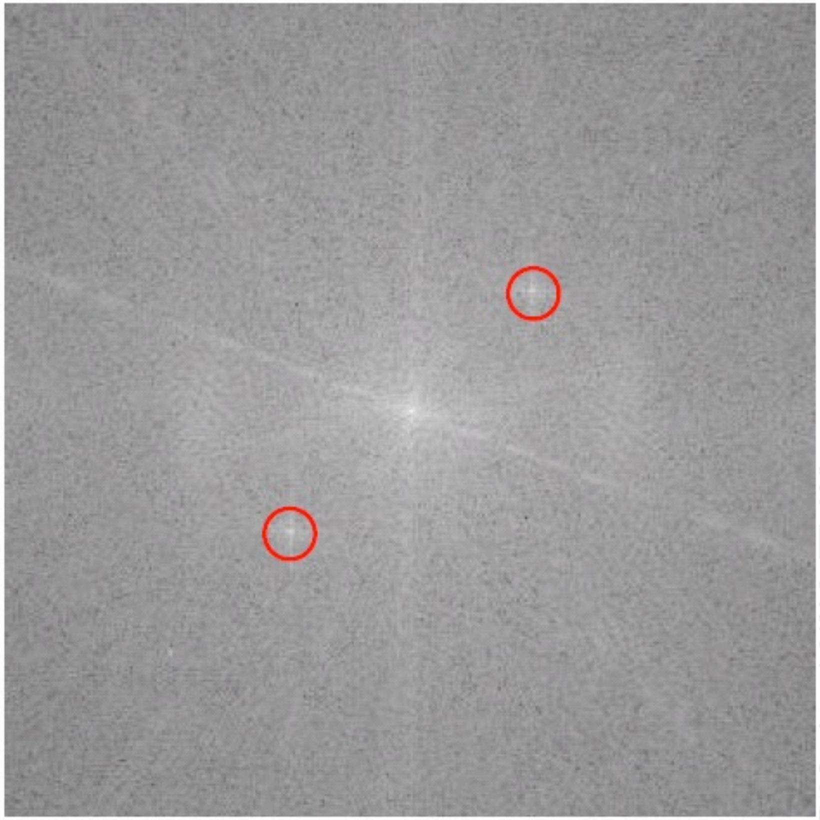

image + noise

Noise Reduction through Notch Filtering

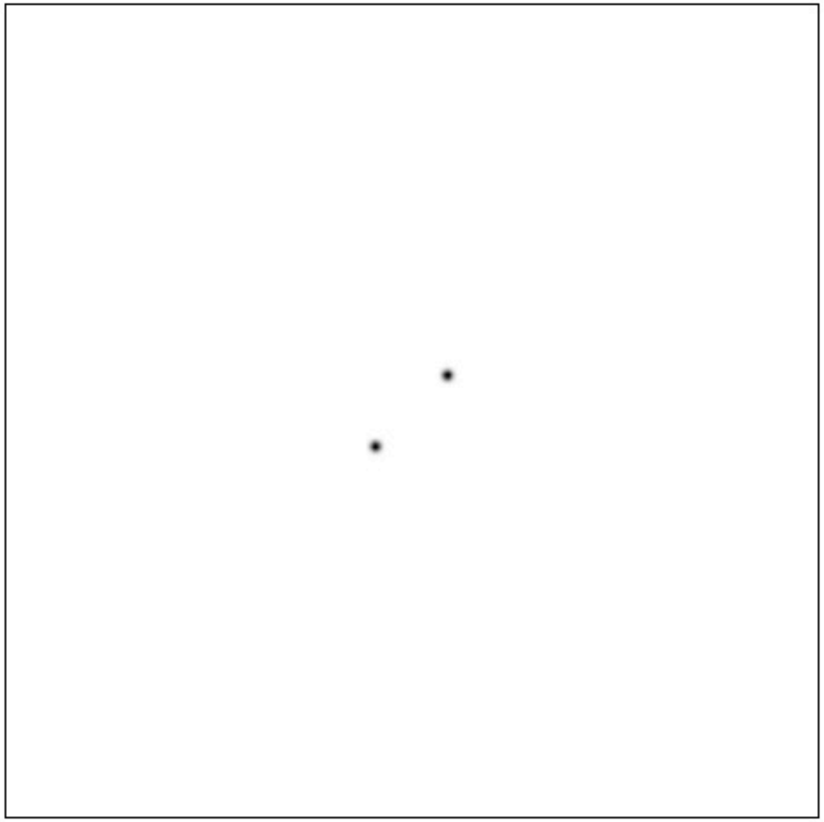

noise mask

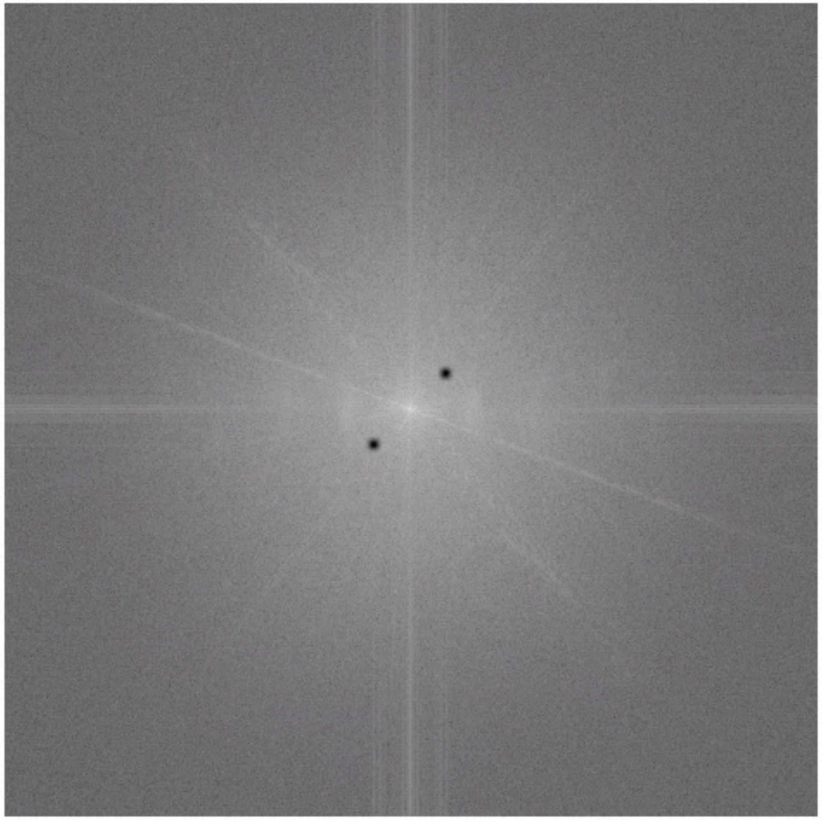

masked power spectrum

Inverse of Masked Fourier Transform

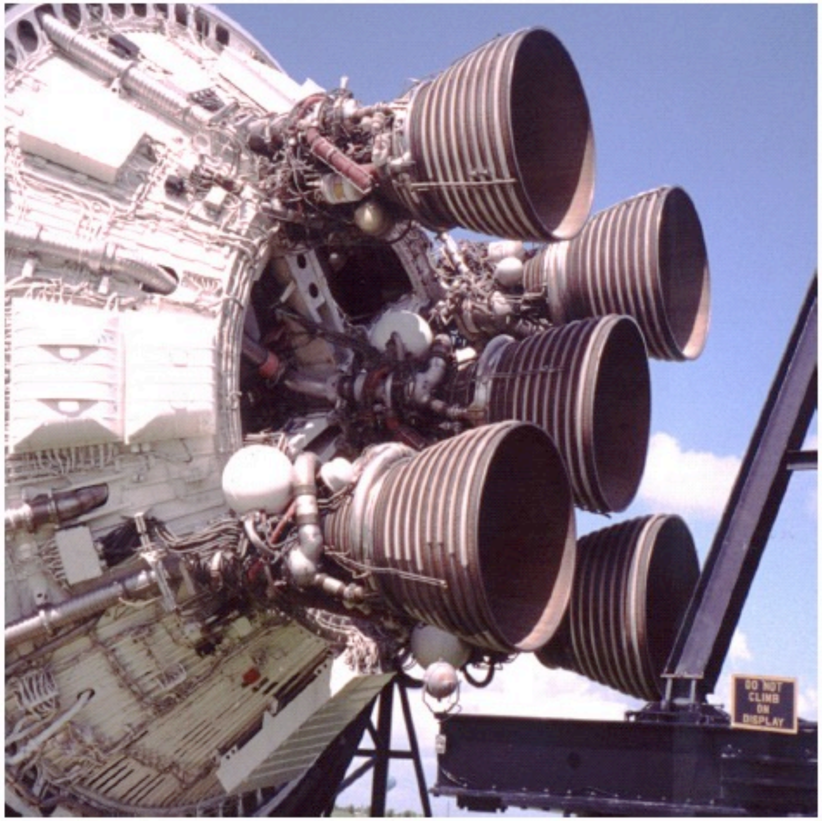

original image

noise reduced image

Halftone Dots

Image scanned (600 dpi) from a magazine

Detail: Circular patterns, the rosettes, are the result of the halftone dots.

Filtering Out Halftone Dot Distortion

original

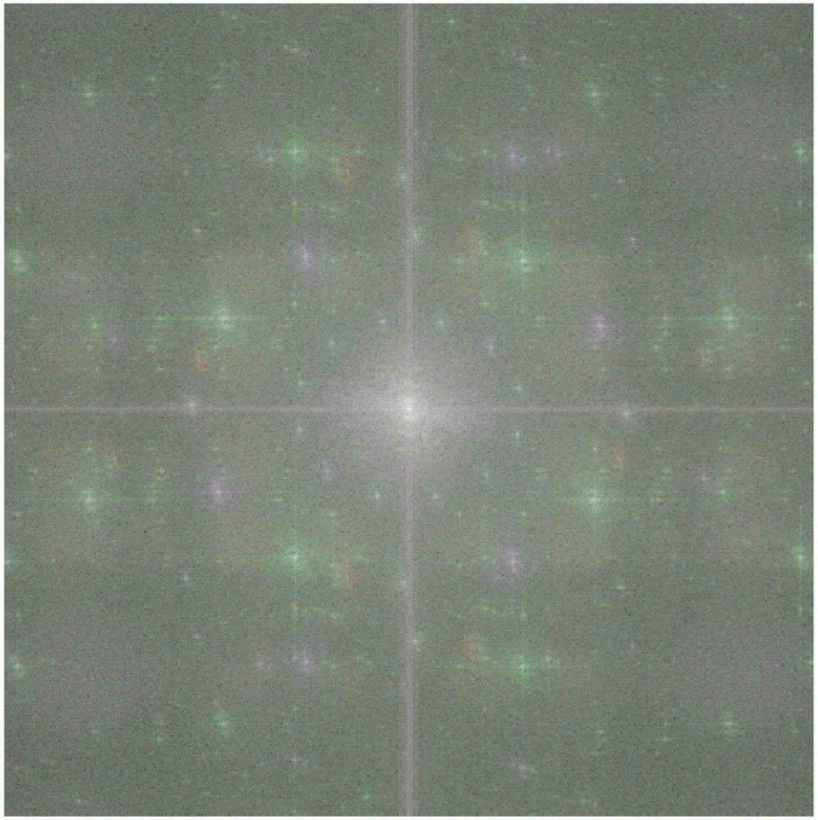

log power spectrum

Notch Filtering

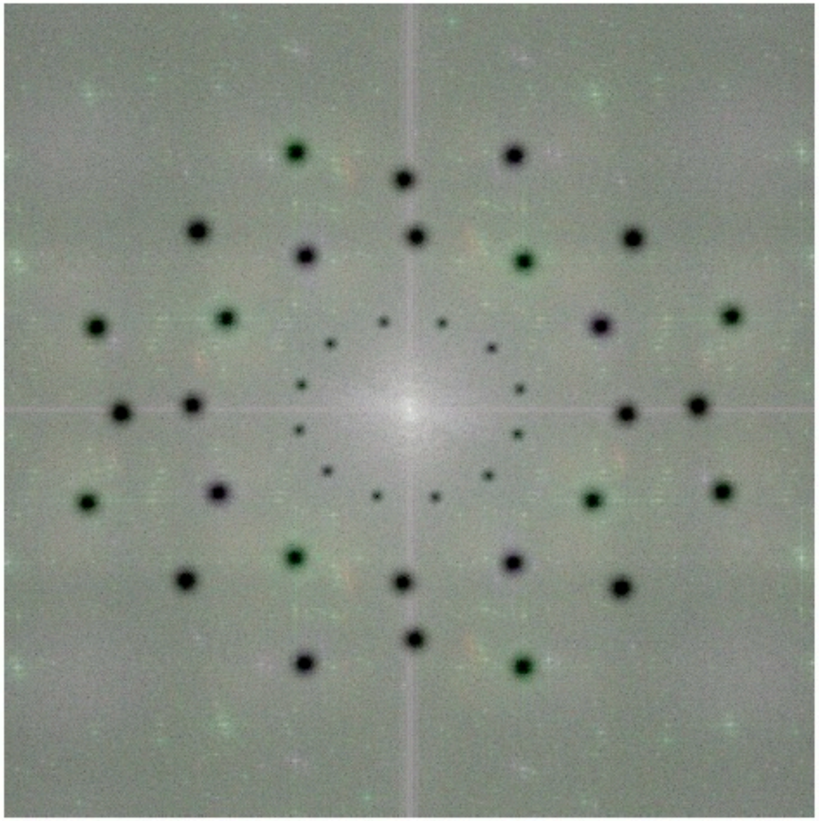

log power spectrum

frequency masked 3

Alias-Effekte und Moiré

Eine eng mit dem Rauschen verwandte Bildstörung sind Alias-Effekte., die sich oft in Form von Moiré-Mustern zeigen. Sie treten vor allem auf, wenn das signaltheoretische Abtasttheorem verletzt ist, d.h. wenn Signale mit einer nicht ausreichenden Abtastfrequenz abgetastet werden.

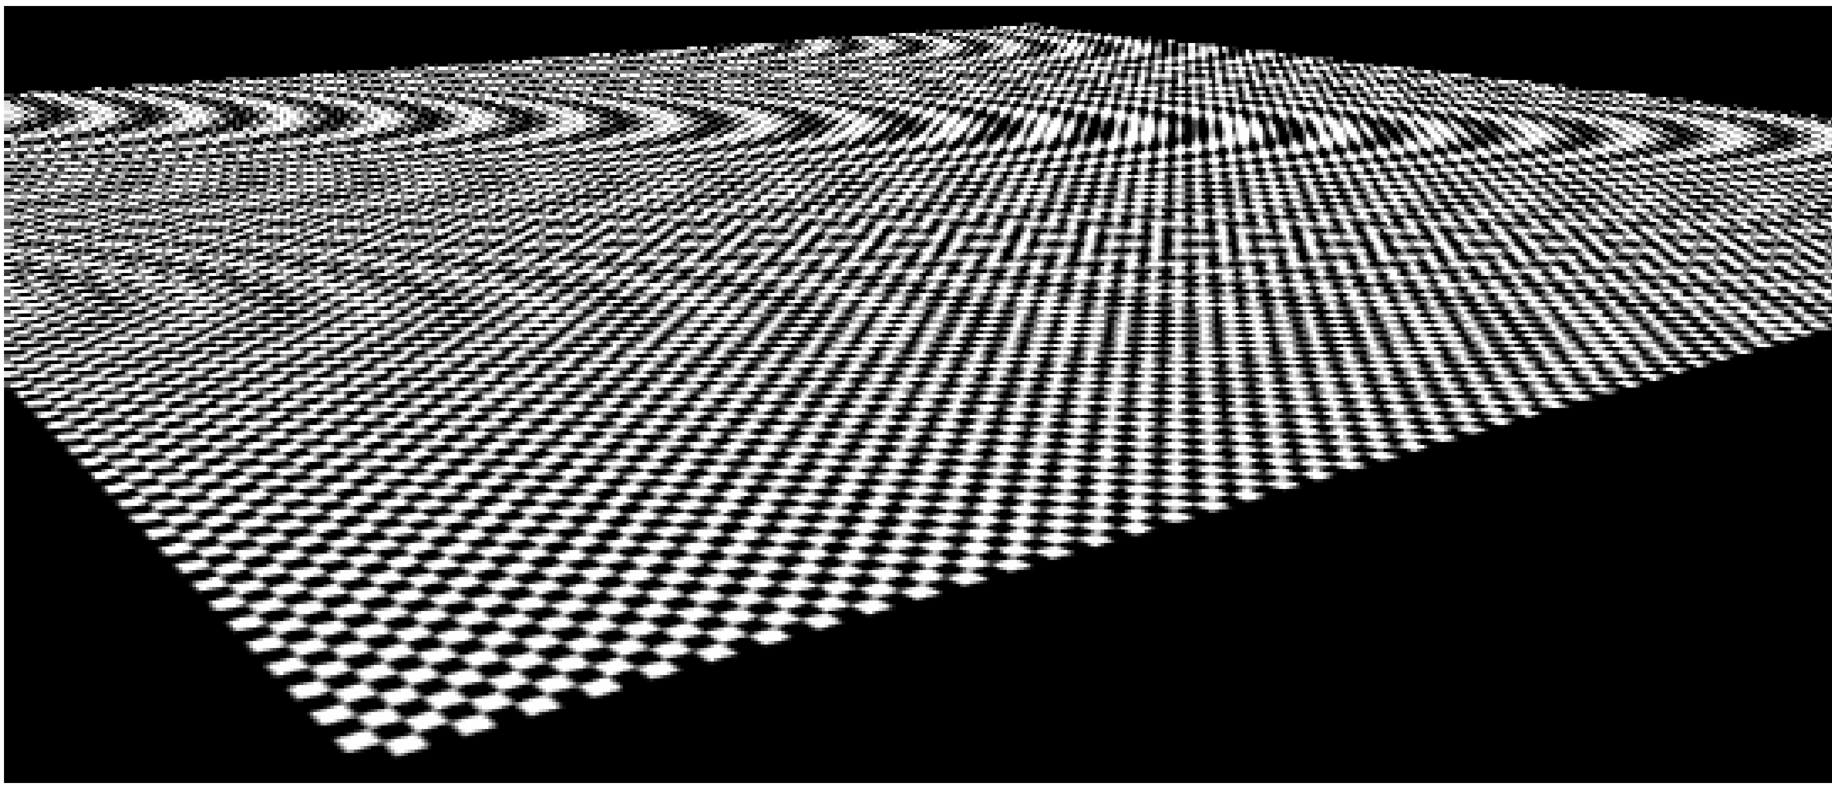

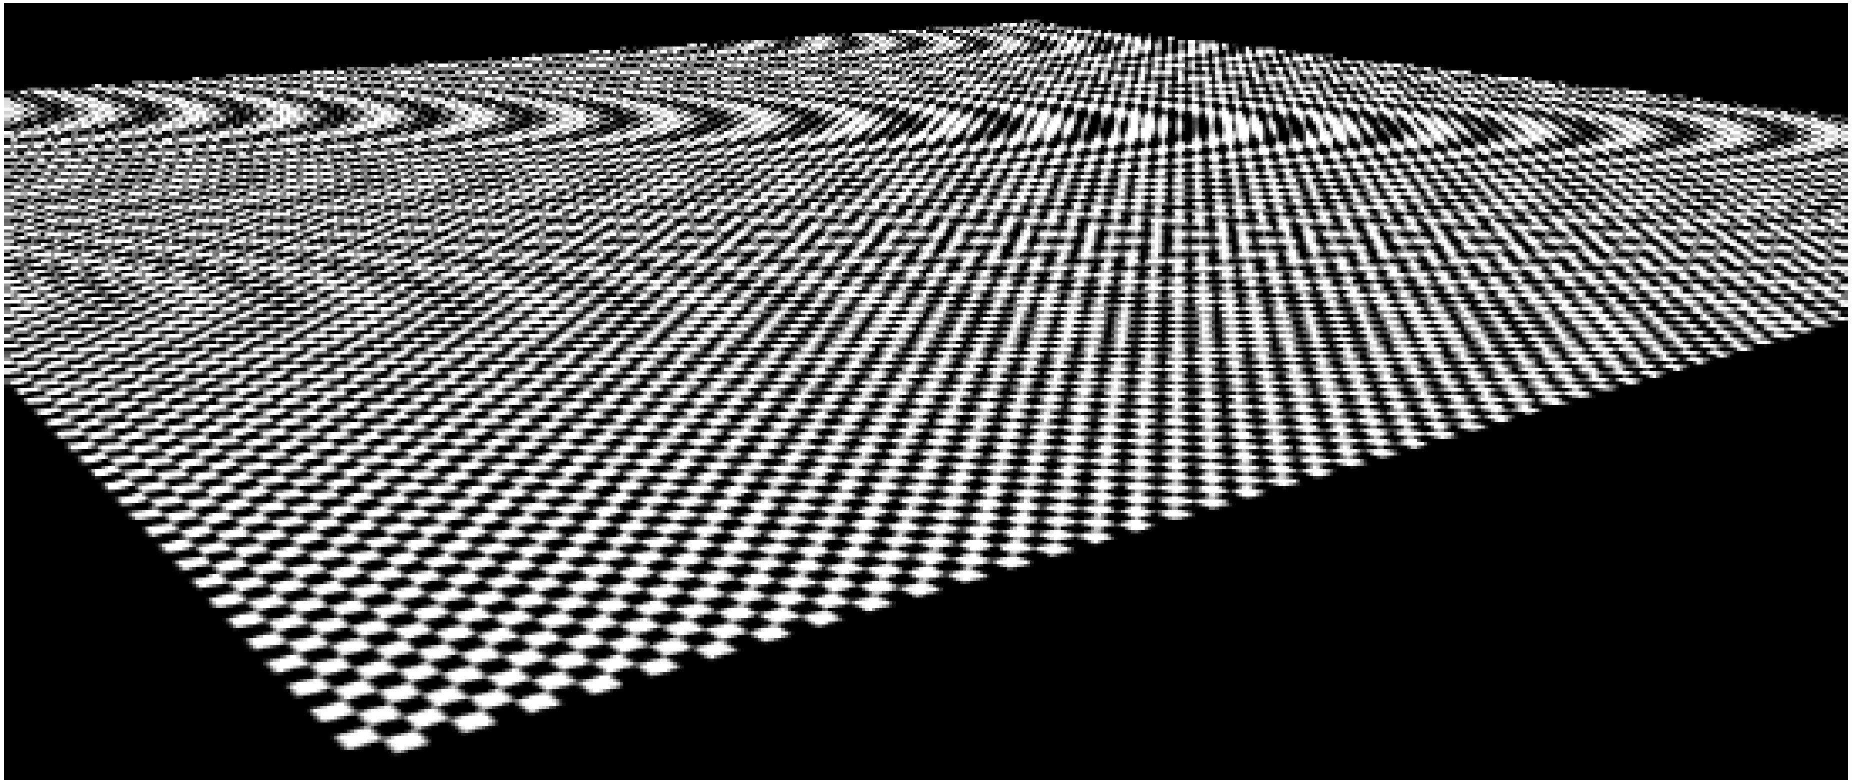

Bei zu starker Unterabtastung geht der Alias-Effekt über in ein örtlich unkorreliertes Rausch-Signal

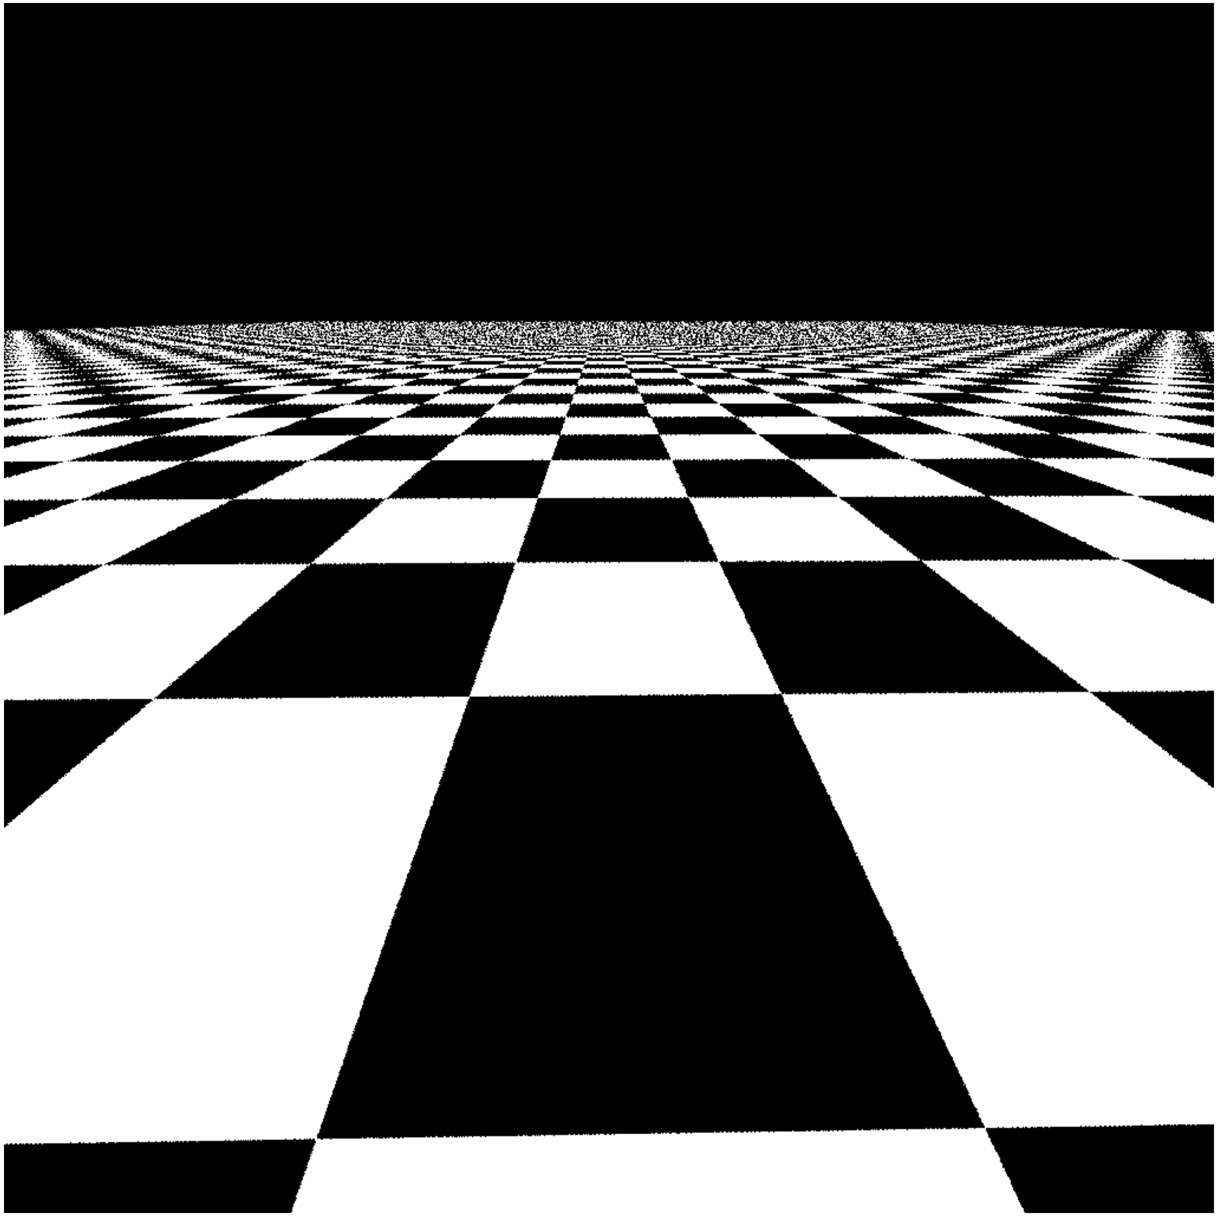

Eine signaltheoretisch korrekte Bandbegrenzung kann zu deutlich sichbarer Unschärfe führen.

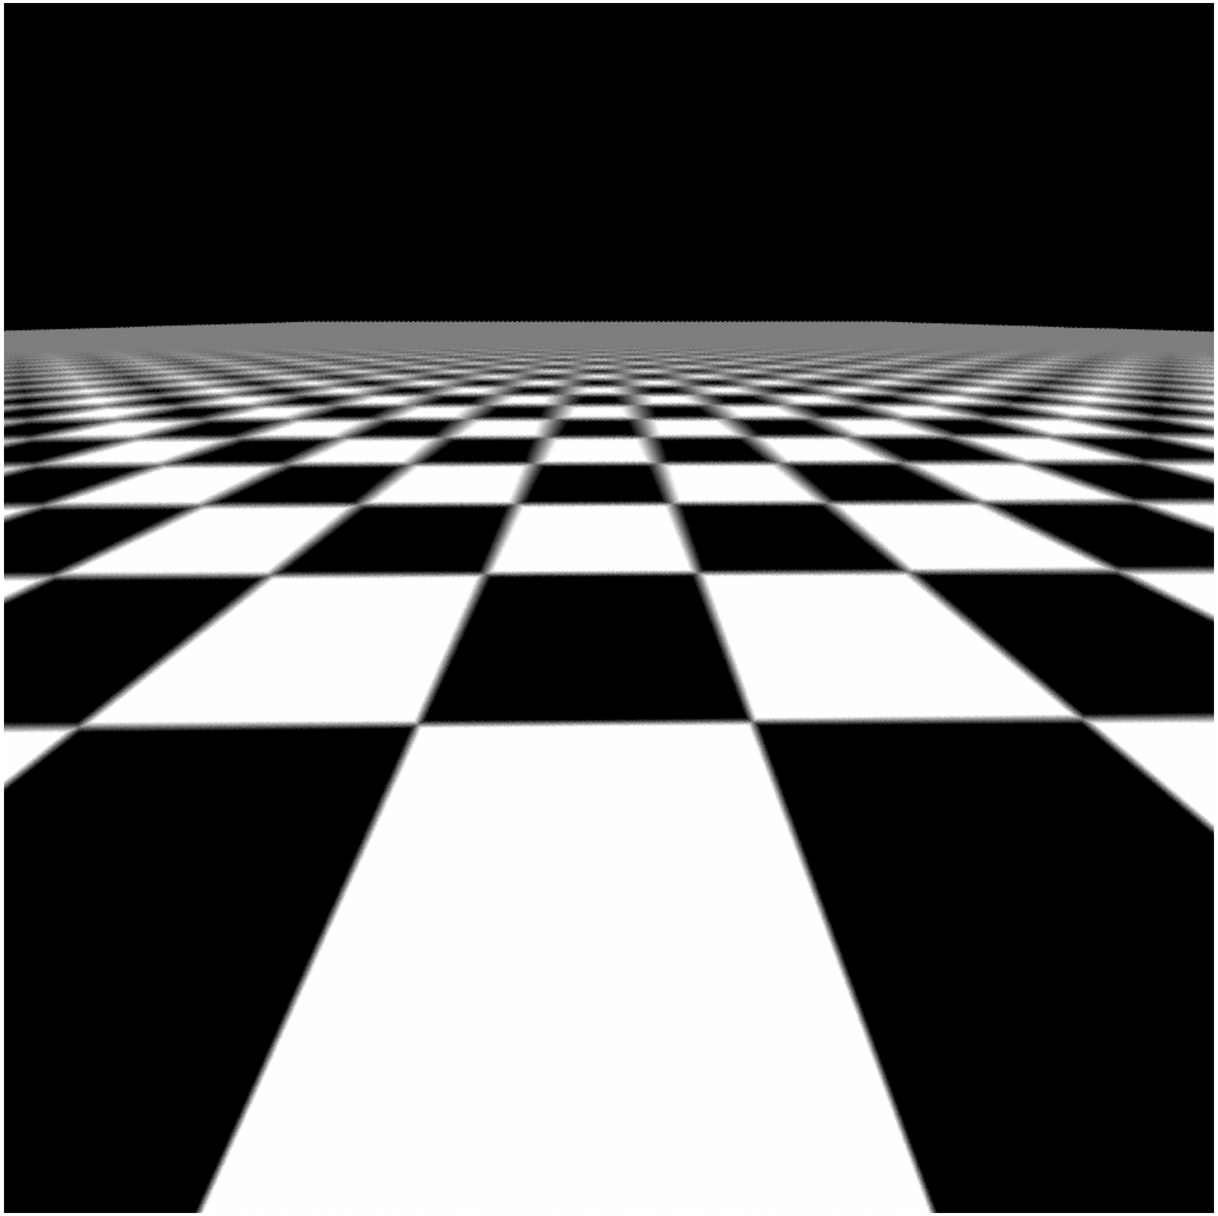

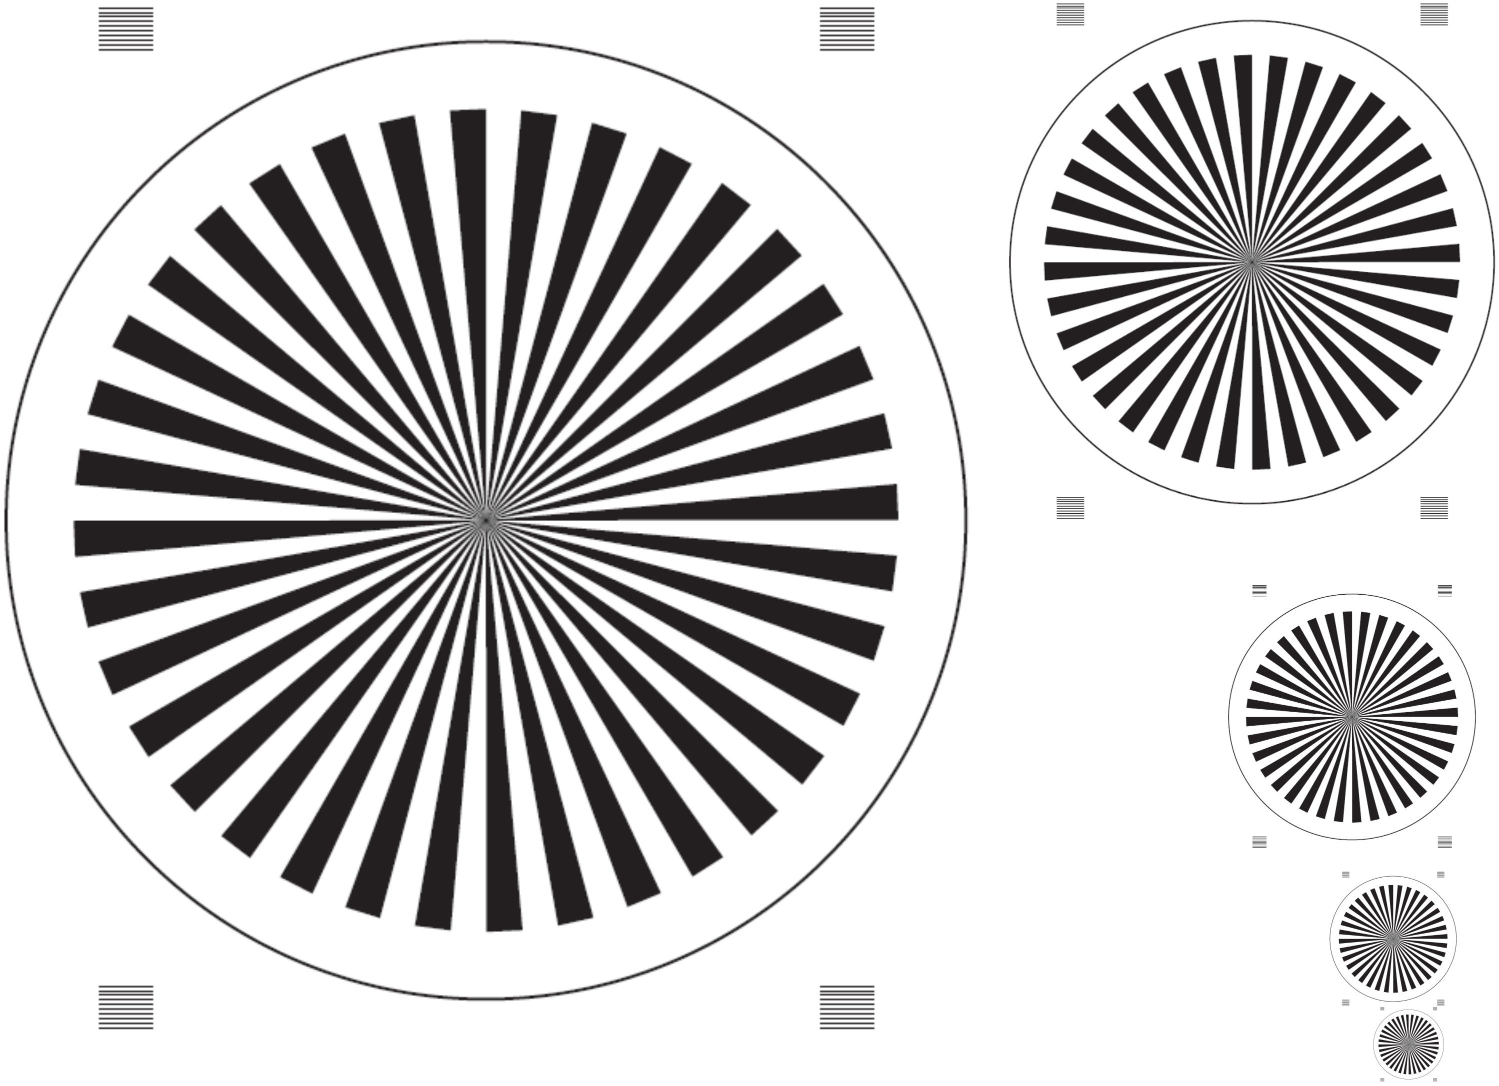

Gerade bei der Aufnahme von Mauerwerk aus der Entfernung treten oft starke **Moiré-Muster** auf.

Auch der *Treppen-Effekt* an schrägen Linien ist ein Alias-Effekt.

Durch zueinander versetzte Farbsensoren können auch *Farb-Alias*-Effekte auftreten.