Image Processing and Computer Vision

Image Processing and Computer Vision

- **Processing of continuous images**
  - **linear filtering**
  - **Fourier transformation**
- **Wiener filtering**
- **Nonlinear diffusion**

Computer Vision

What is computer vision? **interpreting images!**

The computer sees 1001110100101010000000001110101...

Image Processing

What is image processing? **restoring images without extraction of semantic information!**

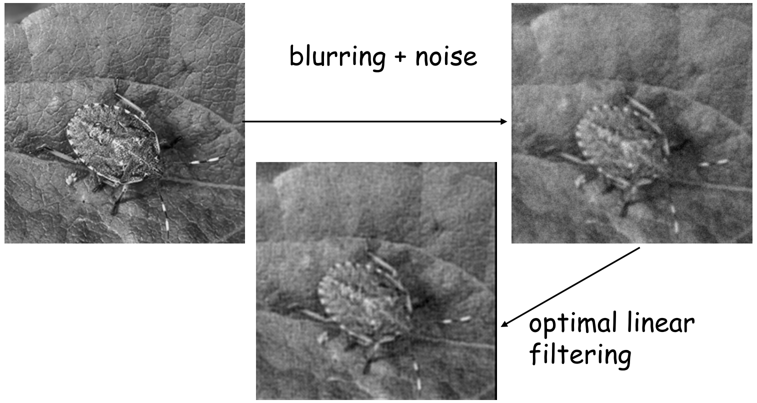

The Image Formation Process

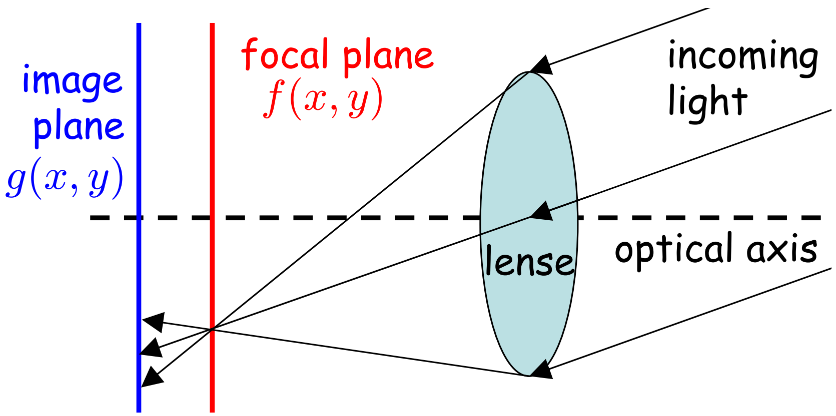

Mathematical Modelling of Image Processing

Def.: An image is a continuous, two-dimensional function of the light intensity

$$f : \begin{array}{ccc} \mathbb{R}^2 & \to & \mathbb{R}_+ \\ (x,y) & \mapsto & f(x,y) \end{array}$$

Question: How can we compensate an image deformation, e.g., defocussing?

Goal: reconstruct $f(x,y)$ from $g(x,y)$ in the presence of noise!

Model assumption:

1) When $f(x,y)$ is shifted then $g(x,y)$ is shifted as well.

2) Doubling the incoming light intensity will double the brightness $g(x,y)$.

Linear Shift-Invariant Systems

Strategy for restauration: invert the transformation which has mapped the original image $f(x,y)$ to the defocussed image $g(x,y)$.

Linearity: (assumption)

$$f_1 \longrightarrow \boxed{\text{transform}} \longrightarrow g_1$$

$$f_2 \longrightarrow \boxed{\text{transform}} \longrightarrow g_2$$

$$\alpha f_1 + \beta f_2 \longrightarrow \boxed{\text{transform}} \longrightarrow \alpha g_1 + \beta g_2 \; \forall \alpha, \beta \in \mathbb{R}$$

- Linearity is typically only in the low intensity range fulfilled since physical systems tend to saturate.
- f_i, g_i are intensities $\equiv$ power per area with $f_i, g_i \geq 0$ in the full domain.
- Often we experience non-linear imaging errors!

Shift invariance: (assumption)

$$f(x,y) \longrightarrow \boxed{\text{transform}} \longrightarrow g(x,y)$$

$$f(x-a, y-b) \longrightarrow \boxed{\text{transform}} \longrightarrow g(x-a, y-b)$$

- Shift invariance holds only in a limited range since images are finite objects.

Remarks: The assumption of linearity is a significant limitation but it gives the advantage that the linear filter theory is completely developed.

- An analogous one-dimensional theory applies to passive electrical circuits, although there time is the essential dimension and causality constraints the signal.

How Can We Identify a Transformation?

Dirac's δ-function (1D): $\int_{-\infty}^{\infty} \delta(x-a)f(x)dx = f(a)$

- Integration with the δ-function "samples" the function $f(x)$ at the position $x_0 = a$.
- The δ-function is a "generalized function".
- Regularization:

$$\delta(x) = \lim_{\epsilon \to 0} \begin{cases} \frac{1}{\epsilon} & |x| \leq \frac{\epsilon}{2} \\ 0 & \text{else} \end{cases}$$

or

$$\delta(x) = \lim_{\epsilon \to 0} \frac{1}{\sqrt{2\pi}\epsilon} \exp(-\frac{x^2}{2\epsilon^2})$$

Convolution and the Point Spread Function

Assumption: $\delta(x,y) \longrightarrow \boxed{\mathcal{T}} \longrightarrow h(x,y)$

With linearity and shift invariance it holds:

$$\begin{aligned} g(x,y) &= \mathcal{T} f(x,y) \\ &= \mathcal{T} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\xi,\eta)\delta(x-\xi,y-\eta)d\xi d\eta \\ &\overset{\text{linearity}}{=} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\xi,\eta) \underbrace{[\mathcal{T}\delta(x-\xi,y-\eta)]}_{\substack{h(x-\xi,y-\eta) \\ \text{[shift inv.]}}} d\xi d\eta \\ &= (f * h)(x,y) \end{aligned}$$

Linear, shift invariant systems can be written as convolutions!

Identification of the Kernel

Let $f(x,y) = \delta(x,y)$, i.e., the image is a white dot with "infinite" intensity. Then the measured image $g(x,y)$ is given by

$$\begin{aligned} g(x,y) &= (\delta * h)(x,y) \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \delta(\xi,\eta)h(x-\xi,y-\eta)d\xi d\eta \\ &= h(x,y) \end{aligned}$$

$$\Rightarrow \quad \mathcal{T}\delta(x,y) = h(x,y)$$

$\Rightarrow$ testing the linear shift-invariant system with a δ-peak will reveal the convolution kernel $h(x,y)$ of the system.

Schematic View of a Convolution

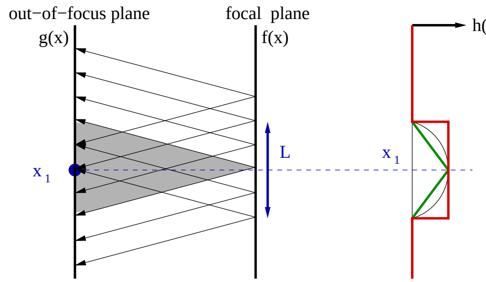

- $g(x_1)$ depends on $f(x)$ for all $x \in [x_1 - \frac{L}{2}, x_1 + \frac{L}{2}]$.
- convolution kernel $h_{x_1}(x)$ describes the influence of $f(x)$ onto $g(x_1)$.
- shift invariance of $h_{x_1}(x)$ results in cumulative influence:

$$\begin{aligned} g(x_1) &= \int_{-L/2}^{L/2} f(x)h(x_1-x)dx = \int_{-L/2}^{L/2} f(x_1-x)h(x)dx \\ &\approx f(0)h(x_1)\Delta + f(\Delta)h(x_1-\Delta)\Delta + f(2\Delta)h(x_1-2\Delta)\Delta + \ldots \\ &\quad + f(-\Delta)h(x_1+\Delta)\Delta + f(-2\Delta)h(x_1+2\Delta)\Delta + \ldots \end{aligned}$$

Convolution: 1D-Example

$$y(t) = (x * h)(t) = \int x(\tau)h(t-\tau)d\tau$$

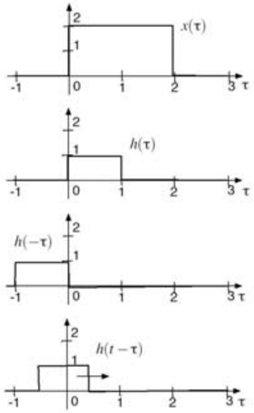

Convolution: 1D-Example (cont'd)

$$y(t) = (x * h)(t) = \int x(\tau)h(t-\tau)d\tau$$

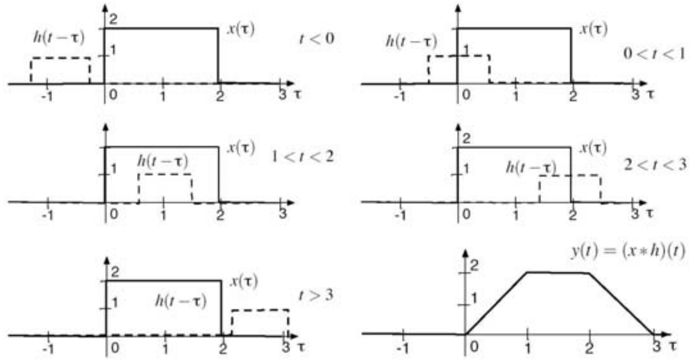

Facts about Convolution

- Linear shift-invariant (LSI) systems can be written as convolutions.
- The convolution kernel h characterizes the LSI system uniquely.
- Cascades of LSI systems: the convolution is commutative and associative:

$$\begin{aligned} g * h &= h * g \\ (f * g) * h &= f * (g * h) \end{aligned}$$

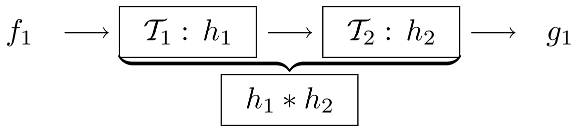

⇒ one of the most important operations in signal processing

Convolution Kernel for Image Defocussing

Defocussing an image amounts to convolving it with a 'pillbox':

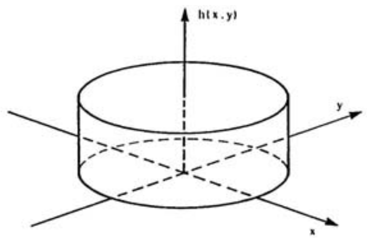

$$h(x,y) = \begin{cases} \frac{1}{\pi R^2} & x^2 + y^2 \le R^2 \\ 0 & \text{otherwise} \end{cases}$$

Note: this convolution kernel is normalized: $\int\int h(x,y)dxdy = 1$

Convolution Kernel for Image Defocussing

original image

convolved with pillbox kernel

A Motion Kernel

Each light dot is transformed into a short line along the x-axis:

$$h(x, y) = \frac{1}{2l} \left[\theta(x + l) - \theta(x - l)\right] \delta(y)$$

Denoising Time Series

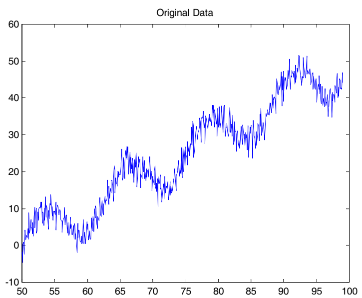

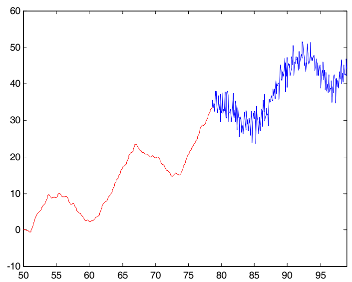

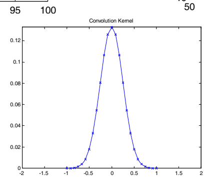

Lena with Gaussian Blurring and Noise

Gaussian blurring kernel:

$$h(x, y) = \frac{1}{2\pi\sigma^2} \exp(-\frac{x^2 + y^2}{2\sigma^2})$$

Lena Convolved with a Laplacian Filter

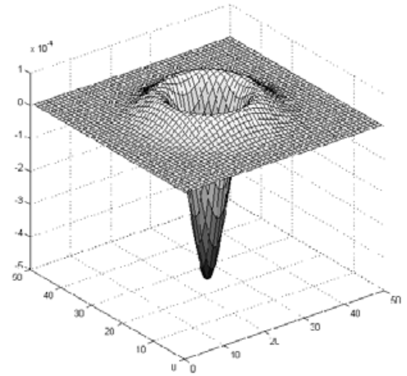

Laplacian filter: $$h(x, y) = \nabla^2 \exp(-\frac{x^2 + y^2}{2\sigma^2})$$

Note: here the normalization is $\int\int h(x, y)dxdy = 0$.

The Fourier Transformation

Def.: Let f be an absolutely integrable function over $\mathbb{R}$. The Fourier transformation of f is defined as

$$\hat{f}(u) \equiv \mathcal{F}[f(x)] = \int_{-\infty}^{+\infty} f(x) \exp(-\imath 2\pi ux)dx.$$

The inverse Fourier transformation is given by the formula

$$f(x) \equiv \mathcal{F}^{-1}[\hat{f}(u)] = \int_{-\infty}^{+\infty} \hat{f}(u) \exp(\imath 2\pi ux)du.$$

Note: while $f(x)$ is always real, $\hat{f}(u)$ is typically complex.

- $\hat{f}(u)$ is also called the **continuous spectrum** of $f(x)$.
- If x is a space coordinate, then u is called the **spatial frequency**.

Inversion formula: $f(x)$ is represented as a continuous superposition of waves with amplitude $\hat{f}(u)$.

Example of an odd function approximated by sinus waves (Remember: $\exp(\imath x) = \cos(x) + \imath \sin(x)$):

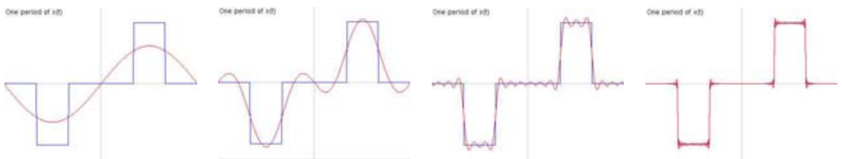

$$f(x) \approx \hat{f}(u_0)\sin(2\pi u_0 x) + \hat{f}(u_1)\sin(2\pi u_1 x) + \hat{f}(u_2)\sin(2\pi u_2 x) + \ldots$$

Fourier Transformation: Example 1 (box)

Given the box function

$$f(x) = \frac{1}{2l}(\theta(x + l) - \theta(x - l)) = \begin{cases} \frac{1}{2l} & \text{if} |x| \leq l \\ 0 & \text{otherwise} \end{cases}$$

the Fourier transform is

$$\begin{aligned} \hat{f}(u) \equiv \mathcal{F}[f(x)] &= \int_{-\infty}^{+\infty} f(x) \exp(-\imath 2\pi ux)dx \\ &= \int_{-l}^{l} \frac{1}{2l} \cdot (\cos(2\pi ux) - \imath \underbrace{\sin(2\pi ux)}_{\int \to 0})dx \\ &= \frac{\sin(2\pi ul)}{2\pi ul} \equiv \text{sinc}(2\pi ul) \end{aligned}$$

Fourier Transformation: Example 1 (box)

Graphs of box and sinc-function for $l = \frac{1}{2}$:

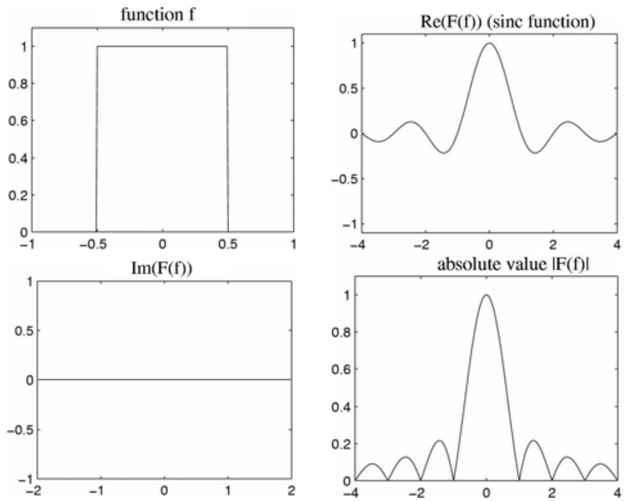

Fourier Transformation: Example 2 (Gauss)

Given the function

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma_x} \exp(-\frac{x^2}{2\sigma_x^2})$$

the Fourier transform is

$$\begin{aligned} \hat{f}(u) &\equiv \mathcal{F}[f(x)] = \int_{-\infty}^{+\infty} f(x) \exp(-\imath 2\pi u x) dx \\ &= \frac{1}{\sqrt{2\pi}\sigma_x} \int_{-\infty}^{\infty} \exp(-\frac{x^2}{2\sigma_x^2}) \cdot (\cos(2\pi u x) - \imath \underbrace{\sin(2\pi u x)}_{\int \to 0}) \, dx \\ &=^{\dagger} \exp(-\frac{u^2}{2\sigma_u^2}) \qquad \text{where} \quad \sigma_u = \frac{1}{2\pi\sigma_x} \end{aligned}$$

† [Abramowitz, Stegun: Handbook of Mathematical Functions, 1972]

⇒ the Fourier transform of a Gaussian is a (unnormalized) Gaussian!

The larger the variance σ_x^2, the smaller the variance σ_u^2: $\sigma_x \cdot \sigma_u = \frac{1}{2\pi}$

Fourier Transformation: Example 3 (Dirac's δ)

The Fourier transform of Dirac's δ-function is

$$\begin{aligned} \hat{\delta}(u) \equiv \mathcal{F}[\delta(x)] &= \int_{-\infty}^{+\infty} \delta(x) \exp(-\imath 2\pi u x) dx \\ &= \exp(-\imath 2\pi u \cdot 0) \\ &= 1 \end{aligned}$$

⇒ the Fourier transform of the δ-function equals 1 for *all* frequencies u.

Properties of the Fourier Transformation

Linearity: If $\mathcal{F}[f(x)] = \hat{f}(u)$ and $\mathcal{F}[g(x)] = \hat{g}(u)$ then it holds for all complex numbers $a, b \in \mathbb{C}$

$$\mathcal{F}[af(x) + bg(x)] = a\hat{f}(u) + b\hat{g}(u)$$

Shift: If $\mathcal{F}[f(x)] = \hat{f}(u)$ then it holds for $c \in \mathbb{R}$

$$\mathcal{F}[f(x - c)] = \hat{f}(u) \exp(-\imath 2\pi c u)$$

Modulation: If $\mathcal{F}[f(x)] = \hat{f}(u)$ then it holds for $c \in \mathbb{R}$

$$\mathcal{F}[f(x) \exp(\imath 2\pi c x)] = \hat{f}(u - c)$$

Scaling: If $\mathcal{F}[f(x)] = \hat{f}(u)$ and $c > 0$

$$\mathcal{F}[f(cx)] = \frac{1}{c}\hat{f}(\frac{u}{c})$$

Differentiation: Let f be piecewise continuous and absolutely integrable. If the function $xf(x)$ is absolutely integrable then the Fourier transform $\hat{f}$ is continuous and differentiable. It holds

$$\begin{aligned} \mathcal{F}[xf(x)] &= \frac{\imath}{2\pi}\frac{d}{du}\hat{f}(u) \\ \mathcal{F}[\frac{d}{dx}f(x)] &= \imath 2\pi u \hat{f}(u) \end{aligned}$$

Parseval's Equality: Let f be piecewise continuous and absolutely integrable. Then the Fourier transform $\hat{f}(u) = \mathcal{F}[f(x)]$ satisfies:

$$\int_{-\infty}^{\infty} |f(x)|^2 dx = \int_{-\infty}^{\infty} |\hat{f}(u)|^2 du$$

Power Spectrum: Considering the auto-correlation function $\Phi_{ff}(x)$ of a complex function f for $x \in \mathbb{R}$,

$$\Phi_{ff}(x) = \int_{-\infty}^{\infty} \bar{f}(\xi - x) f(\xi) d\xi \ .$$

The Fourier transform is given by

$$\hat{\Phi}_{ff}(u) \equiv \mathcal{F}[\Phi_{ff}(x)] = |\hat{f}(u)|^2.$$

($\bar{f}(x)$ is the conjugate complex function of $f(x)$)

Fourier Transform of Convolution

Given: convolution $g(x) = (f * h)(x) = \int f(\xi)h(x-\xi)d\xi$

Calculate Fourier transform of g:

$$\begin{aligned} \hat{g}(u) &\equiv \mathcal{F}[g(x)] = \int_{-\infty}^{+\infty} \left[\int_{-\infty}^{+\infty} f(\xi)h(x-\xi)d\xi\right] \exp(-\imath 2\pi ux)dx \\ &= \int_{-\infty}^{+\infty} f(\xi) \left[\int_{-\infty}^{+\infty} h(x-\xi)\exp(-\imath 2\pi ux)dx\right] d\xi \\ &= \int_{-\infty}^{+\infty} \hat{h}(u) f(\xi) \exp(-\imath 2\pi u\xi) d\xi \\ &= \hat{h}(u)\hat{f}(u) \end{aligned}$$

⇒ Convolution in spatial domain becomes multiplication in Fourier space.

Modulation Transfer Function

System Behavior in Fourier Space: How is a harmonic oscillation transformed by convolution kernel h?

⇒ amplitude modulation $A(u)$:

$\exp(\imath 2\pi ux)$ ⟶ [kernel $h(x)$] ⟶ $A(u)\exp(\imath 2\pi ux)$

Eigenfunction of the convolution with eigenvalue $A(u)$ is the oscillation $f(x) = \exp(\imath 2\pi ux)$.

$$\begin{aligned} \text{Output} \quad g(x) &= (f*h)(x) = \int \exp(\imath 2\pi u\xi)h(x-\xi)d\xi \\ &= \exp(\imath 2\pi ux)\int \exp(-\imath 2\pi u\xi)h(\xi)d\xi \\ &= \hat{h}(u)\exp(\imath 2\pi ux) \end{aligned}$$

Note: the eigenvalue $A(u)$ equals $\hat{h}(u) = \mathcal{F}[h](u)$.

Image Filtering in the Frequency Domain

2D Fourier transformation of an image $f(x, y)$:

$$\hat{f}(u,v) \equiv \mathcal{F}[f(x,y)] = \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} f(x,y)\exp(-\imath 2\pi(ux+vy))dxdy$$

High-pass filtering: remove low frequencies, for example choose maximum value B:

$$\hat{f}_{\text{hp}}(u,v) = \begin{cases} \hat{f}(u,v) & \text{if } u^2+v^2 > B^2 \\ 0 & \text{otherwise} \end{cases}$$

Inverse Fourier transformation yields high-pass-filtered image $f_{\text{hp}}(x,y) = \mathcal{F}^{-1}[\hat{f}_{\text{hp}}(u,v)]$

Example of Image Filtering


original image


high-pass-filtered

⇒ edge detection

Low-pass filtering: analogous to high-pass filter, but remove high frequencies

Example:

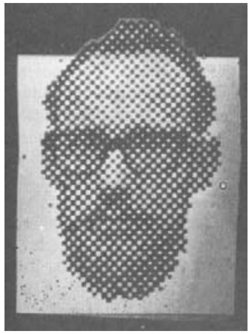
original image

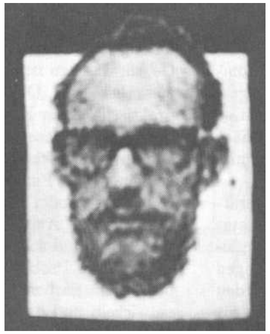
low-pass-filtered

⇒ removing noise

The Image Restoration Problem

$$f(x,y) \longrightarrow \boxed{h(x,y)} \longrightarrow g(x,y) \longrightarrow \boxed{\tilde{h}(x,y)} \longrightarrow f(x,y)$$

The 'inverse' kernel $\tilde{h}(x,y)$ should compensate the effect of the image degradation $h(x,y)$, i.e.,

$$(\tilde{h} * h)(x,y) = \delta(x,y)$$

$\tilde{h}$ may be determined more easily in Fourier space:

$$\mathcal{F}[\tilde{h}](u,v) \cdot \mathcal{F}[h](u,v) = 1$$

To determine $\mathcal{F}[\tilde{h}]$ we need to estimate

1. the distortion model $h(x,y)$ (point spread function) or $\mathcal{F}[h](u,v)$ (modulation transfer function)
2. the parameters of $h(x,y)$, e.g. r for defocussing.

Image Restoration: Example

Example: motion blur $h(x,y) = \frac{1}{2l}\big(\theta(x+l) - \theta(x-l)\big)\delta(y)$

(a light dot is transformed into a small line in x direction).

Fourier transformation:

$$\mathcal{F}[h](u,v) = \frac{1}{2l}\int_{-l}^{+l} \exp(-\imath 2\pi u x) \underbrace{\int_{-\infty}^{+\infty} \delta(y)\exp(-\imath 2\pi v y)dy}_{=1}\, dx$$

$$= \frac{\sin(2\pi u l)}{2\pi u l} =: \mathrm{sinc}(2\pi u l)$$

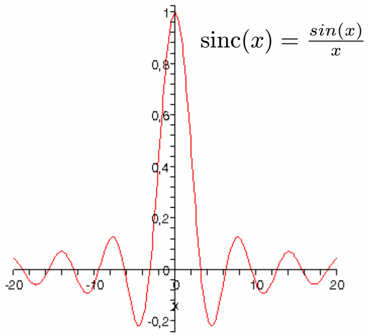

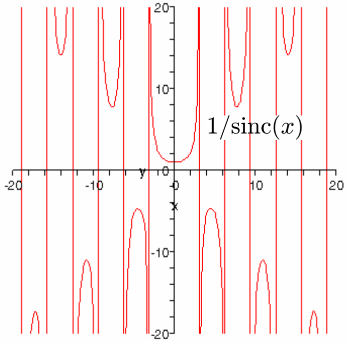

$\hat{h}(u) = \mathcal{F}[h](u) = \mathrm{sinc}(2\pi u l)$ $\mathcal{F}[\tilde{h}](u) = 1/\hat{h}(u)$

Problems:

- Convolution with the kernel h completely cancels the frequencies $\frac{\nu}{2l}$ for $\nu \in \mathcal{Z}$. Frequencies which disappear cannot be recovered!
- Noise amplification for $\mathcal{F}[h](u,v) \ll 1$.

Avoiding Noise Amplification

Regularized reconstruction filter:

$$\tilde{\mathcal{F}}[\tilde{h}](u,v) = \frac{\mathcal{F}[h]}{|\mathcal{F}[h]|^2 + \epsilon^2}$$

Singularities are avoided by the regularization ϵ^2.

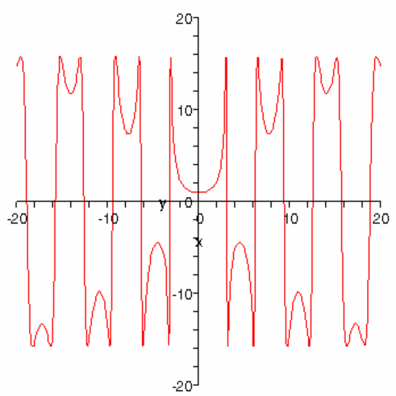

The Wiener Filter: Optimal Linear Filtering for Noise Suppression and Image Reconstruction

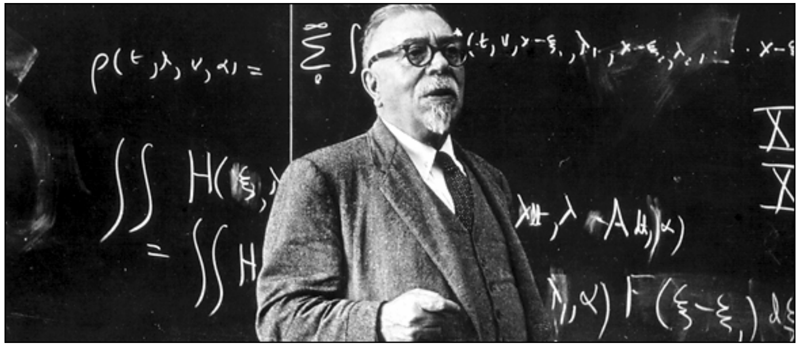

American mathematician who developed the theory of Brownian motion; $\Rightarrow$ Wiener measure, numerical PDE solutions, *linear filter theory*.

Cybernetics as the new science for systems design and control. Norbert Wiener broke new ground in robotics, computer control, and automation.

Optimal Linear Filtering and Noise Suppression

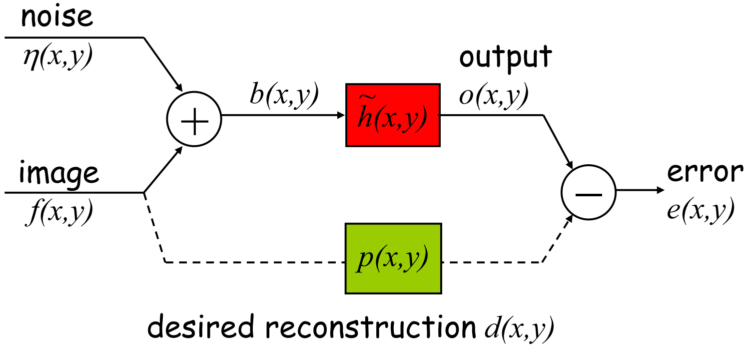

Given: $b(x,y) = f(x,y) + \eta(x,y)$

image = signal + noise

Goal: Reconstruct signal $f(x,y)$ as "good" as possible.

Noise model: we assume that the signal and the noise are uncorrelated, i.e. the cross-correlation is zero:

$$\Phi_{f\eta}(a,b) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \bar{f}(x-a, y-b)\eta(x,y)\, dx\, dy = 0.$$

Task: find $o(x,y)$ which reconstructs the original image $f(x,y)$ *as good as possible* from the observed image $b(x,y)$!
(setting $d(x,y) = f(x,y)$ with $p(x,y)$ being the identity map)

In some situations the desired reconstruction $d(x,y)$ might differ from $f(x,y)$ since we might prefer a smoothed or sharpened version (given by the transformation $p(x,y)$) of the original image.

Assumption: use a *linear* filter $\tilde{h}(x,y)$ for reconstruction, i.e.,

$$o(x,y) = (b * \tilde{h})(x,y).$$

Goal of the Wiener Filter

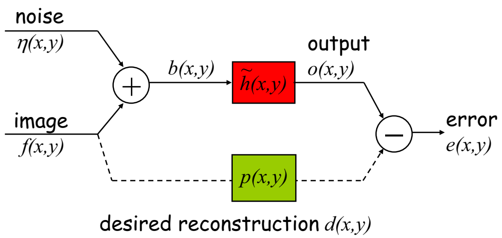

Quality measure for image restoration:
What does "as good as possible" actually mean?
⇒ **average quadratic error**

$$\text{error} = E = \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} (o(x,y) - d(x,y))^2\, dxdy$$

Goal: find the kernel $\tilde{h}$ that minimizes this error.

Derivation of the Wiener Filter

Error decomposition: $(\mathbf{x} := (x,y)^\top)$

$$E = \int_\Omega (o(\mathbf{x}) - d(\mathbf{x}))^2\, d\mathbf{x} = \int_\Omega \left(o^2 - 2od + d^2\right)\, d\mathbf{x}$$
$$= \underbrace{\int_\Omega o(\mathbf{x})^2\, d\mathbf{x}}_{(1)} - 2\underbrace{\int_\Omega o(\mathbf{x})d(\mathbf{x})\, d\mathbf{x}}_{(2)} + \underbrace{\int_\Omega d(\mathbf{x})^2\, d\mathbf{x}}_{(3)}$$

⇒ simplify each of the three integrals:

integral (3): $\int_\Omega d(\mathbf{x})^2\, d\mathbf{x} = \Phi_{dd}(0,0)$

where $\Phi_{dd}(0,0)$ is d's auto-correlation with no displacement.

integral (2): inserting $o(\mathbf{x}) = (b * \tilde{h})(\mathbf{x})$ yields

$$\begin{aligned}\int_\Omega o(\mathbf{x})d(\mathbf{x})d\mathbf{x} &= \int_\Omega \left[\int_\Omega b(\mathbf{x}-\boldsymbol{\xi})\tilde{h}(\boldsymbol{\xi})\, d\boldsymbol{\xi}\right] d(\mathbf{x})\, d\mathbf{x} \\ &= \int_\Omega \underbrace{\left[\int_\Omega b(\mathbf{x}-\boldsymbol{\xi})d(\mathbf{x})\, d\mathbf{x}\right]}_{\Phi_{bd}(\boldsymbol{\xi})} \tilde{h}(\boldsymbol{\xi})d\boldsymbol{\xi} \\ &= \int_\Omega \Phi_{bd}(\boldsymbol{\xi})\tilde{h}(\boldsymbol{\xi})d\boldsymbol{\xi}\end{aligned}$$

where $\Phi_{bd}(\boldsymbol{\xi})$ is the cross-correlation of b and d with displacement $\boldsymbol{\xi} = (\xi_1, \xi_2)^\top$.

integral (1): inserting $o(\mathbf{x}) = (b * \tilde{h})(\mathbf{x})$ yields

$$\begin{aligned}\int_\Omega o^2 d\mathbf{x} &= \int_\Omega \left((b*\tilde{h})(\mathbf{x})\right)^2 d\mathbf{x} \\ &= \int_\Omega \left(\int_\Omega\int_\Omega b(\mathbf{x}-\boldsymbol{\xi})b(\mathbf{x}-\boldsymbol{\alpha})\tilde{h}(\boldsymbol{\xi})\tilde{h}(\boldsymbol{\alpha})\, d\boldsymbol{\alpha}\, d\boldsymbol{\xi}\right) d\mathbf{x} \\ &\overset{\mathbf{x}'=\mathbf{x}-\boldsymbol{\alpha}}{=} \int_\Omega\int_\Omega \underbrace{\left[\int_\Omega b(\mathbf{x}'-\boldsymbol{\xi}+\boldsymbol{\alpha})b(\mathbf{x}')\, d\mathbf{x}'\right]}_{\Phi_{bb}(\boldsymbol{\xi}-\boldsymbol{\alpha})} \tilde{h}(\boldsymbol{\xi})\tilde{h}(\boldsymbol{\alpha})\, d\boldsymbol{\alpha}\, d\boldsymbol{\xi} \\ &= \int_\Omega\int_\Omega \Phi_{bb}(\boldsymbol{\xi}-\boldsymbol{\alpha})\tilde{h}(\boldsymbol{\xi})\tilde{h}(\boldsymbol{\alpha})\, d\boldsymbol{\alpha}\, d\boldsymbol{\xi}\end{aligned}$$

The first term in the average error defines a quadratic form of the auto-correlation $\Phi_{bb}(\boldsymbol{\xi}-\boldsymbol{\alpha})$ with $\tilde{h}$.

Wiener Filter Defined by Correlations

The average quadratic error can now be rewritten in terms of various auto/cross correlations:

$$\begin{aligned}E(\tilde{h}) &= \underbrace{\int_\Omega\int_\Omega \Phi_{bb}(\boldsymbol{\xi}-\boldsymbol{\alpha})\tilde{h}(\boldsymbol{\xi})\tilde{h}(\boldsymbol{\alpha})\, d\boldsymbol{\alpha}\, d\boldsymbol{\xi}}_{(1)} \\ &\quad -2\underbrace{\int_\Omega \Phi_{bd}(\boldsymbol{\xi})\tilde{h}(\boldsymbol{\xi})\, d\boldsymbol{\xi}}_{(2)} + \underbrace{\Phi_{dd}(0,0)}_{(3)}\end{aligned}$$

To minimize $E(\tilde{h})$ w.r.t. the reconstructing filter $\tilde{h}$ is a problem of variational calculus, e.g., $\min_{\tilde{h}} \int f(\tilde{h}(\mathbf{x}))d\mathbf{x}$.

Variation of the Wiener Filter

Next, we find the filter $\tilde{h}$ that minimizes the error function $E(\tilde{h})$, using the *variational calculus*:

- we assume that the kernel $\tilde{h}(x,y)$ minimizes $E(\tilde{h})$.
- we choose an *arbitrary* function $\delta\tilde{h}(x,y)$;
  ($\delta\tilde{h}(x,y) = 0$ on the boundary of the image)
- then $\tilde{h}(x,y) + \epsilon \cdot \delta\tilde{h}(x,y)$ is also a valid kernel ($\epsilon \geq 0$).
- **Minimality Condition:** since $\tilde{h}(x,y)$ minimizes $E(\tilde{h})$, it has to be a minimum of $E(\tilde{h})$ with the condition:

$$\frac{\partial}{\partial\epsilon} E(\tilde{h} + \epsilon\cdot\delta\tilde{h})\Big|_{\epsilon=0} = 0 \quad \forall\, \delta\tilde{h}(x,y) \in \mathcal{C}^0$$

Replace $\tilde{h}$ by $\tilde{h} + \epsilon \cdot \delta\tilde{h}$ to obtain $E(\tilde{h} + \epsilon \cdot \delta\tilde{h})$:

$$
\begin{aligned}
E(\tilde{h} + \epsilon \cdot \delta\tilde{h}) &= \int_\Omega\int_\Omega \Phi_{bb}(\boldsymbol{\xi} - \boldsymbol{\alpha})\Big(\tilde{h}(\boldsymbol{\xi}) + \epsilon\delta\tilde{h}(\boldsymbol{\xi})\Big) \times \\
&\qquad \Big(\tilde{h}(\boldsymbol{\alpha}) + \epsilon\delta\tilde{h}(\boldsymbol{\alpha})\Big)\, d\boldsymbol{\alpha}\, d\boldsymbol{\xi} \\
&\quad -2\int_\Omega \Phi_{bd}(\boldsymbol{\xi})\Big(\tilde{h}(\boldsymbol{\xi}) + \epsilon\delta\tilde{h}(\boldsymbol{\xi})\Big)\, d\boldsymbol{\xi} + \Phi_{dd}(0,0) \\
&= E(\tilde{h}) + 2\epsilon \int_\Omega \underbrace{\int_\Omega \Phi_{bb}(\boldsymbol{\xi} - \boldsymbol{\alpha})\tilde{h}(\boldsymbol{\alpha})\, d\boldsymbol{\alpha}}_{(\Phi_{bb}*\tilde{h})(\boldsymbol{\xi})}\, \delta\tilde{h}(\boldsymbol{\xi})\, d\boldsymbol{\xi} \\
&\quad - 2\epsilon \int \Phi_{bd}(\boldsymbol{\xi})\, \delta\tilde{h}(\boldsymbol{\xi}, \eta)\, d\boldsymbol{\xi} + \mathcal{O}(\epsilon^2)
\end{aligned}
$$

$$
\Rightarrow \quad \frac{\partial}{\partial\epsilon} E(\tilde{h} + \epsilon \cdot \delta\tilde{h})\Big|_{\epsilon=0} = -2\int_\Omega \Big(\Phi_{bd}(\boldsymbol{\xi}) - (\Phi_{bb}*\tilde{h})(\boldsymbol{\xi})\Big)\delta\tilde{h}(\boldsymbol{\xi})\, d\boldsymbol{\xi}
$$

Since $\delta\tilde{h}(\boldsymbol{\xi})$ is an arbitrary function, the equation

$$
\frac{\partial}{\partial\epsilon} E(\tilde{h} + \epsilon \cdot \delta\tilde{h})\Big|_{\epsilon=0} = 0
$$

requires the integrand $(\Phi_{bd} - \Phi_{bb} * \tilde{h})$ to vanish for all values $\mathbf{x} = (x, y)^\top$ (fundamental theorem of variational calculus):

$$
\Phi_{bd}(\mathbf{x}) = (\Phi_{bb} * \tilde{h})(\mathbf{x}) \quad \text{Wiener-Hopf equation}
$$

The convolution kernel (point spread function) $\tilde{h}(\mathbf{x})$ of the optimal linear filter has to satisfy the Wiener-Hopf equation.

Fourier Analysis of the Wiener Filter

In Fourier space the Wiener-Hopf equation yields:
($\hat{f} := \mathcal{F}[f]$ denotes the Fourier transform)

$$
\begin{aligned}
\hat{\Phi}_{bd} &= \hat{\Phi}_{bb} \cdot \mathcal{F}[\tilde{h}] \\
\mathcal{F}[\tilde{h}](u,v) &= \frac{\hat{\Phi}_{bd}(u,v)}{\hat{\Phi}_{bb}(u,v)} = \frac{\hat{\Phi}_{fd}(u,v)}{\hat{\Phi}_{ff}(u,v) + \hat{\Phi}_{\eta\eta}(u,v)}
\end{aligned}
$$

The last equality holds because we assumed that

- $b(x,y) = f(x,y) + \eta(x,y)$,
- the noise η is **not** correlated with the signal f: $\Phi_{f\eta}(x,y) = 0$ for all x, y.
  $\Rightarrow \Phi_{bb} = \Phi_{f+\eta,f+\eta} = \Phi_{ff} + \underbrace{\Phi_{f\eta}}_{=0} + \underbrace{\Phi_{\eta f}}_{=0} + \Phi_{\eta\eta}$

Wiener Filter: Improving a Noisy Image

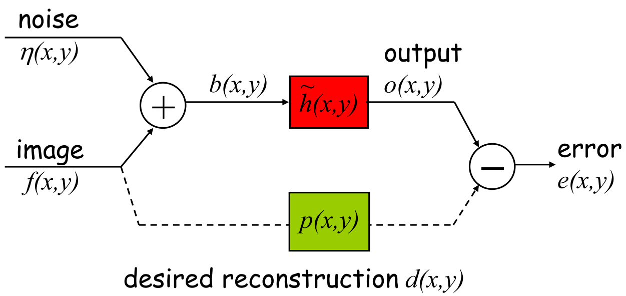

If $d = f$, the Fourier transform of the optimal linear filter for the (unknown) original signal f is

$$
\mathcal{F}[\tilde{h}](u,v) = \frac{\hat{\Phi}_{ff}(u,v)}{\hat{\Phi}_{ff}(u,v) + \hat{\Phi}_{\eta\eta}(u,v)} = \frac{1}{1 + \frac{\hat{\Phi}_{\eta\eta}(u,v)}{\hat{\Phi}_{ff}(u,v)}}
$$

Signal-to-Noise Ratio

Definition: the ratio

$$\text{SNR}(u,v) = \frac{\hat{\Phi}_{ff}(u,v)}{\hat{\Phi}_{\eta\eta}(u,v)}$$

is called the *signal-to-noise ratio* (at the frequencies (u,v)).

SNR(u,v) **large:** the filter behaves almost like the identity map.

SNR(u,v) **small:** the filter is proportional to the SNR.
$\Rightarrow$ damping.

Statistics of Natural Images

Observation [Fields, 1987]: the power spectrum of natural images $f(x,v)$ decays as

$$\begin{aligned}\hat{\Phi}_{ff}(u,v) &= \hat{\Phi}_{ff}(\rho,\theta) \propto \frac{1}{\rho^2}\\ \Rightarrow \hat{\Phi}_{ff}(\rho) &= \int \hat{\Phi}_{ff}(\rho,\theta)\rho\, d\theta \propto \int \frac{1}{\rho^2}\rho\, d\theta \propto \frac{1}{\rho}\end{aligned}$$

Note: in the Fourier space, the polar coordinates ρ, θ are used in place of the Cartesian coordinates u, v (frequencies).

Assumption concerning noise: the noise is spatially uncorrelated, i.e.,

$$\Phi_{\eta\eta}(x,y) = \Phi_0 \cdot \delta(x,y)$$

$$\begin{aligned}\Rightarrow \text{ Fourier transform: } \quad \hat{\Phi}_{\eta\eta}(u,v) &= \Phi_0\\ \Rightarrow \text{ polar coordinates: } \quad \hat{\Phi}_{\eta\eta}(\rho) &= \int \hat{\Phi}_{\eta\eta}(\rho,\theta)\rho\, d\theta\\ &= \int \Phi_0 \rho\, d\theta \propto \Phi_0 \cdot \rho\end{aligned}$$

Wiener Filter: Improving Noisy Natural Images

natural image:

- power spectrum: $\hat{\Phi}_{ff}(\rho) \propto \frac{1}{\rho}$
- noise: $\hat{\Phi}_{\eta\eta}(\rho) \propto \Phi_0 \cdot \rho$

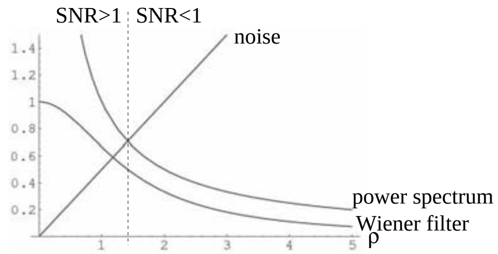

Wiener filter: $\mathcal{F}[\tilde{h}](\rho) = \left(1 + \frac{\hat{\Phi}_{\eta\eta}(\rho)}{\hat{\Phi}_{ff}(\rho)}\right)^{-1}$

Two limiting cases:

- SNR $\gg 1 \quad \Rightarrow \quad \mathcal{F}[\tilde{h}] \approx 1$
  ... no modulation of the low frequencies

- SNR $\ll 1 \quad \Rightarrow \quad \mathcal{F}[\tilde{h}] \approx \hat{\Phi}_{ff}/\hat{\Phi}_{\eta\eta} \ \propto \ 1/\rho^2$
  ... damping of the high frequencies

Optimal Linear Filtering for Image Reconstruction with Simultaneous Noise Suppression

Assumption: There exists a "degradation kernel" h which has transformed the image before noise perturbation!

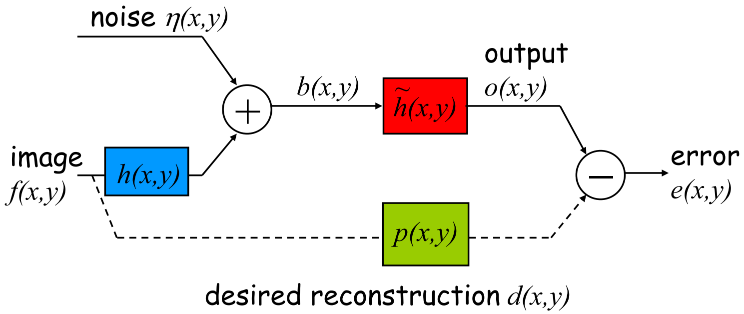

Given:

$b(x,y)$	=	$(f*h)(x,y)$	+	$\eta(x,y)$
image	=	signal f & degradation h	+	noise

Noise: we assume that the signal and the noise are uncorrelated: $\Phi_{f\eta} = 0$.

Task: reconstruct $f(x,y)$ as good as possible from $b(x,y)$!

Assumption: use a *linear* filter $\tilde{h}(x,y)$ to compensate the degradation and to filter out the noise, i.e.,

$$o(x,y) = (b * \tilde{h})(x,y)$$

Derivation of Reconstruction Wiener Filter

Autocorrelation of image $b(x,y)$:

$$\begin{aligned}\hat{\Phi}_{bb} &= \hat{\Phi}_{f*h+\eta, f*h+\eta} = \hat{\Phi}_{f*h,f*h} + 2\hat{\Phi}_{f*h,\eta} + \hat{\Phi}_{\eta\eta} \\ &= \hat{h}^2\hat{\Phi}_{ff} + \hat{h}\underbrace{\hat{\Phi}_{f\eta}}_{=0} + \hat{h}\underbrace{\hat{\Phi}_{\eta f}}_{=0} + \hat{\Phi}_{\eta\eta},\end{aligned}$$

since a correlation of a convolution $f*h$ with a function g is the convolution of the correlation $f*g$ with the kernel h.

Result in Fourier space:

$$\hat{\Phi}_{bd} = \hat{\Phi}_{bb} \cdot \mathcal{F}[\tilde{h}] \quad \text{... as before}$$

$$\mathcal{F}[\tilde{h}](u,v) = \frac{\hat{\Phi}_{bd}(u,v)}{\hat{\Phi}_{bb}(u,v)} = \frac{\hat{h}(u,v)\cdot\hat{\Phi}_{fd}(u,v)}{\hat{h}^2(u,v)\cdot\hat{\Phi}_{ff}(u,v) + \hat{\Phi}_{\eta\eta}(u,v)}$$

Assumption: $d(x,y) = f(x,y)$, i.e., the desired image is the original one:

$$\mathcal{F}[\tilde{h}](u,v) = \frac{\hat{h}(u,v)}{\hat{h}^2(u,v) + \underbrace{\frac{\hat{\Phi}_{\eta\eta}(u,v)}{\hat{\Phi}_{ff}(u,v)}}_{=1/\mathrm{SNR}(u,v)}}$$

Note: this filter corresponds to the heuristic regularization for avoiding noise amplification (slide 39): $\epsilon^2 = 1/\mathrm{SNR}(u,v)$

Two limiting cases:

- SNR $\gg 1 \;\Rightarrow\; \mathcal{F}[\tilde{h}](u,v) \approx \frac{1}{\hat{h}(u,v)}$
  ... cf. direct derivation of image restoration kernel (slide 36)
- SNR $\ll 1 \;\Rightarrow\; \mathcal{F}[\tilde{h}] \approx \hat{h}(u,v)\hat{\Phi}_{ff}/\hat{\Phi}_{\eta\eta} \propto 1/\rho^2$
  ... in natural images $\Rightarrow$ damping of high frequencies.

Wiener Filter: Sharpening and Denoising Natural Images

Assume: blurring with Gaussian kernel (in polar coordinates):

$$h(r) \propto \exp\left(-\frac{r^2}{2\sigma_r^2}\right)$$

$$\Rightarrow \quad \hat{h}(\rho) \propto \exp\left(-\frac{\rho^2}{2\sigma_\rho^2}\right)$$

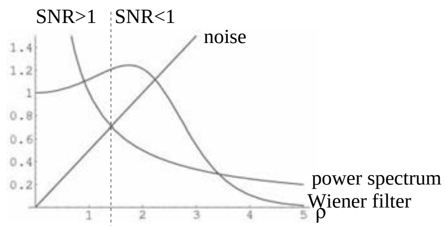

$\Rightarrow$ **Wiener filter:**

$$\mathcal{F}[\tilde{h}](\rho) \propto \frac{\exp\left(-\frac{\rho^2}{2\sigma_\rho^2}\right)}{\left(\exp\left(-\frac{\rho^2}{2\sigma_\rho^2}\right)\right)^2 + \underbrace{\frac{\hat{\Phi}_{\eta\eta}(\rho)}{\hat{\Phi}_{ff}(\rho)}}_{\propto \hat{\Phi}_0 \cdot \rho^2 \text{ in natural images}}}$$

Lena: original & noisy (PSNR=7.2)

Mean Square Error: $\mathrm{MSE}(f,g) = \frac{1}{\Omega}\int_\Omega (f(\mathbf{x}) - g(\mathbf{x}))^2 d\mathbf{x}$

peak SNR: $\mathrm{PSNR}(f,g) := 20\log_{10}\left(\frac{255}{\sqrt{\mathrm{MSE}}}\right)$

Optimal Linear Filter (Lena PSNR=7.2)

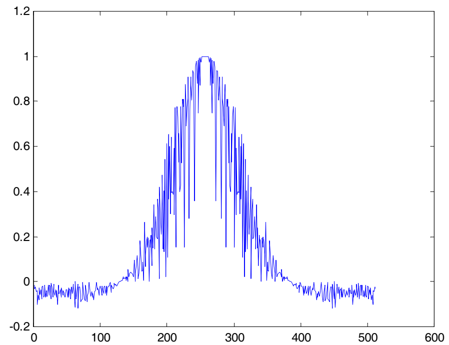

Lena: noisy image & reconstruction

The Lena image has been blurred with a Gaussian kernel of $\sigma_{\text{kernel}} = 3$ and it has been degraded with Gaussian noise ($\sigma_{\text{noise}} = 30$).
Image quality: PSNR=7.2

Peak SNR of reconstruction: PSNR=24.7

Reconstruction Filter