Introduction to Fourier Analysis

Fourier Analysis:

A basic tool with many uses, including the analysis of image sampling, windowing, and filtering.

Suppose $s(t)$ is a bounded, complex-valued signal for $t \in [0,1]$. Then $s(t)$ is represented as a sum of basis functions:

$$s(t) =_{a.e.} \sum_{n=-\infty}^{\infty} a_n b_n(t). \qquad \text{(Fourier Series)}$$

Fourier basis:

$$b_n(t) \equiv e^{i2\pi nt} = cos(2\pi nt) + isin(2\pi nt), n \in \mathcal{Z}.$$

Properties of the Fourier Basis:

- **Orthogonality:**

$$\begin{aligned}\langle b_m(t), b_n(t)\rangle &\equiv \int_0^1 b_m^*(t) b_n(t) dt \\ &= \int_0^1 e^{-i2\pi mt} e^{i2\pi nt} dt = \int_0^1 e^{i2\pi(n-m)t} dt \\ &= \delta_{n,m}.\end{aligned}$$

- **Periodicity:**

$$b_n(t) = b_n(t+1).$$

The signal $s(t)$ has been (effectively) periodically extended for t beyond $[0,1]$.

- **Shift Invariance:**

$$b_n(t+c) = \lambda(c) b_n(t),$$

with $\lambda(c) = e^{i2\pi nc}$.

Fourier Transform

Fourier Series:

$$s(t) = \sum_{n=-\infty}^{\infty} a_n b_n(t).$$

(Finite Interval) Fourier Transform: $\{a_n\}_{n=-\infty}^{\infty}$

$$a_n = \langle b_n(t), s(t) \rangle \equiv \int_0^1 b_n^*(t) s(t) dt.$$

Note the Fourier transform of $s(t+c)$ is given by:

$$\begin{aligned} a_n(c) &= \langle b_n(t), s(t+c) \rangle, \\ &= \langle b_n(t-c), s(t) \rangle, && \text{(by periodicitiy)} \\ &= e^{i2\pi nc} \langle b_n(t), s(t) \rangle, && \text{(by shift invariance)} \\ &= e^{i2\pi nc} a_n(0). \end{aligned}$$

We find that the *amplitude spectrum*, namely $\{|a_n|\}_{n=-\infty}^{\infty}$, is shift invariant, but the *phase spectrum*, $\{\angle a_n\}_{n=-\infty}^{\infty}$, is not.

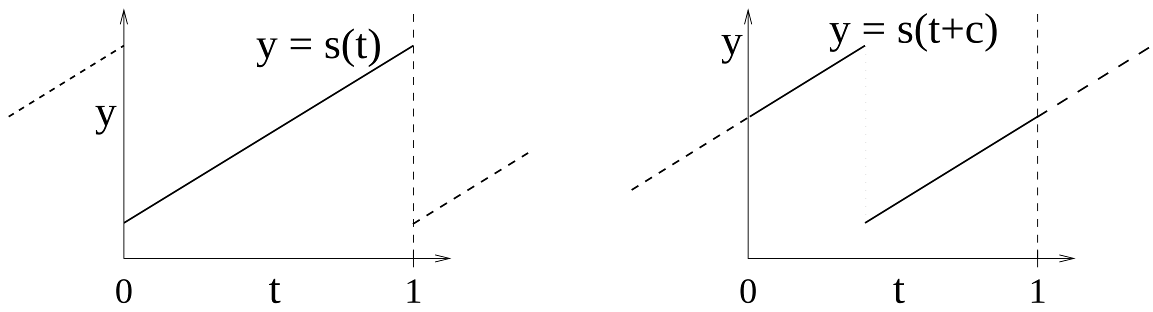

Fourier Basis Functions in 2D

For $\vec{x} \equiv (x, y) \in [0, 1] \times [0, 1]$, and any integers n, m, define:

$$b_{m,n}(x, y) \equiv e^{i2\pi mx} e^{i2\pi ny} = cos(\vec{k}_{m,n} \cdot \vec{x}) + isin(\vec{k}_{m,n} \cdot \vec{x}),$$

Here $\vec{k}_{m,n} \equiv 2\pi(m, n)^T$.

Grating for (m,n) = (1,-3)

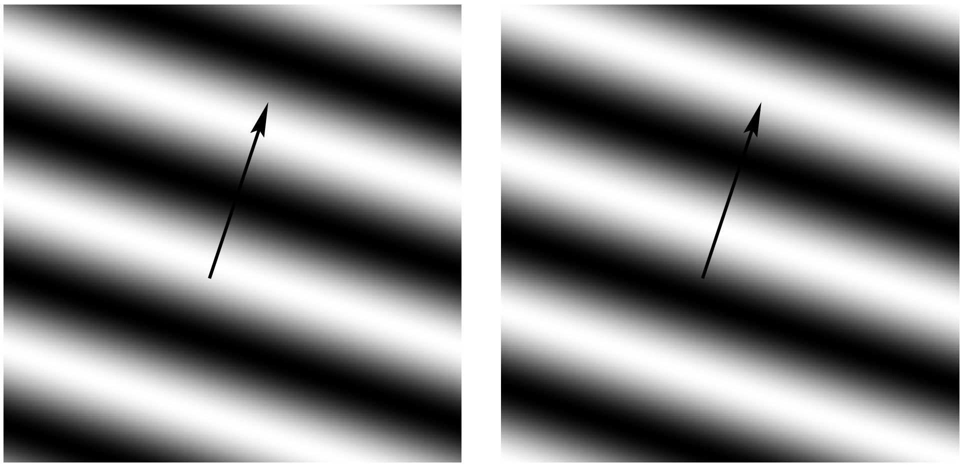

Real Imag

Grating for (m,n) = (7,1)

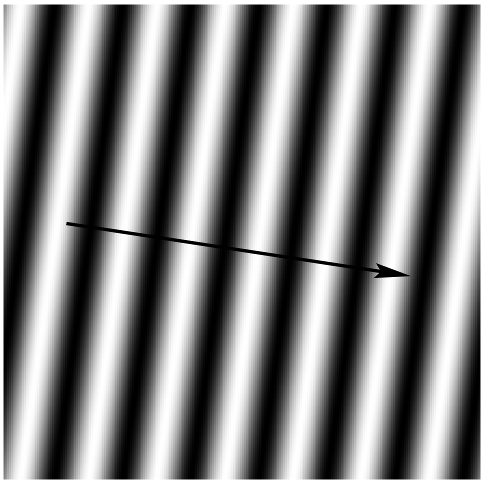

Real

Fourier Transform in 2D

Consider the 2D signal $s(x, y) \in \mathcal{C}$ for $(x, y) \in [0, 1] \times [0, 1]$.

The 2D basis functions are orthonormal, that is,

$$\langle b_{m,n}(\vec{x}), b_{k,j}(\vec{x}) \rangle \equiv \int_0^1 \int_0^1 b^*_{m,n}(x, y) b_{k,j}(x, y) dx dy$$
$$= \delta_{m,k} \delta_{n,j}.$$

The 2D Fourier transform is given by

$$a_{m,n} = \langle b_{m,n}(\vec{x}), s(\vec{x}) \rangle.$$

Example

A blocks image and it's amplitude spectrum (i.e. $|a_{m,n}|$):

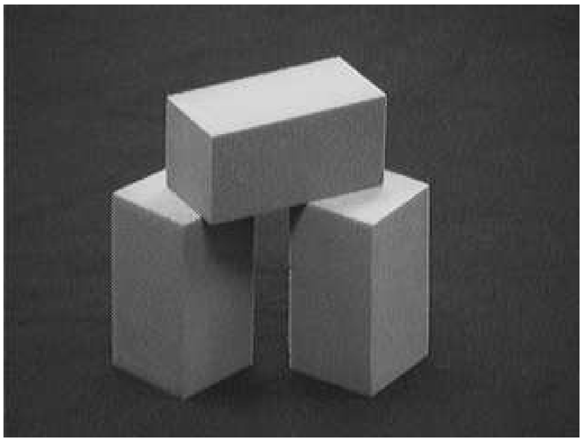

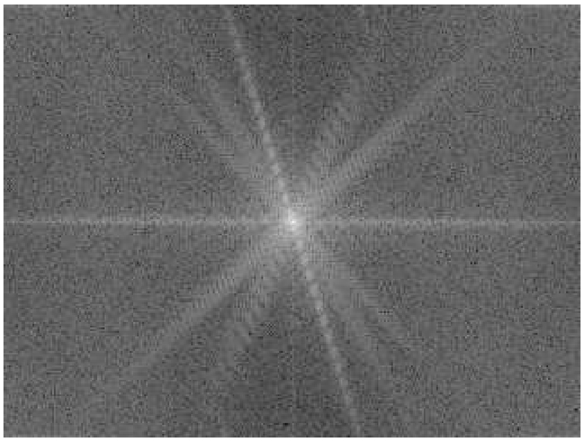

Note the amplitude spectrum satisfies $|a_{m,n}| = |a_{-m,-n}|$ for any real-valued input.

Sampling Images

Consider sampling the continuous distribution of irradiance falling on the image plane.

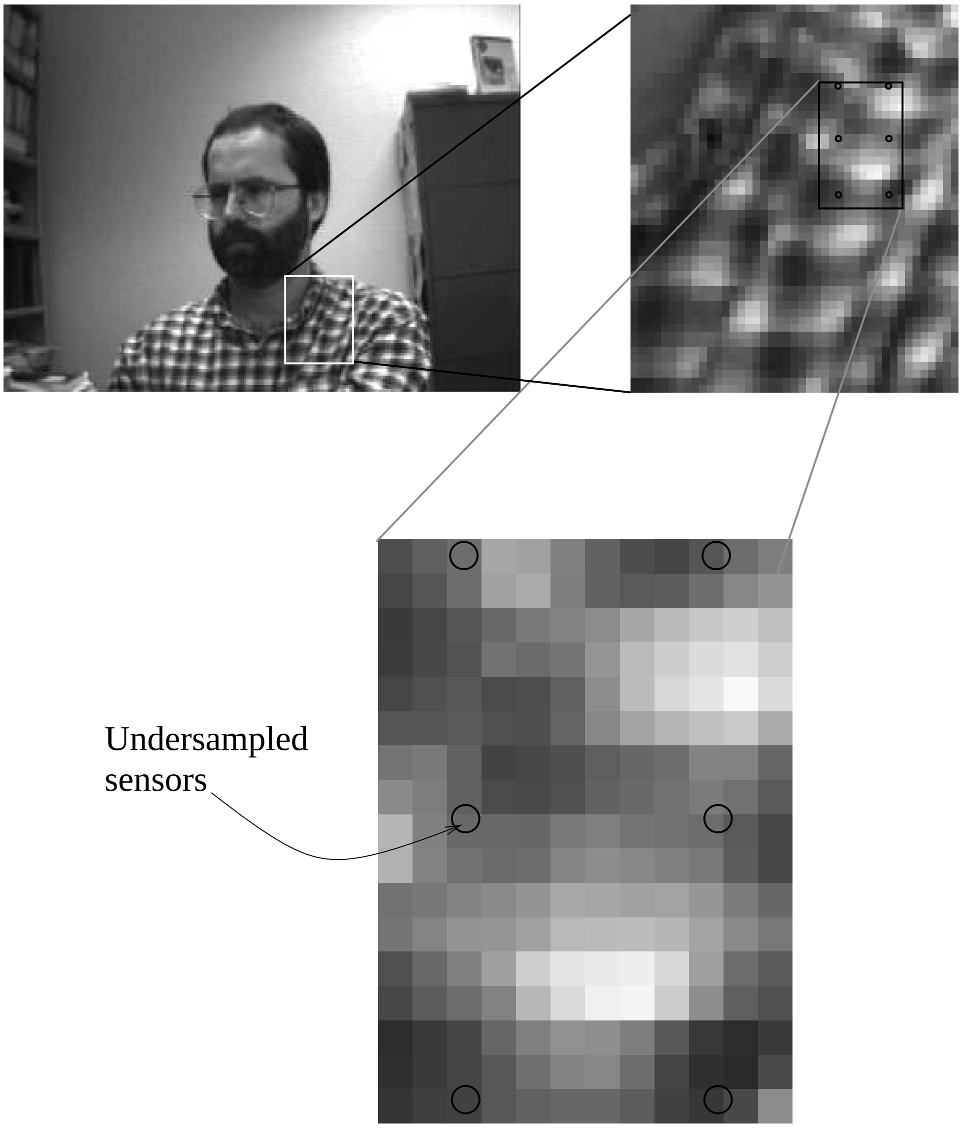

If we were given only the sparse samples, could you approximately reconstruct the original signal?

Sampling Continuous Signals

Consider the simple 1D signal $s(t)$, for $t \in [0, 1]$.

Sample Points:

$$t_j = j/N, \quad \text{for } j \in \{0, \ldots, N-1\}.$$

Discrete Sampling:

$$\begin{aligned}\vec{s} \equiv& (s(t_0), s(t_1), \ldots, s(t_{N-1}))^T = D_N s(t)\\ \vec{b}_n \equiv& (b(t_0), b(t_1), \ldots, b(t_{N-1}))^T = D_N b(t).\end{aligned}$$

Now, from the Fourier series for $s(t)$, we find

$$\begin{aligned}\vec{s} =& D_N s(t) = D_N \sum_{n=-\infty}^{\infty} a_n b_n(t) = \sum_{n=-\infty}^{\infty} a_n D_N b_n(t)\\ =& \sum_{n=-\infty}^{\infty} a_n \vec{b}_n.\end{aligned}$$

Aliasing:

- Note $\vec{b}_n \in \mathcal{C}^N$, so at most N of these can be linearly independent.
- Any maximal linearly independent set can serve as a basis.
- Often choose $n \in (-N/2, N/2]$ (low-pass) for the basis.
- The other sampled basis functions, say $\vec{b}_m$, must be linear combinations of the chosen basis, say $\{\vec{b}_n\}_{n=-\lfloor (N-1)/2 \rfloor}^{\lfloor N/2 \rfloor}$. These other $\vec{b}_m$'s, not included in the basis, are said to be *aliased.*

Details of Aliasing

Sampled Fourier Basis Functions:

$$\vec{b}_n = D_N b_n(t) == (b_n(t_0), b_n(t_1), \ldots, b_n(t_{N-1}))^T.$$

Note that

$$\begin{aligned} b_n(t_j) =& e^{i2\pi n t_j} = e^{i2\pi nj/N} e^{i(2\pi n/N)j} \\ =& e^{i\omega_n j} = (z_n)^j. \end{aligned}$$

Here $\omega_n = 2\pi n/N$ is called the *frequency* of $b_n(t)$. Also,

$$\begin{aligned} z_n =& z_n e^{i2\pi k}, \qquad\qquad (\text{since } e^{i2\pi} = 1) \\ =& e^{i2\pi(n+kN)/N} = z_{n+kN}. \end{aligned}$$

Therefore, for all integers n and k, $\vec{b}_n = \vec{b}_{n+kN}$. The Fourier series becomes

$$\vec{s} = \sum_{n=-\lfloor (N-1)/2 \rfloor}^{\lfloor N/2 \rfloor} \left[\sum_{k=-\infty}^{\infty} a_{n+kN} \right] \vec{b}_n = \sum_{n=-\lfloor (N-1)/2 \rfloor}^{\lfloor N/2 \rfloor} \hat{a}_n \vec{b}_n.$$

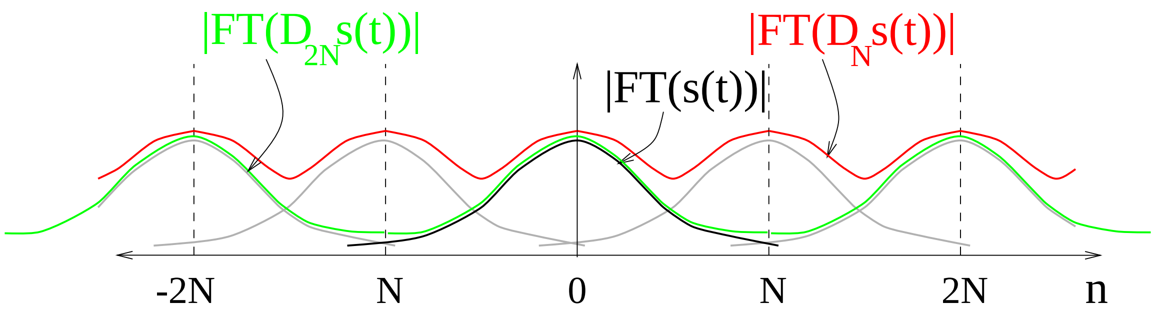

Aliasing *destroys* information about the original continuous signal whenever there are significant nonzero coefficients, a_n, for any frequency larger than the Nyquist *frequency* $\omega_{N/2} = \pi$.

Example of Aliasing

Consider the perspective image of an infinite checkerboard. The signal is dominated by high frequency components for image pixels near the horizon.

A properly designed camera blurs the continuous signal before sampling. This is a consequence of:

- the PSF due to diffraction (i.e. the Airy disk),
- imperfect focus,
- averaging the signal over each CCD element.

These operations all attenuate high frequency components in the signal (effectively by spatial averaging).

Without this (physical) preprocessing, the sampled image can be severely aliased:

Aliased perspective image of a checker board.

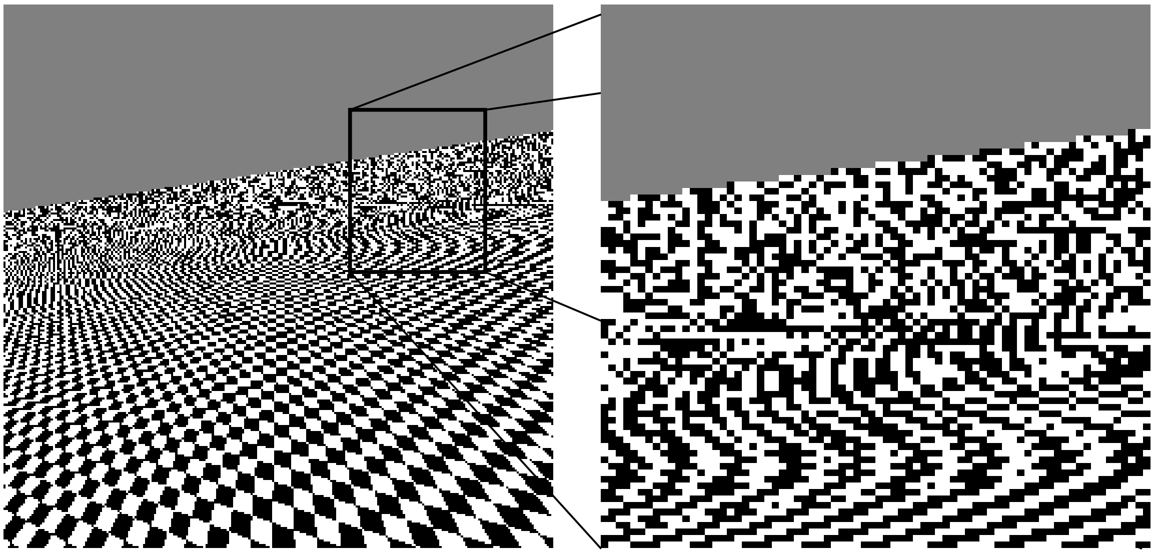

Note: Aliasing must also be controlled while warping discretely sampled images, a process which is often used in stereo and motion estimation algorithms.

Discrete Fourier Transform

Consider the discretely sampled signal, $\vec{s} = (s_0, s_1, \ldots, s_{N-1})^T \in \mathcal{C}^N$. The Fourier basis vectors are

$$\vec{b}_n = (1, z_n, \ldots, z_n^{N-1})^T, \quad z_n = e^{i2\pi n/N}.$$

These vectors can be collected into the matrix

$$B = \left(\vec{b}_{-n_0}, \ldots, \vec{b}_{-1}, \vec{b}_0, \vec{b}_1, \ldots, \vec{b}_{n_1}\right) \in \mathcal{C}^{N \times N},$$

where $n_0 = \lfloor (N-1)/2 \rfloor$ and $n_1 = \lfloor N/2 \rfloor$. It then follows that $B^* B = NI$, so the matrix $\frac{1}{\sqrt{N}} B$ is unitary.

Fourier Transform: The transform of $\vec{s}$ is

$$\vec{a} = B^* \vec{s}.$$

Inverse Fourier Transform: The original signal, $\vec{s}$, is recovered from the Fourier transform, $\vec{a}$, using

$$\vec{s} = \frac{1}{N} B \vec{a}.$$

The Fast Fourier Transform can evaluate both the forward and inverse transforms in $O(N \log N)$ time.

Real-Valued Signals

Suppose $\vec{s} = (s_0, s_1, \ldots, s_{N-1})^T \in \Re^N$. Then $\vec{s}$ is the same as the complex conjugate, $\bar{\vec{s}}$. By construction $\vec{b}_n = \bar{\vec{b}}_{-n}$. And it follows by the uniqueness of the Fourier transform that

$$a_n = \bar{a}_{-n}.$$

In particular, for real-valued signals:

- There are N degrees of freedom in the Fourier transform representation (same as for the signal vector).
- The transform and the inverse can both be done in real-valued arithmetic using sines and cosines in place of complex exponentials.
- The corresponding matrix, B, for the transform is orthogonal, up to a scale factor, (i.e. $B^T B = NI$).

Discrete Cosine Transform

Another variation is the Discrete Cosine Transform. Here, instead of assuming the signal is periodic with period N, it is extended by successively reversing the signal from left to right so that it has the period $2N$. This avoids the large step discontinuities at the borders, which can dominate the Fourier transform.

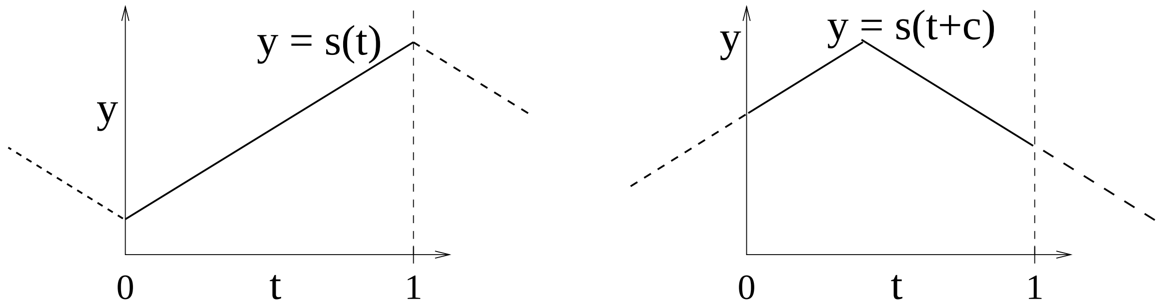

Here the signal is twice as long, but it is flip symmetric around the middle. As a consequence:

- The anti-symmetric sine terms cannot contribute, and only the cosine terms remain in the Fourier transform.
- The matrix for the transform is orthogonal and can be evaluated in $O(N \log N)$ time.
- The advantage of using the DCT is that the artifacts due to the borders of the image are decreased.
- JPEG compression uses the DCT within small image regions.

Example: Discrete Cosine Transform

The standard Fourier transform provides an amplitude spectrum consistent with the periodic signal:

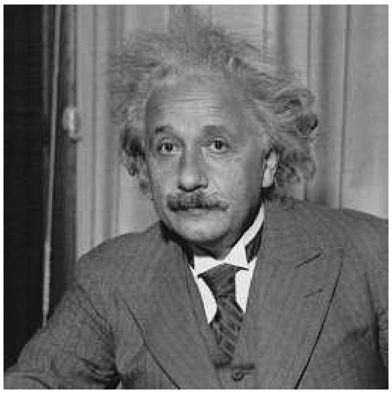
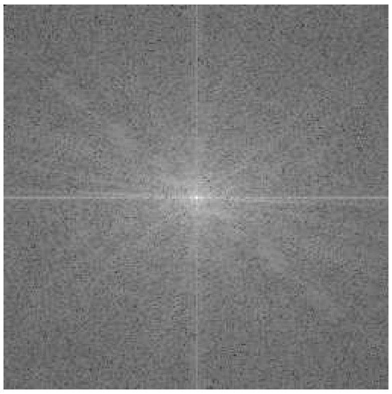
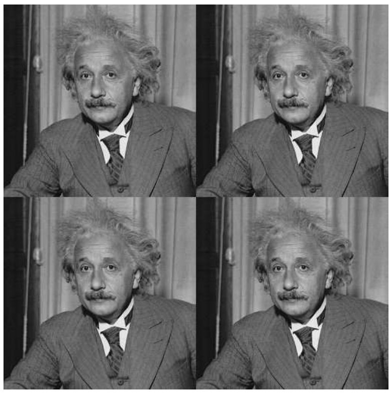

Note the significant response for horizontal and vertical gratings, which are artifacts due to the image borders.

The DCT transform tiles the image in a doubly periodic fashion which helps to reduce the large steps around the borders.

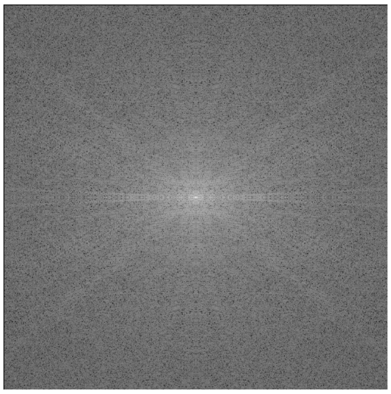
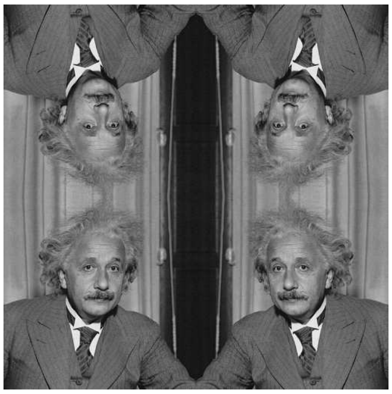

The DCT requires only one quadrant in this response. (We show all four for comparison purposes.) Note the responses for horizontal and vertical gratings have been signifcantly reduced.