Sampling and Filtering Tutorial: Sept 27, '04

Sampling and Filtering Images — representing a continuous signal with a discrete set of values taken at a particular interval.

References:

- linSysTutorial.m
- samplingTutorial.m
- imageTutorial.m
- ~jepson/pub/matlab/iseToolbox/tutorials/
- Your lecture notes!
- http://www.cis.rit.edu/people/faculty/montag/vandplite/pages/chap_9/ch9p1.html
- http://www.uwlax.edu/faculty/will/svd/

Questions

samplingTutorial.m is roughly organized to answer the following sequence of questions and more!

- What does an impulse train, and a sampled signal look like?
- What the Fourier transform of a sampled signal looks like?
- What is aliasing?
- How can we diminish the effect of aliasing?
- How can we reconstruct a signal from its sampled version?

Questions

imageTutorial.m is roughly organized to answer the following sequence of questions and more!

- What does the DFT of 2-D signal look like?
- What does the DFT of an image looks like?
- What is the frequency response of typical filters?
- Why is aliasing such a big problem in images?

Impulse Trains

Sampling a continuous signal $s(t)$ is equivalent to multiplying the signal with an impulse train (also known as the comb function) $comb_T(x) = \sum_{n=-\infty}^{\infty} \delta(x - nT)$. Here T is the sampling period and determines the spacing between samples in the sampled signal $s(n) = s(t) \cdot comb_T(x)$.

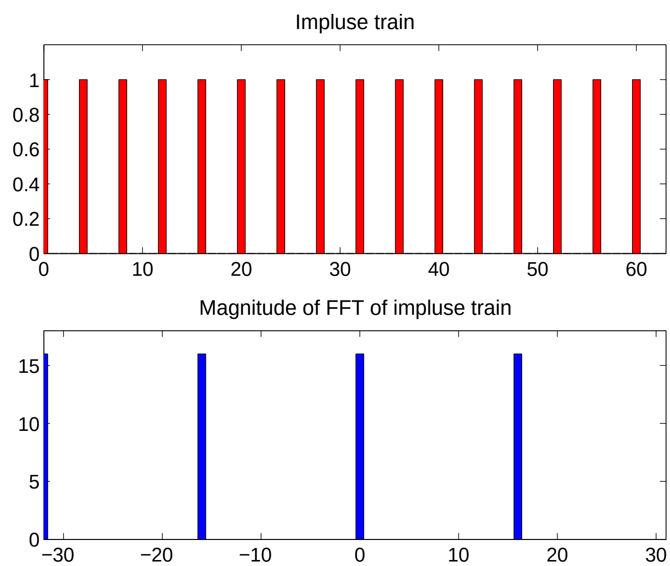

Figure 1: An impulse train and its DFT

The figure above shows an impulse train and its DFT (which is just another impulse train!). The important thing to remember is that the spacing of the impulses in the time (or space) domain and the frequency domain is inversely related.

Sampling a Signal

Sampling involves multiplication in the time or space domain, this corresponds to convolution in the frequency domain. Hence, the FT of a sampled signal can be obtained by convolving the FT of the original signal with the FT of the impulse train.

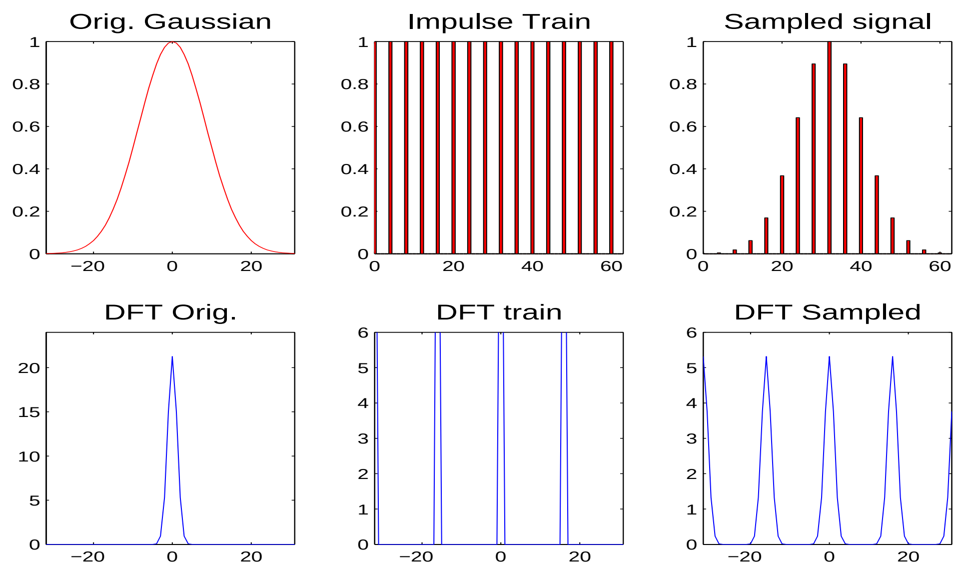

Figure 2: Original signal, impulse train, sampled signal, and their DFTs

You can see that we get replicas of the Fourier spectrum of the original signal, the spacing of these replicas corresponds to the spacing between impulses in the DFT of the original impulse train.

Aliasing

Aliasing occurs when the replicas of the Fourier spectrum of the original signal overlap, this means that the sampled signal is 'masquerading' as a different signal.

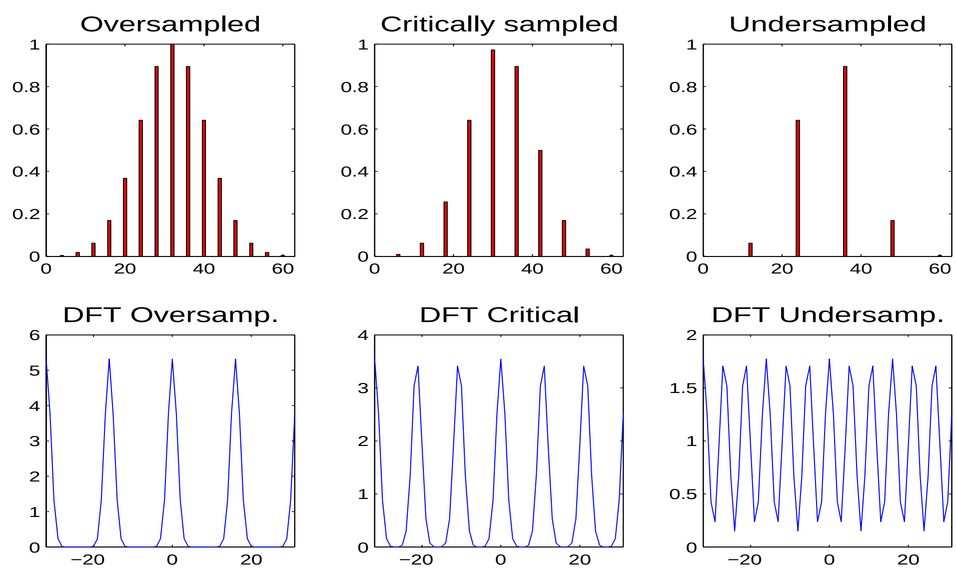

Figure 3: Same Gaussian sampled at different intervals, and corresponding DFTs

As the space between samples increases, the spacing between the replicas of the spectrum of the original signal decreases until the replicas overlap, when the replicas overlap we get aliasing, frequency information necessary for reconstructing the original signal is lost.

Sampling considerations

Sampling pattern in the human eye depends on the arrangement of rods and cones in the retina.

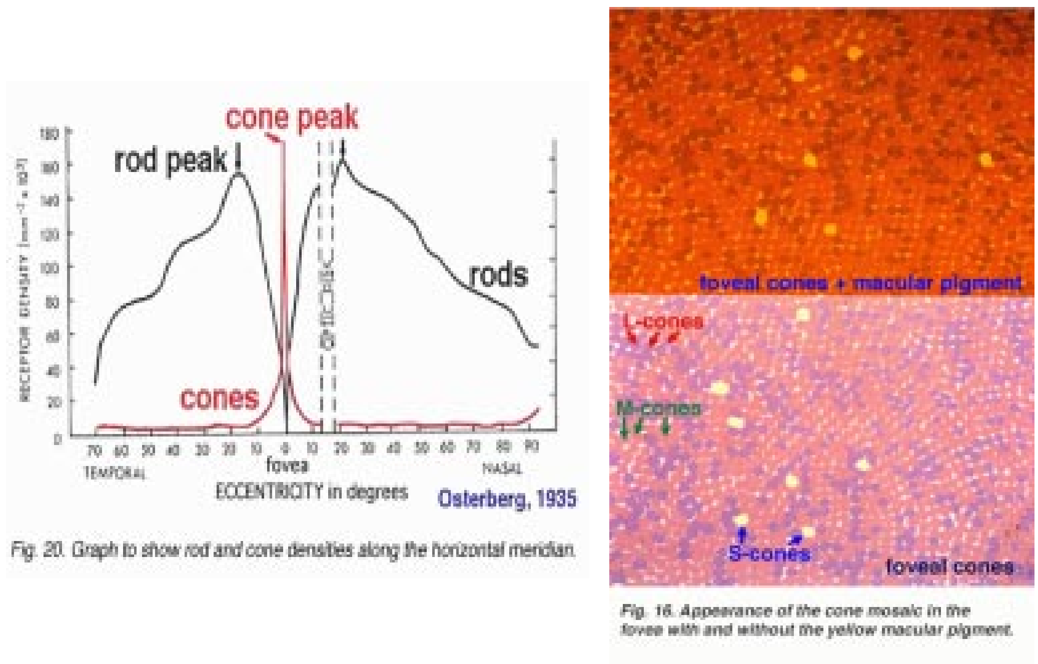

Figure 4: a) Distribution of rods/cones across the retina (from Osterberg, 1935), b) Cone mosaic in the fovea (from Lall and Cone, 1996)

How does the sampling pattern affect the perception of images? what about the sampling pattern in digital cameras?

Singular Value Decomposition

To understand SVD, first recall that any $n \times m$ matrix A defines a linear transformation from $\Re^m$ to $\Re^n$. In other words, the operation

$$\vec{b} = A\vec{x}$$

takes a vector $\vec{x}$ in $\Re^m$ and maps it to a vector $\vec{b}$ in $\Re^n$. To understand the mapping performed by A, we should consider what happens to vector x when left-multiplied (hit) by A.

First, notice that any vector $\vec{b}$ produced by the matrix can be expressed as a linear combination of the columns of A as shown in the following example:

$$\begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ 1 & 0 & -1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}$$

$$x_1 \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} + x_2 \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix} + x_3 \begin{bmatrix} 0 \\ -1 \\ -1 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}$$

In general, $Ax = x_1 A_{*,1} + ... x_m A_{*,m}$ so we see that $\vec{b} = A\vec{x}$ must be in the span of the columns of A. There can be at most m linearly independent columns in A, so when $m < n$ this restricts $\vec{b}$ to a subspace of $\Re^n$. Let the **column space** of A be $colspace(A) = span(A_{*,1}, \ldots, A_{*,m})$.

Secondly, notice that each element $b_i, 1 \leq i \leq n$ is the projection of vector $\vec{x}$ onto the i^{th} row of A. We define the **row space** of A as $rowspace(A) = colspace(A^t) = span(A^T_{*,1}, ... A^T_{*,n})$. A basic theorem in Linear Algebra states that for any $n \times m$ matrix, $rank(A) = dim(rowspace(A)) = dim(colspace(A))$.

Since $\vec{x}$ is first projected onto the row space of A, any $\vec{x}$ that is perpendicular to this space will yield $A\vec{x} = \vec{0}$. In this case, we say that $\vec{x}$ is a member of the (right) **null space** of A, the $(\text{right null space of A} \equiv \{\vec{x} | A\vec{x} = \vec{0}\}$, and it is the orthogonal complement of the row space of the matrix. Similarly, the orthogonal complement of the column space of A is known as the left null space of A.

Now, given any $(n \times m)$ matrix A, we can express A as a product of three matrices

$$A = U\Sigma V^T$$

where U is an $(n \times n)$ matrix that contains an orthonormal basis for the column space of A, Σ is an $(n \times m)$ diagonal matrix that contains the **singular values** of A, the number of non-zero diagonal elements in Σ is equal to the rank of A, and V is an $(m \times m)$ matrix that contains an orthonormal basis for the row space of A.

The columns of U whose corresponding singular value is non-zero constitute the orthogonal basis for the column space of A, the remaining columns span the left null space of the matrix. Similarly, the columns of V whose corresponding singular value is non-zero form an orthonormal basis for the row space of the matrix, while the remaining columns span the right null space of A.

Every matrix has an SVD of the form shown above, this means that any matrix, rectangular or not, is the product of three matrices two of which are orthonormal coordinate transformations, and the third is a rescaling of each axis separately (e.g. a square matrix S is given by a rotation V^T followed by a rescaling of each axis, followed by a rotation/reflection U). That is all that a matrix can do to $\vec{x}$!

More about SVD

- The **singular values** of A are the positive square roots of the **eigenvalues** of $A^T A$.

- The singular vectors are sorted by convention in decreasing order of magnitude of the associated singular value.

- $A_k = U_k \Sigma_k V_k^T$ is the closest rank k approximation of A. A_k minimizes the sum of the squares of the elements of $A - A_k$. Here U_k is a matrix formed with the first k columns of U, Σ_k is a diagonal matrix that contains the first k singular values, and V_k contains the first k columns of V.

- SVD is commonly used to determine the **principal components** of a data distribution as shown in the following figure

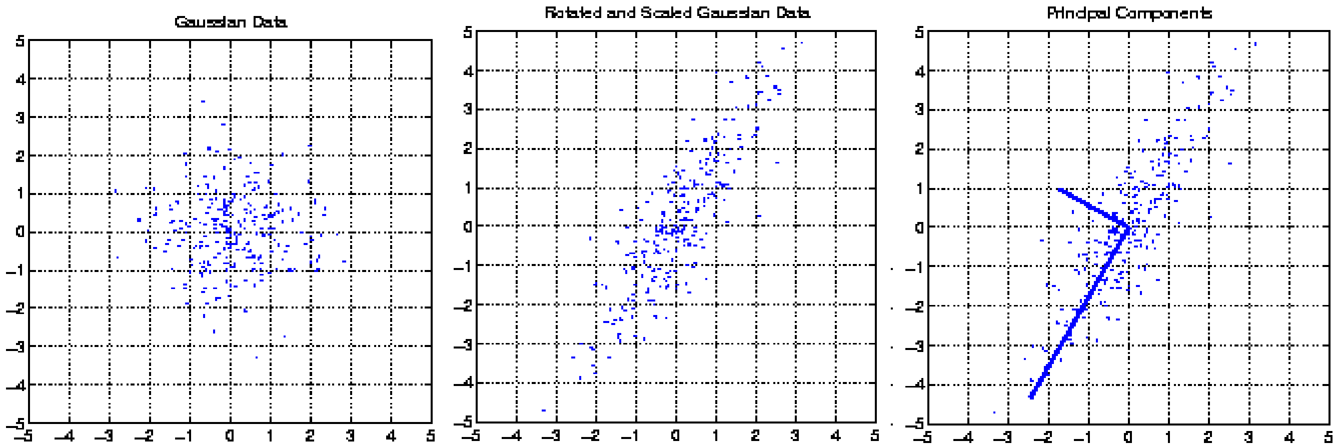

Figure 5: a) 2-D Gaussian distribution, b) Scaled and rotated distribution, c) Principal components from SVD

- Another use is to approximate the inverse matrix for matrices that are either non-square, or not full rank. The pseudo inverse of a non-invertible matrix M is given by

$$pinvM = VWU^T$$

where W is a diagonal matrix, and $W_{i,i} = 1/\sigma_i$ for $\sigma_i > \tau$ and 0 otherwise.