

Machine learning for vision

Winter 2015

Roland Memisevic

Lecture 4, February 3, 2015




Sampling a continuous signal

The Fourier transform of the Dirac comb with spacing D is the Dirac comb with spacing $\frac{1}{D}$.

- ▶ Recall that elementwise multiplication in the time domain is convolution in the frequency domain.
- ▶ Sampling therefore generates replicates of the original spectrum known as **(spectral) alias**:



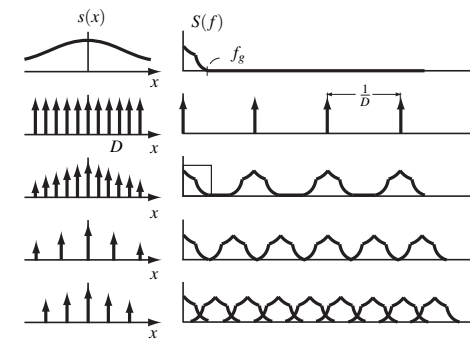
Sampling a continuous signal



- ▶ Sampling (= A/D conversion) can be formalized as multiplication of a continuous signal by a grid of Dirac deltas = "Dirac comb".



Sampling generates spectral aliases




- ▶ The spectrum of the digitized signal is the infinite superposition of translated copies of the original spectrum.




Sampling generates spectral aliases

- ▶ In theory, we can perfectly undo this effect and recover the original signal by low-pass filtering, if the spectral aliases don't overlap.
- ▶ Otherwise, sampling will lose information.



Sampling theorem in 2d




Sampling theorem and Nyquist frequency


The sampling theorem (wikipedia)

If a function $s(t)$ contains no frequencies higher than D hertz, it is completely determined by giving its ordinates at a series of points spaced $1/(2D)$ seconds apart. The frequency $1/(2D)$ is called Nyquist frequency.

Sampling theorem in 2d




Aliasing example




Navigation icons: back, forward, search, etc.

Aliasing in 2d




Navigation icons: back, forward, search, etc.

Aliasing in 2d




Navigation icons: back, forward, search, etc.

Aliasing in 2d



Navigation icons: back, forward, search, etc.


Aliasing in 2d



Roland Memisevic

Machine learning for vision


Aliasing in 2d



Roland Memisevic

Machine learning for vision

Aliasing in 2d



Roland Memisevic

Machine learning for vision

Aliasing in 2d



Roland Memisevic

Machine learning for vision

Aliasing in 2d



Aliasing in 2d



Smoothing

- ▶ An immediate practical implication of the sampling theorem:

If you want to subsample an image, do smoothing first.



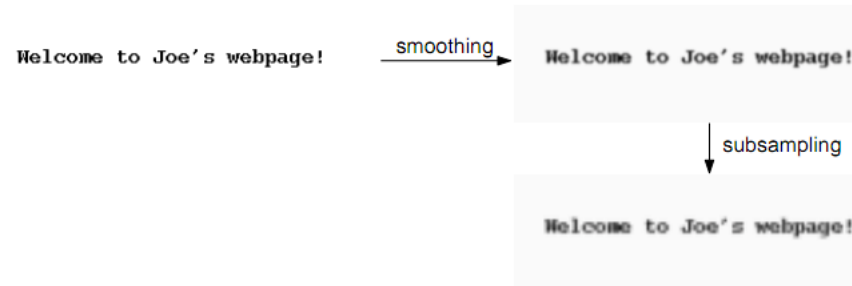
Example

Welcome to Joe's webpage! $\xrightarrow{\text{subsampling}}$ Welcome to Joe's webpage!

example from
<http://redwood.berkeley.edu/bruno/psc129/handouts/aliasing.pdf>



Example



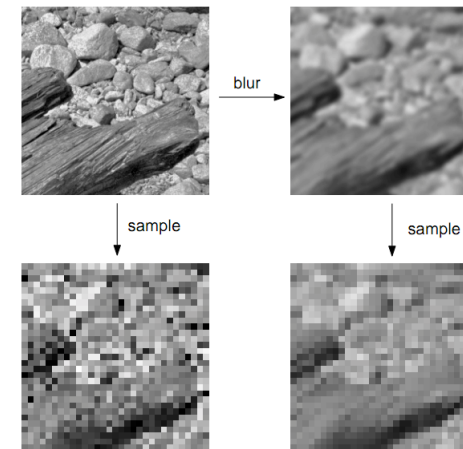
example from
<http://redwood.berkeley.edu/bruno/psc129/handouts/aliasing.pdf>

Ideal interpolation

The diagram shows the process of ideal interpolation. It starts with a discrete signal $S(f)$ represented by a series of impulses at frequencies $0, 1/d, 2/d$. This is convolved with a sinc function, which is the Fourier transform of a rectangular window. The result is a continuous signal that is a sinc function centered at each impulse location. A box below the diagram contains the inequality $f_g \leq \frac{1}{2} \cdot \frac{1}{d}$, representing the Nyquist frequency.

- ▶ To generate the original (continuous) signal from the sampled signal, do low-pass filtering.
- ▶ Applying an ideal low-pass filter amounts to convolving the signal with the Fourier transform of the box function,
- ▶ Thus, **ideal interpolation is convolution of the discrete signal with the sinc-function.**

Example




example from
<http://redwood.berkeley.edu/bruno/psc129/handouts/aliasing.pdf>

DFT leakage


- ▶ We can think of the DFT of a finite signal as the DFT of a periodic signal after multiplying it by a rectangular window.
- ▶ The DFT spectrum you get can be thought of as the spectrum of the periodic signal convolved with a sinc-function.
- ▶ Because of the zero-crossings of the sinc-function the convolution will have no effect on signal components whose frequencies are integer multiples of the window length.
- ▶ For any other components, the convolution will generate additional components in the spectrum.
- ▶ This effect is known as **leakage**.

Leakage example




Navigation icons: back, forward, search, etc.

Leakage example




Navigation icons: back, forward, search, etc.

Leakage example




Navigation icons: back, forward, search, etc.

Leakage example



Navigation icons: back, forward, search, etc.

Leakage example




Navigation icons

Windowing

- ▶ It is best to think of leakage as *leaking into*: frequency components or noise can leak into the frequency bins that you actually want to detect.
- ▶ Leakage cannot be avoided.
- ▶ But a window other than the box-window may lead to different, potentially less undesirable, leakage properties.
- ▶ This may allow us to re-distribute the effect of leakage across the bins such that it causes the least harm for the application at hand.


Navigation icons

Leakage with box window




Navigation icons

Leakage with box window




Navigation icons

Leakage with Gaussian window




Navigation icons: back, forward, search, etc.

Leakage with Gaussian window




Navigation icons: back, forward, search, etc.

Leakage with small Gaussian window



Navigation icons: back, forward, search, etc.

Leakage with small Gaussian window



Navigation icons: back, forward, search, etc.

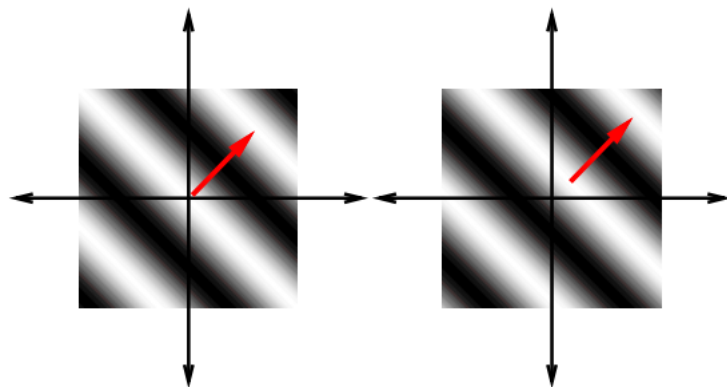
Windowing and Short Time Fourier Transform

- ▶ The **Short-Time Fourier Transform (STFT)** is an application of window functions:
- ▶ Fourier-transform the signal *locally*, then view the resulting set of spectra as a function of time or space.
- ▶ In 1d, the result is called *spectrogram*.
- ▶ The choice of window is important for the reasons just discussed.
- ▶ An STFT using a Gaussian window is also called **Gabor transform**.



Example

Wave:



Gabor features (2d)

- ▶ 2d Gabor functions are highly common in visual recognition tasks.
- ▶ They can be defined as the product of two exponential functions (one for the window, one for the phasor):


$$g(K, \sigma, x_0, y_0, \gamma, u, v, P) = K \exp\left(-\frac{1}{\sigma^2}((x - x_0)^2 + \gamma^2(y - y_0)^2)\right) \cdot \exp(i2\pi(ux + vy) + P)$$

- ▶ Many variations are possible (for example, Gabor filters with oriented rather than axis parallel windows, etc.).




Example

Window:



Example

Gabor feature:




Frequency channels

- ▶ In many applications, local Gabor features are used as *filters*, ie. they are scanned across the image.
- ▶ This naturally raises the question:
- ▶ What is the frequency response of a Gabor filter?
- ▶ It is a *localized blob* in the frequency domain:
- ▶ The Fourier transform of a phasor times a Gaussian will be a delta-peak convolved with a Gaussian.
- ▶ Thus, Gabor filters are **oriented bandpass filters**.

Frequency channels

- ▶ In many applications, local Gabor features are used as *filters*, ie. they are scanned across the image.
- ▶ This naturally raises the question:
- ▶ What is the frequency response of a Gabor filter?

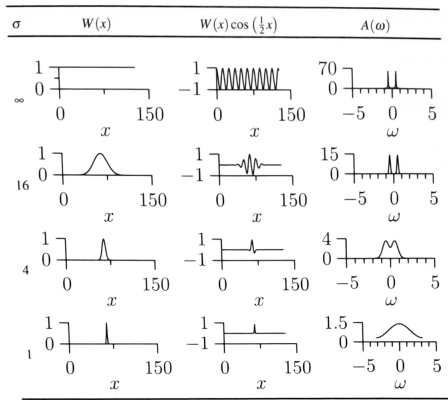
A spectrogram (top) of an utterance



| b | ey | z | th | ih | er | em |
| Bayes' | Theorem |

- ▶ (from Bishop, 2006)
- ▶ The visual analog of the spectrogram is 3-dimensional feature map.


The uncertainty principle



from: Natural Image Statistics (Hyvarinen, Hurri, Hoyer; 2009)



In 2d: orientation uncertainty



from: Natural Image Statistics (Hyvarinen, Hurri, Hoyer; 2009)

