Sampling and Quantization

Luca Guarnera, Ph.D.

Research Fellow

luca.guarnera@unict.it

University of Catania

Dipartimento di Matematica e Informatica

Sampling and Quantization

- Given a continuous signal, a **finite number of "samples"** representative of the signal must be chosen.

- The value at each individual point in the signal is a real number, **discrete values must be chosen to represent the signal correctly**.

Sampling and Quantization

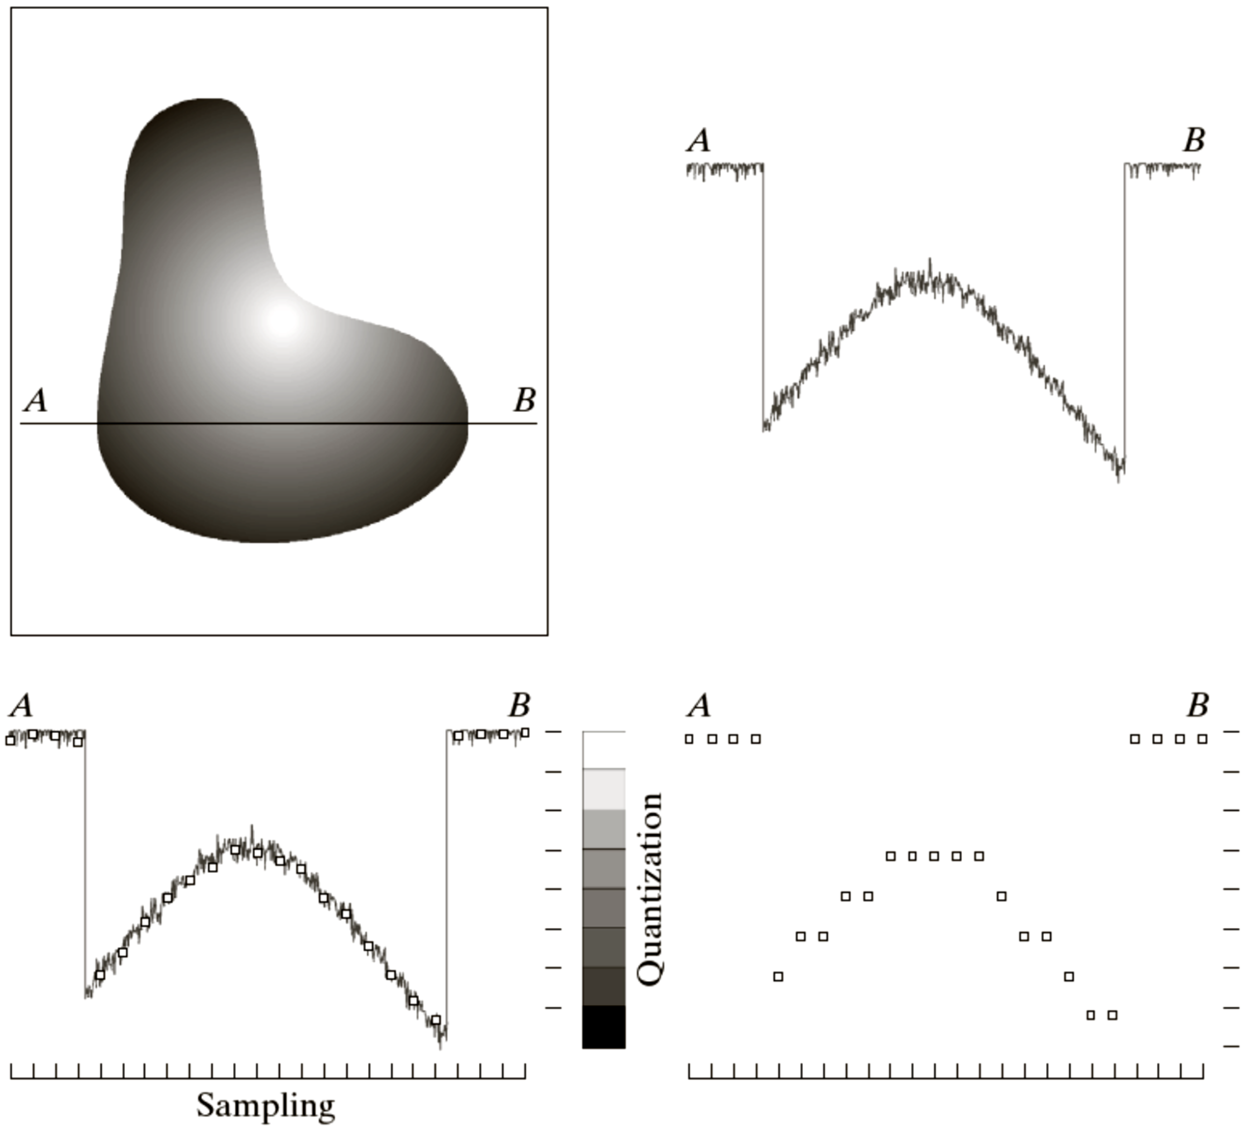

Università di Catania

Sampling

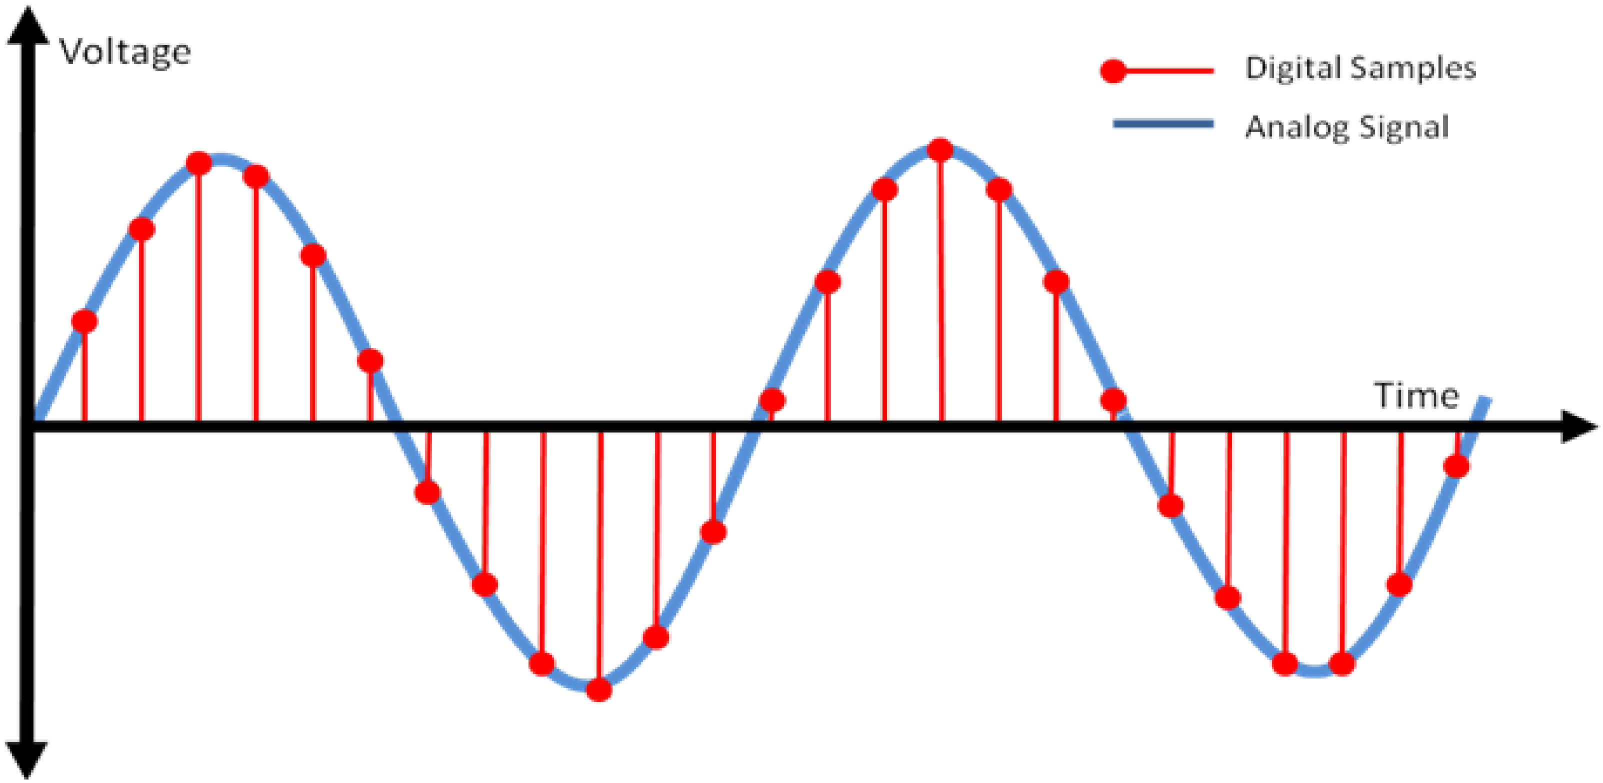

An error in the sampling rate can distort a signal in two ways:

- **Too low a sampling causes details and information to be lost; although serious such a loss is often a necessity:**

  we cannot keep millions of samples and are content to lose information as long as we keep the database of obtained measurements at a manageable size.

- **Too low a sampling can cause details NOT PRESENT in the original to appear in the image.**

  The signal is "altered" and changed into something "other". This is referred to as "**aliasing**." Aliasing is a subtle phenomenon but because it is unpredictable it requires attention.

How to choose the right sampling?

- A fundamental theorem is used to choose the right sampling value: the **Shannon theorem**.
- This theorem is based on the **Nyquist rate measure**.
- But what is it?

Nyquist rate (Harry Nyquist, 1928)

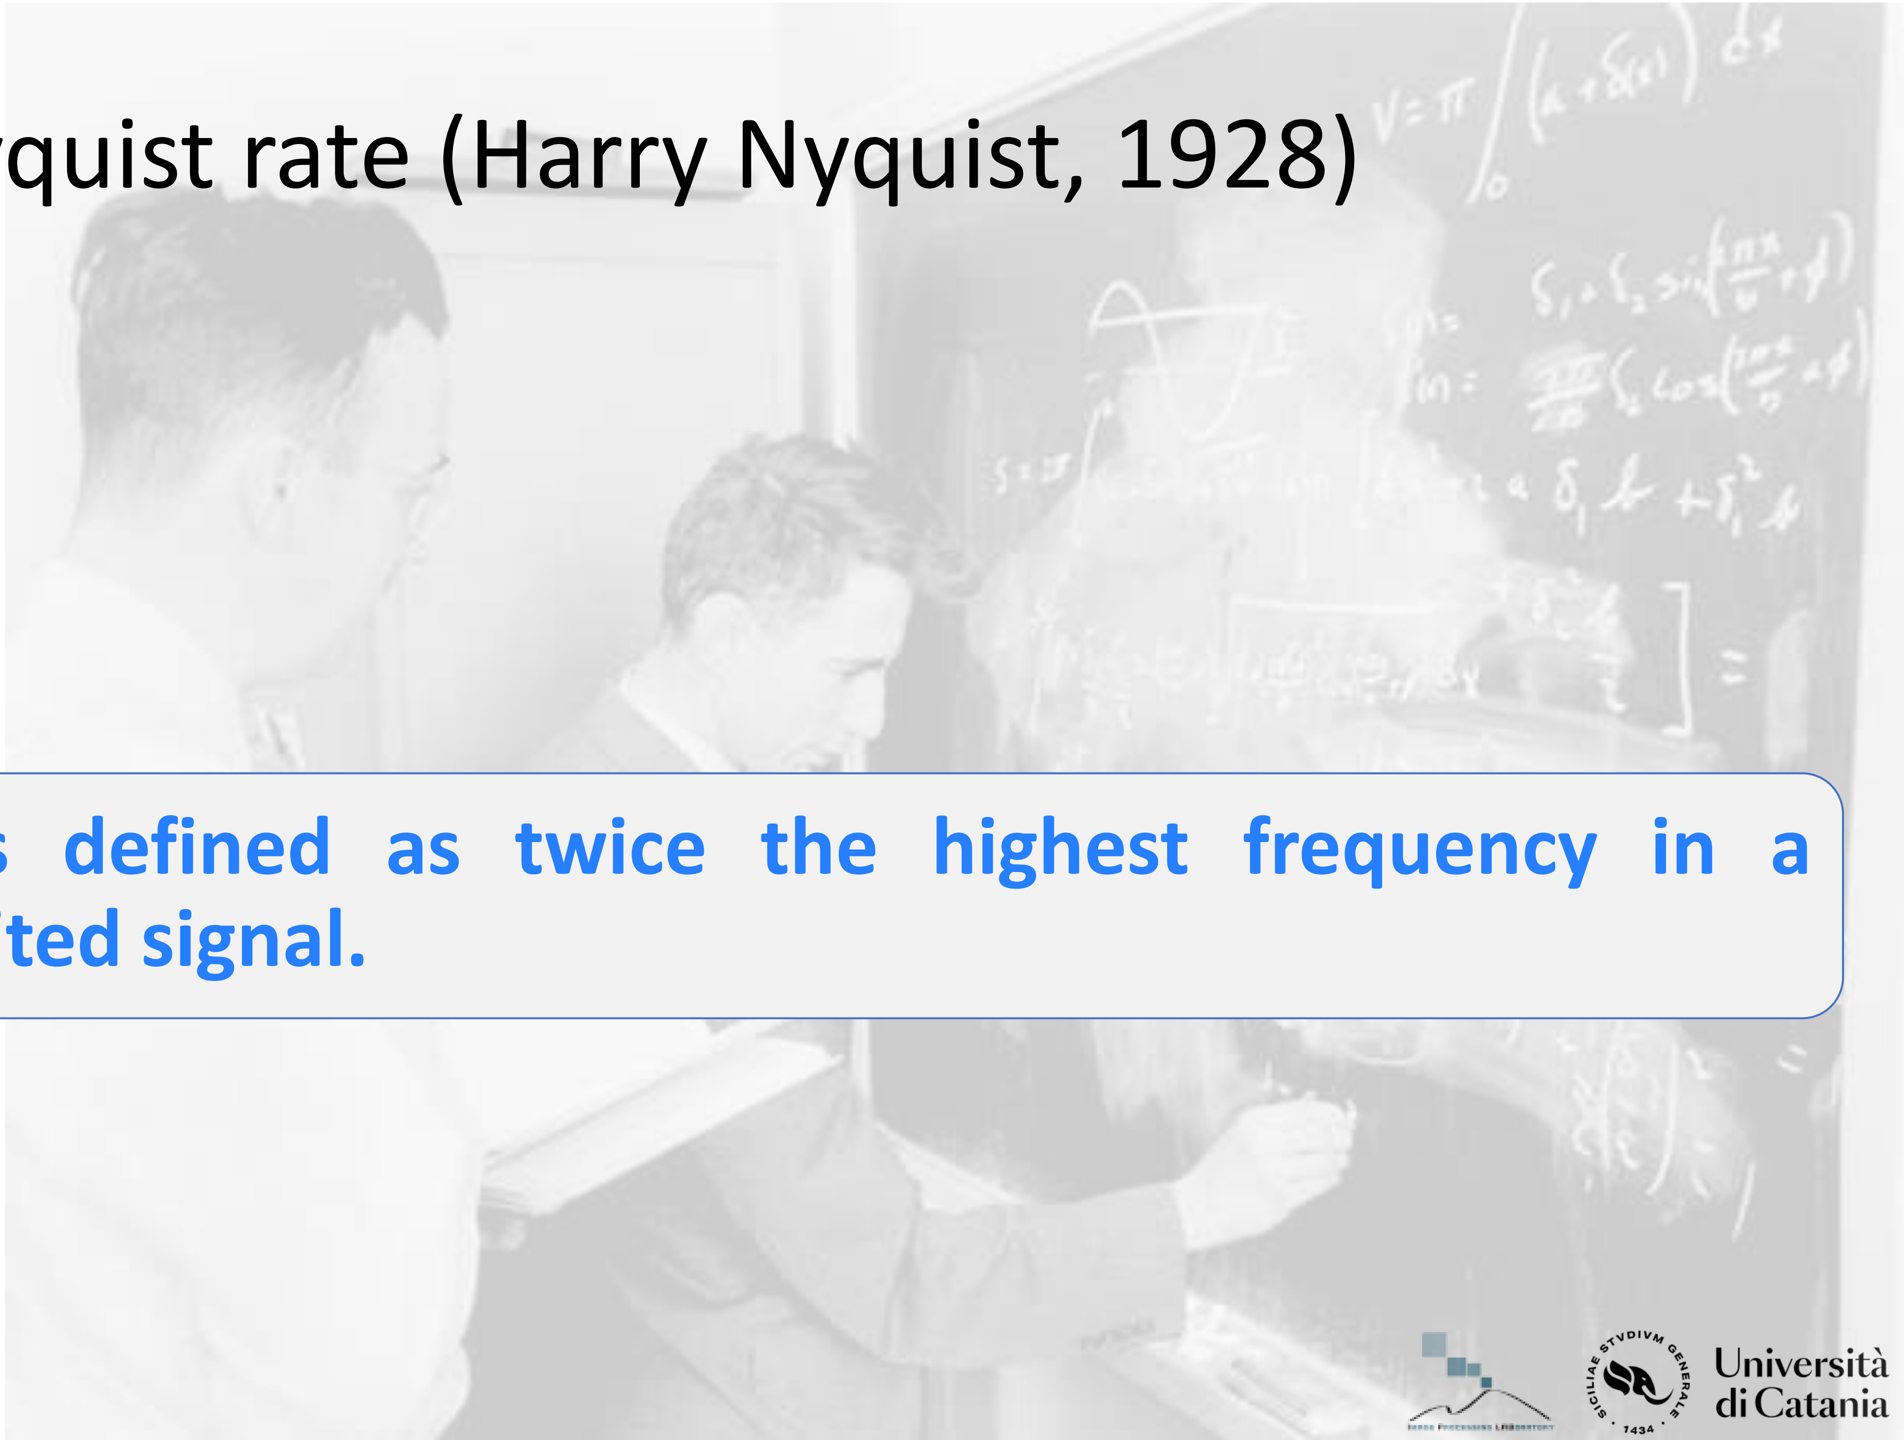

Nyquist rate is defined as twice the highest frequency in a continuous, limited signal.

Time domain...

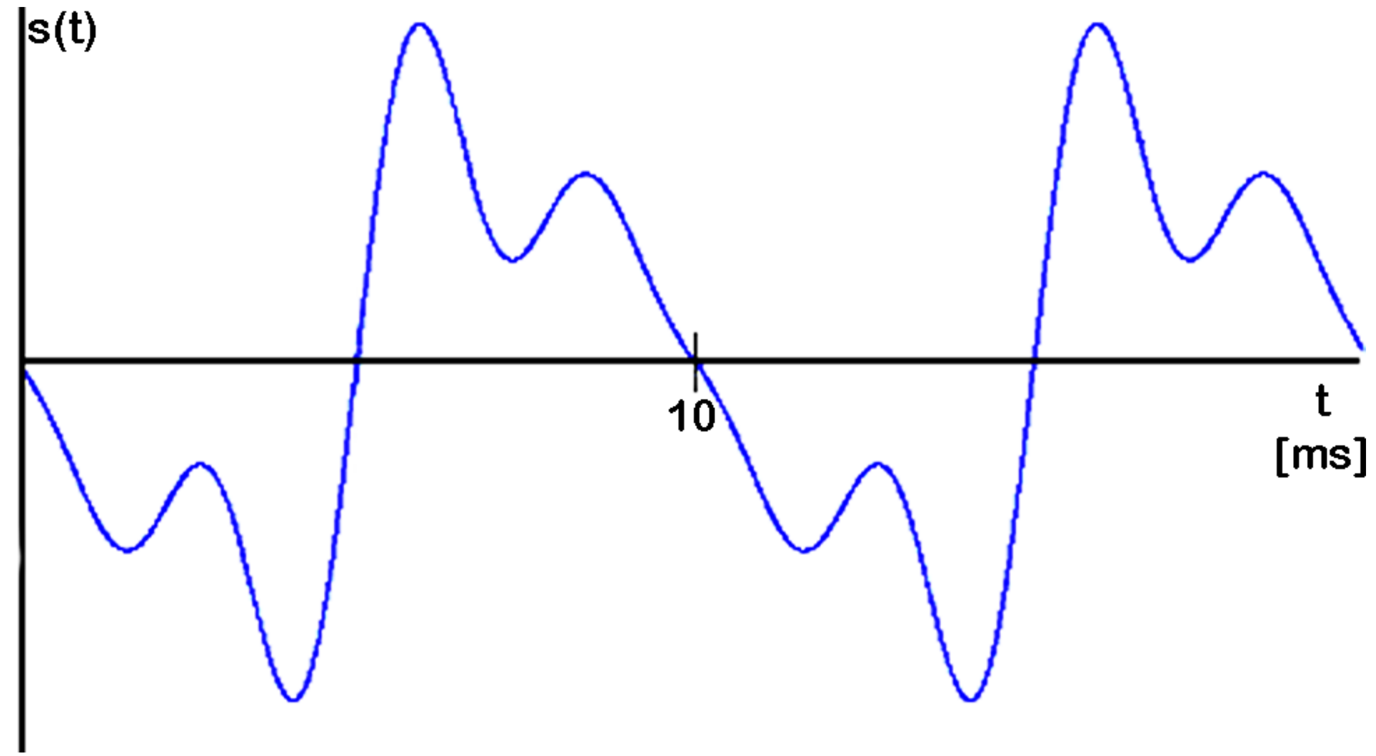

Through the operation of **Fourier Series**...

... and frequency

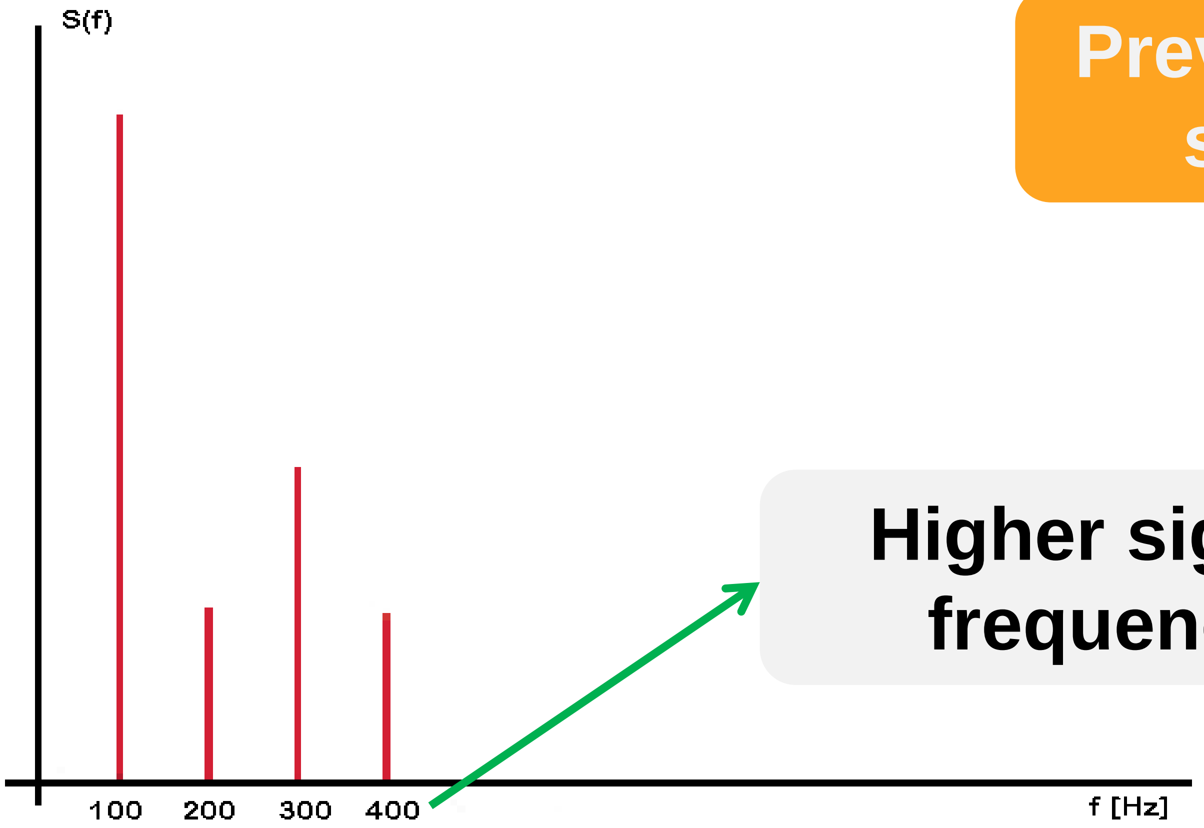

Previous signal spectrum

Higher signal frequency

... we move to a different domain: **frequencies**.

Shannon sampling theorem (Claude E. Shannon, 1949)

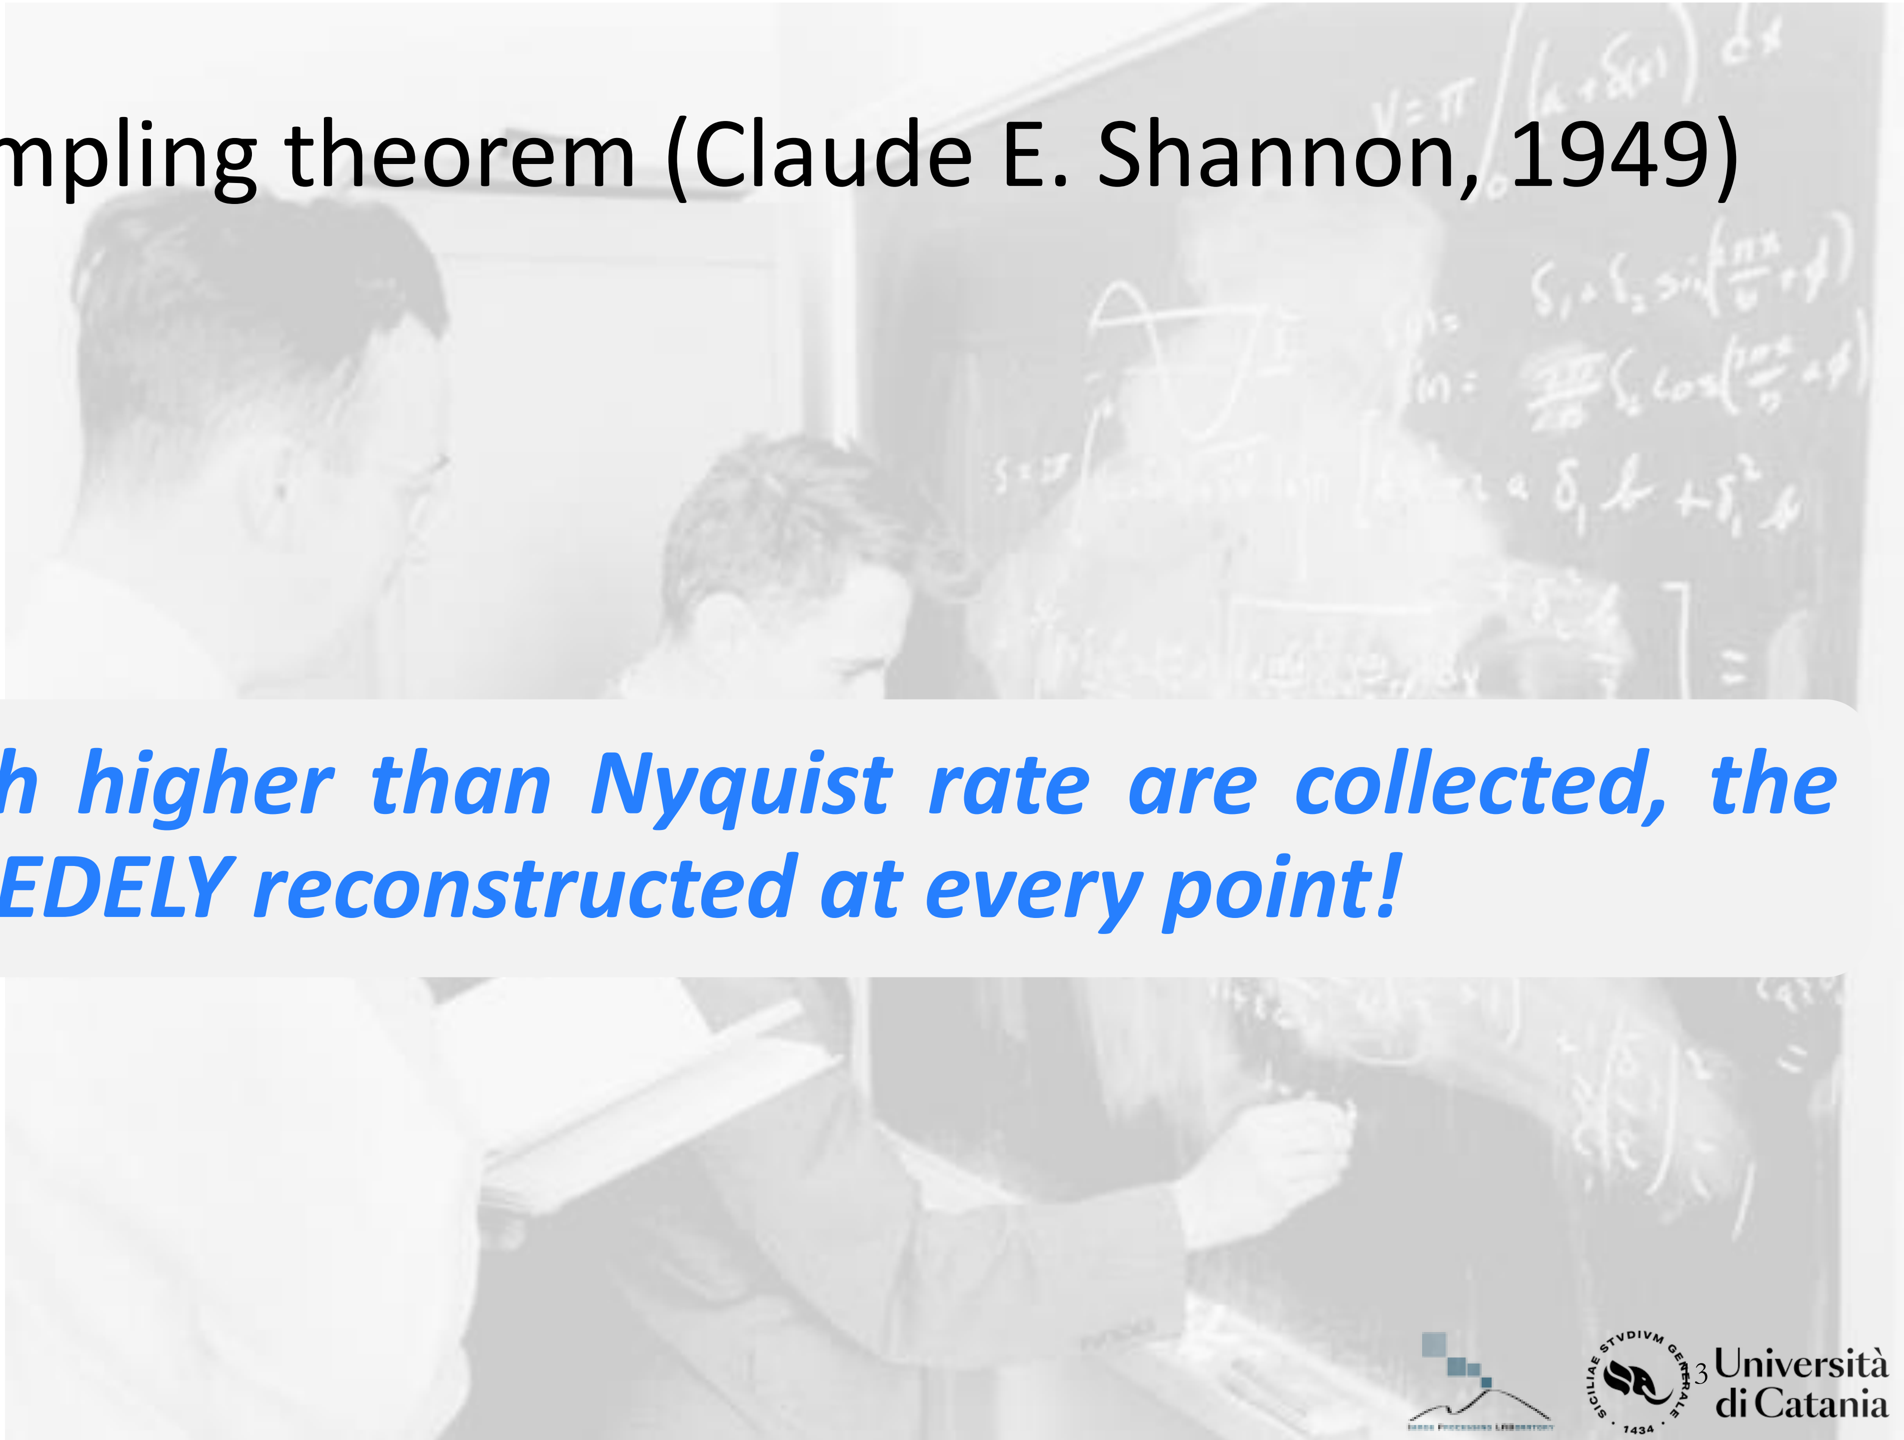

If samples with higher than Nyquist rate are collected, the signal can be FEDELY reconstructed at every point!

Application to images

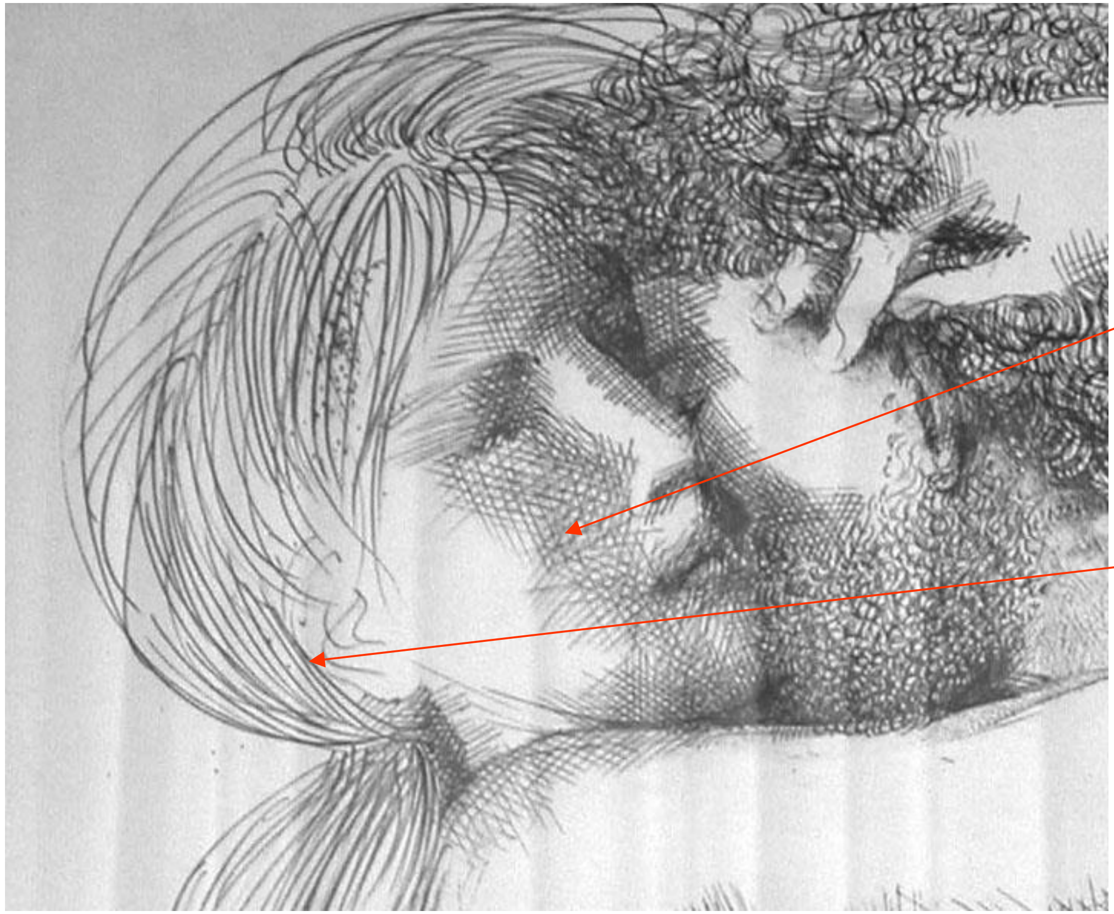

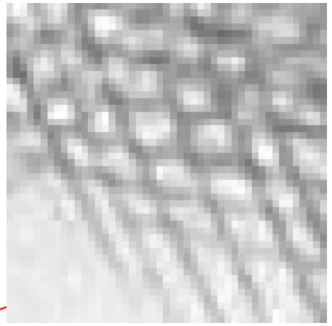

Diametro tratto: 4 pixel

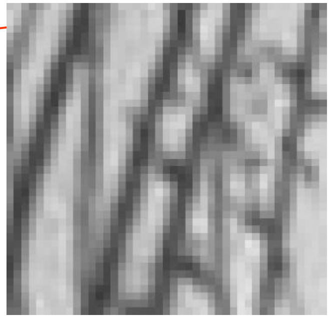

Diametro tratto: 6 pixel

Correct sampling

We use the fine lines. If we preserve these, then we have also preserved the others. Nyquist rate turns out to be: 360. We will then take only 360 samples and reconstruct with bilinear interpolation the image.

Original with 720 x 720 samples

Sampled with 360 x 360 samples

Correct sampling

- I rewrite the image so that the sampled one has the same size as the original one (I use interpolation algorithm).

Original with 720 x 720 samples

Reconstructed with 360 x 360 samples

Wrong sampling

We decide we want to ignore fine lines.

Original with 720 x 720 samples

Sampled with 240 x 240 samples

Wrong sampling

- I rewrite the image so that the sampled one has the same size as the original one (I use interpolation algorithm).

Original with 720 x 720 samples

Reconstructed with 240 x 240 samples

Wrong sampling

- We sample at a frequency below the Nyquist rate.

Original with 720 x 720 samples

Sampled with 120 x 120 samples

Wrong sampling

- I rewrite the image so that the sampled one has the same size as the original one (I use interpolation algorithm).

Original with 720 x 720 samples

Reconstructed with 120 x 120 samples


720 campioni


360 campioni


240 campioni

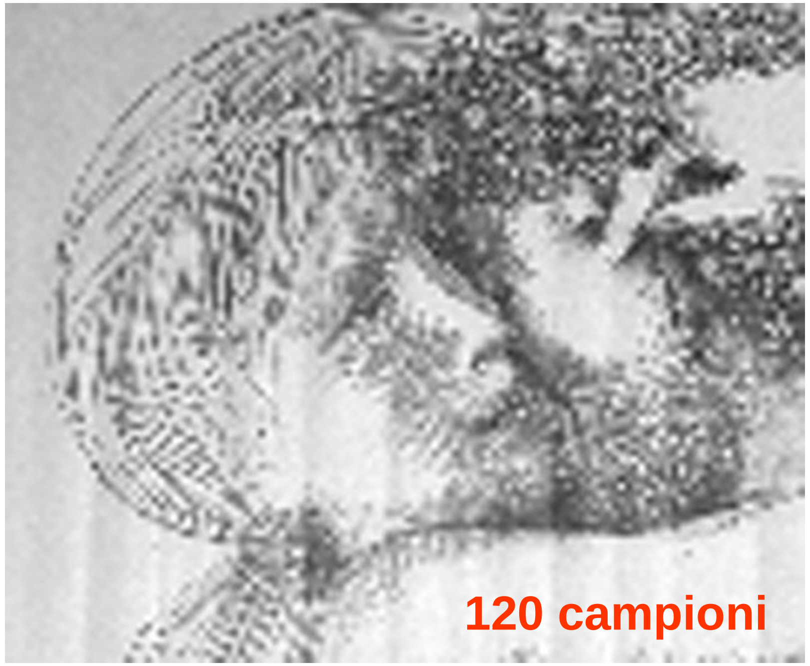
120 campioni

Subsampling

- But what exactly happens if you sample at a **lower frequency** than at the **Nyquist rate**?

Significant details are lost and often new details are introduced that are not present in reality.

Aliasing

- This phenomenon is called **frequency aliasing** or simply **aliasing**.
- With aliasing, high frequencies are "masked" as low frequencies and processed as such in the sampling phase.
- Aliasing comes from Alias i.e. false identity!

Loss of detail

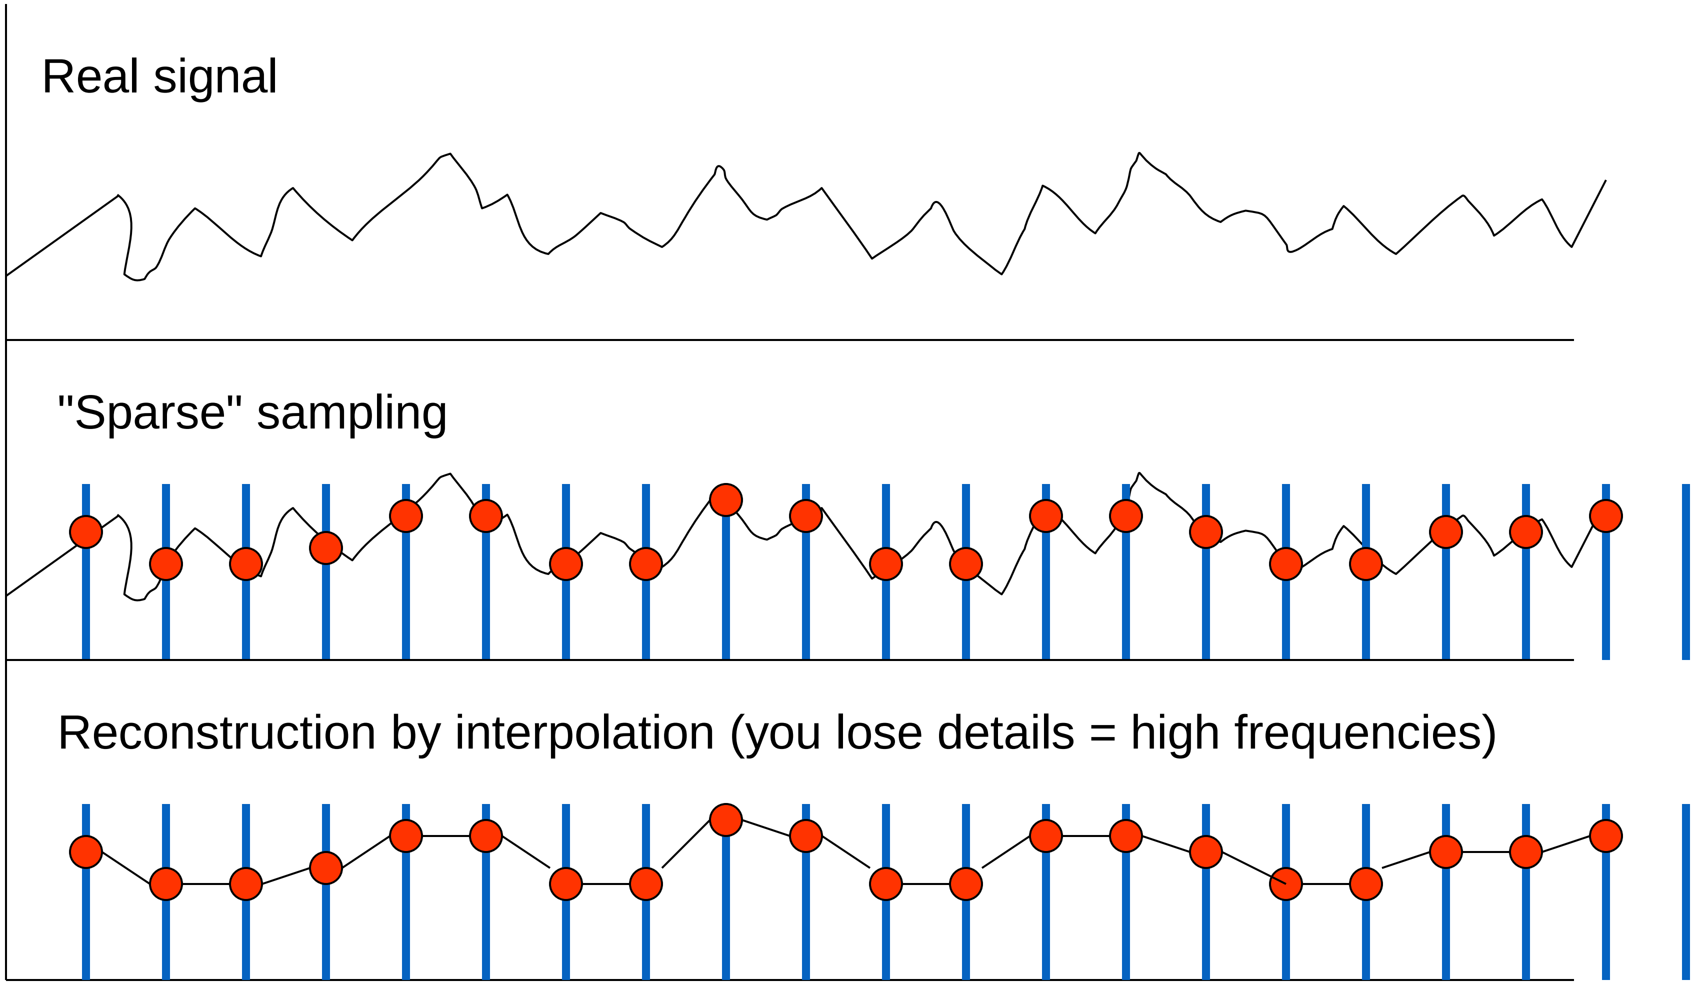

Aliasing

Sampling
One pixel per 256

Reconstruction with
bicubic interpolation

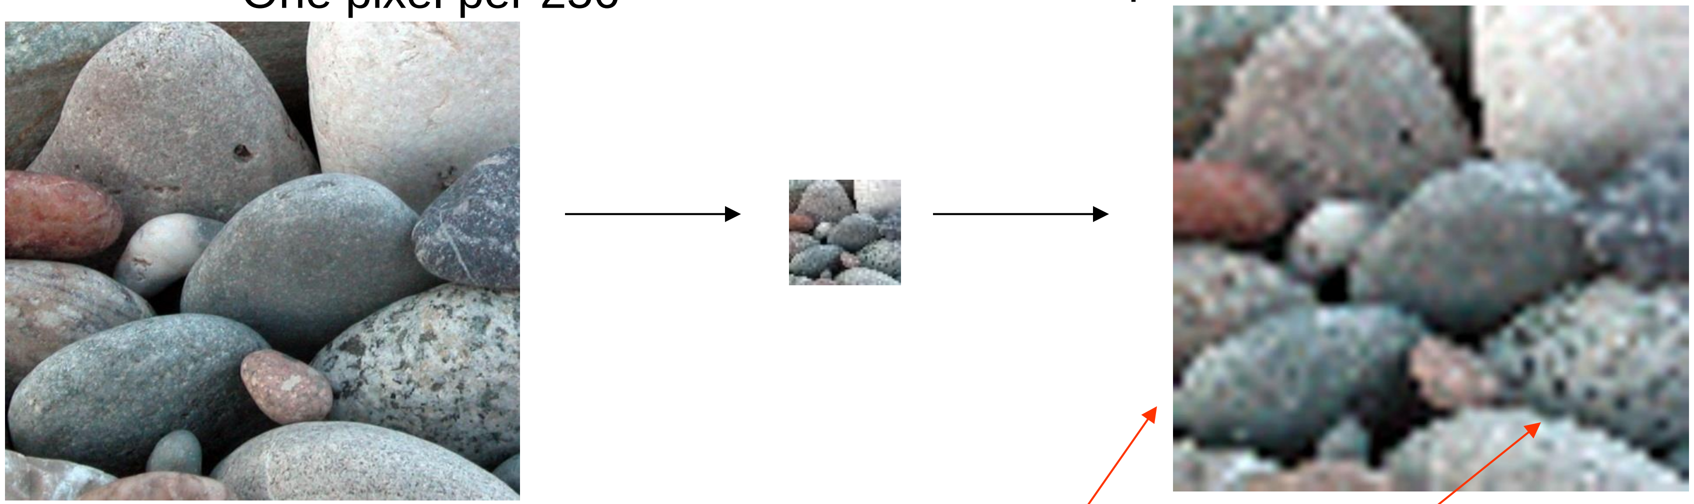

Details are lost, scratches and patterns on the rocks have become indistinguishable and NEW details have appeared!

- Obvious scaling on the edges of the rocks.
- Holes that were not present in the original!

Aliasing


original

New nonexistent
details ("artifacts")


reconstructed

Università
di Catania

Aliasing

- In real life, aliasing is always present even under minimal conditions.
- It is introduced when the signal is required to be limited in order to be sampled.
- Aliasing can be reduced by applying a smoothing function on the original signal before sampling (anti-aliasing).

Aliasing

a b c

FIGURE 4.17 Illustration of aliasing on resampled images. (a) A digital image with negligible visual aliasing. (b) Result of resizing the image to 50% of its original size by pixel deletion. Aliasing is clearly visible. (c) Result of blurring the image in (a) with a 3 × 3 averaging filter prior to resizing. The image is slightly more blurred than (b), but aliasing is not longer objectionable. (Original image courtesy of the Signal Compression Laboratory, University of California, Santa Barbara.)

Real pictures from Facebook

Università di Catania

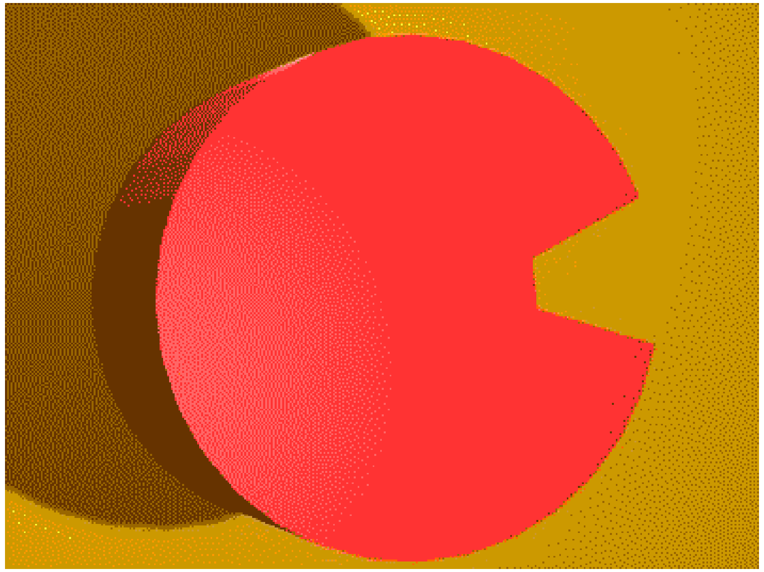

Animation built with all available samples

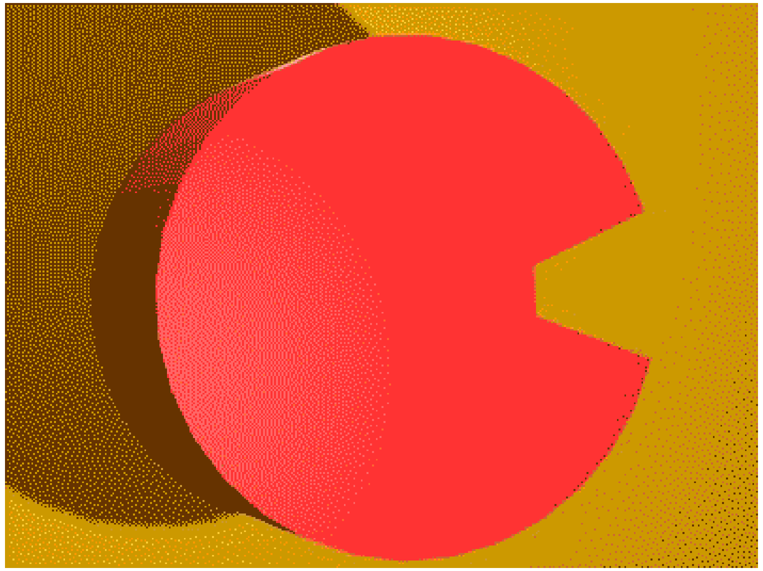

Animation built using one sample every 4

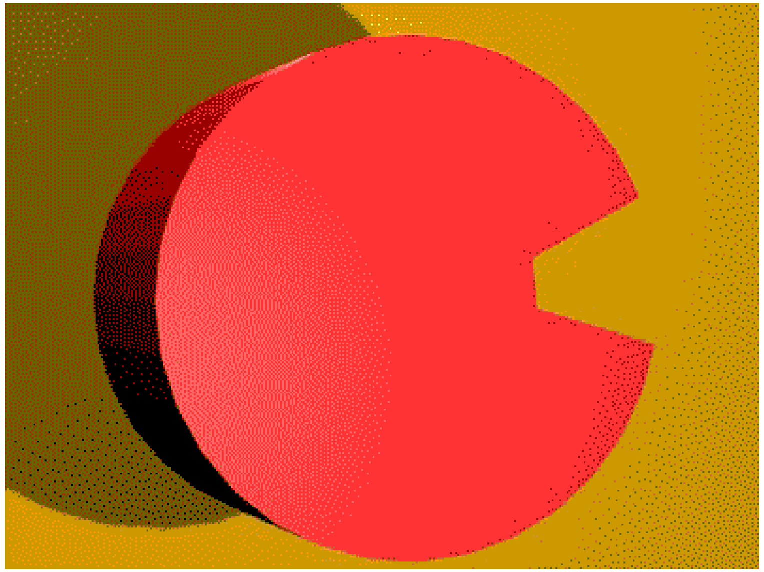

Animation built using one sample every 25

Università di Catania

Quantization

Università di Catania

Quantization

- **Sensors are analog equipment: they provide brightness measurements as REAL numbers.**
- **It is useful to round these values and keep them in a certain range.**
- **This process is called QUANTIZATION**

Quantization: general procedure

If the values to be quantized are real numbers in the range [a, b] and you want to quantize on n levels:

- You set n+1 numbers (**$t_0, t_1, \dots, t_n$**) in **[a, b]** such that:

$$t_0 = a < t_1 < t_2 < \dots < t_n = b$$

- The number **x** in **[a,b]** will be assigned to the quantization level **k** if it turns out:

$$t_k <= x < t_{k+1}$$

- **(b is assigned to level k.)**

- Each level is assigned a representative Q_i value

Quantizzazione: diagramma

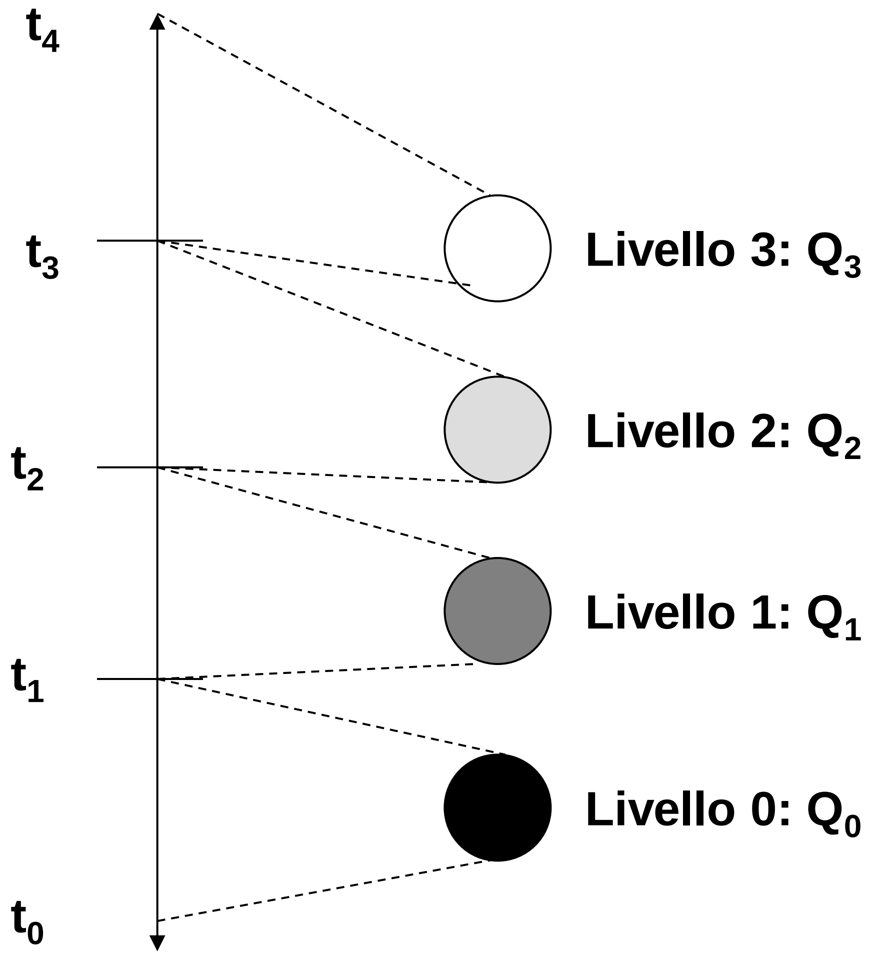

Having fixed the number of quantization levels, the problem arises of how to represent these levels in memory. Obviously we will use numerical labels.

How many bits are needed to remember which level of brightness is measured at a point?

In the example 2 = log(4) will suffice.

In general if there are N levels we need to represent N numeric labels and we will need a number of bits equal to:

B = log (N)

Quantization: effects on images

2 levels
1 bit

4 levels
2 bit

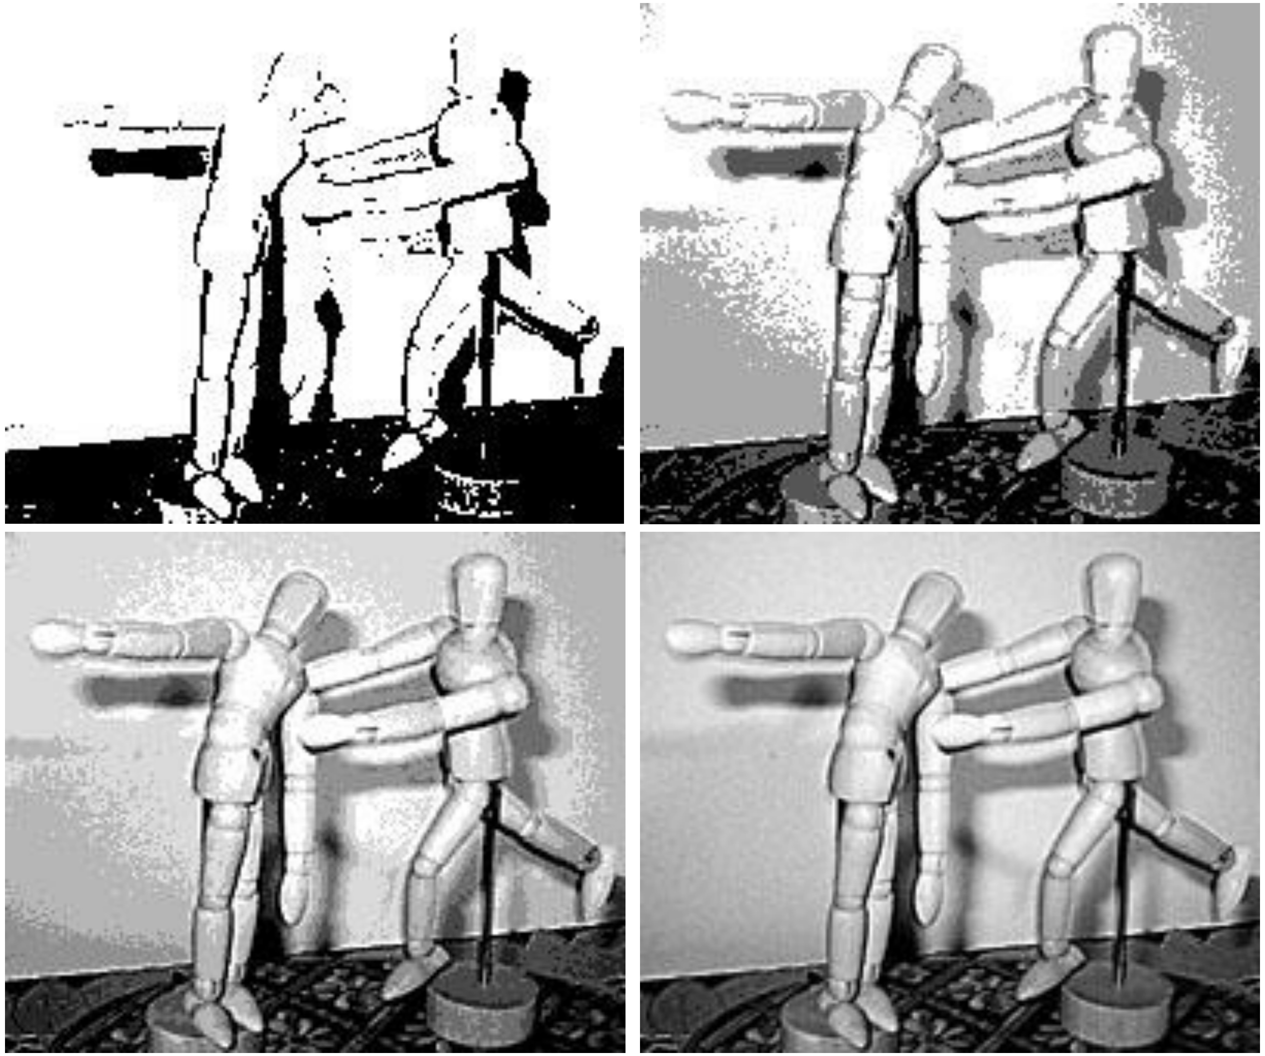

8 levels
3 bit

256
levels
8 bit

Sampling and Quantization

Luca Guarnera, Ph.D.

Research Fellow

luca.guarnera@unict.it

University of Catania
Dipartimento di Matematica e Informatica