




# Machine Learning for Physical Sensing & Perception

## CMSC 818V : Fall 2026

## Lecture 2: Signals as Measurements of Physical Phenomena


Nirupam Roy

COMPUTER SCIENCE
UNIVERSITY OF MARYLAND


**Tu-Th 12:30 – 1:45pm**

**IRB 2107**







# The Forward Measurement Model

*Assembling x, $\vartheta$, F, and n*

# What is a forward model?

$$y = F(x; \theta) + n$$

**A model of the physical process that maps the state of the world to the measurements produced by a sensor.**

$F$ maps the latent physical state $x$ into the expected sensor measurement $y$, given system/environment parameters $\theta$.

---

*Notice what F does not have to be: it does not need to be the physical state itself, or the measurement itself — it is the mapping between them.*

# The forward question vs. the inverse question

**The forward question**

*"If I know where the target is, what signal should the radar receive?"*

**The inverse question**

*"Given the received signal, where is the target?"*

---

An instance — a target at distance d:

$$y(t) = \alpha\, s(t - 2\ \tau) + n(t)$$

$$y(t) = \alpha\, s(t - 2\ x/c) + n(t)$$

$$x \rightarrow y$$

*forward model*

$$y \rightarrow \hat{x}$$

*inverse problem*

# The forward model, across sensing systems

*In every system, the forward model predicts the measurement from the physical state.*

| System | Physical state x | Forward model predicts |
|---|---|---|
| **Camera** | 3D scene, materials, lighting | Pixel intensities |
| **Microphone** | Sound source + room | Recorded waveform |
| **Radar** | Target position / velocity | RF echo / IQ samples |
| **LiDAR** | Surface geometry | Return times / intensities |
| **IMU** | Body motion | Acceleration and gyro readings |
| **CT** | Internal attenuation map | X-ray projections |

# θ: what does it actually mean?

$$y = F(x; \theta) + n$$

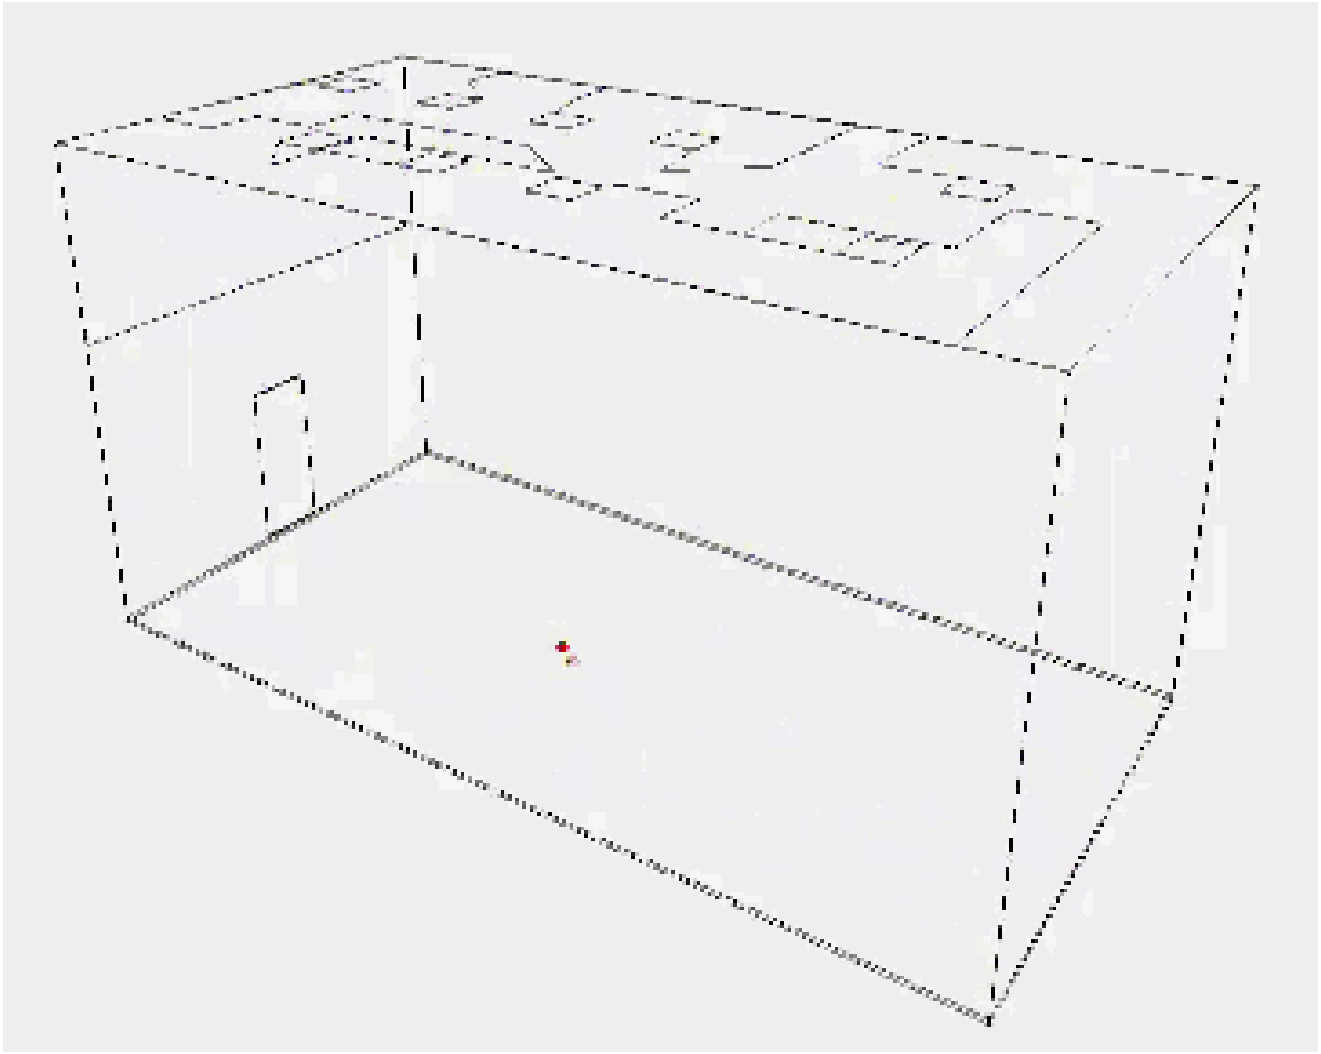
**Parameters describing the sensing system, environment, and propagation conditions that determine the forward process.**

$\theta$ is not what we are trying to infer — it is what the forward model needs to know to relate $x$ to $y$.

---

*$\vartheta$ is a parameter of the model, not a label — some $\vartheta$ are known and chosen, others are unknown and must be estimated alongside x.*

# θ: what does it actually mean?



# θ across five sensing systems

| System | x — quantity of interest | θ — forward-model parameters |
|---|---|---|
| **Camera** | Object geometry / pose | Illumination, camera pose, focal length, exposure |
| **Microphone array** | Source location | Room geometry/materials, mic locations, speed of sound |
| **Radar** | Target position / velocity | Antenna geometry, carrier frequency, transmitted waveform |
| **WiFi localization** | User location | AP position, room geometry, furniture, antenna response |
| **IMU** | Motion / trajectory | Sensor bias, scale factor, orientation, gravity |

# The x / θ boundary depends on the task, not the quantity

**Task: speaker localization**

$X$ = speaker location

$\theta$ = room geometry / materials

**Task: room reconstruction**

$X$ = room geometry / materials

$\theta$ = everything else (source signal, mic array geometry, …)

*The same physical quantity — room geometry — is ϑ in one task and x in the other. The split is set by what you are trying to infer, not by the physics itself.*

# F vs. θ: model structure vs. instantiation

$F$ specifies the structure of the forward model; $\theta$ specifies the parameters that instantiate it for a particular system or environment.

$$y(t) = (1/r)\ s(t - r/c) + n(t)$$

$x = r$ (distance) $\quad \theta = c$ (propagation speed)

---

For a richer room-acoustic model:

$\theta$ = { room geometry, wall reflection coefficients, microphone positions, $c$, … }

| $F$ | $\theta$ | $x$ |
|---|---|---|
| model structure / physical laws | parameters that instantiate that model | physical state we want to infer |

# Why not just call θ "nuisance parameters"?

A **nuisance parameter** is a parameter that affects the sensor measurement but is not the quantity you actually want to estimate.

**Some θ are known, deliberately chosen design parameters:**

- Antenna spacing
- Radar carrier frequency
- Camera focal length

**Others are genuinely unknown nuisance parameters:**

- Room materials — and other unmodeled environmental conditions

---

**θ mixes both known design choices and unknown nuisances — so "nuisance parameters" is too narrow a name for what it actually captures.**

# F and n: process and uncertainty

## *F*: the forward model

The known — or partially known — physical process transforming x into measurements. Sometimes physics-based, sometimes learned, often both.

## *n*: measurement noise

The measurement uncertainty not captured by the deterministic part of the model — sensor electronics, thermal noise, unmodeled effects.

# From a deterministic map to a probabilistic model

Once we admit noise, y is no longer a deterministic function of x — it is a random variable given x.

$$p(y \mid x)$$

*the probabilistic measurement model*

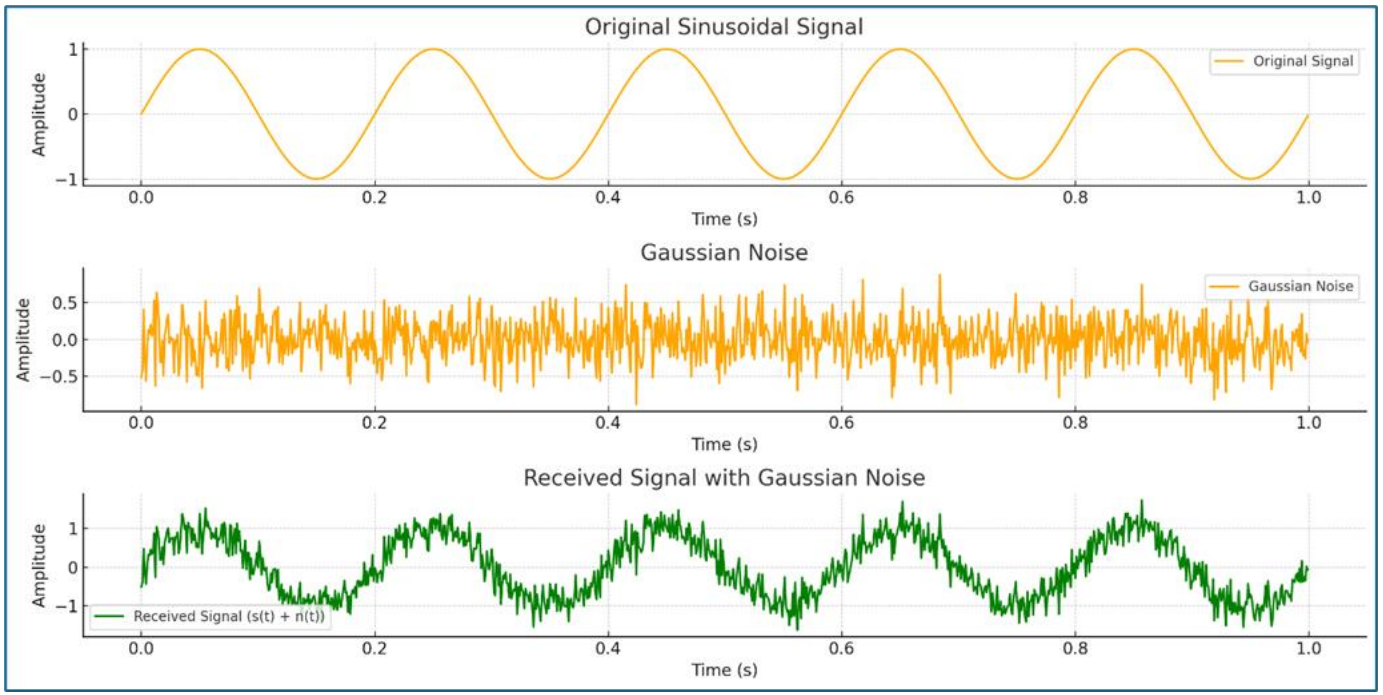
---

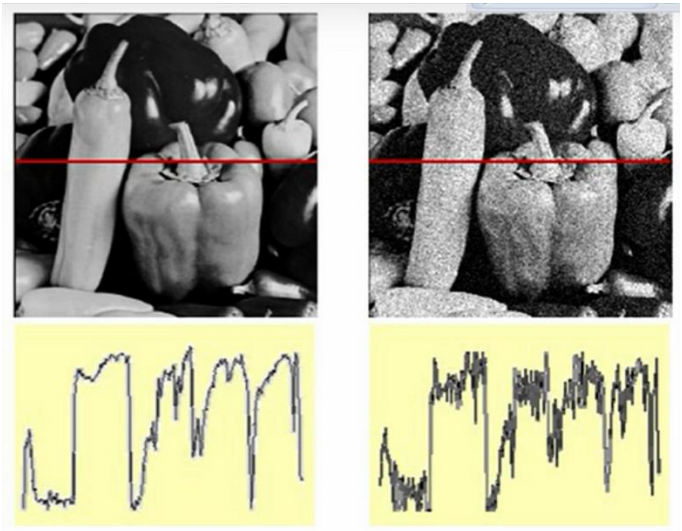
*Bayesian inference (later lectures) are built on this model.*

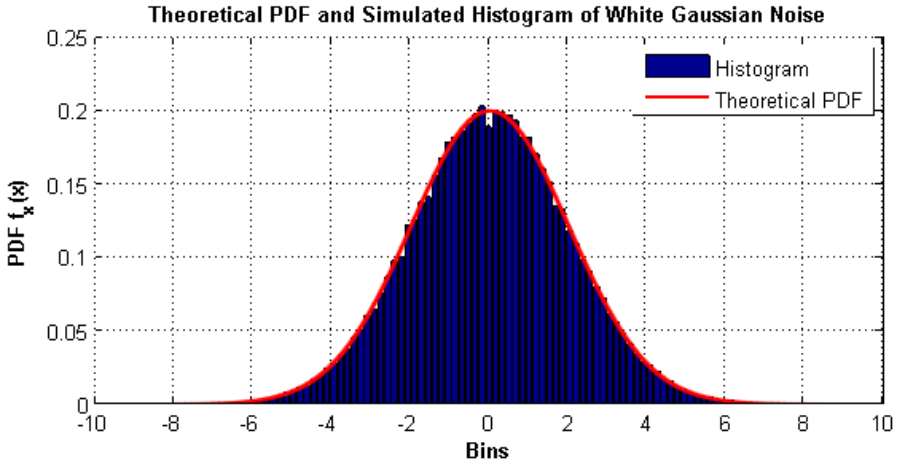
# Deterministic Model vs. Stochastic Measurement

*Why y = F(x) is almost never literally true*

# Gaussian noise





# Why Gaussian?

**The Central Limit Theorem (CLT) (roughly):**
When many independent random contributions are added together, their sum tends to have a Gaussian (normal) distribution, even if the individual contributions are not Gaussian.



Galton board

# Why y = F(x) is almost never literally true

Instead, we write y = F(x) + n. For Gaussian noise:

$$n \sim \mathcal{N}(0, \sigma^2) \longrightarrow p(y|x) \propto \exp\left[-\|y - F(x)\|^2 / 2\sigma^2\right]$$

---

- This is the origin of a very common pattern in estimation: under Gaussian measurement noise, maximizing p(y|x) becomes minimizing squared measurement error.
- We will not derive it in detail today — just recognize the shape when it returns in Lecture 3.

$$n \sim \mathrm{N}(0, \sigma^2 I). \longrightarrow y = \mathrm{F}(x) + n \longrightarrow \boxed{y \mid x \sim \mathrm{N}(\mathrm{F}(x), \sigma^2 I)}$$

# What Information Is Lost During Sensing?

*From forward models to inverse problems*

# What information is lost during sensing?

**A sensor does not capture the complete physical state of the world.**

The forward process transforms, mixes, compresses, and often irreversibly discards information before producing the measurement.

---

$$y = F(x; \theta) + n$$

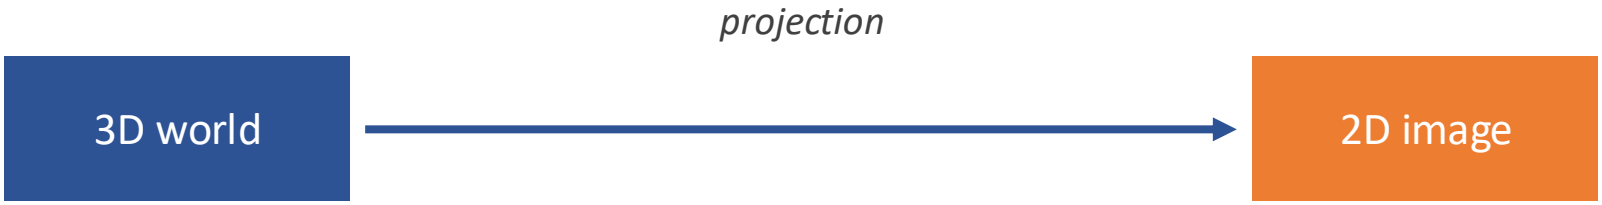
*even with no noise:*

$$y = F(x; \theta)$$

**We generally cannot uniquely recover x from y.**

# 1. Dimensionality reduction

The physical world often has many more degrees of freedom than the measurement — the clearest example is a camera.




A pixel tells us where something projects onto the image plane, but not its depth.

$$u = fX/Z \qquad v = fY/Z$$

---

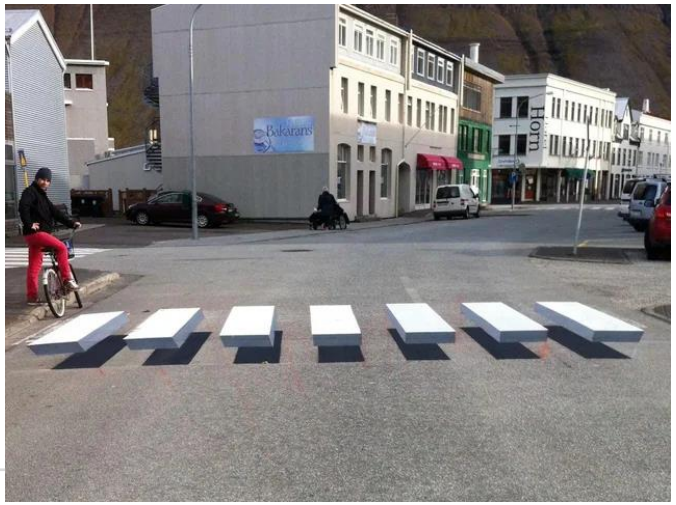
Many 3D points can produce the same (u, v) — e.g. (X, Y, Z) and (2X, 2Y, 2Z) project to the same image location.

*Depth information has been partially lost through projection.*

# 2. Multiple causes, one measurement

**Suppose a camera measures a dark pixel. Why is it dark?**

- Dark material
- Weak illumination
- Surface orientation
- Shadow



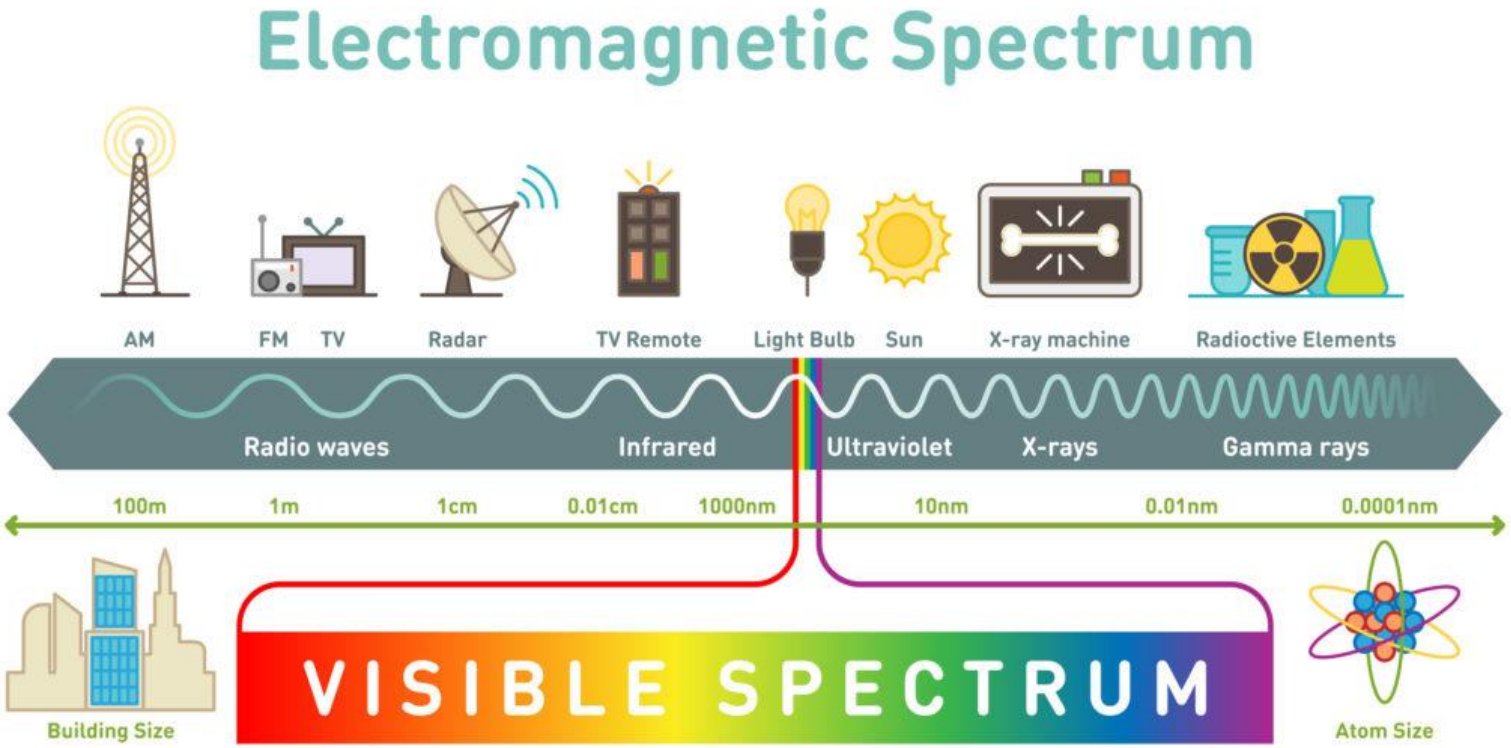
$$y \propto \text{illumination} \times \text{reflectance}$$

*Illumination is governed by ϑ; reflectance depends on x (and sometimes ϑ too).*

$$x_1 \neq x_2 \quad \text{but} \quad F(x_1; \theta_1) = F(x_2; \theta_2)$$

**Different combinations of physical quantities can give the same y — the fundamental many-to-one nature of many forward models.**

# 3. The sensor is an information bottleneck



# 3. The sensor is an information bottleneck

**Every sensor has limited sensitivity.**

$$L(\lambda) \longrightarrow (R, G, B)$$

*A rich electromagnetic spectrum is reduced to approximately three channels.*

---

Two different spectral distributions can therefore produce the same RGB values.

Similarly, a microphone measures pressure variation over a limited frequency range — it does not capture every property of the acoustic field.

**The sensor itself acts as an information bottleneck.**

# 4. Finite resolution destroys detail

If two objects are closer than the radar's range resolution, their reflections may merge into one return.

$$\text{continuous scene} \rightarrow \text{finite samples}$$

$$x(t) \rightarrow x[n]$$

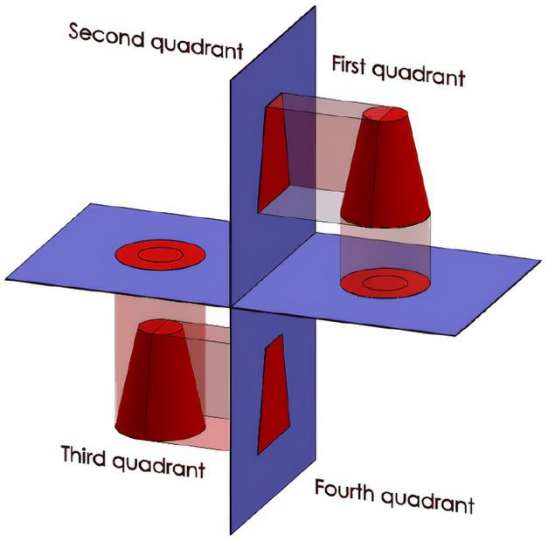
Fine-scale spatial or temporal information can disappear.



# 5. Occlusion: some states are never measured

A camera observes the front surface of a box, but not the back.

$$\text{complete 3D object} \rightarrow \text{visible-surface projection}$$




No amount of processing of a single image can directly recover measurements that were never acquired. Any reconstruction of the back must come from assumptions or prior knowledge.

**The same issue occurs in radar shadowing, ultrasound, LiDAR, and many other sensing systems.**

# 6. Mixing: contributions become entangled

$$y(t) = s(t) * h(t) + n(t)$$

*A microphone measurement contains both the source s(t) and the environment h(t) — they are not separately observed.*

Likewise, wireless CSI may contain the superposition of many propagation paths:

$$H(f) = \Sigma_k \ \alpha_k \exp(-j2\pi f \tau_k)$$

*The receiver measures the combined channel, not necessarily each individual reflector/path.*

**Information may not be completely absent — it may be entangled.**

**lost information ≠ mixed/latent information — a sophisticated inverse method can sometimes disentangle the latter.**

# 7. Noise limits distinguishability (Signal-to-noise ratio or SNR)

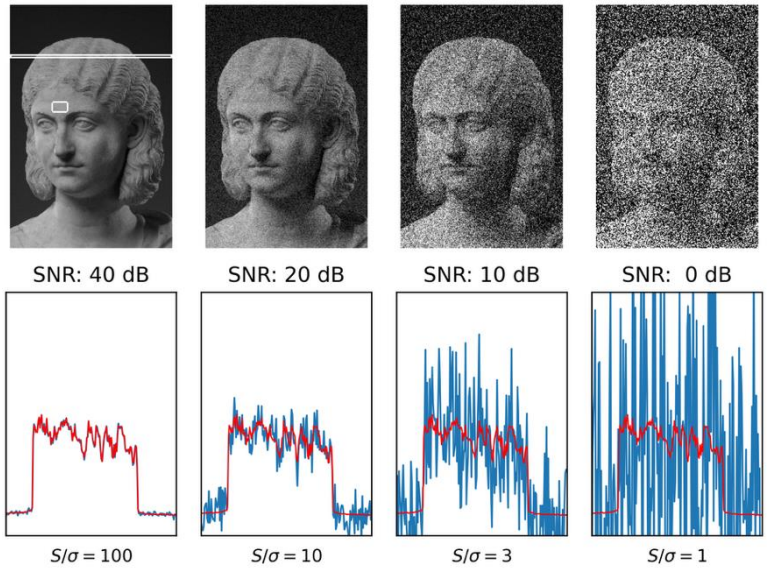
$$y = F(x; \theta) + n$$

Suppose two physical states produce slightly different ideal measurements:

$$F(x_1) \approx F(x_2)$$

---

If the difference is smaller than the measurement noise, the sensor may not be able to distinguish them.




**Even a theoretically invertible mapping may be practically unrecoverable.**

# Three concepts this motivates

**Ambiguity**

Can different physical states explain the same measurement?

**Identifiability**

Can the physical state of interest be uniquely determined from the available measurements?

**Observability**

Do the measurements contain enough information about the state to recover it?

*Especially important for dynamical systems, later in the course.*

$$x_1 \neq x_2, \quad F(x_1) \approx F(x_2)$$

---

*For Lecture 1 we emphasize ambiguity and identifiability, and only briefly mention observability — formal observability fits better with state estimation.*

# What information was lost?

**The sensor gives us *y*. We actually want *x*.**

*How do we get from y back to x?*

# The bridge to inverse problems

Measurements + Physics + Priors + Learning

$\longrightarrow$ Physical inference

---

Much of physical sensing ML exists because the measurement alone is insufficient.

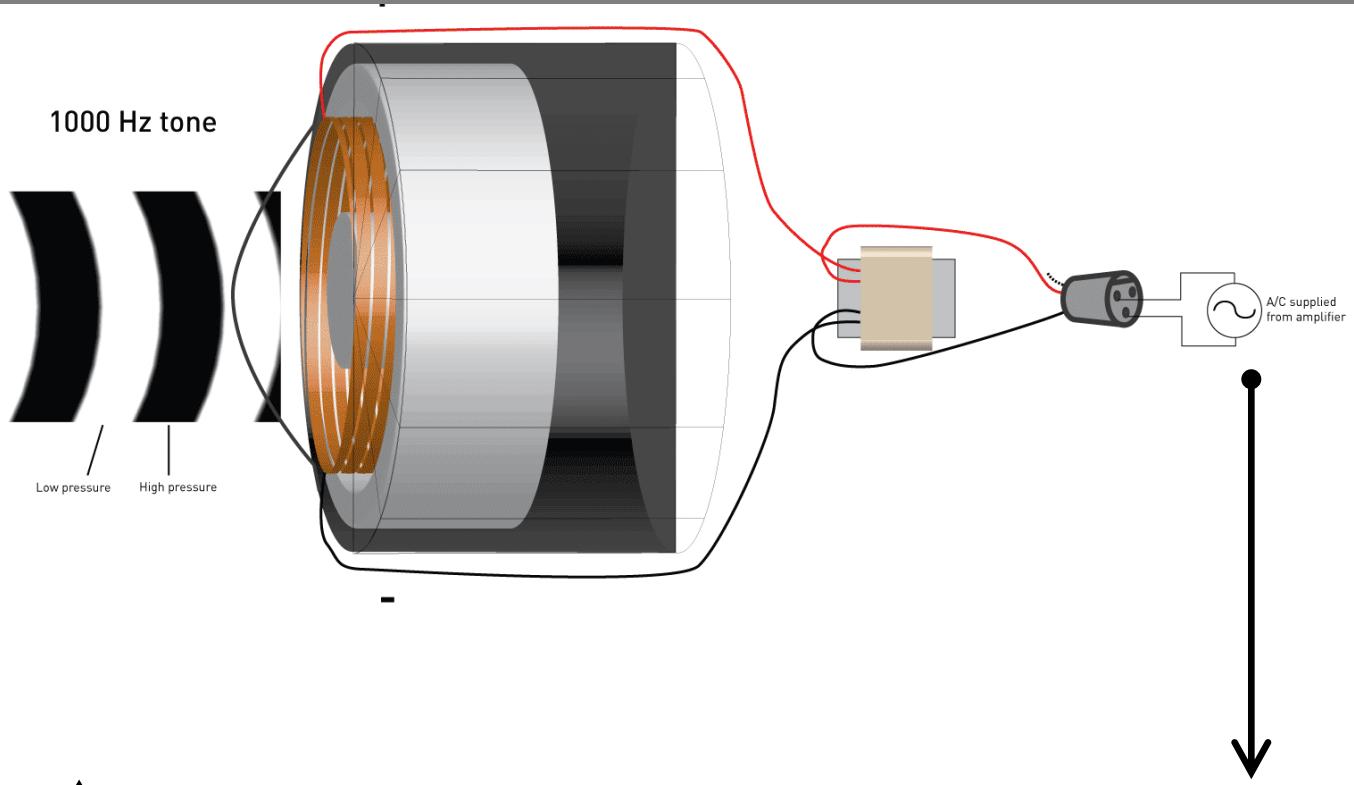
We use physical structure, multiple measurements, priors, and learned regularities to resolve what the forward process made ambiguous.

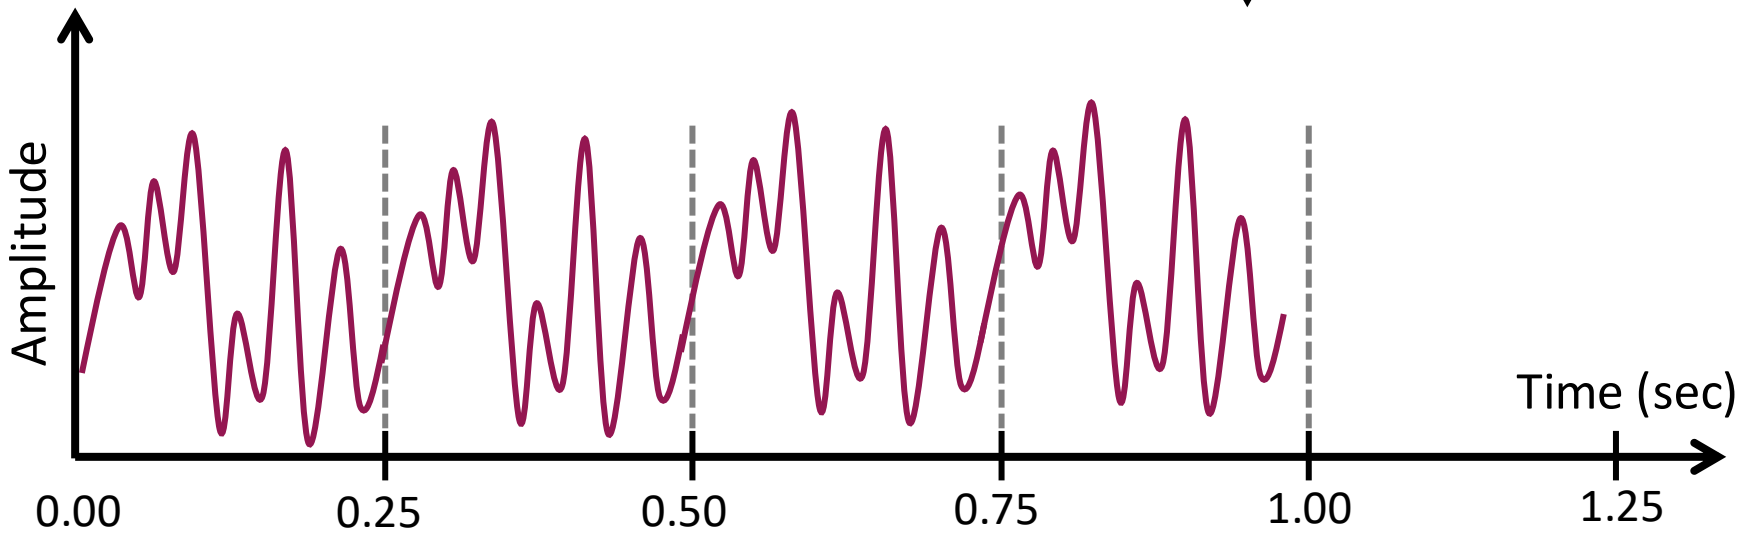
**How does information about the physical world become encoded in “Signals” ?**

**(amplitude, time, frequency, and phase)**

# 1. What is a signal?

# Modeling a Signal



# Define “Signal”

A **signal** is any physical quantity that **varies with one or more independent variables** (such as time, space, or both) and conveys information about a system or phenomenon.

# Define “Signal”

A **signal** is any physical quantity that **varies with one or more independent variables** (such as time, space, or both) and conveys information about a system or phenomenon.

A signal is a **mathematical function** of one or more independent variables (like time or space).

# Define “Signal”

A **signal** is any physical quantity that **varies with one or more independent variables** (such as time, space, or both) and conveys information about a system or phenomenon.

A signal is a **mathematical function** of one or more independent variables (like time or space).

$$\boxed{\text{Signal} = \text{a function that assigns a value to each point in some domain}}$$

$$\boxed{s : \mathcal{D} \rightarrow \mathcal{V}}$$

where:

- $\mathcal{D}$ = **domain**: what variables index the signal;
- $\mathcal{V}$ = **value space** (often loosely called the range): what value the signal takes at each domain point.

# Define “Dimension” of a signal

The number of independent variable of the function.

**1D signals:**

Audio waveform x(t) -- (time).

Stock price x[n] -- (discrete index).

**2D signals:**

Grayscale image x(m,n) -- (row, column)

**3D signals:**

Video x(m,n,t) (space + time)

# Define “Dimension” of a signal

The number of independent variable of the function.

**1D signals:**

Audio waveform x(t) -- (time).

Stock price x[n] -- (discrete index).

**2D signals:**

Grayscale image x(m,n) -- (row, column)

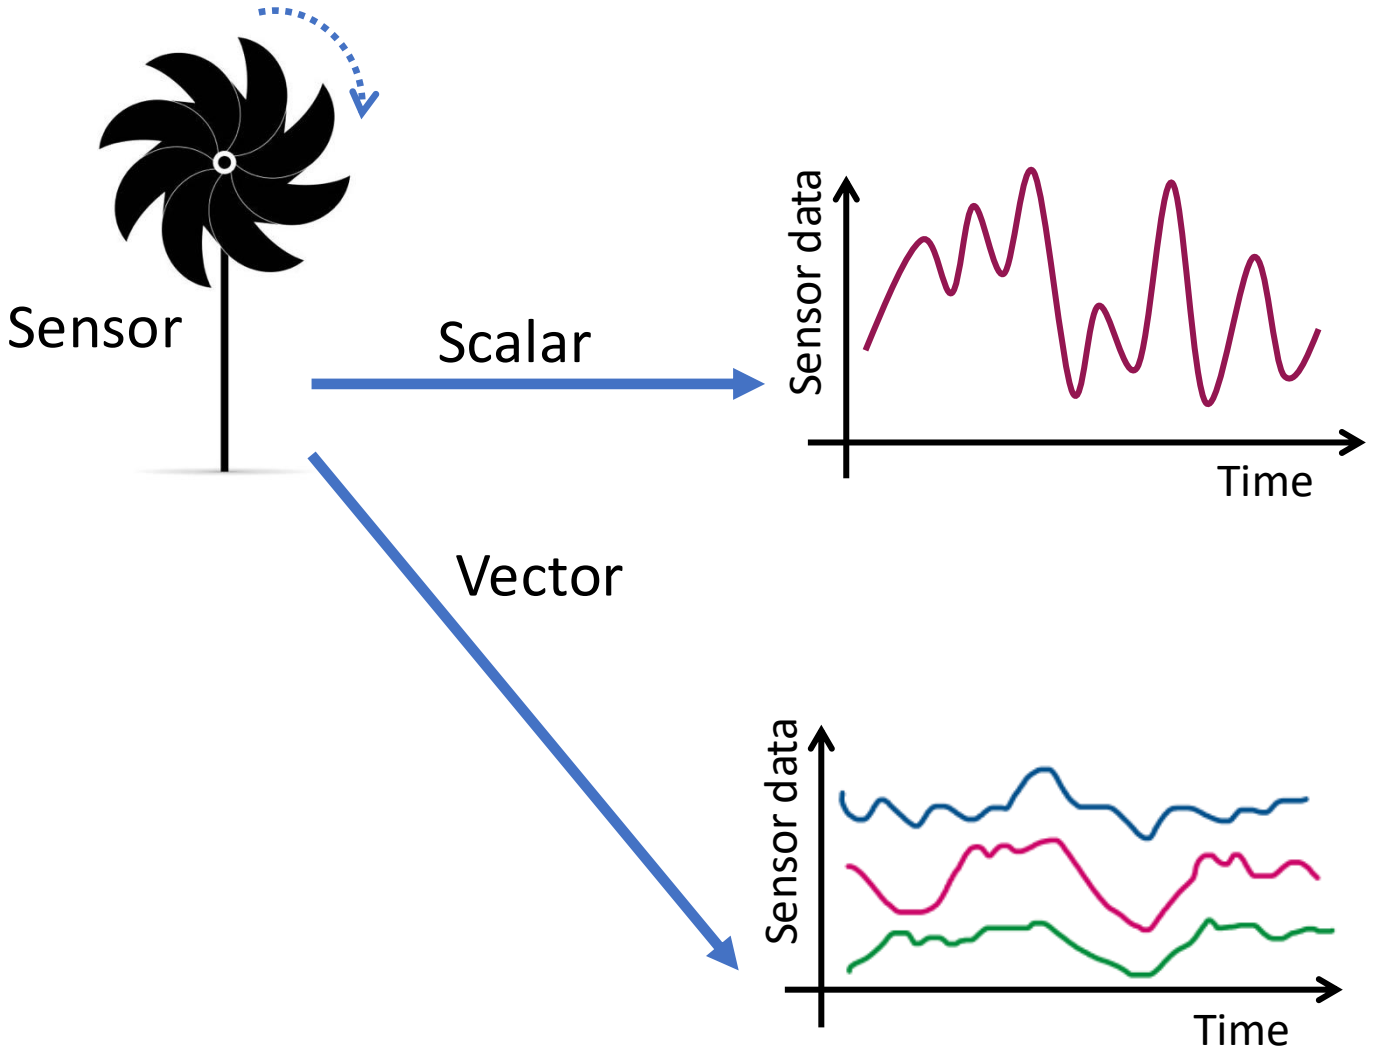
**3D signals:**

Video x(m,n,t) (space + time)

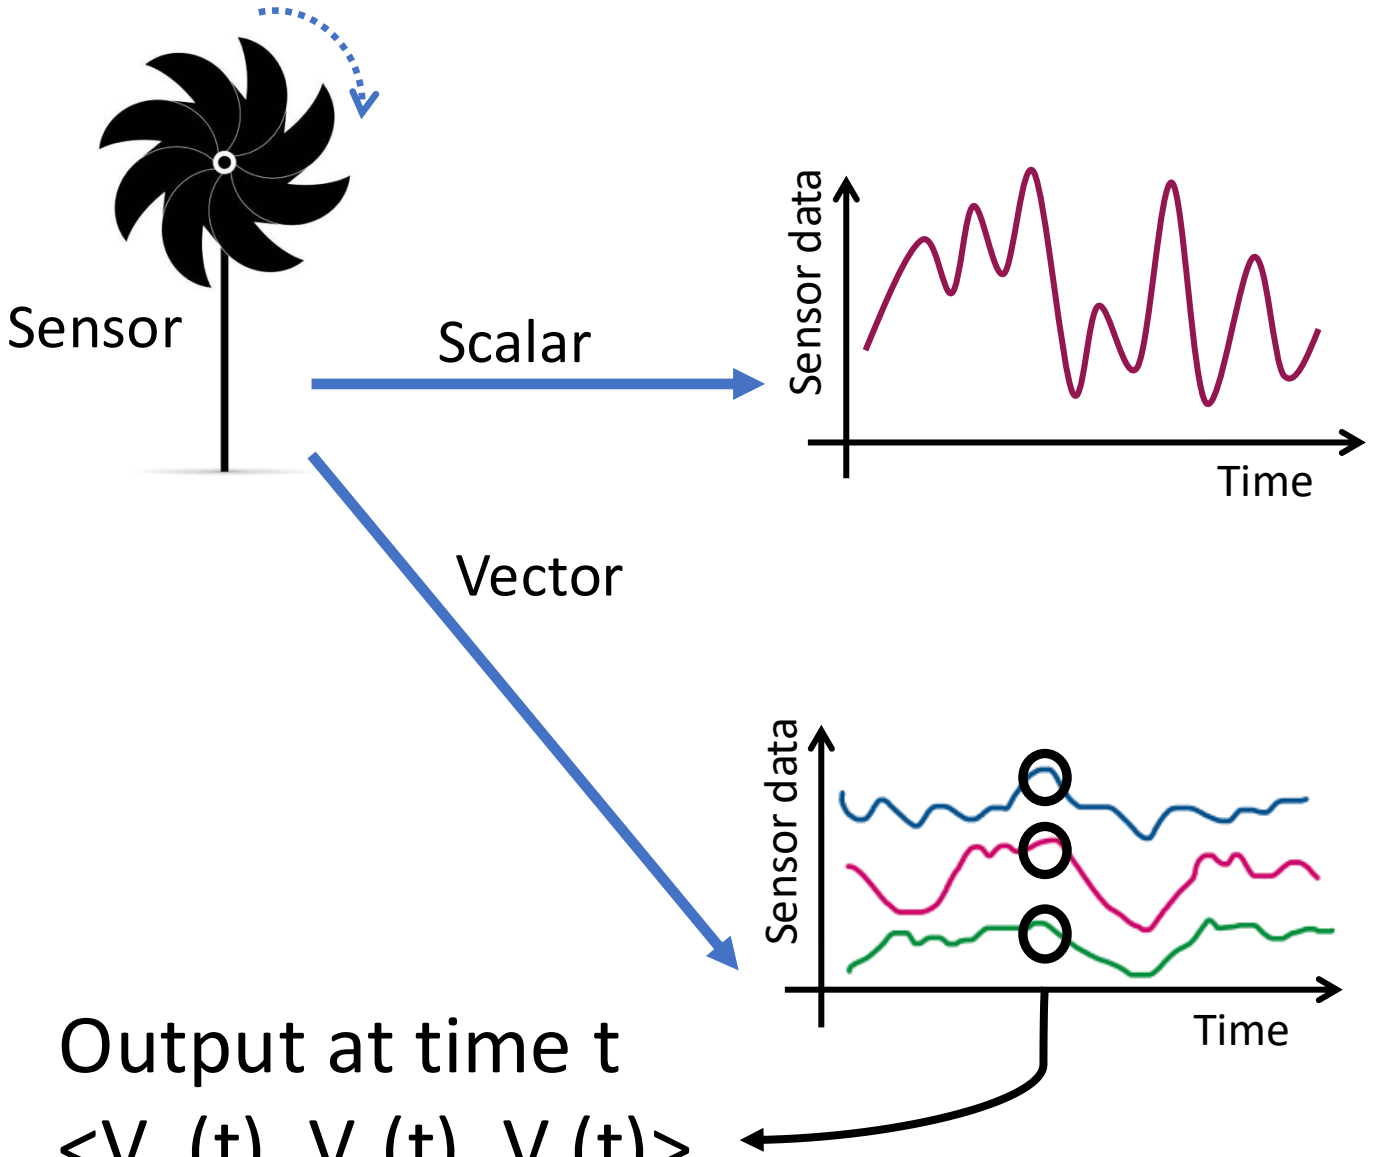
What about a color image?

# Scalar-valued vs. Vector-valued Output of a Sensor

# Sensor data



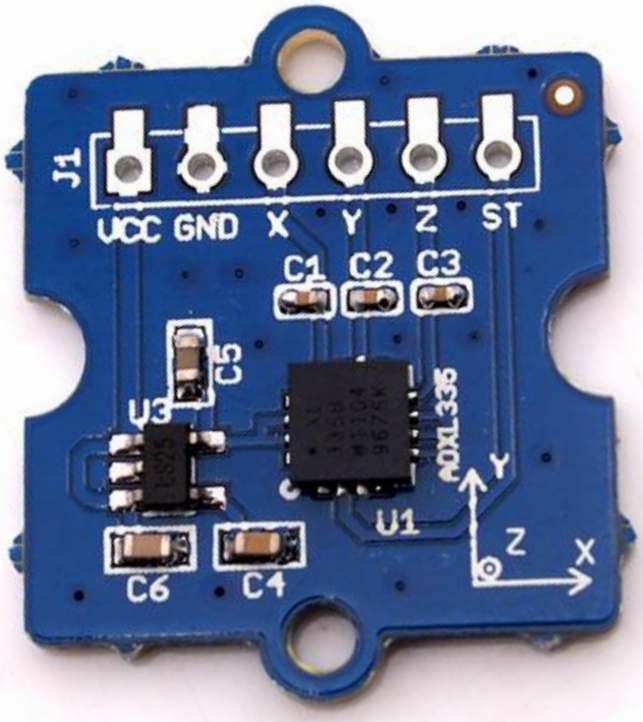
# Sensor data

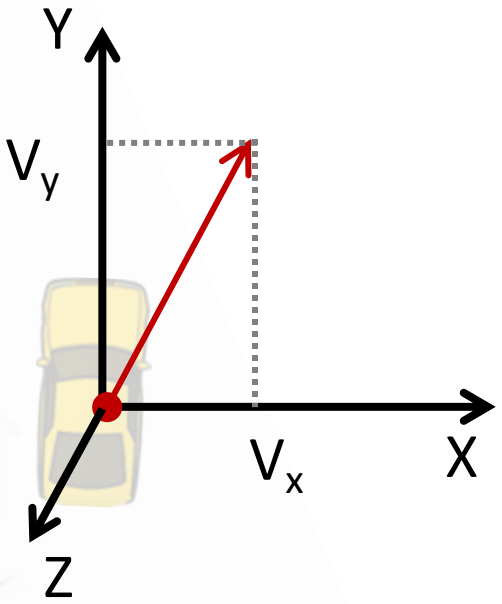

Sensor
Scalar
Sensor data
Time
Vector
Sensor data
Time
Output at time t
<$V_x(t)$, $V_y(t)$, $V_z(t)$>

1D domain, vector output (time → multi-sensor readings):

$$\mathbf{x}(t) = \begin{bmatrix} a_x(t) \\ a_y(t) \\ a_z(t) \end{bmatrix}, \quad n = 1,\ m = 3$$

(Accelerometer signal in 3D).




Output at time t

$\langle a_x(t), a_y(t), a_z(t) \rangle$

3-axis accelerometer sensor

**1D domain, vector output (time → multi-sensor readings):**

$$\mathbf{x}(t) = \begin{bmatrix} a_x(t) \\ a_y(t) \\ a_z(t) \end{bmatrix}, \quad \text{domain=1, value space=3}$$

(Accelerometer signal in 3D).

**2D domain, vector output (images):**

$$\mathbf{x}(m,n) = \begin{bmatrix} R(m,n) \\ G(m,n) \\ B(m,n) \end{bmatrix}, \quad \text{domain=2, value space=3}$$

(Color image, RGB at each pixel).

**3D domain, vector output (video with color):**

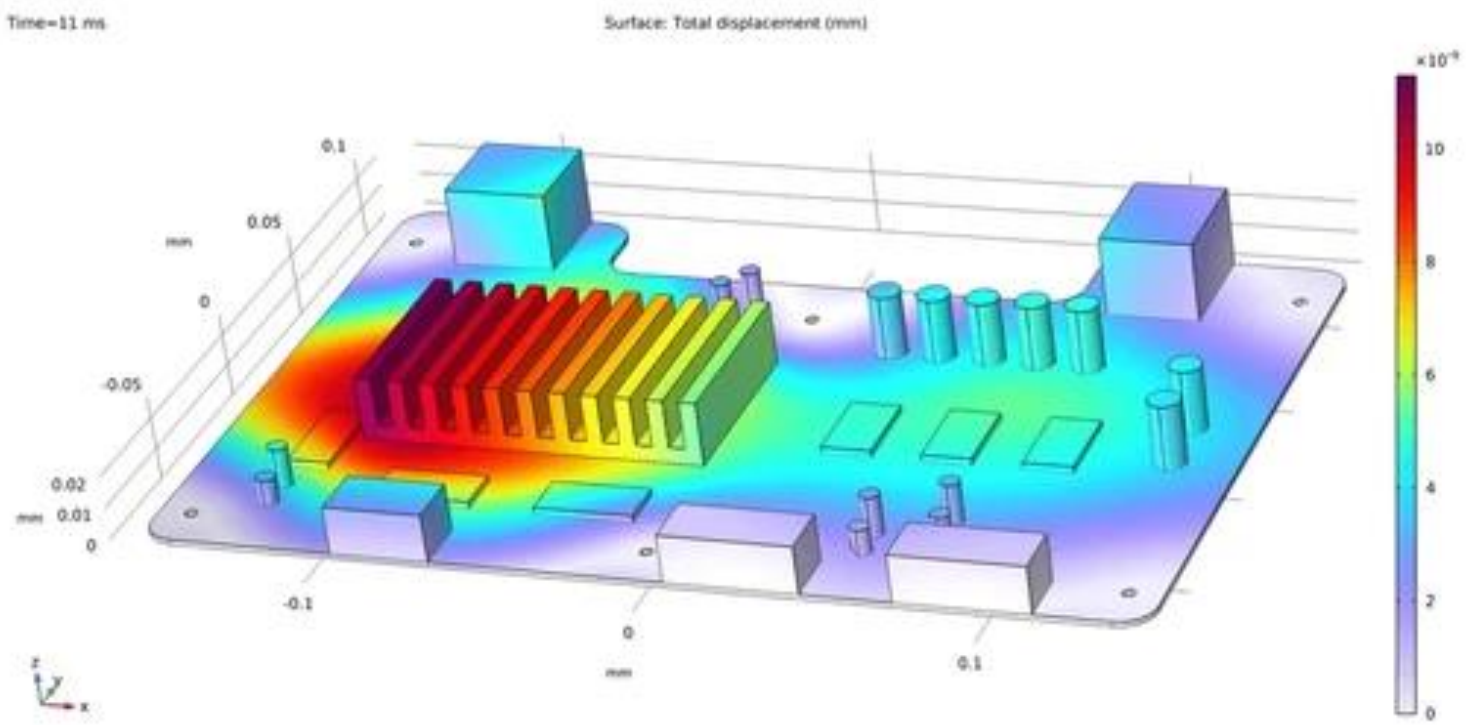
$$\mathbf{x}(m,n,t) = \begin{bmatrix} R(m,n,t) \\ G(m,n,t) \\ B(m,n,t) \end{bmatrix}, \quad \text{domain=3, value space=3}$$

**Higher-dimensional sensing (hyperspectral image):**

$$\mathbf{x}(m,n) = \begin{bmatrix} I_1(m,n) \\ I_2(m,n) \\ \vdots \\ I_k(m,n) \end{bmatrix}, \quad \text{domain=2, value space=3}$$

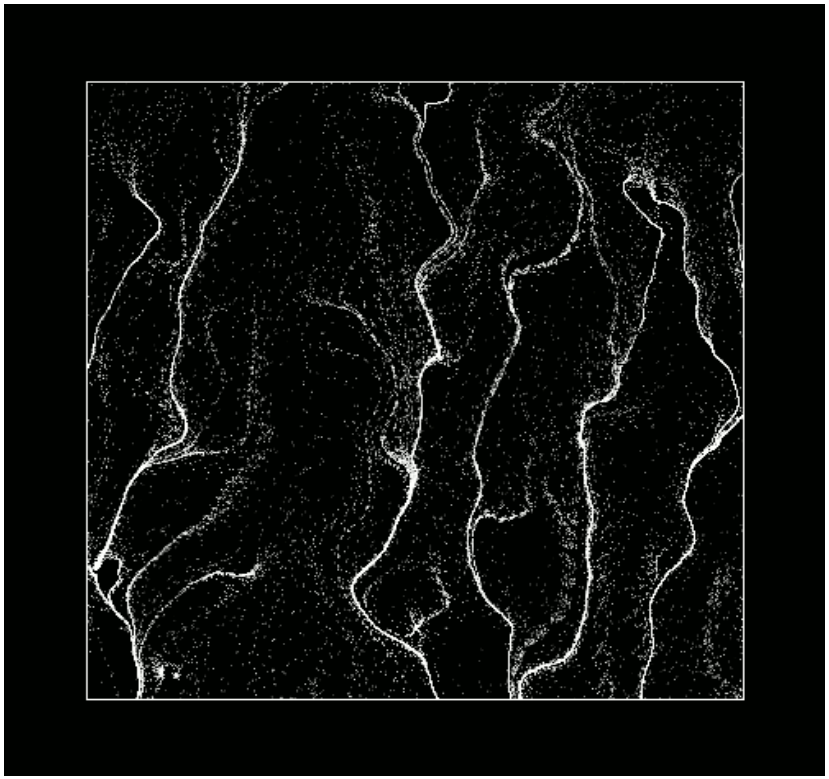
where $k$ could be dozens or hundreds of spectral bands.

# Scalar and Vector fields

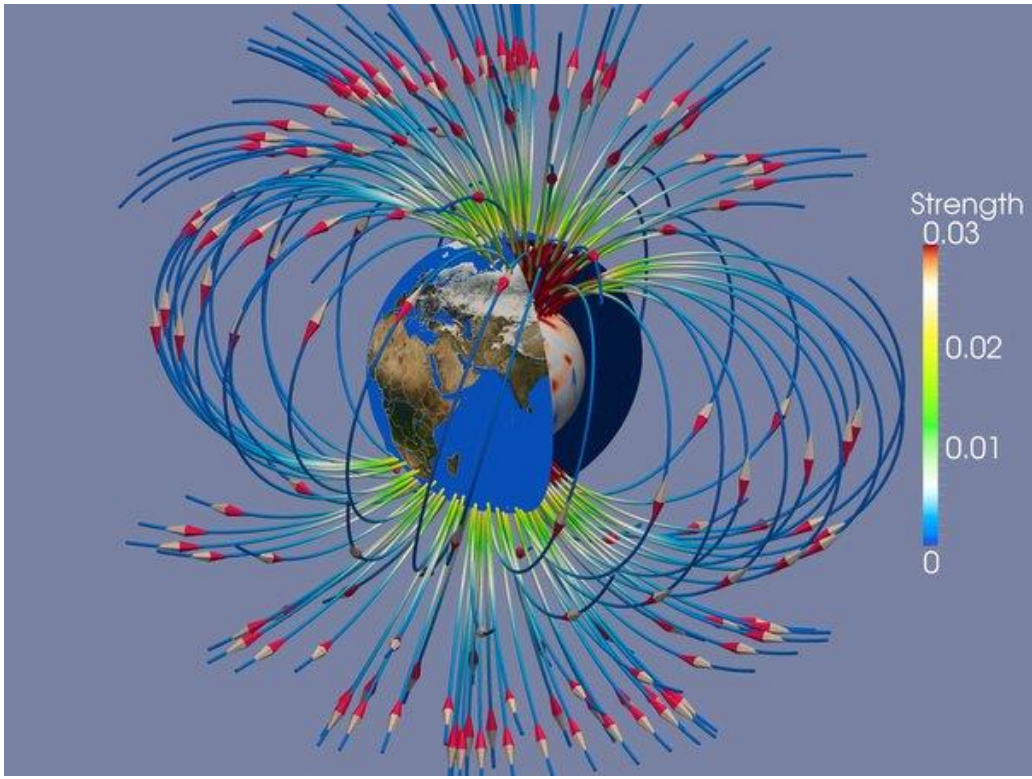



Temperature distribution (scalar field)

# Scalar and Vector fields



Particle flow
(vector field)




Magnetic field
(vector field)

# 2. How Signals ”encodes” physical information

## Amplitude, Frequency, and Phase

# Amplitude: how strongly did the world interact with the signal?

# Path-loss model

$$y(t) = \boldsymbol{A}\, s(t)$$

What determines $\boldsymbol{A}$ ?

- propagation distance
- absorption
- reflectivity
- scattering strength
- source strength
- sensor gain

**Acoustic example — ideal spherical spreading**

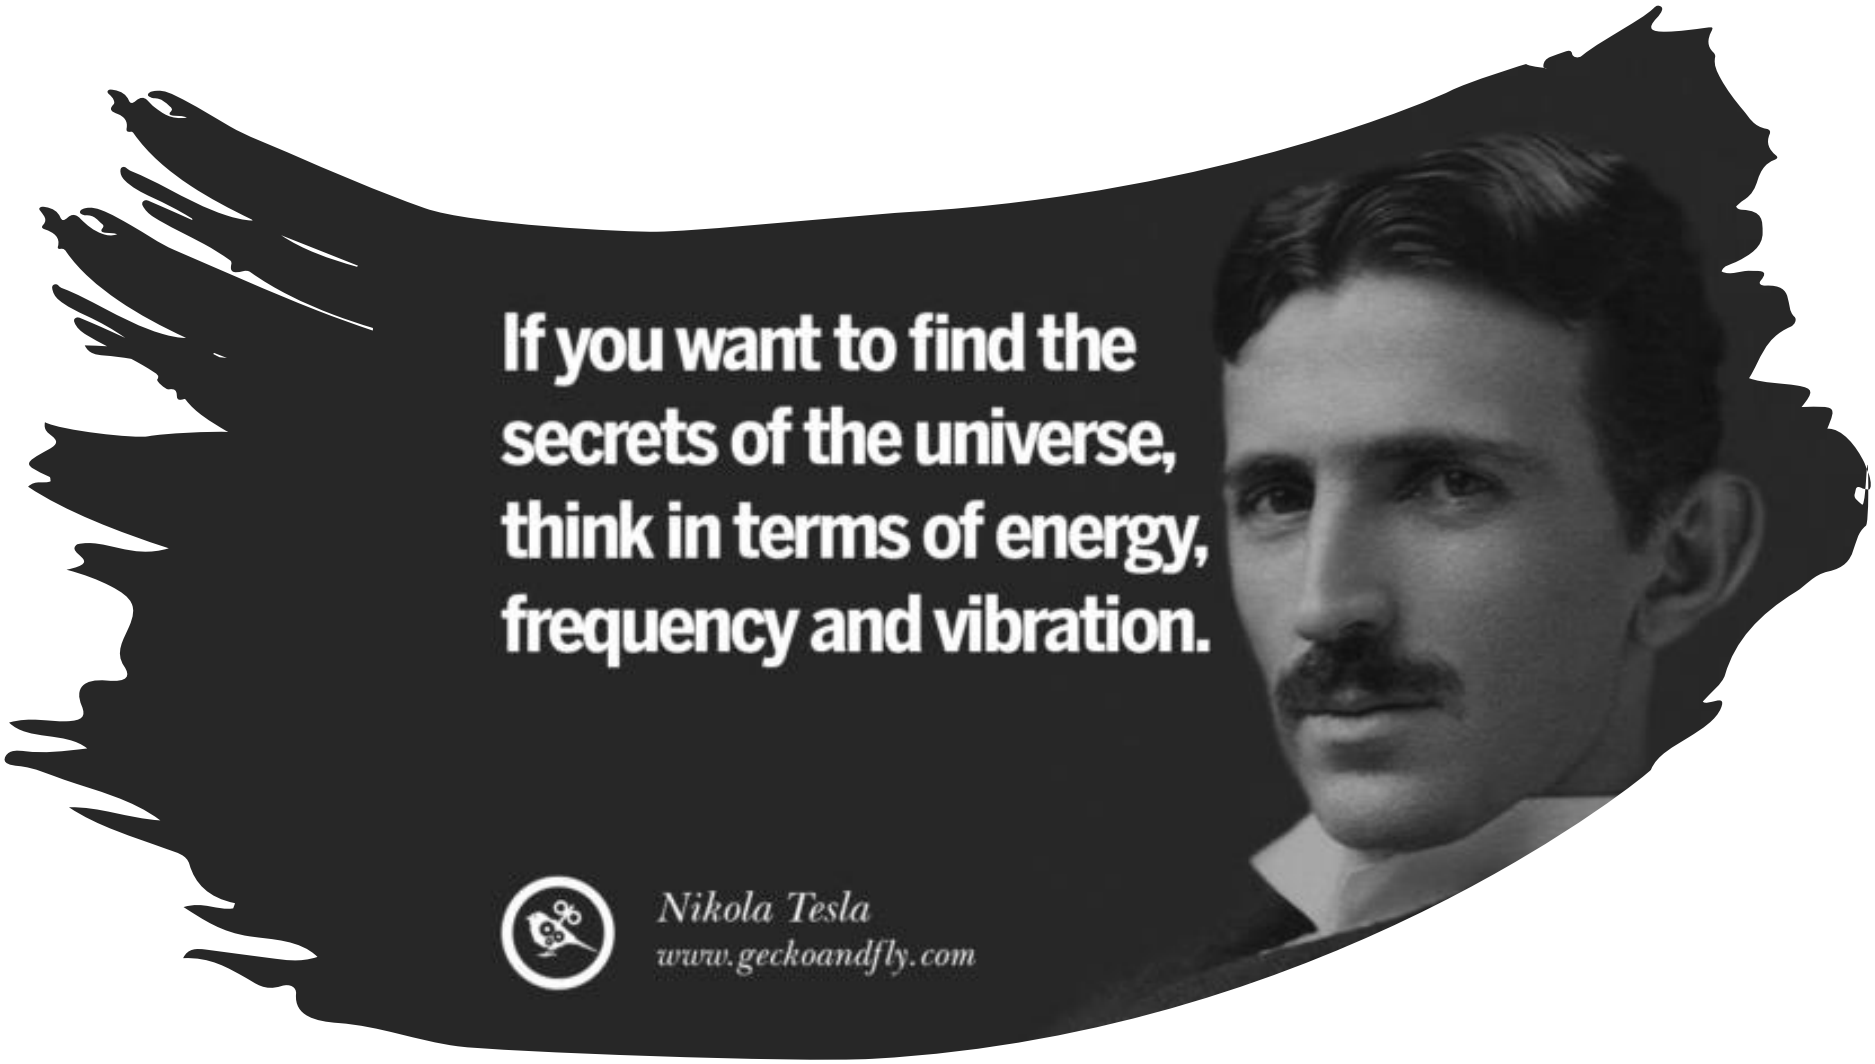
$$p(r) \propto \frac{1}{r}.$$

$$\boxed{y(t) = \alpha s(t) + n(t)}$$

$$\boxed{\mathcal{F}(x;\theta) = \alpha(x;\theta)s(t)}$$

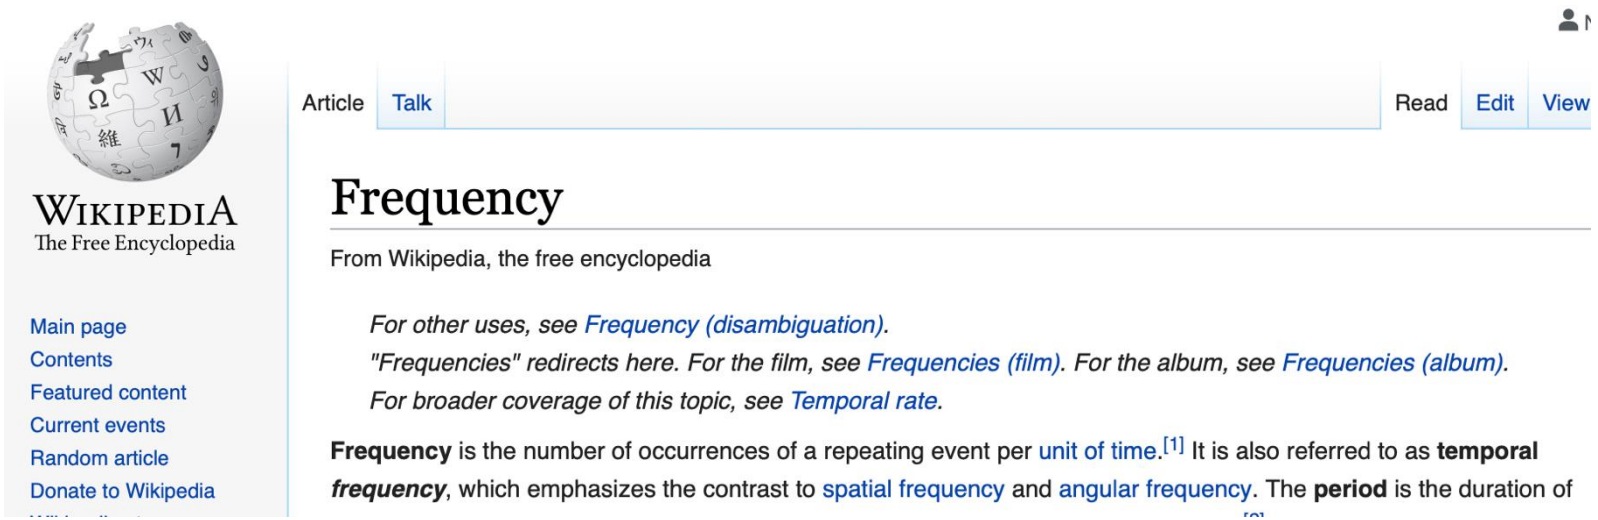
**Amplitude contains physical information — but it is usually confounded.**

*A weak radar return could mean: a distant target · a poorly reflecting target · propagation loss · antenna orientation.*




# What is frequency?

# Definition of Frequency



WIKIPEDIA
The Free Encyclopedia

Main page
Contents
Featured content
Current events
Random article
Donate to Wikipedia

Article Talk Read Edit View

## Frequency

From Wikipedia, the free encyclopedia

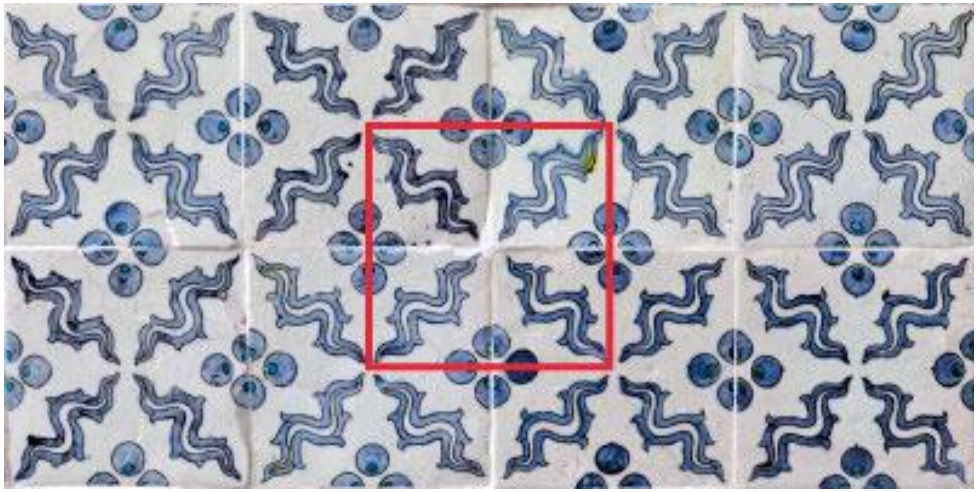
*For other uses, see Frequency (disambiguation).*
*"Frequencies" redirects here. For the film, see Frequencies (film). For the album, see Frequencies (album).*
*For broader coverage of this topic, see Temporal rate.*

**Frequency** is the number of occurrences of a repeating event per unit time.[1] It is also referred to as **temporal *frequency***, which emphasizes the contrast to spatial frequency and angular frequency. The **period** is the duration of

**Frequency** is the number of occurrences of a repeating event per unit time (independent variable).



1D function



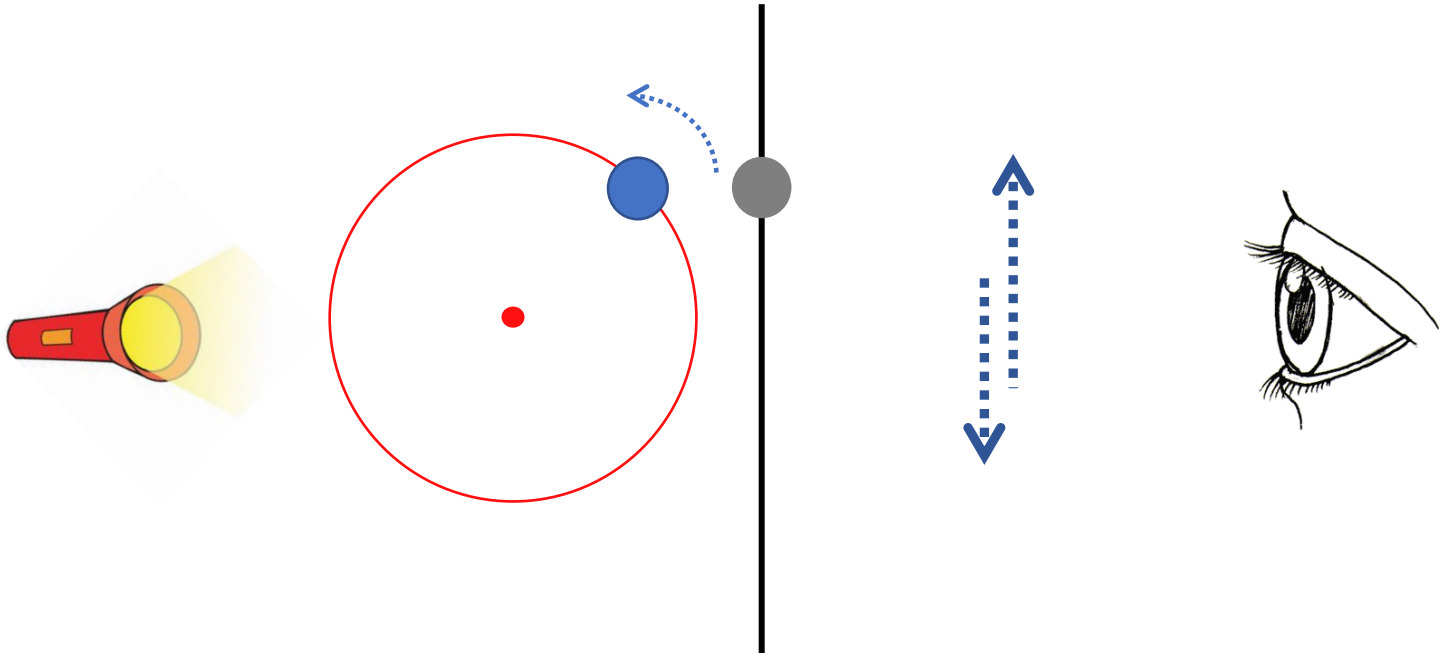
2D function

What can be a simple repeating
“function” == signal?

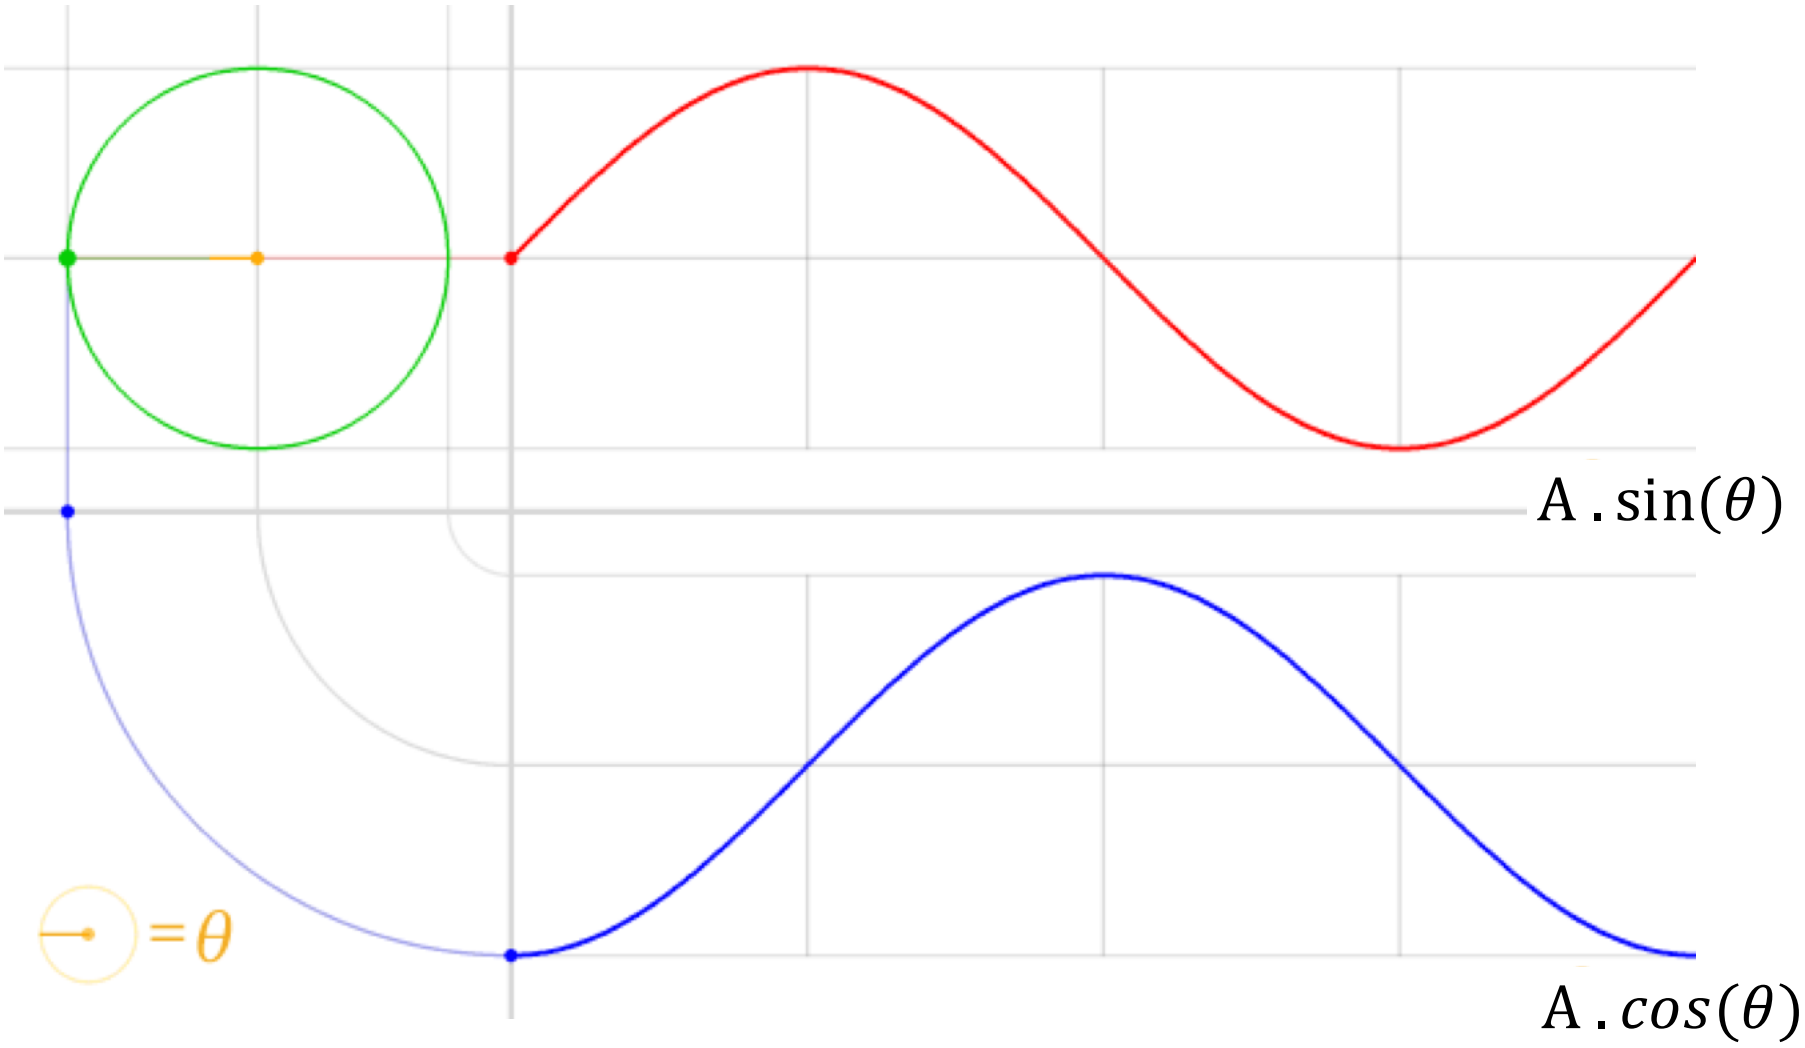
# Frequencies of an arbitrary signal



# Frequency, Amplitude, and Phase

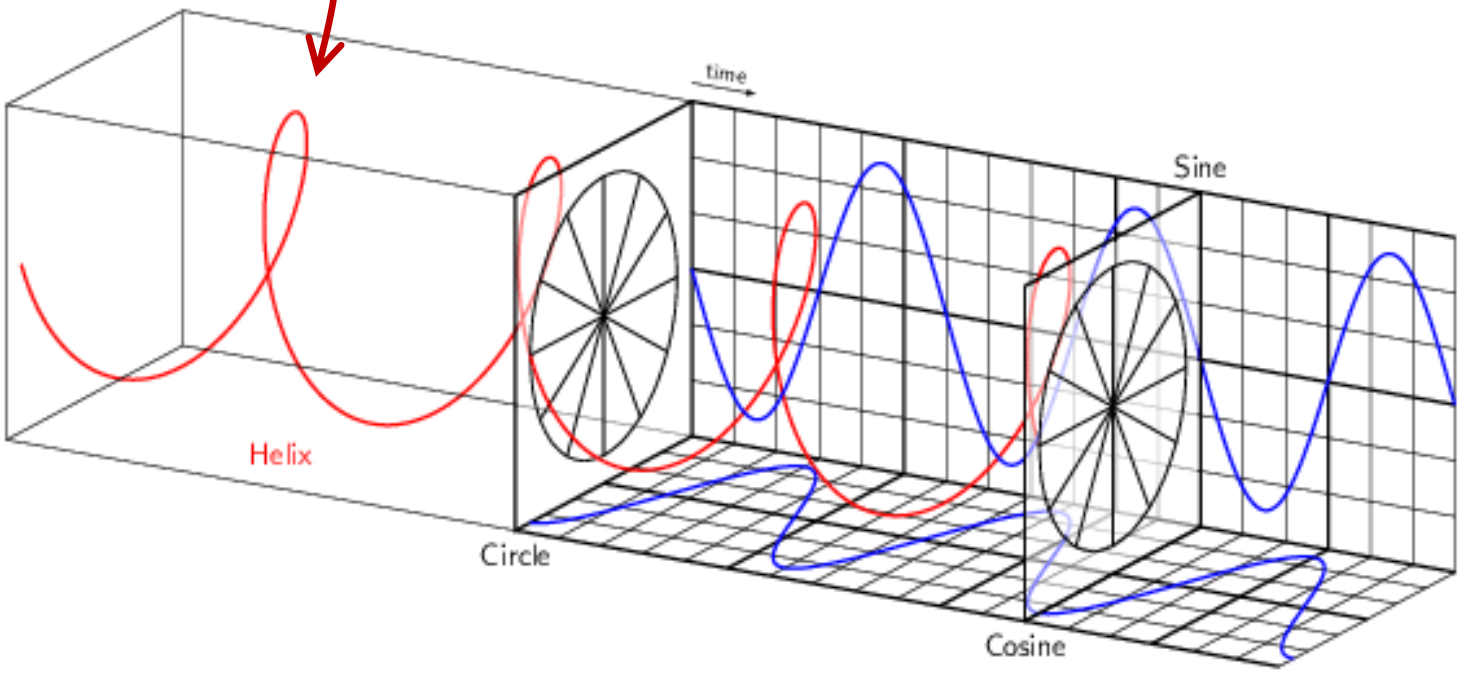


# Model for a signal (frequency, amplitude, and phase)

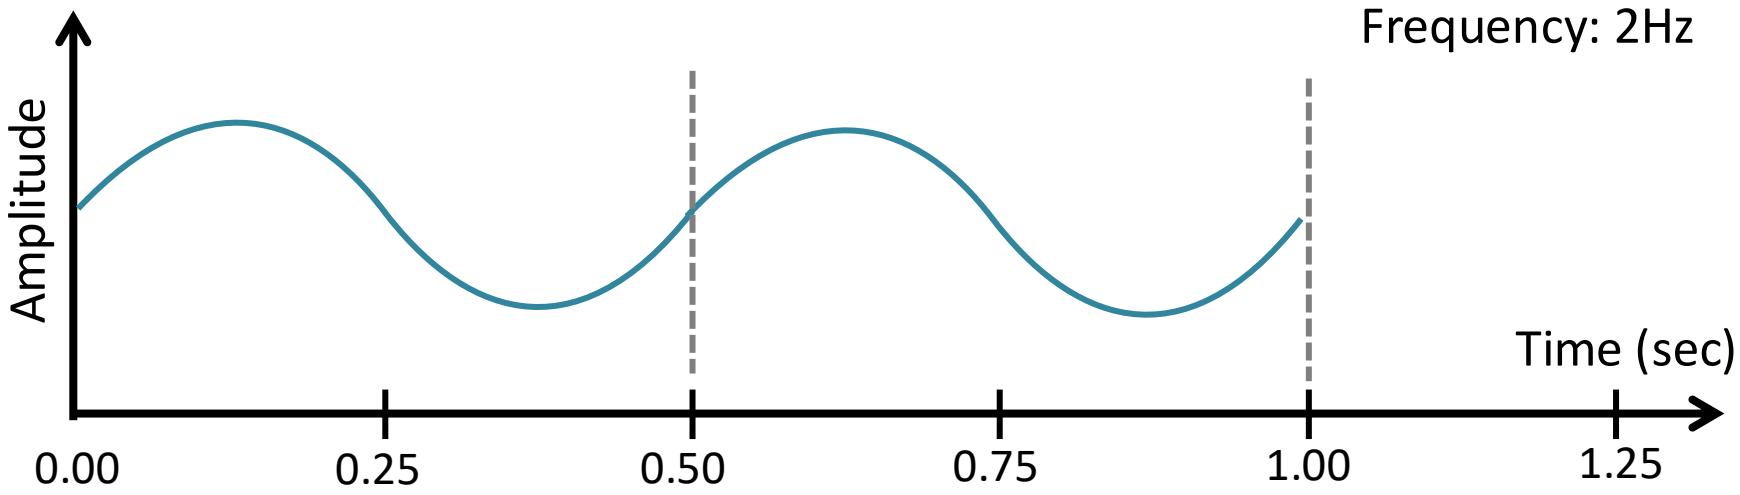


# Model for a signal (frequency, amplitude, and phase)

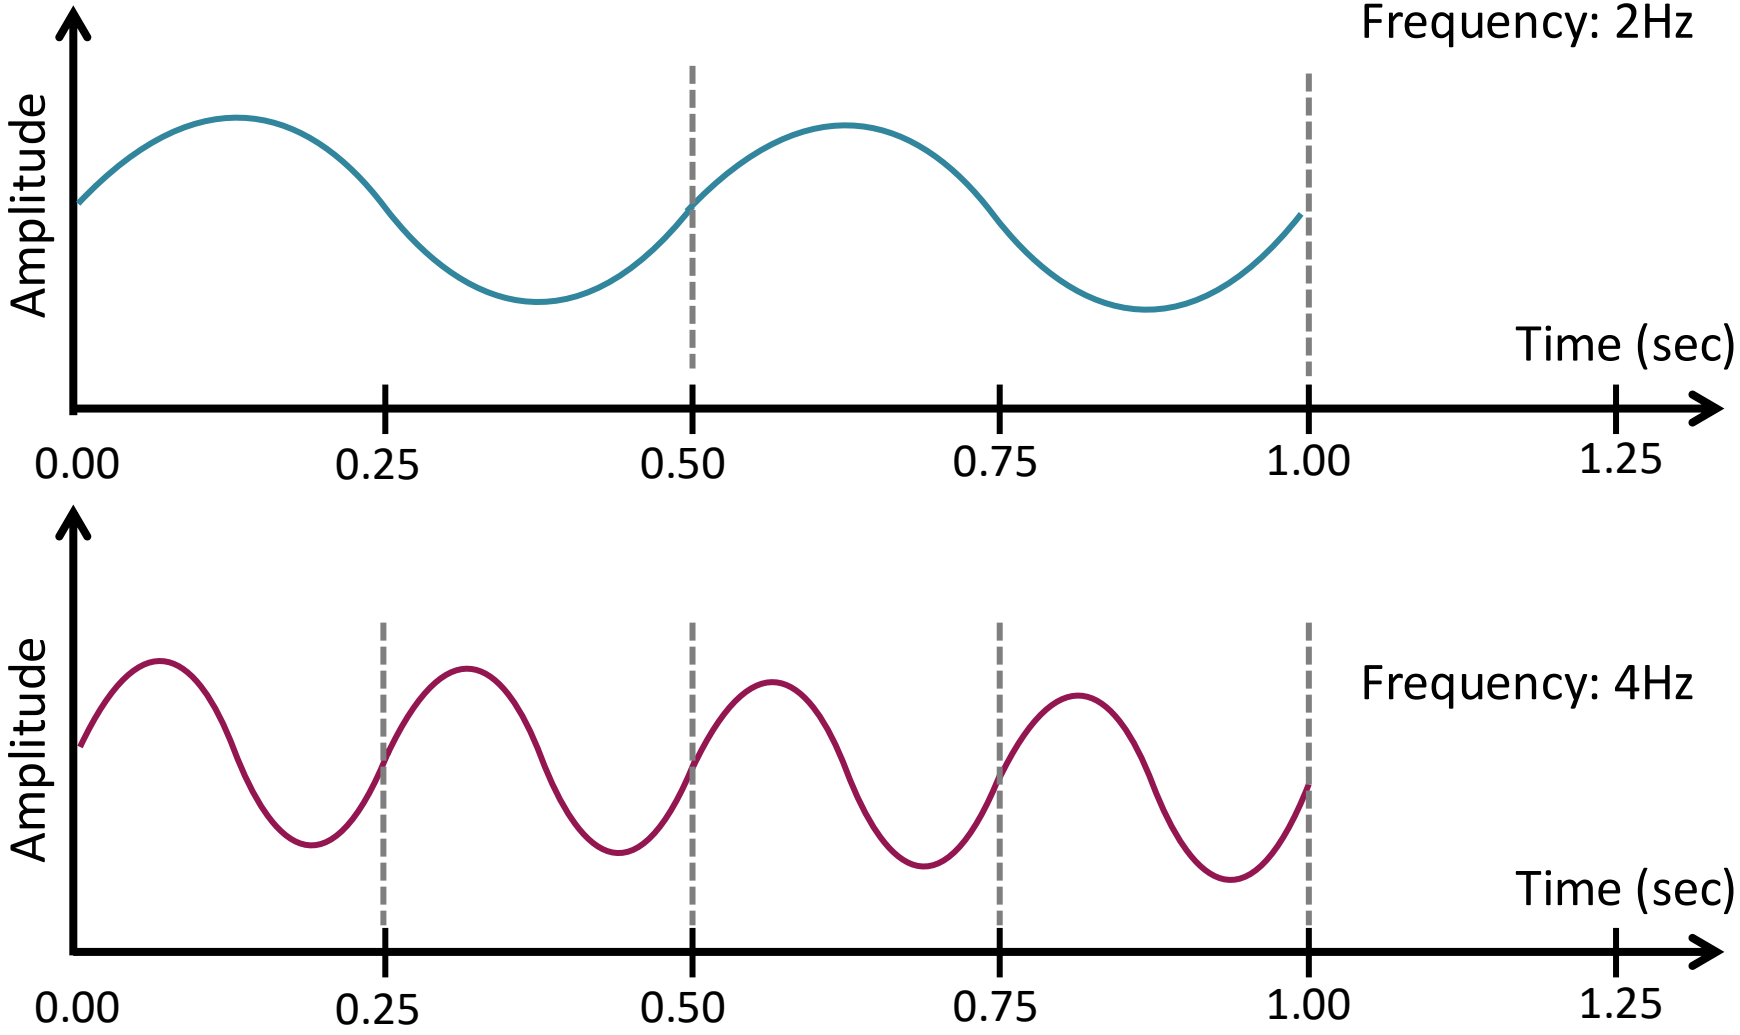
## Simple signal model



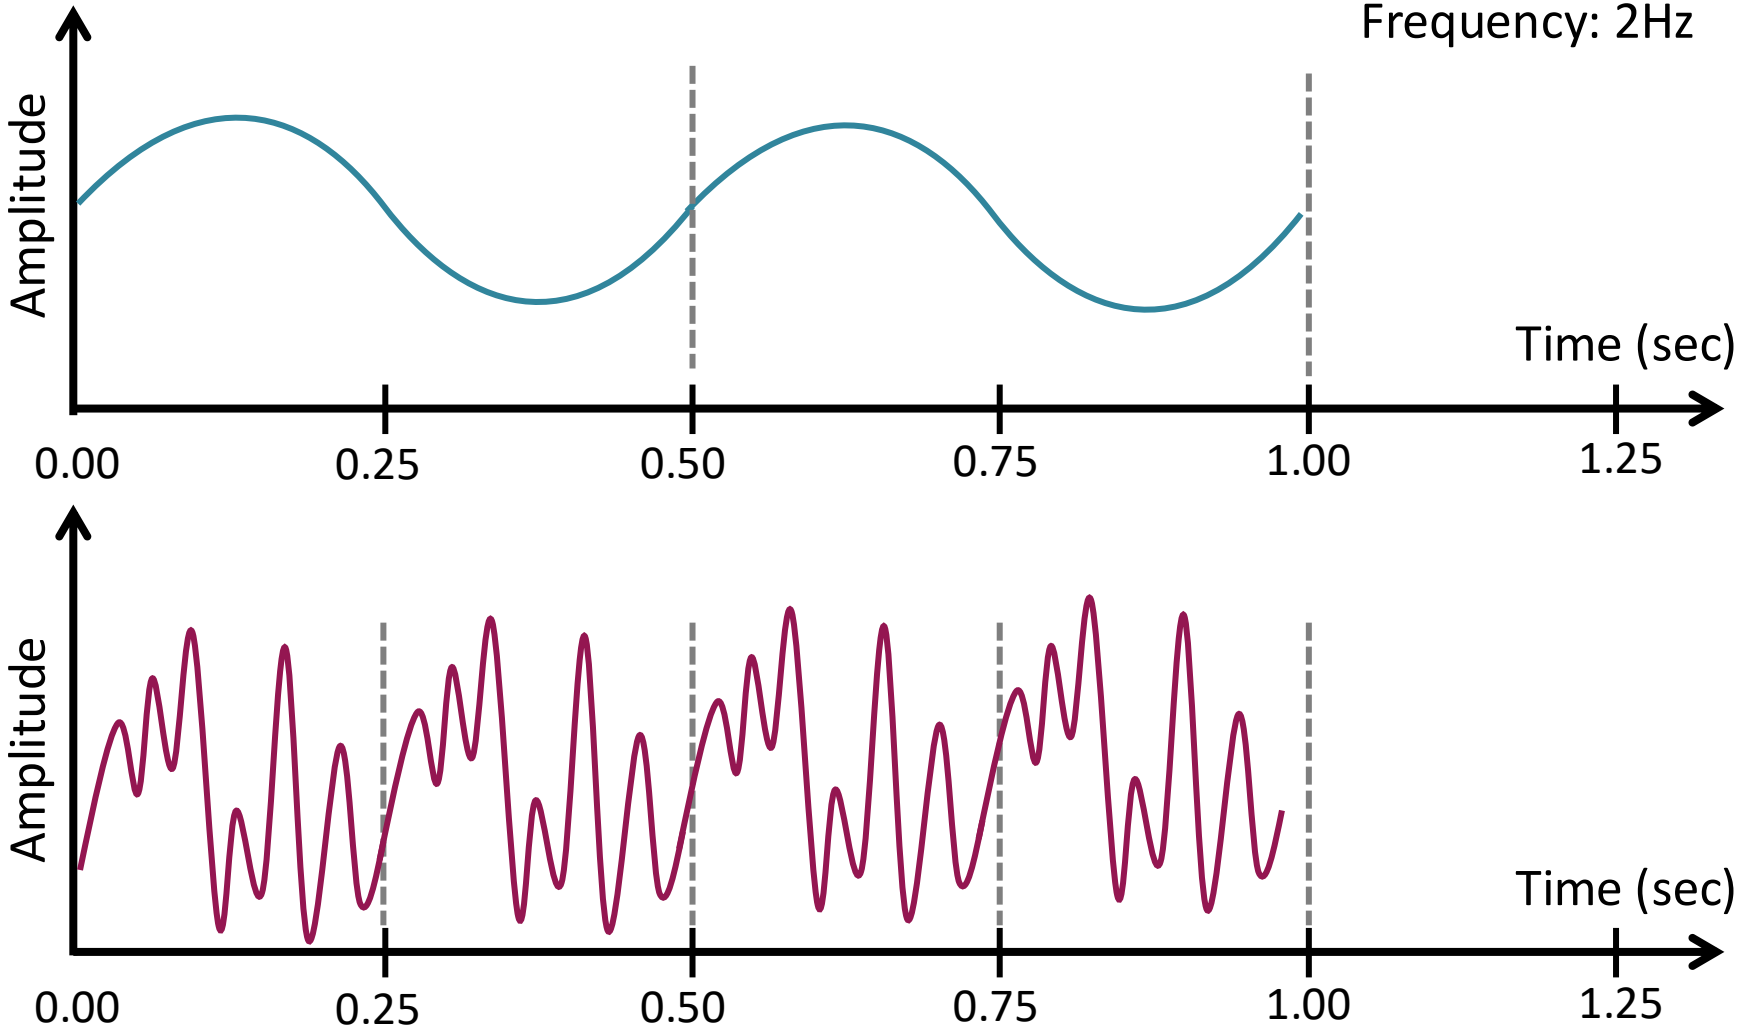
# Frequency, Amplitude, and Phase



# Frequency, Amplitude, and Phase



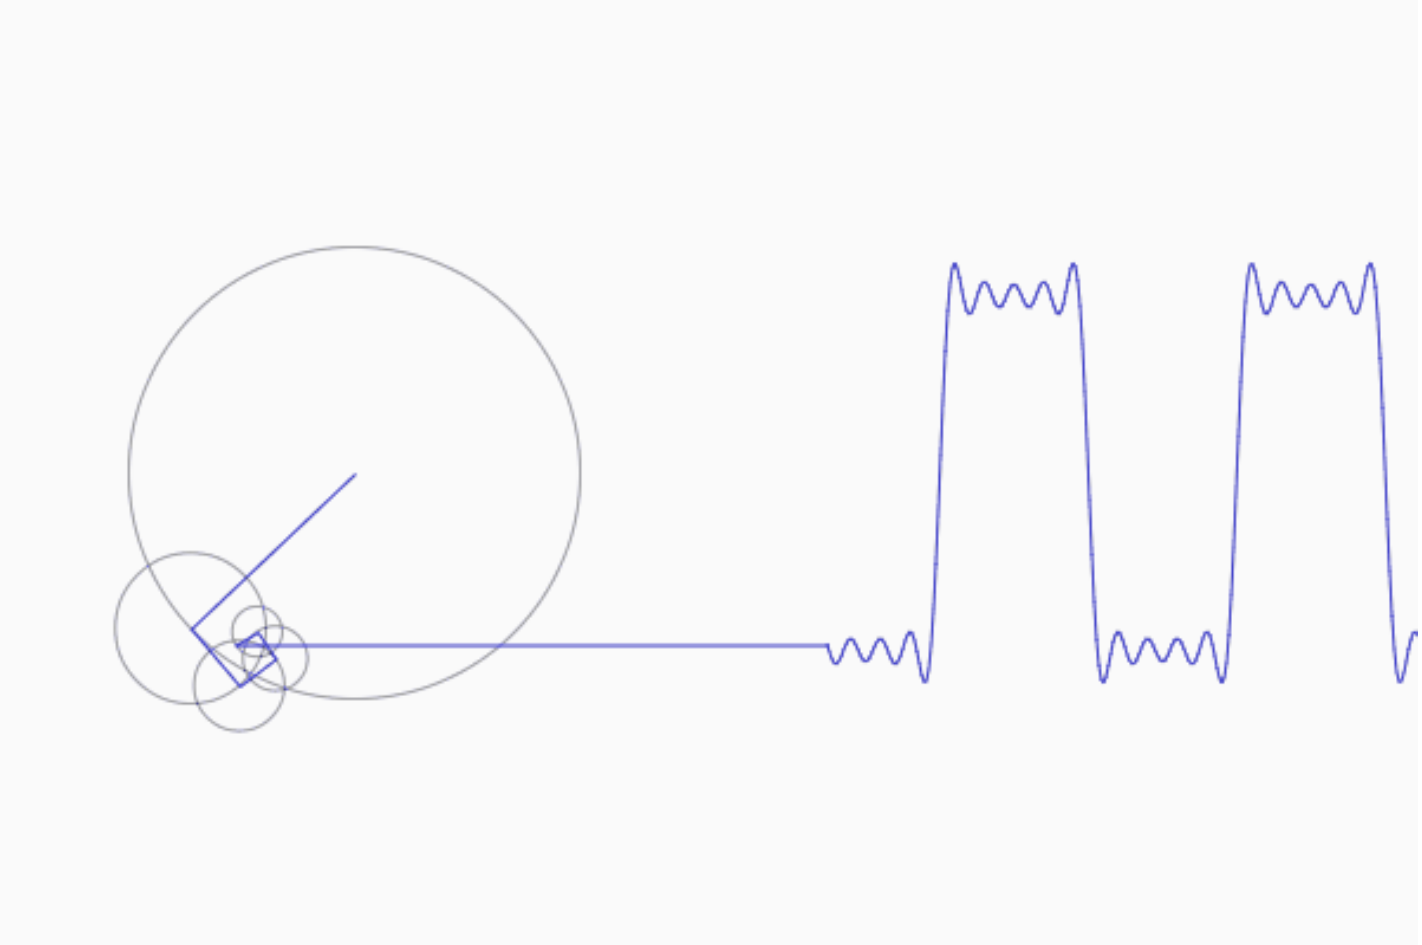
# Frequencies of an arbitrary signal



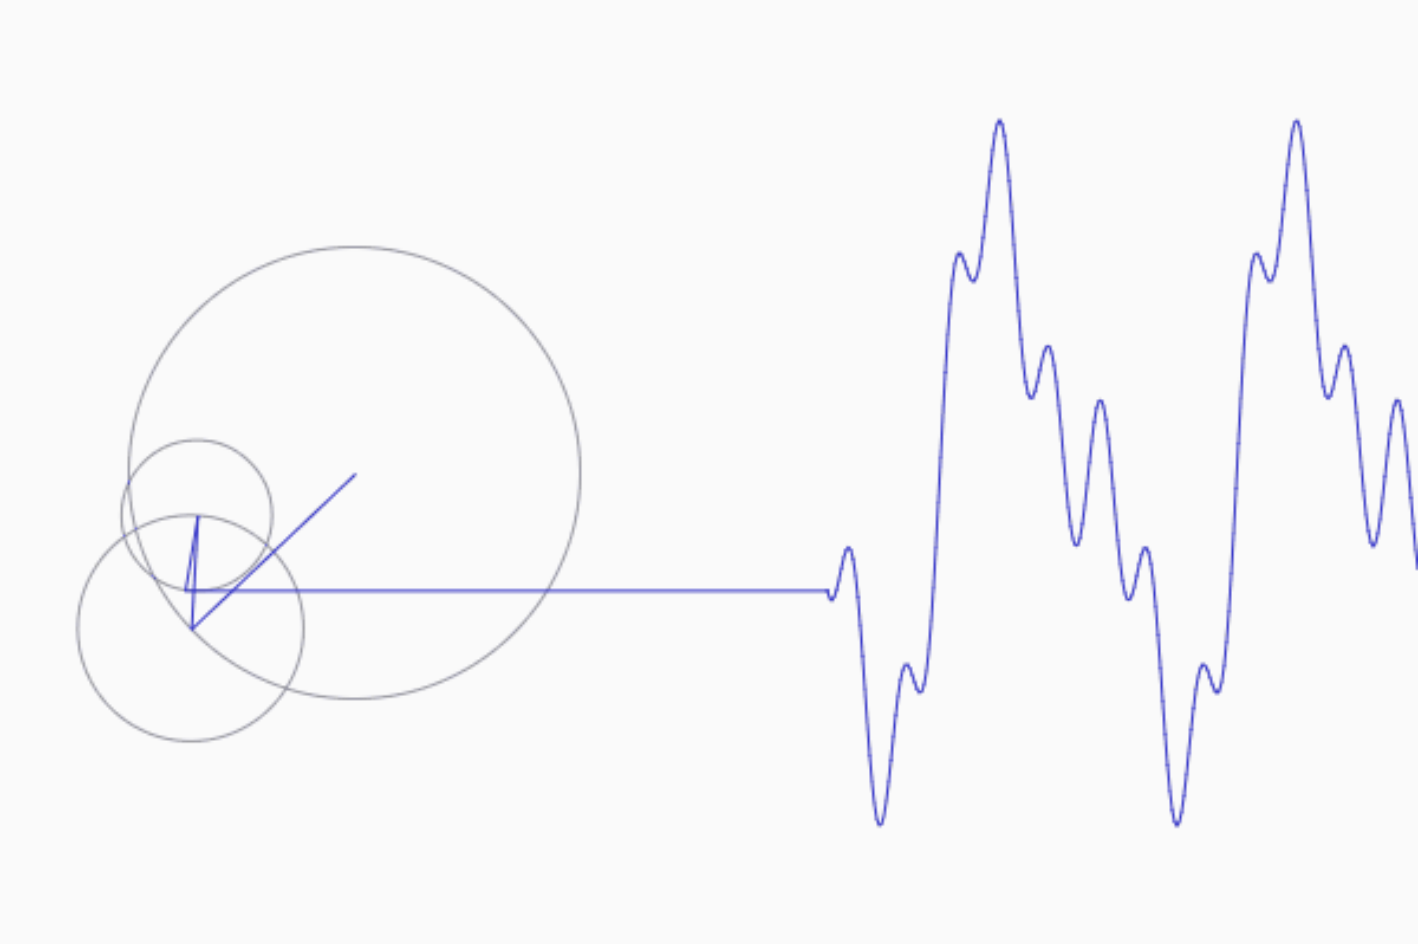
# Frequencies of an arbitrary signal



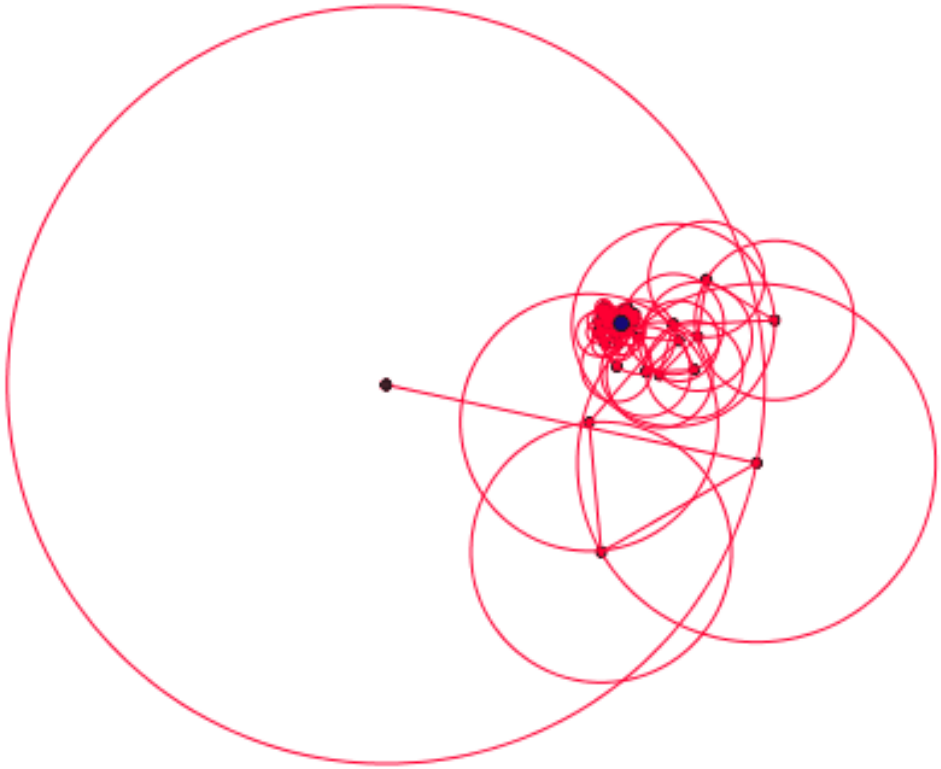
# Frequencies of an arbitrary signal



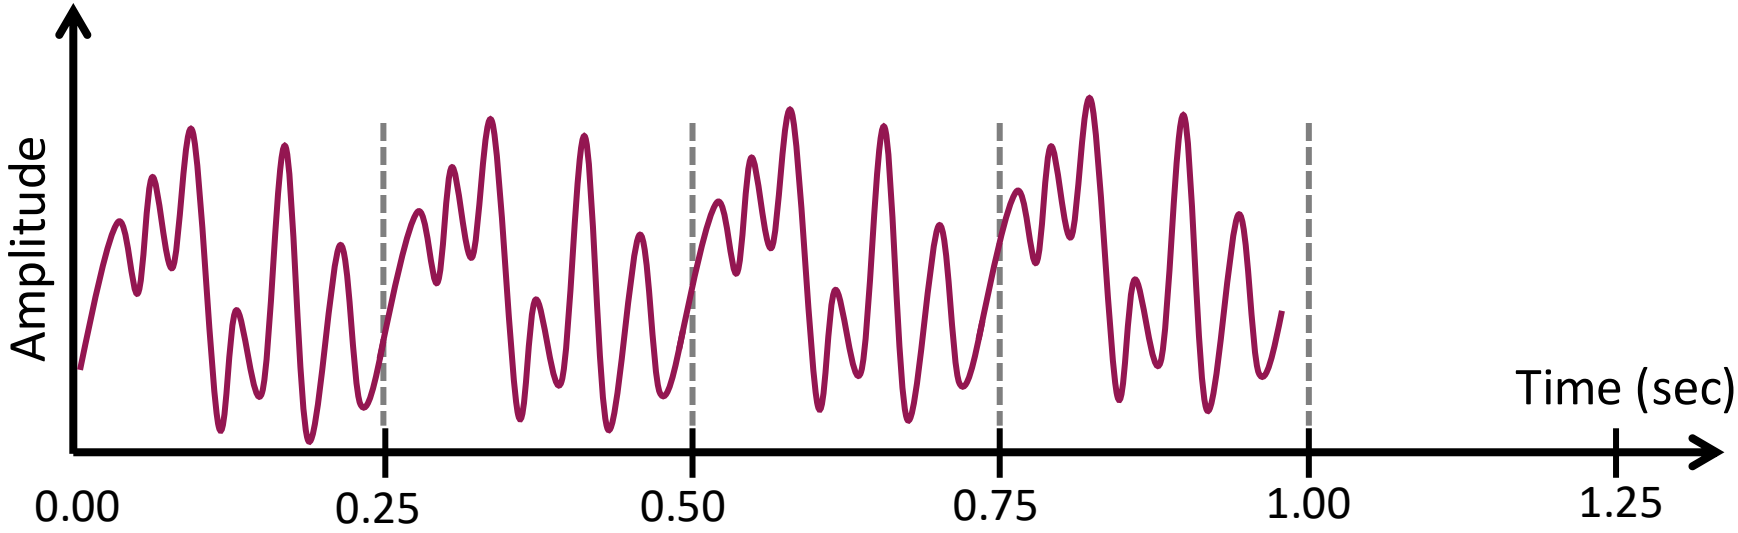
# Frequencies of an arbitrary signal



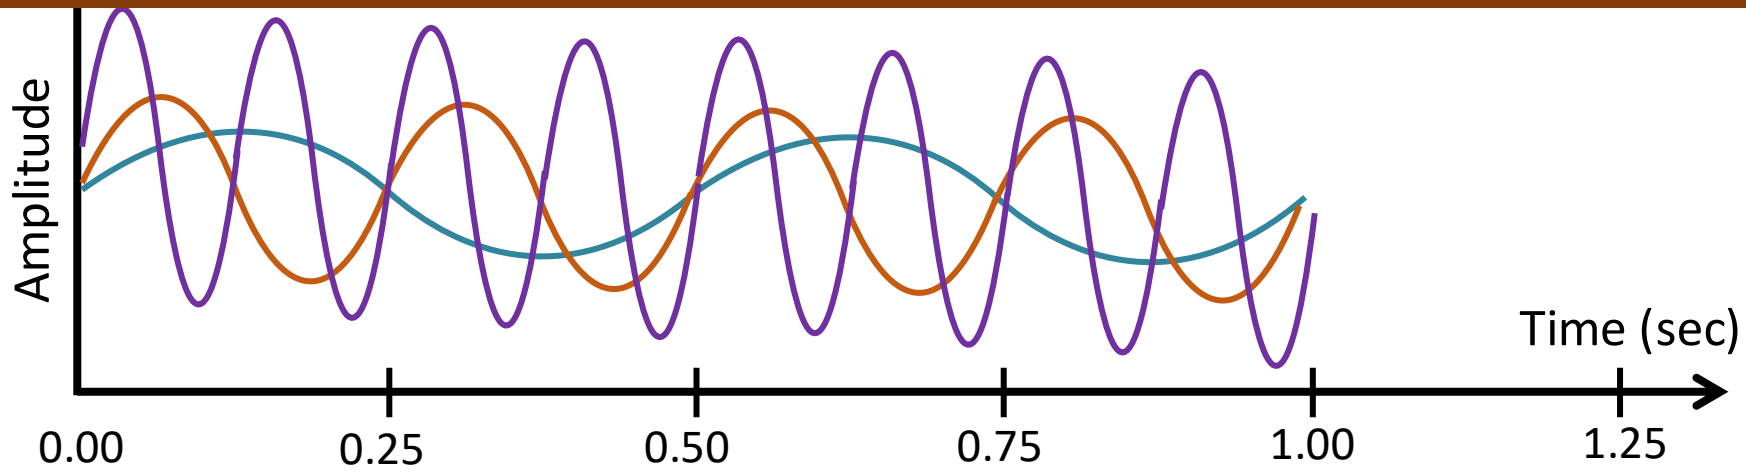
# Frequencies of an arbitrary signal



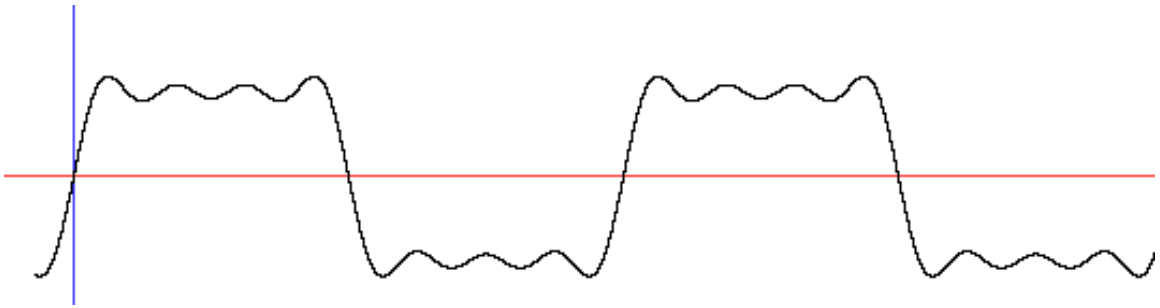
# The concept of the Fourier series




**The key idea: Any function can be written as the sum of sinusoidal functions.**



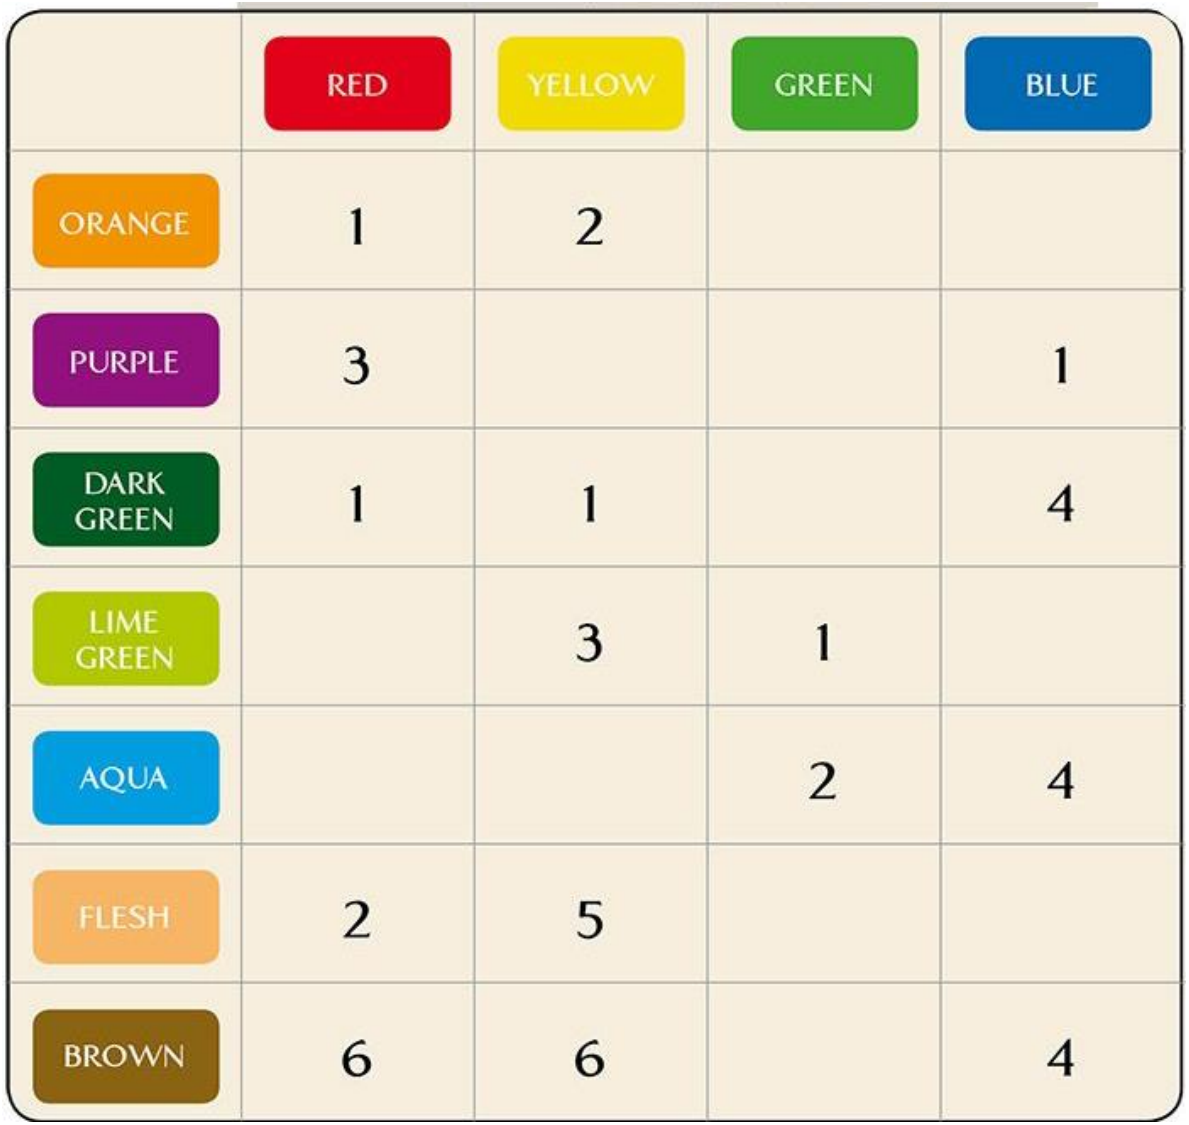
# Time Domain and Frequency Domain



Approx. square wave

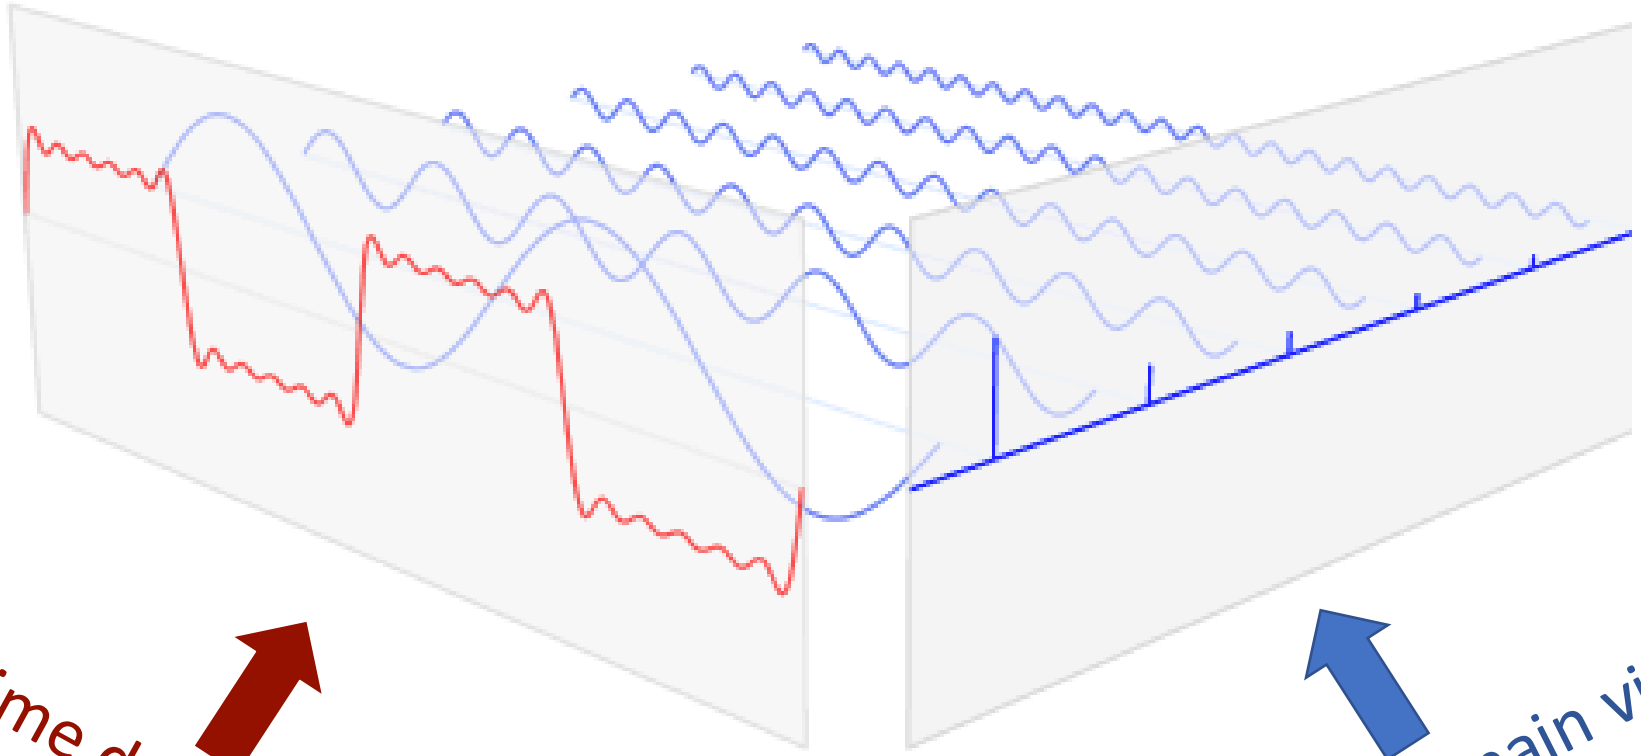
# Analogy: Food coloring chart

Basis for food colors



| | RED | YELLOW | GREEN | BLUE |
|---|---|---|---|---|
| ORANGE | 1 | 2 | | |
| PURPLE | 3 | | | 1 |
| DARK GREEN | 1 | 1 | | 4 |
| LIME GREEN | | 3 | 1 | |
| AQUA | | | 2 | 4 |
| FLESH | 2 | 5 | | |
| BROWN | 6 | 6 | | 4 |

# Time Domain and Frequency Domain



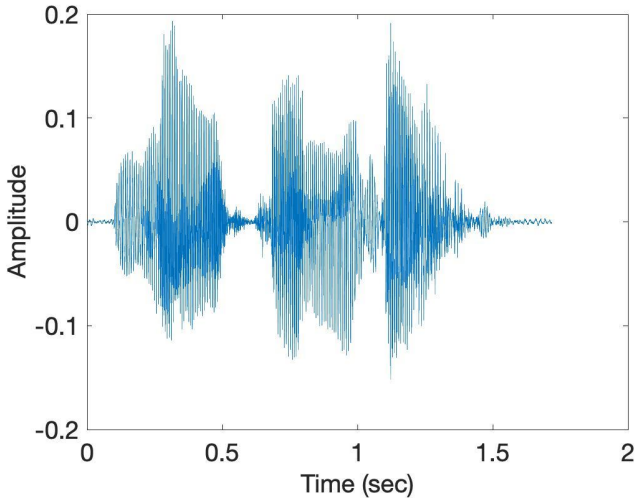
Time domain view

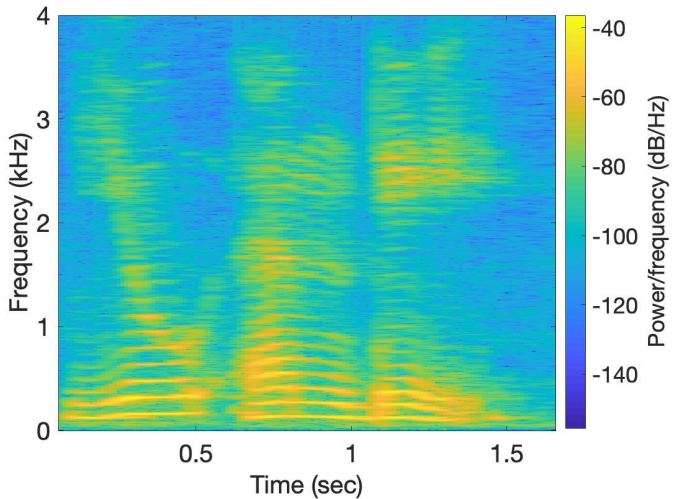
Frequency domain view

# Why do we need the frequency domain view?

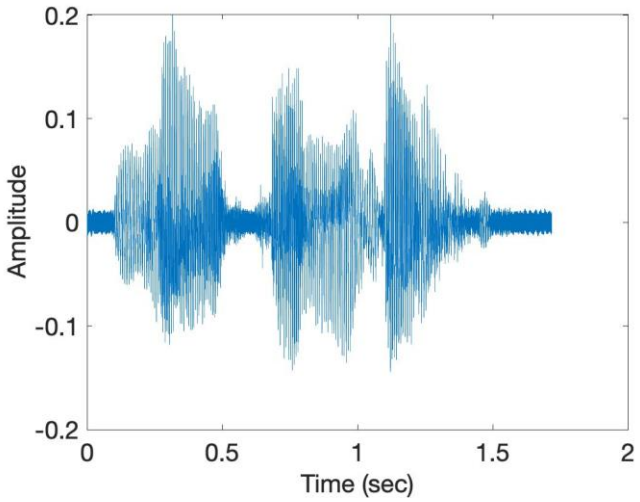
# Why do we need the frequency domain view?

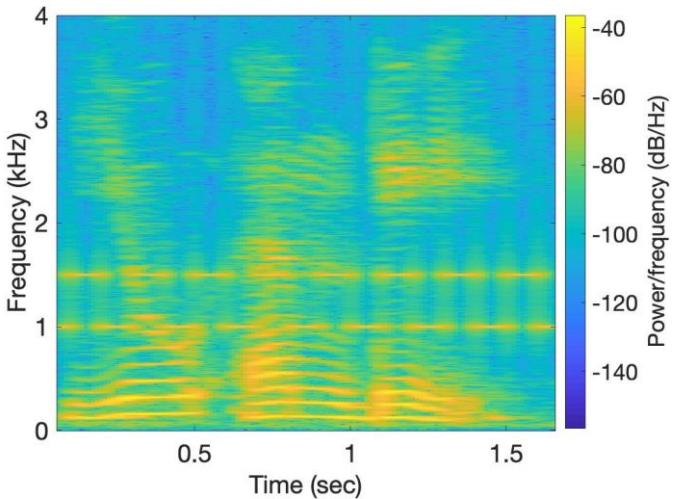
## Sound




## Sound + Noise

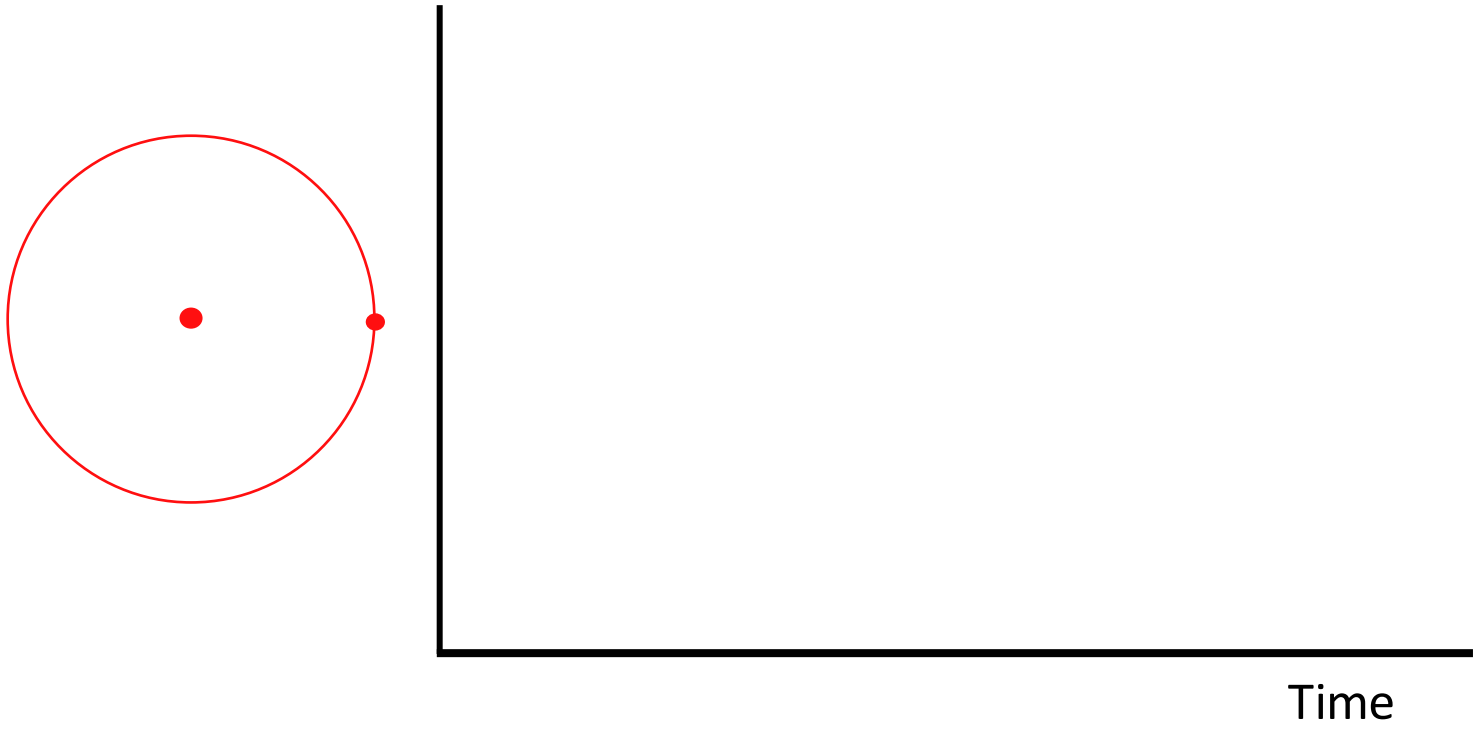



Time domain

Freq. domain

# Model for a signal (frequency, amplitude, and phase)



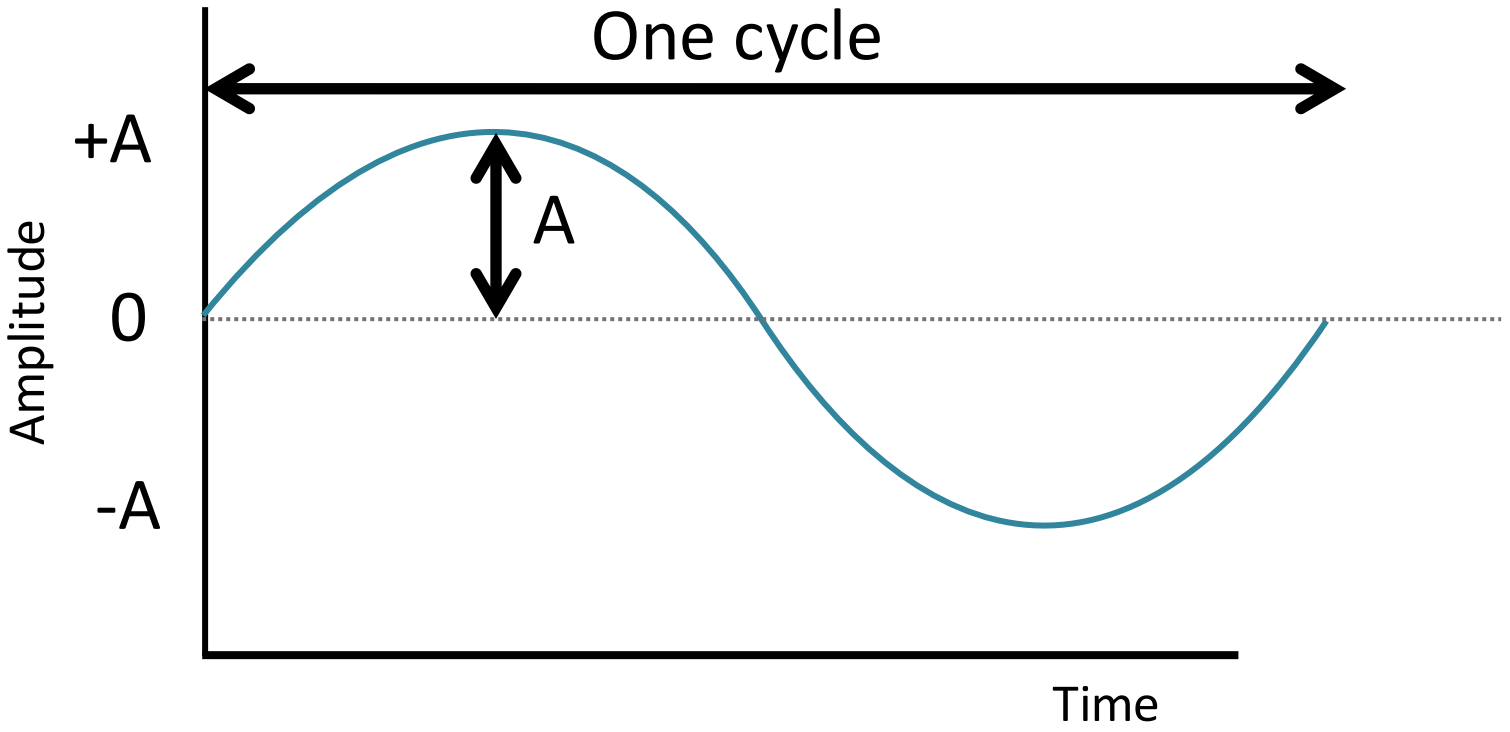
# Model for a signal (frequency, amplitude, and phase)



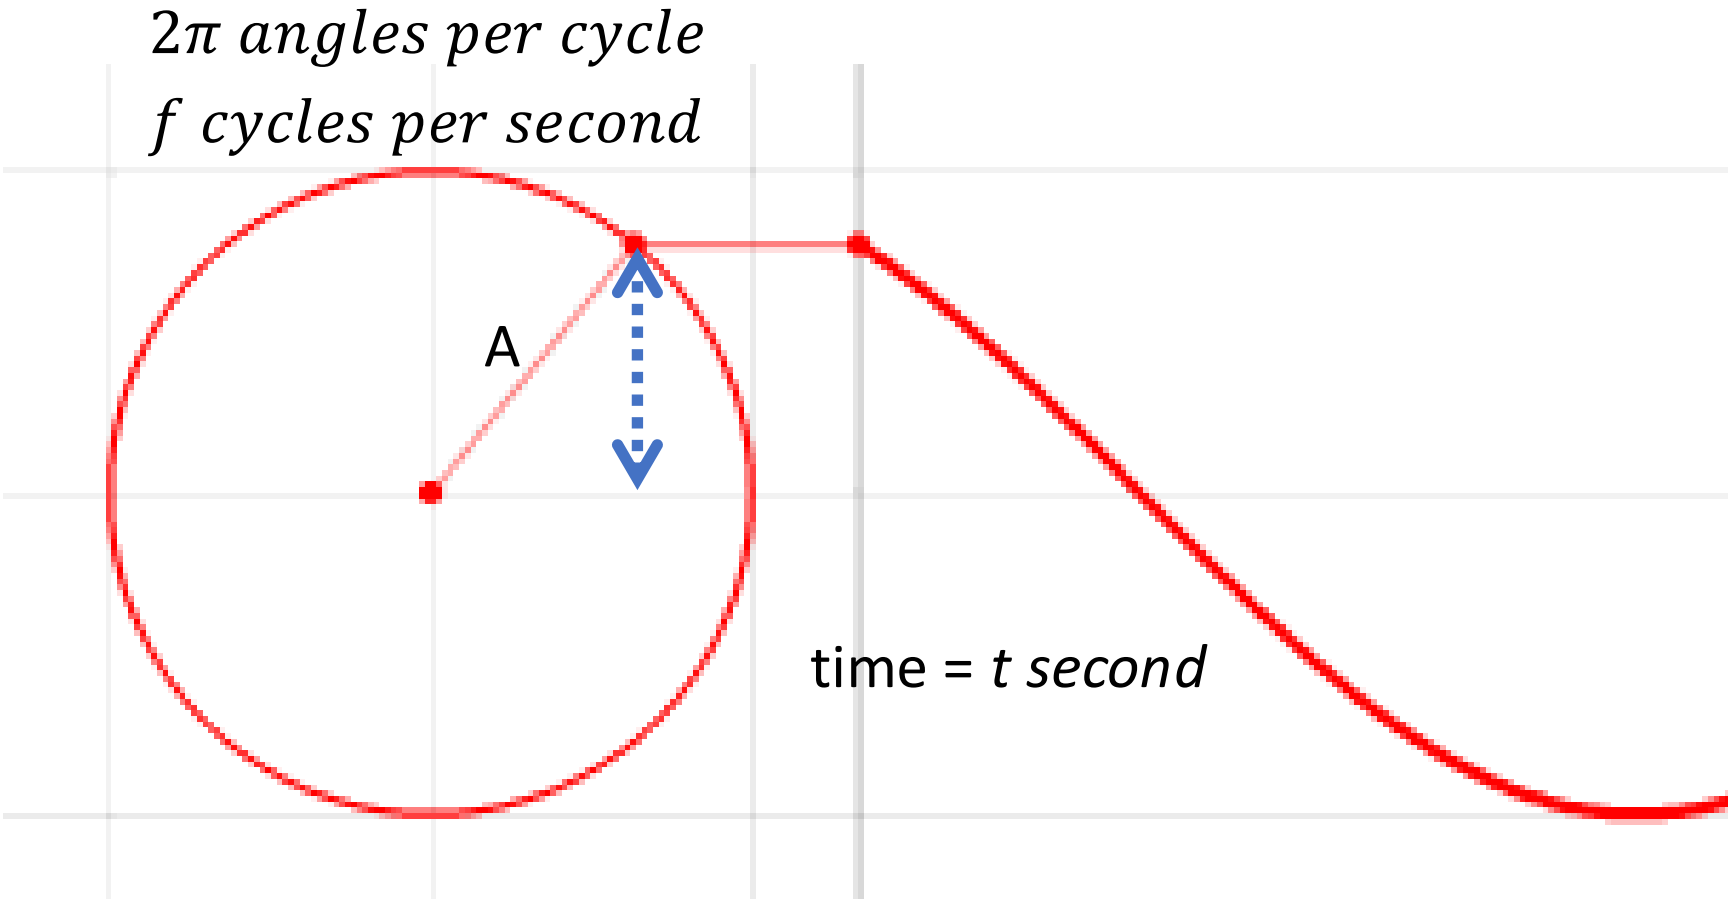

$f$ *cycles per second*

$2\pi$ *angles per cycle*

# Model for a signal (frequency, amplitude, and phase)

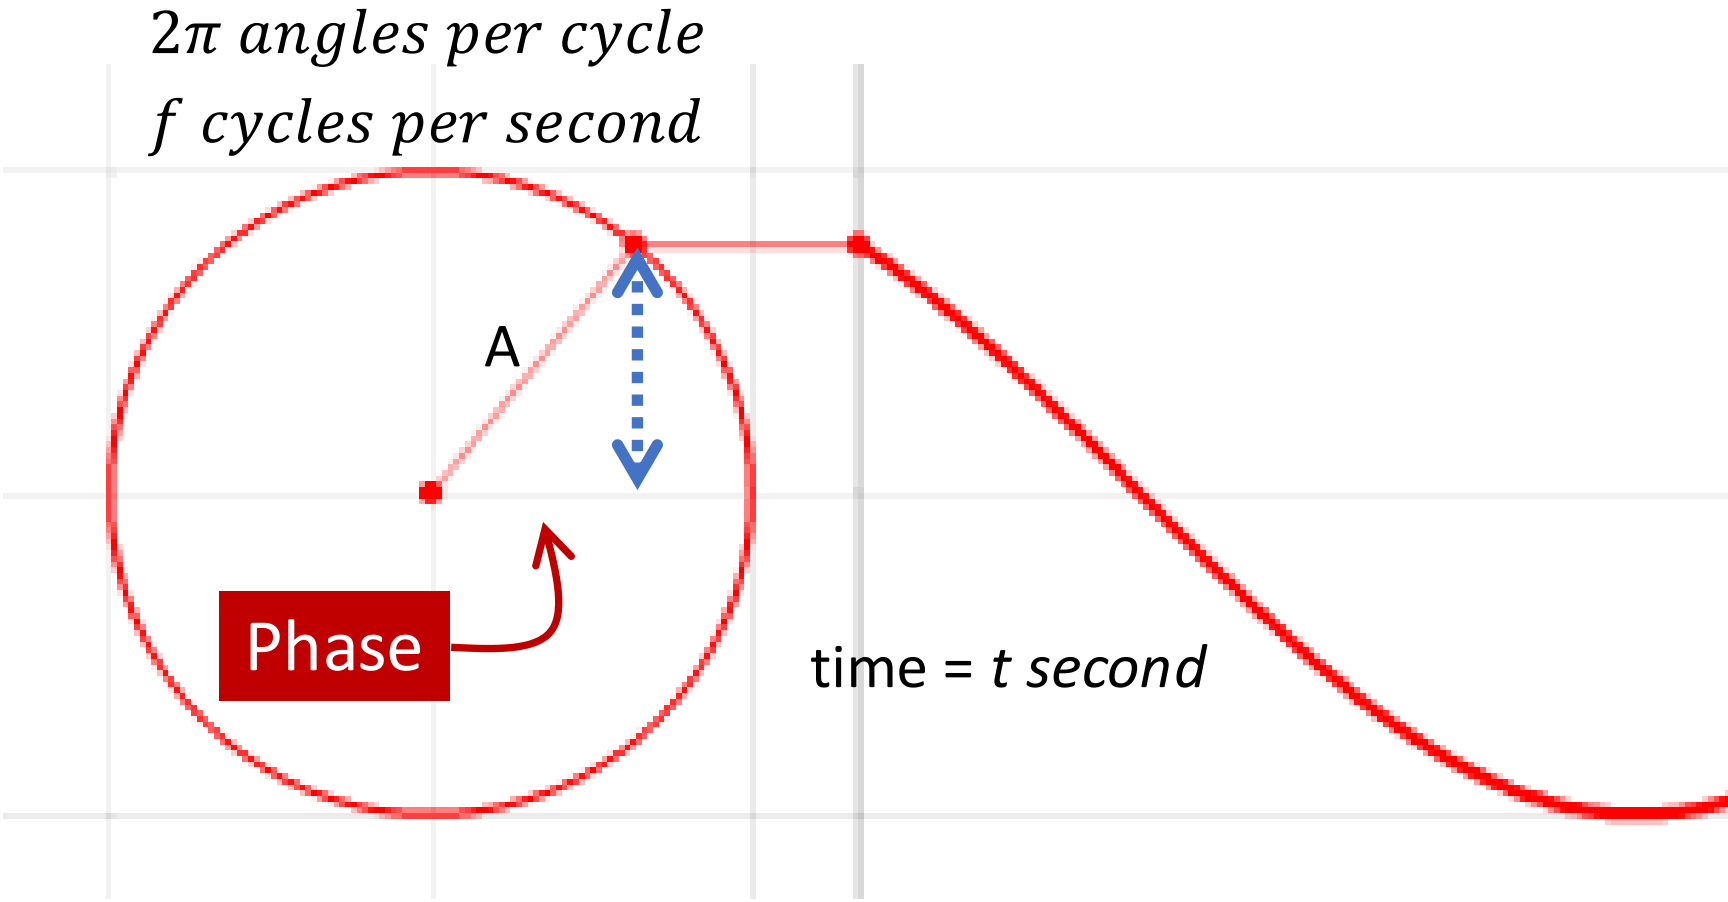


# Model for a signal (frequency, amplitude, and phase)



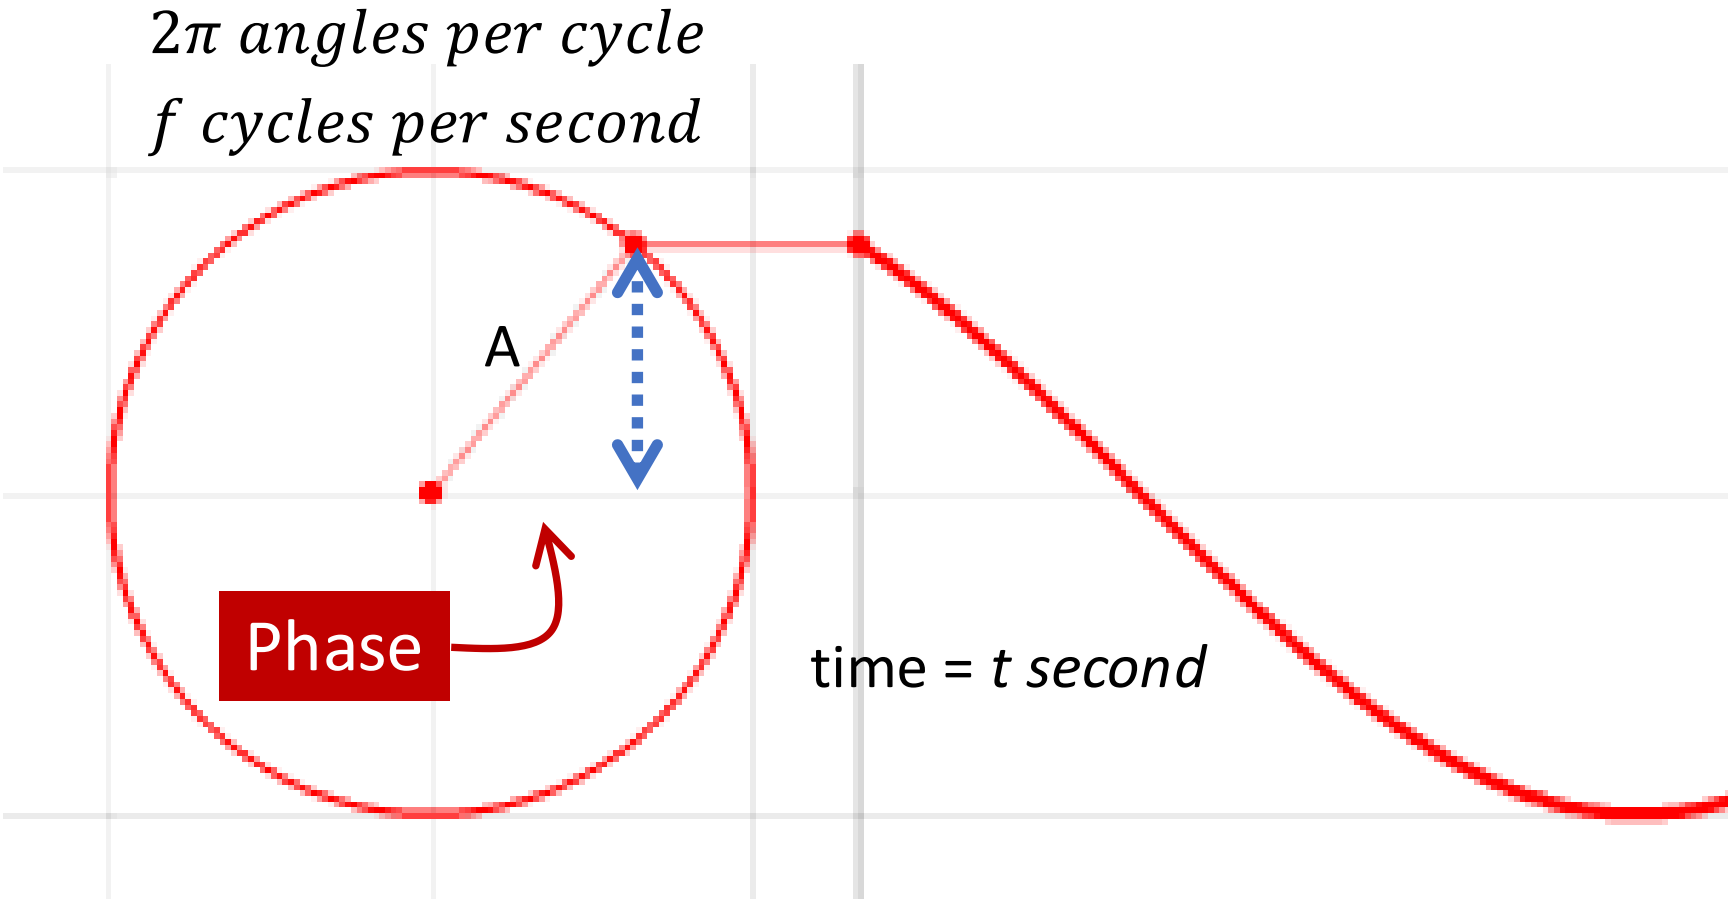

$$= \mathrm{A}\,.\sin(\theta) = \mathrm{A}\,.\sin(2\pi f t)$$

# Model for a signal (frequency, amplitude, and phase)




$$= \mathrm{A}\,.\sin(\theta) = \mathrm{A}\,.\sin(2\pi f t)$$

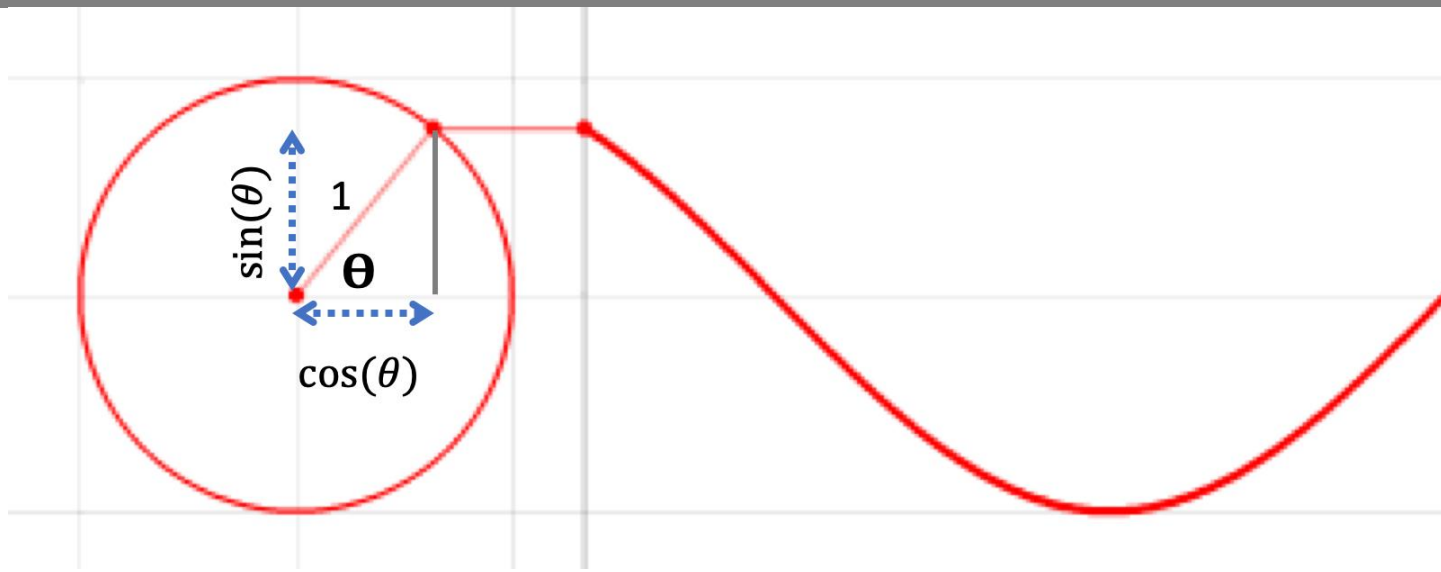
# Model for a signal (frequency, amplitude, and phase)




$$= \quad \mathrm{A}\,.\sin(\theta) \quad = \quad \mathrm{A}\,.\sin(2\pi f t)$$

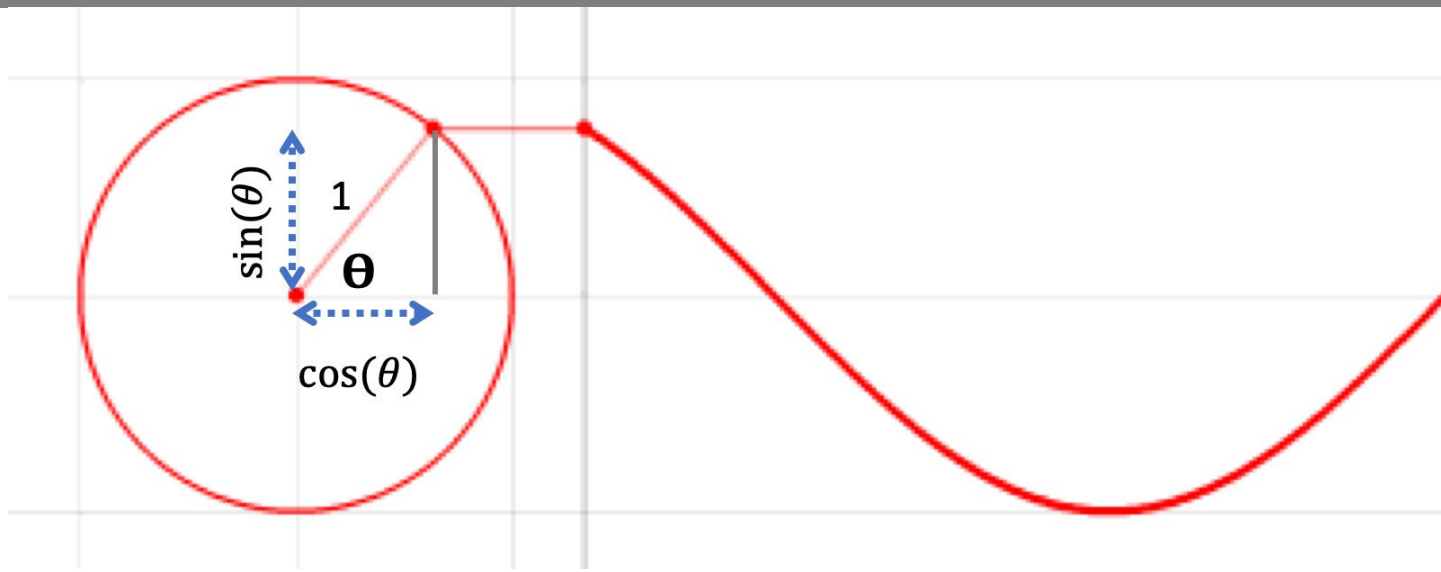
$\mathrm{A}\,.\sin(2\pi f t + \phi)$ -- with initial/additional phase φ

# Model for a signal (frequency, amplitude, and phase)




$$cos(\theta) + j\ \sin(\theta) = e^{j\theta}$$

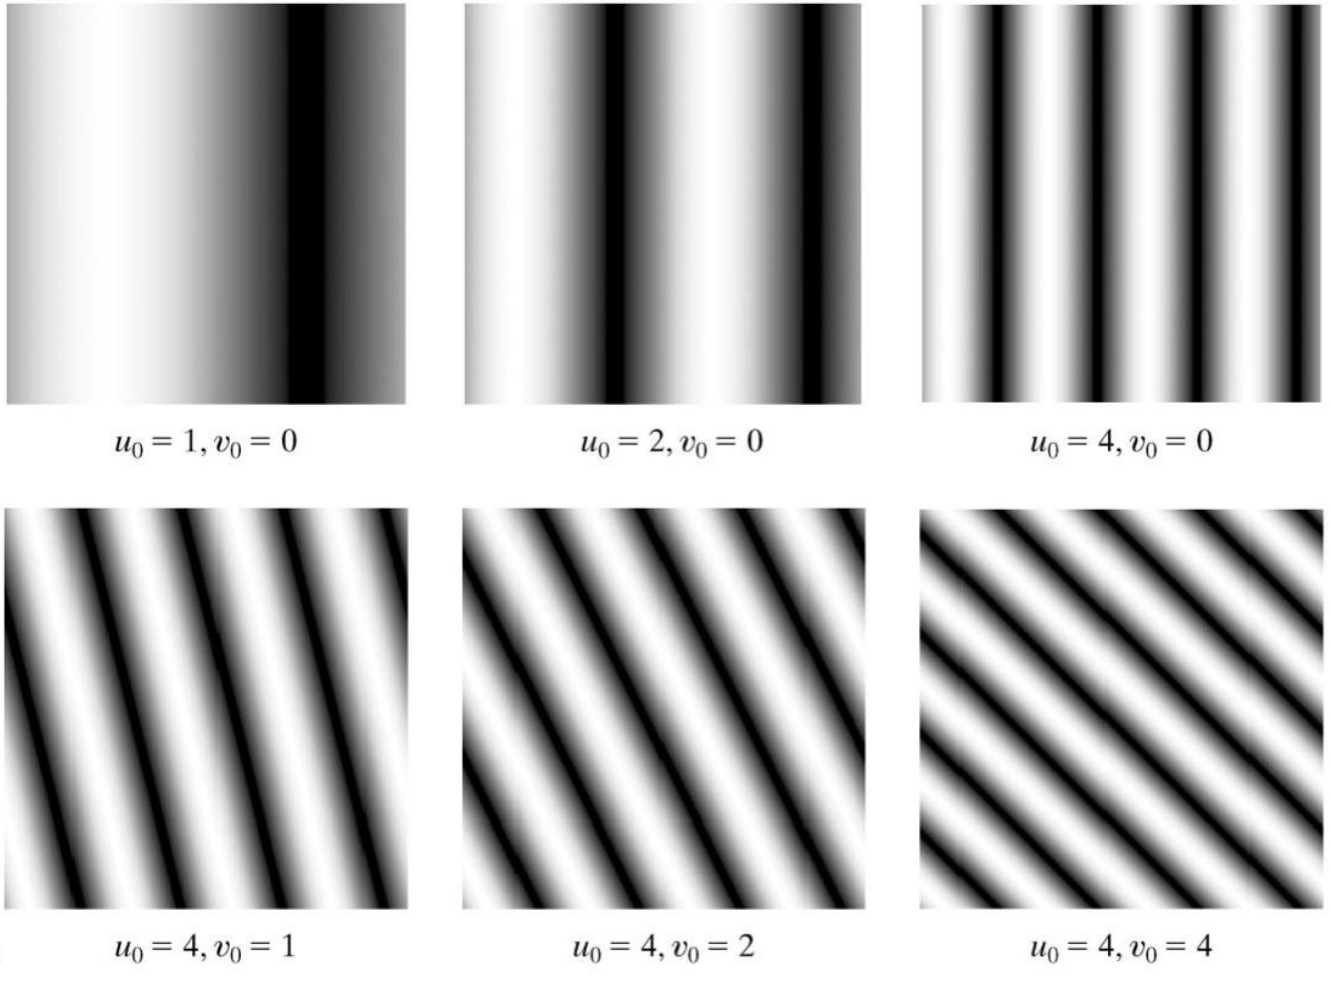
# Model for a signal (frequency, amplitude, and phase)




$$cos(\theta) + j\ \sin(\theta) = e^{j\theta} = e^{j\ 2\pi ft}$$

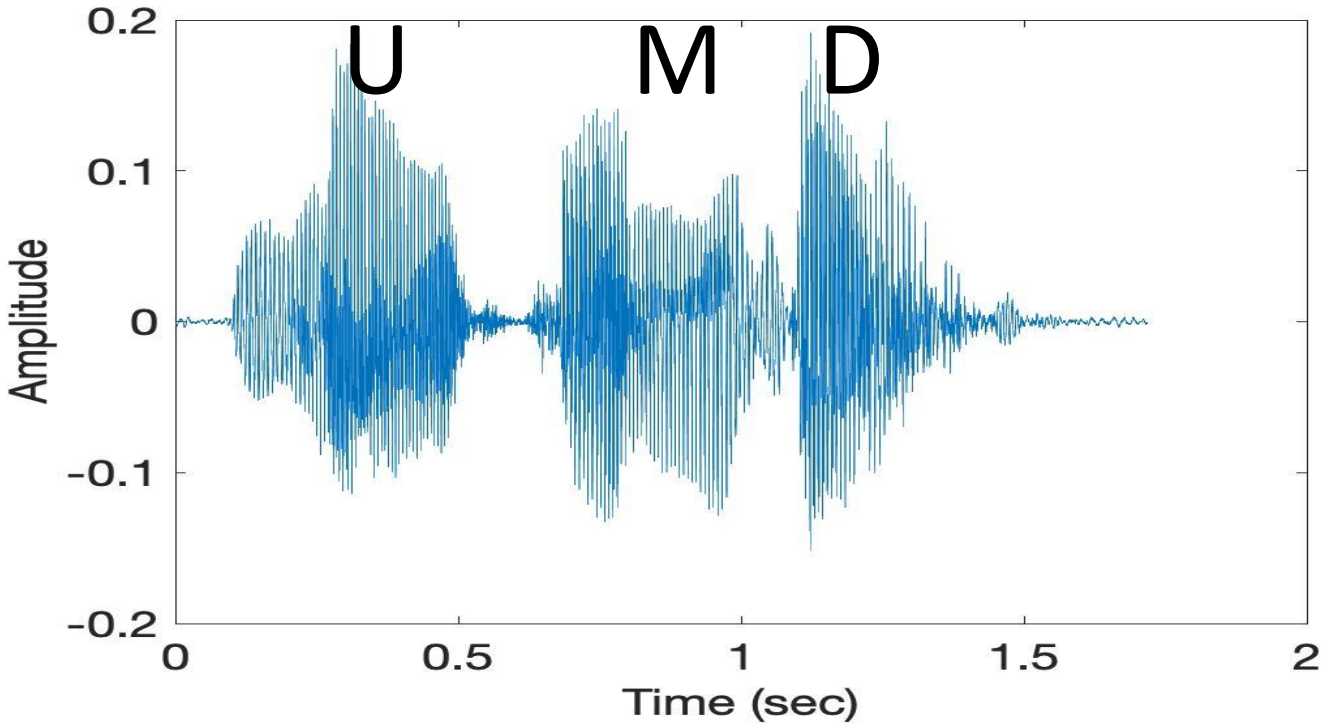
# Holds true for multidimension signals: $e^{j\,2\pi\vec{k}\vec{x}}$

- Intensity images for $s(x,y)=\sin\left[2\pi\left(u_0x+v_0y\right)\right]$

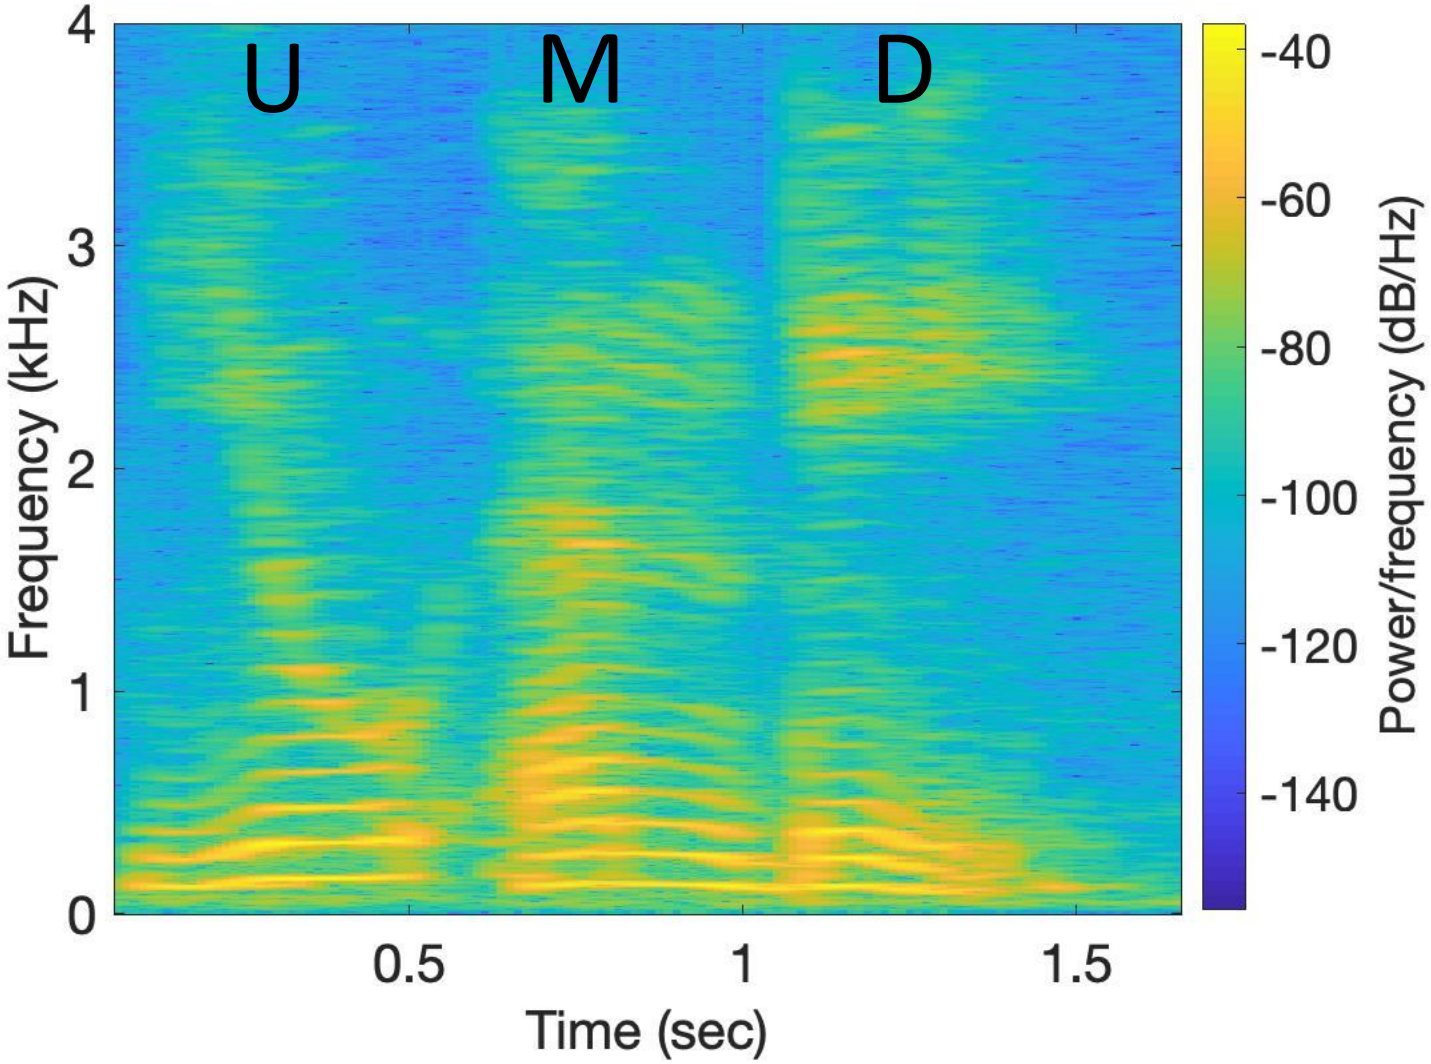


# Visualizing Transient Frequencies

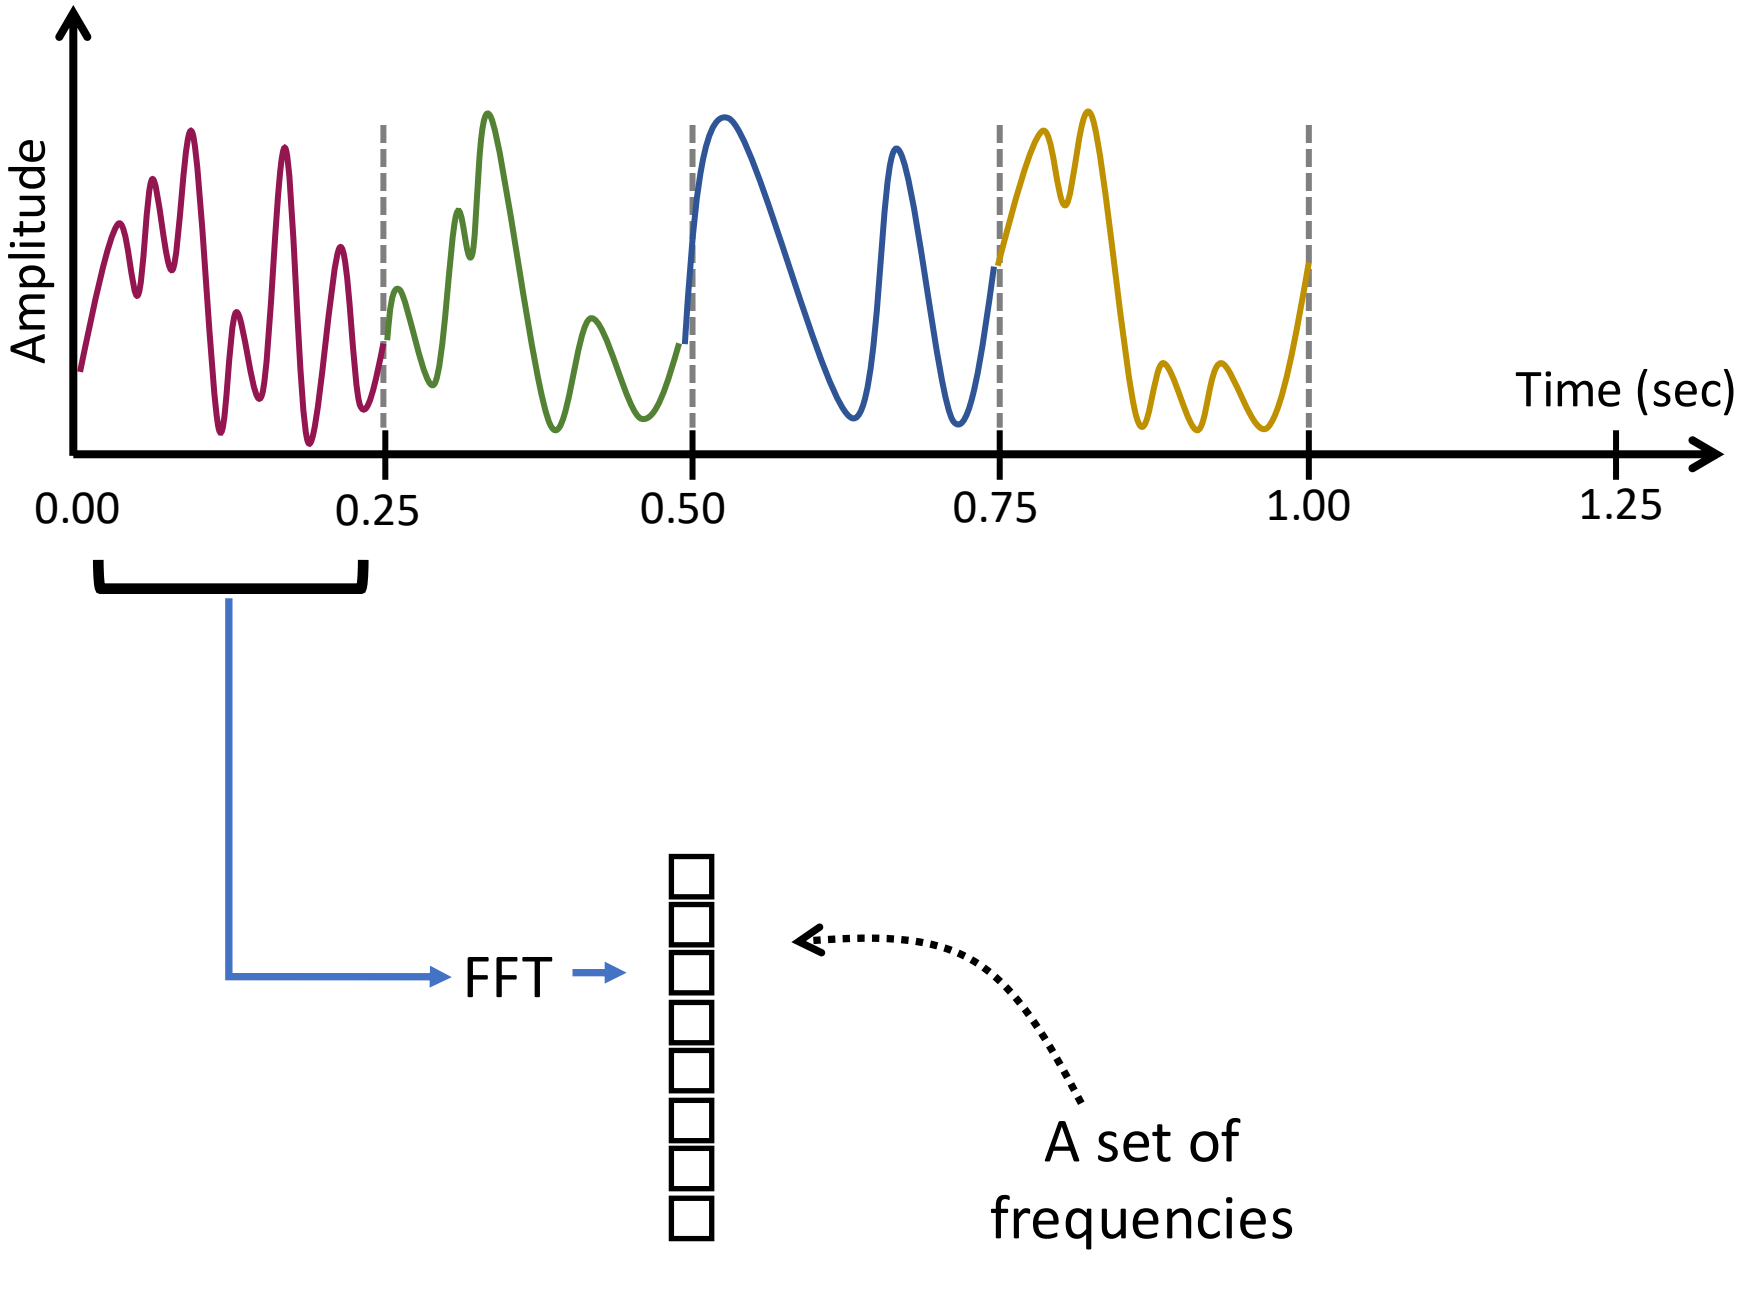
# Visualizing transient frequencies



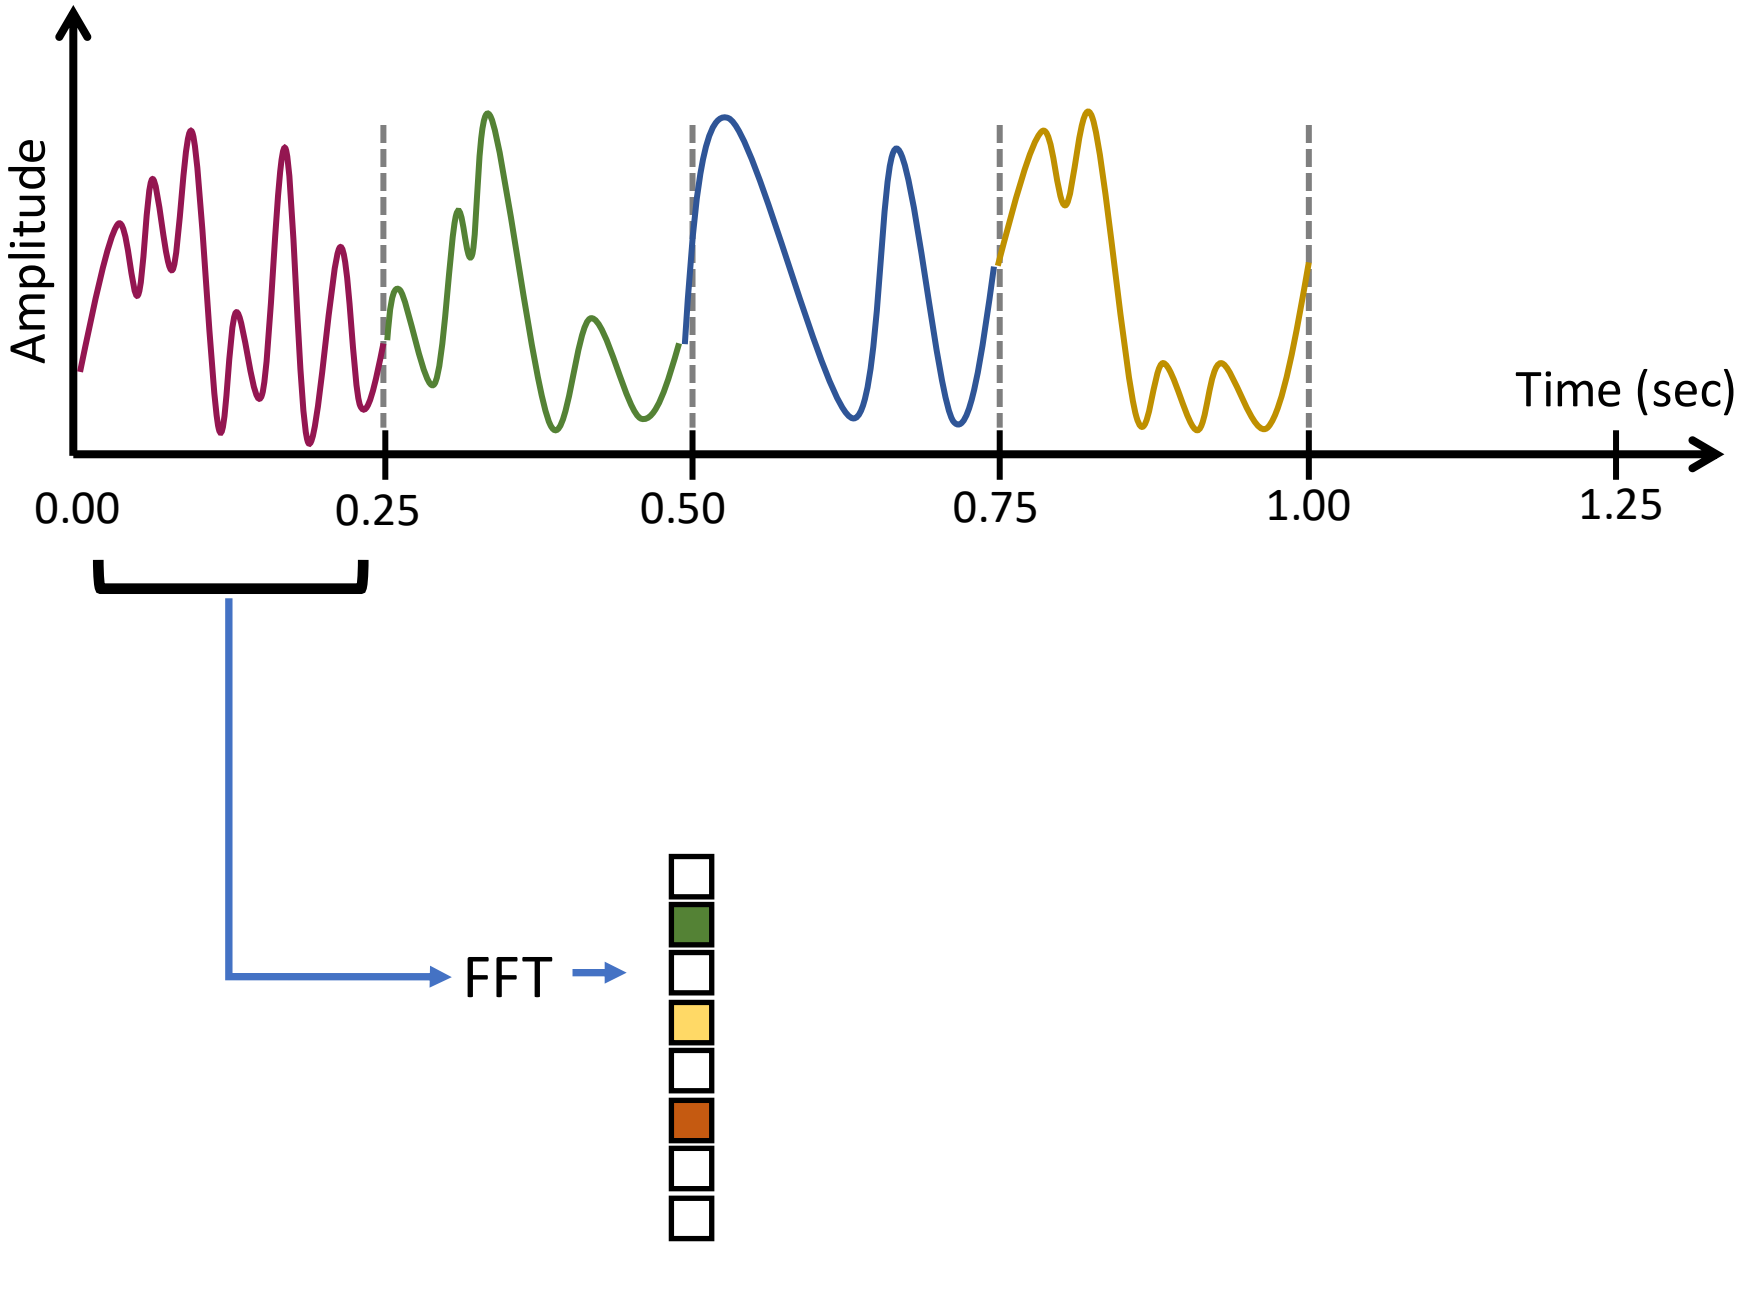
# Spectrogram



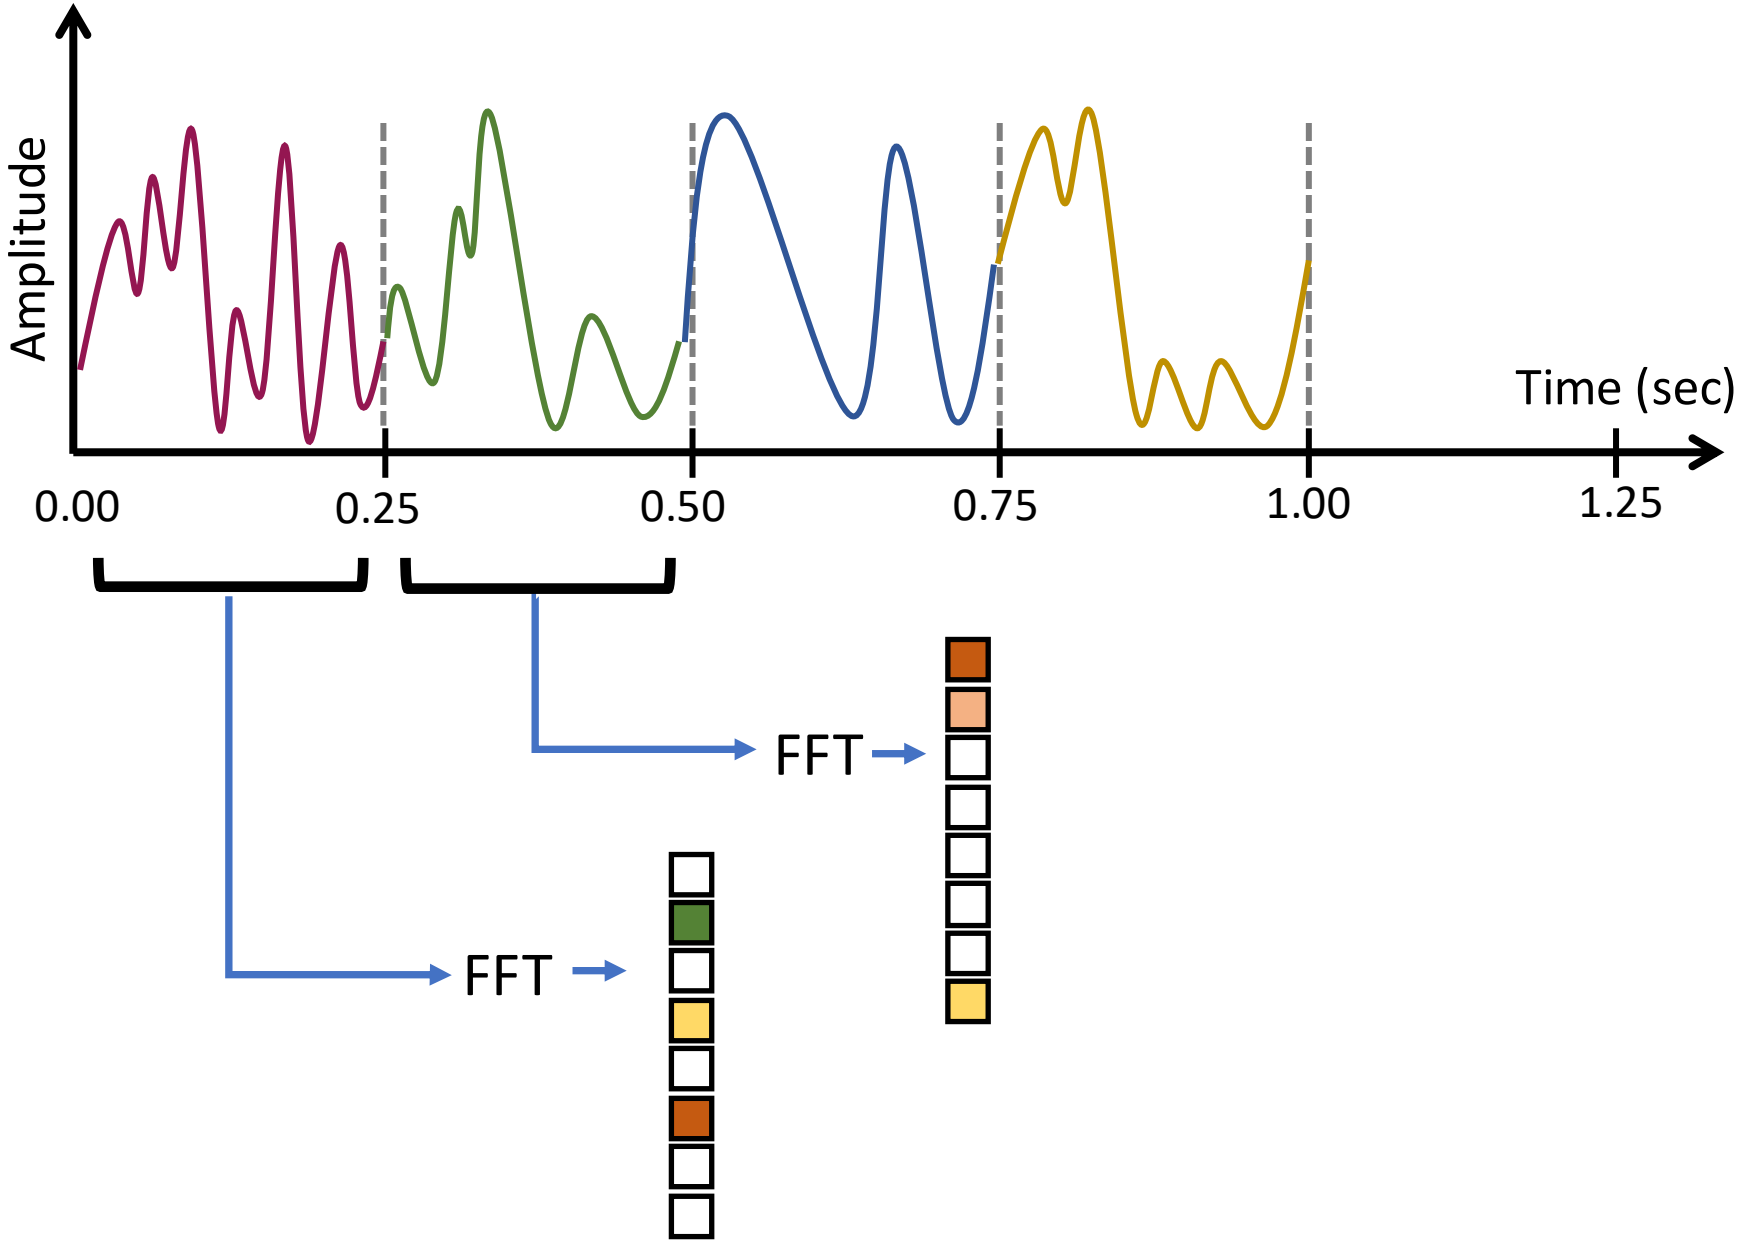
# Frequency components change over time



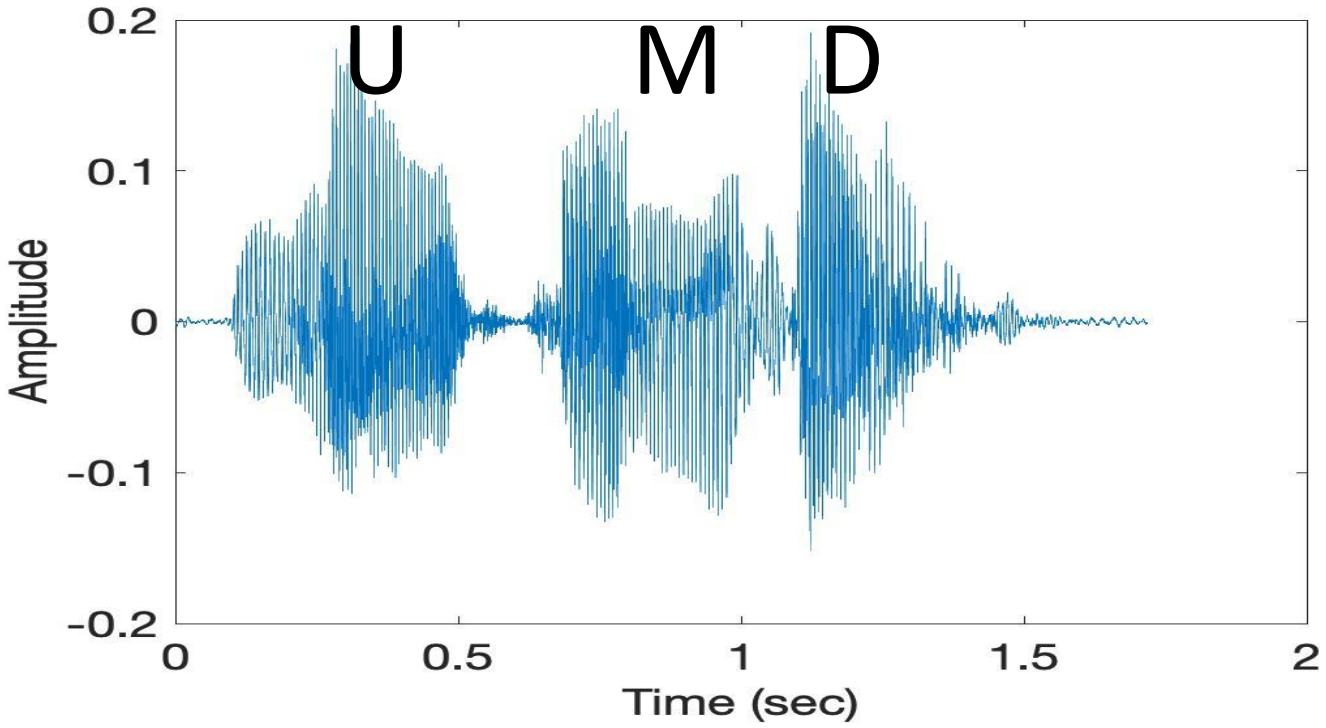
# Frequency components change over time



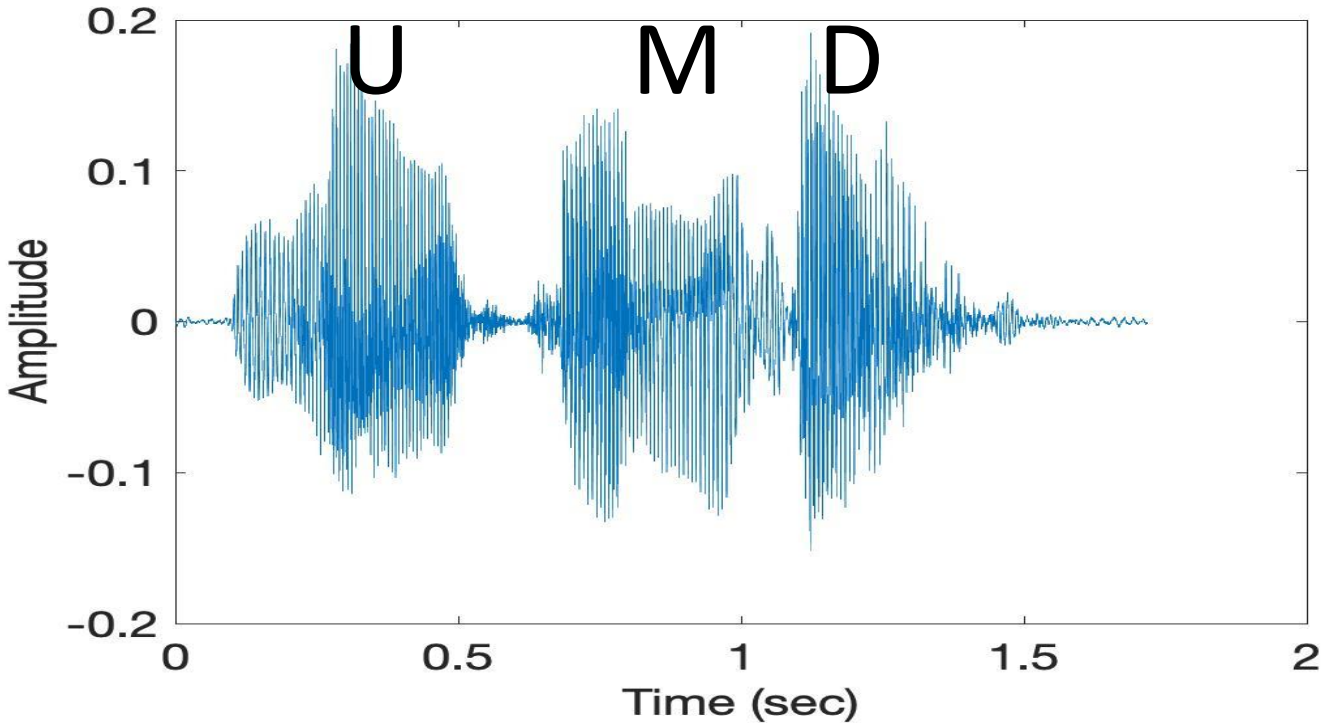
# Frequency components change over time

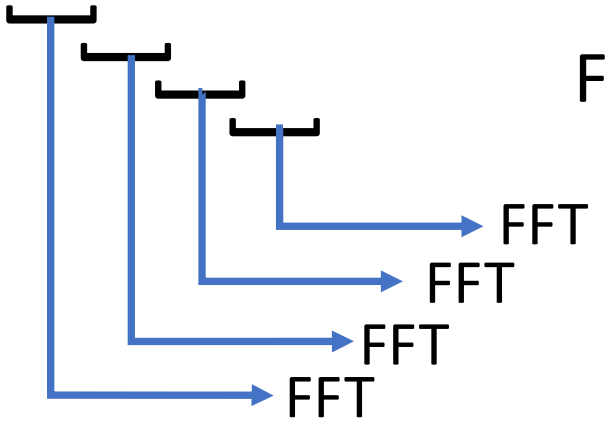


# Frequency components change over time



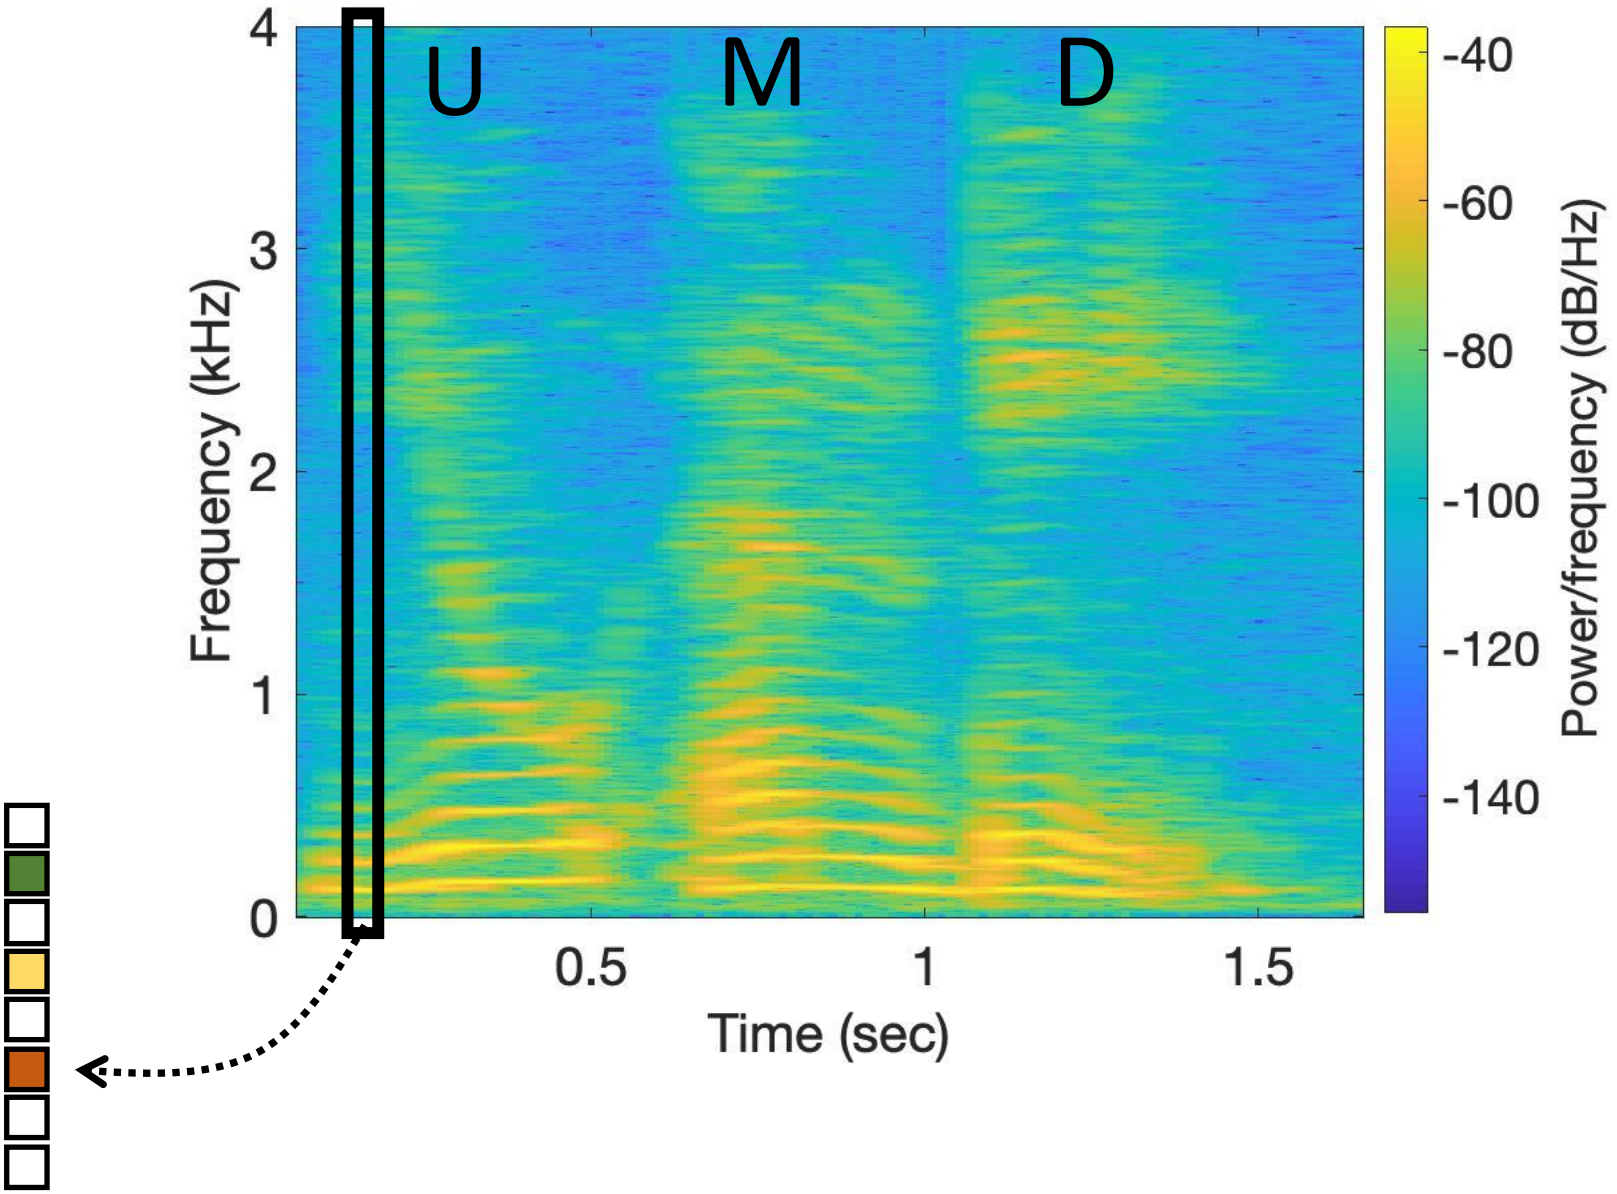
# Spectrogram



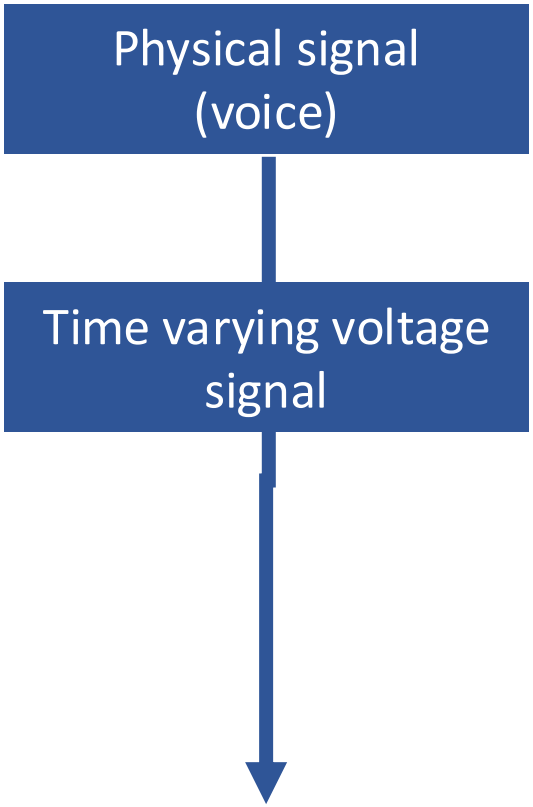




FFT of overlapping windows of samples (Spectrogram)

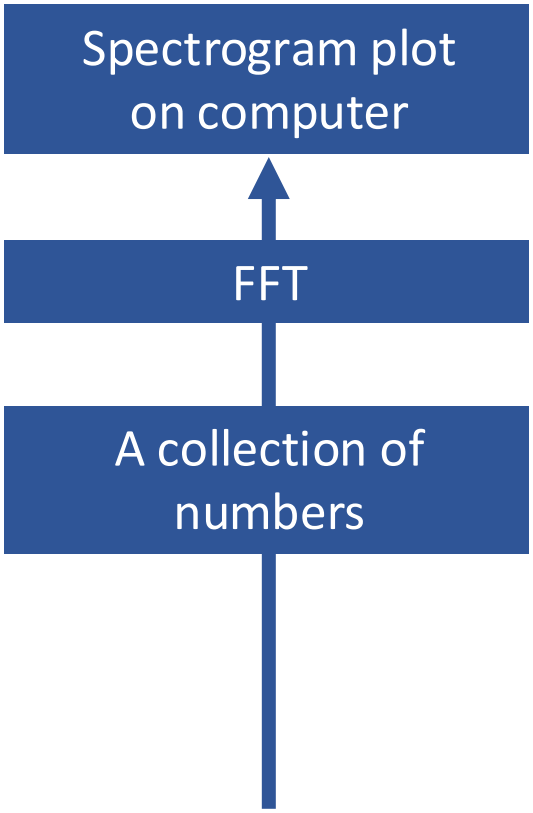
# Spectrogram



# Analog




# Digital



# Sampling

Analog



A physical phenomenon
We want to measure

SENSOR

Analog

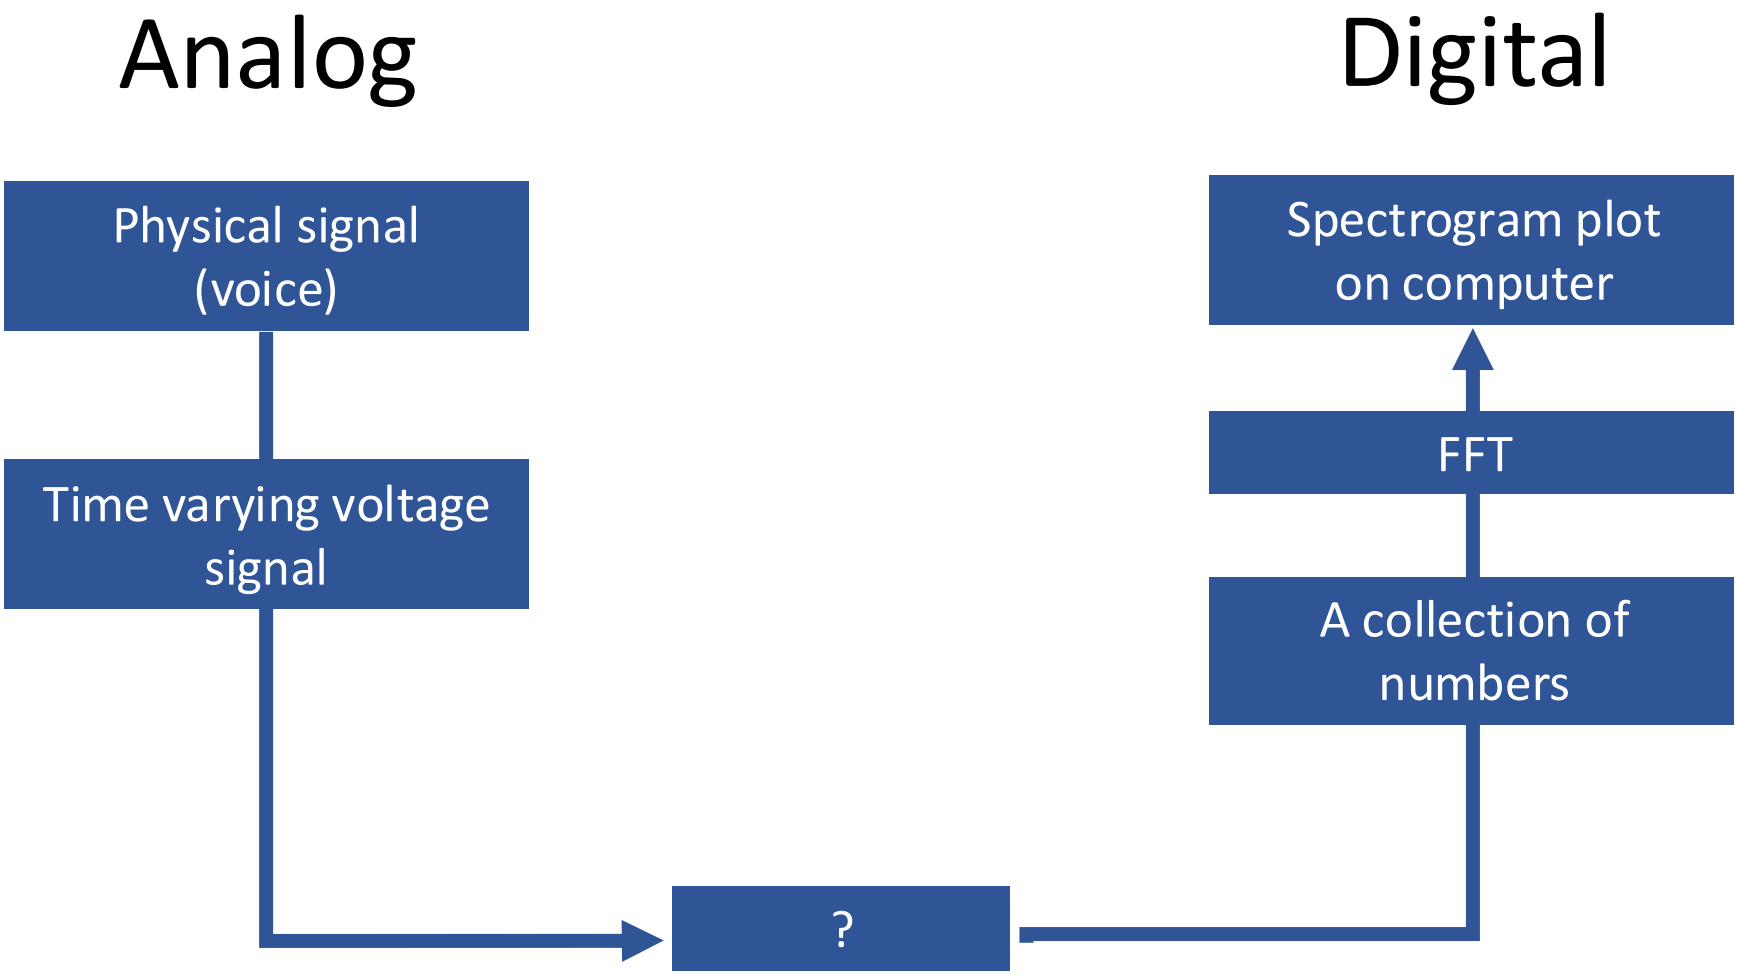
A structure or material That reacts to the physical phenomenon

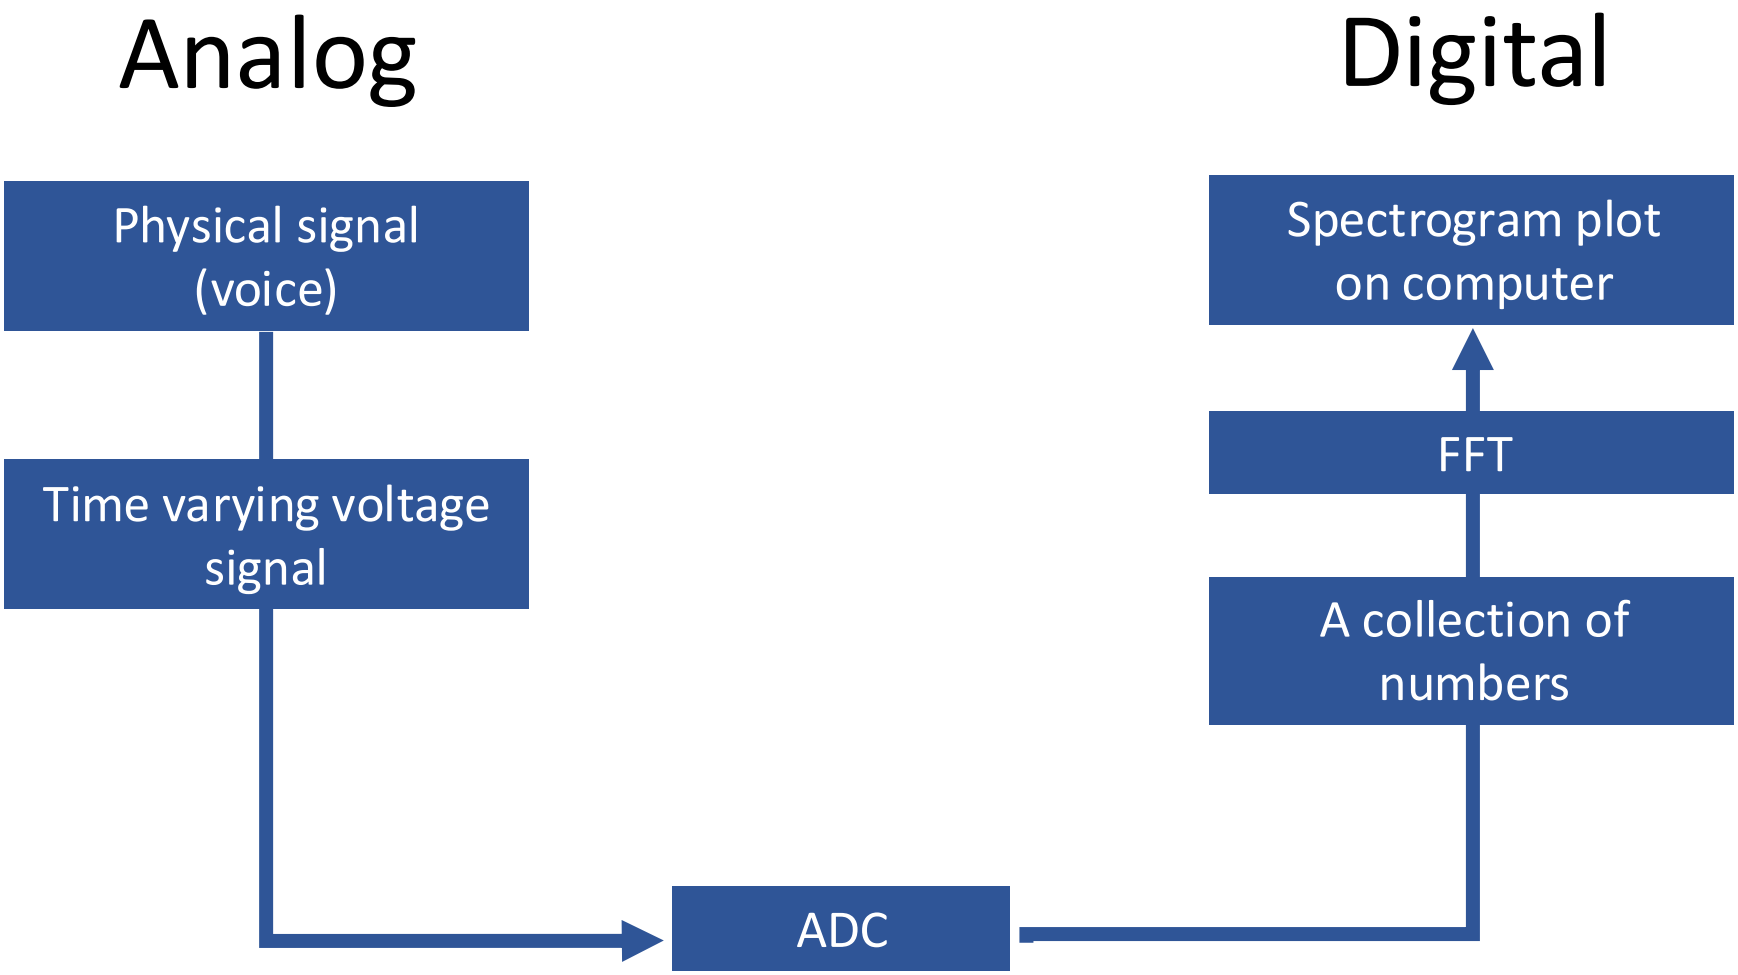
A circuit that converts the reaction to measurable electrical signal (e.g. voltage)

Interpretation of the signal, storing the data

Digital

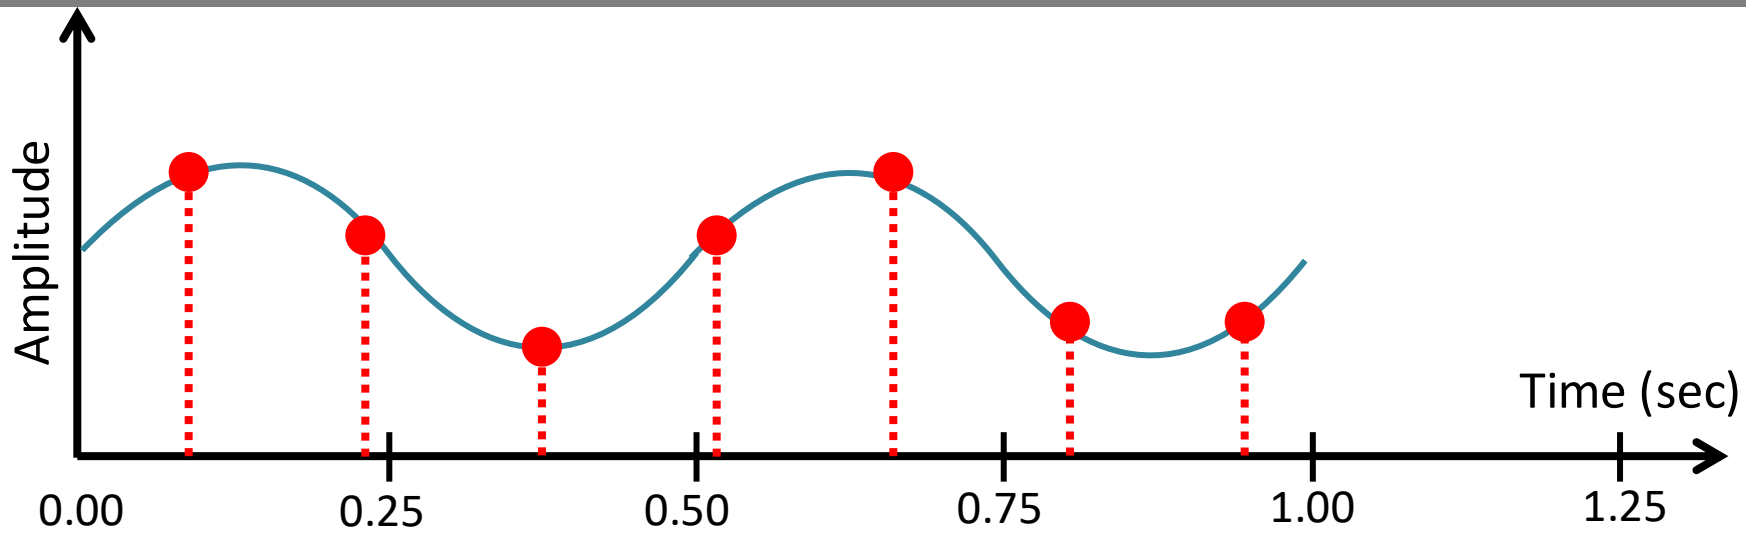



Analog
Digital
Physical signal (voice)
Time varying voltage signal
?
A collection of numbers
FFT
Spectrogram plot on computer

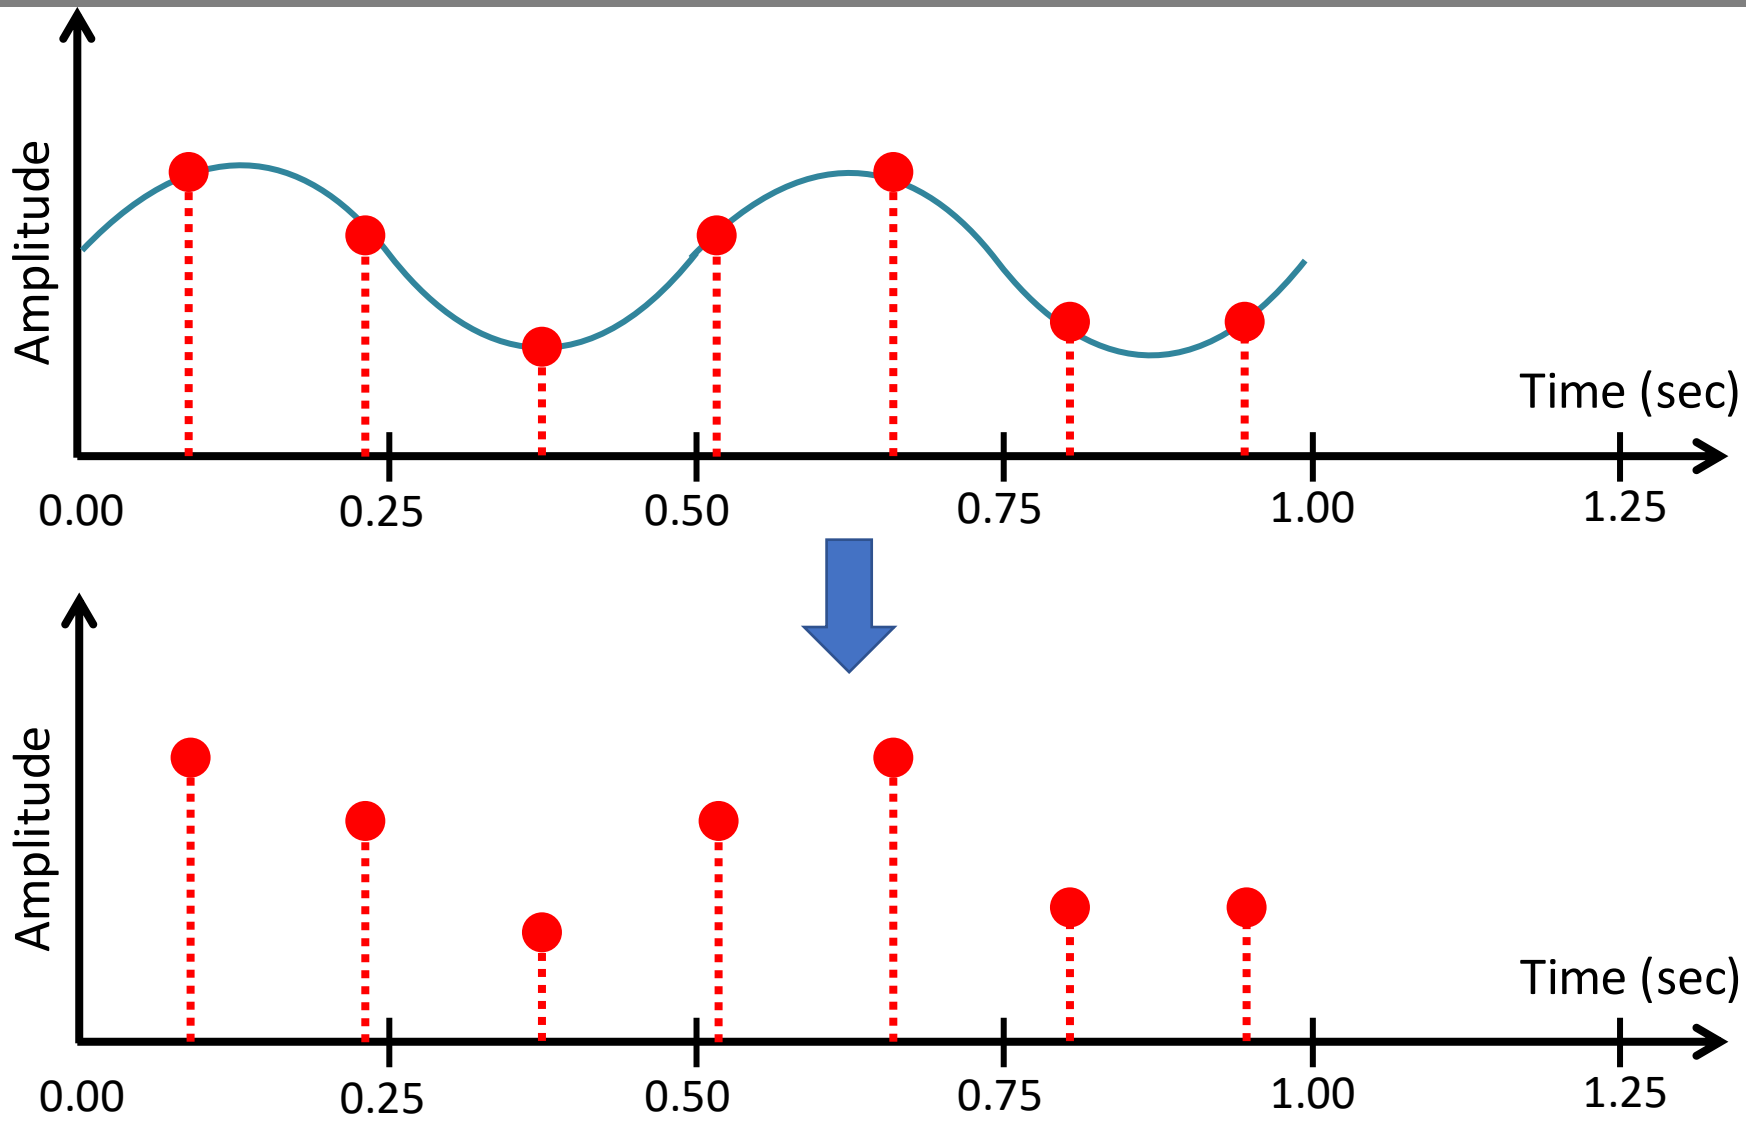



Analog-to-Digital Converter

# Sampling

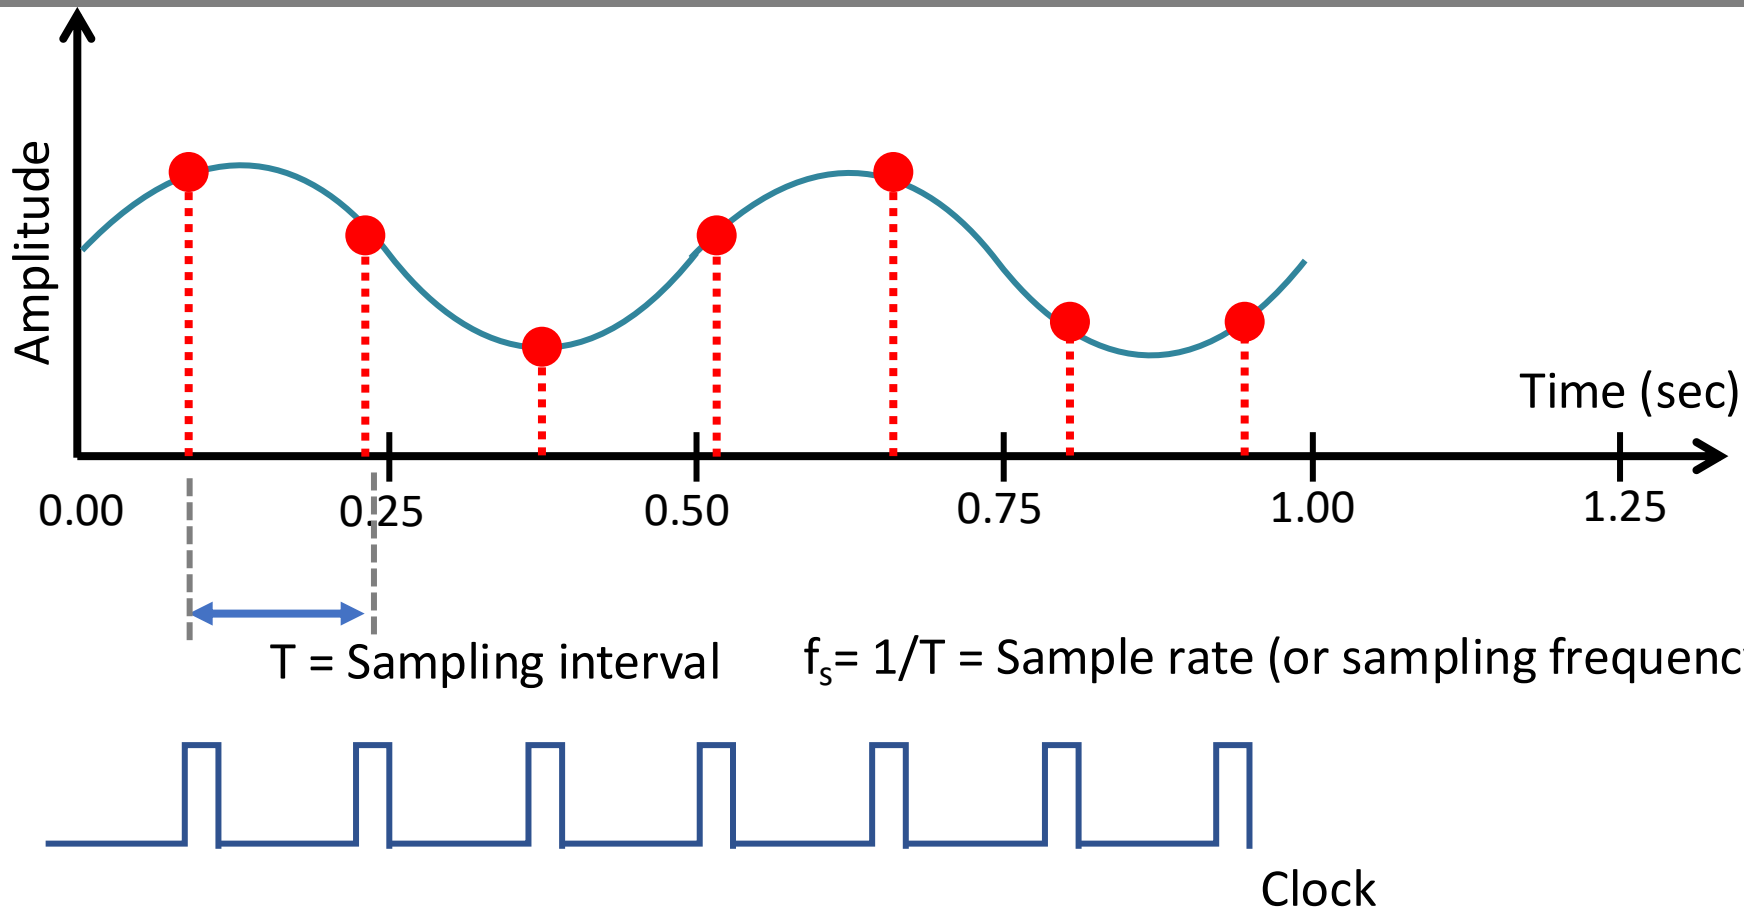


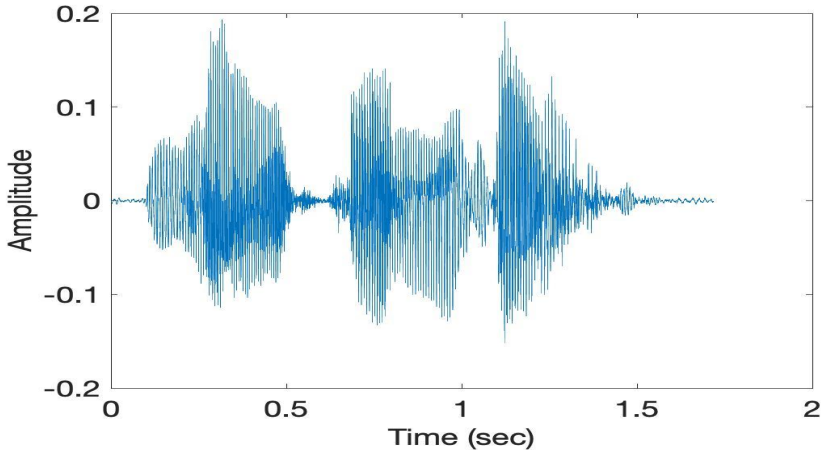
# Sampling



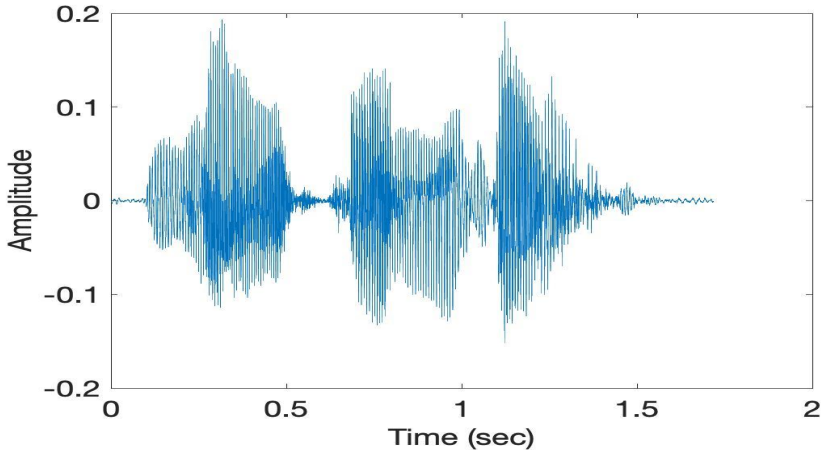

[ 0.34 0.22 0.09 0.21 0.30 0.08 0.09 ]

# Sampling

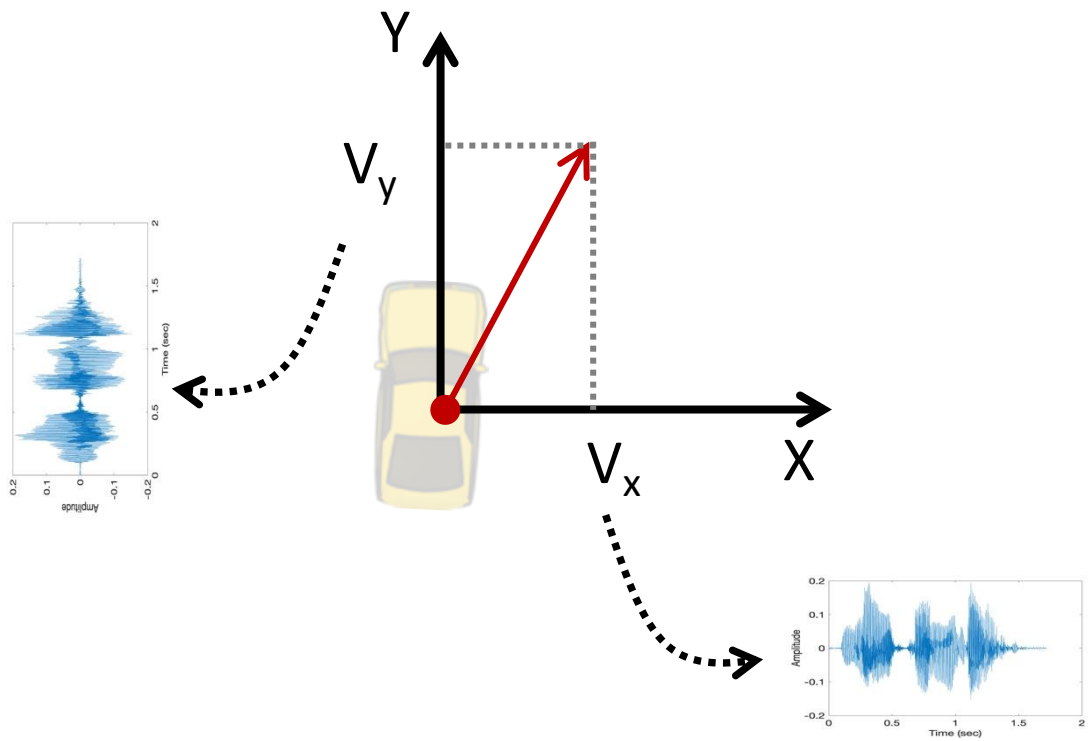

Amplitude
Time (sec)
0.00
0.25
0.50
0.75
1.00
1.25
T = Sampling interval
$f_s$= 1/T = Sample rate (or sampling frequency)
Clock

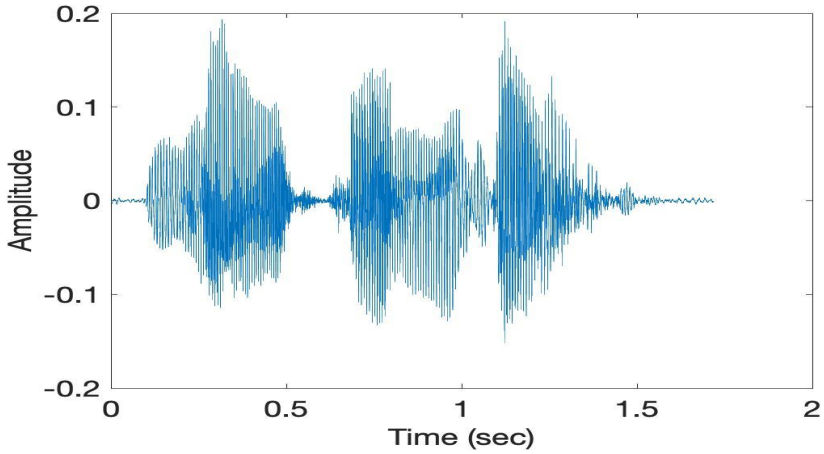



# 1-dimensional sampling

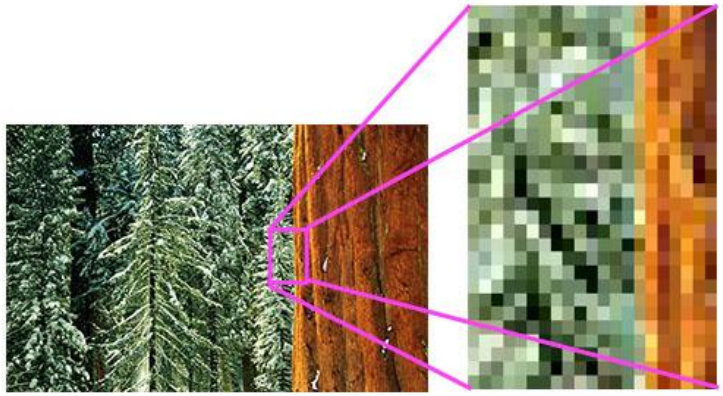



1-dimensional sampling (we sample the domain of the function, not the value-sapce)

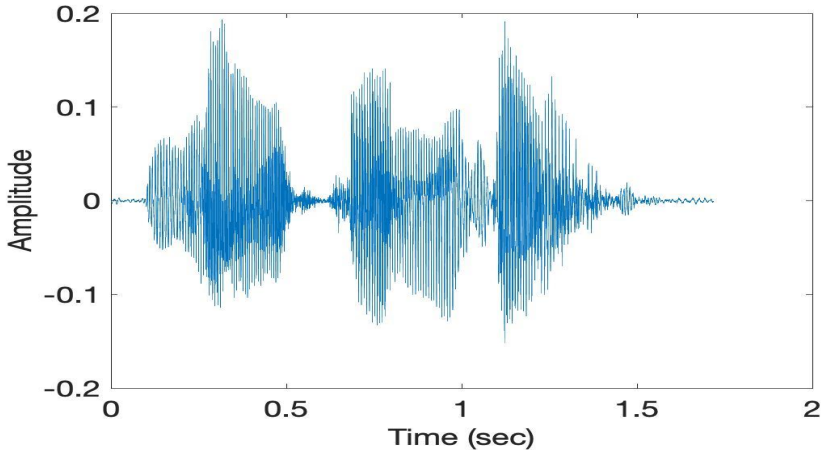




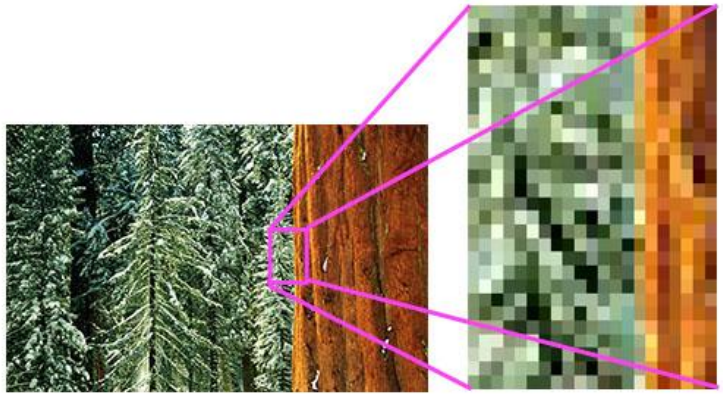

1-dimensional sampling



2-dimensional sampling



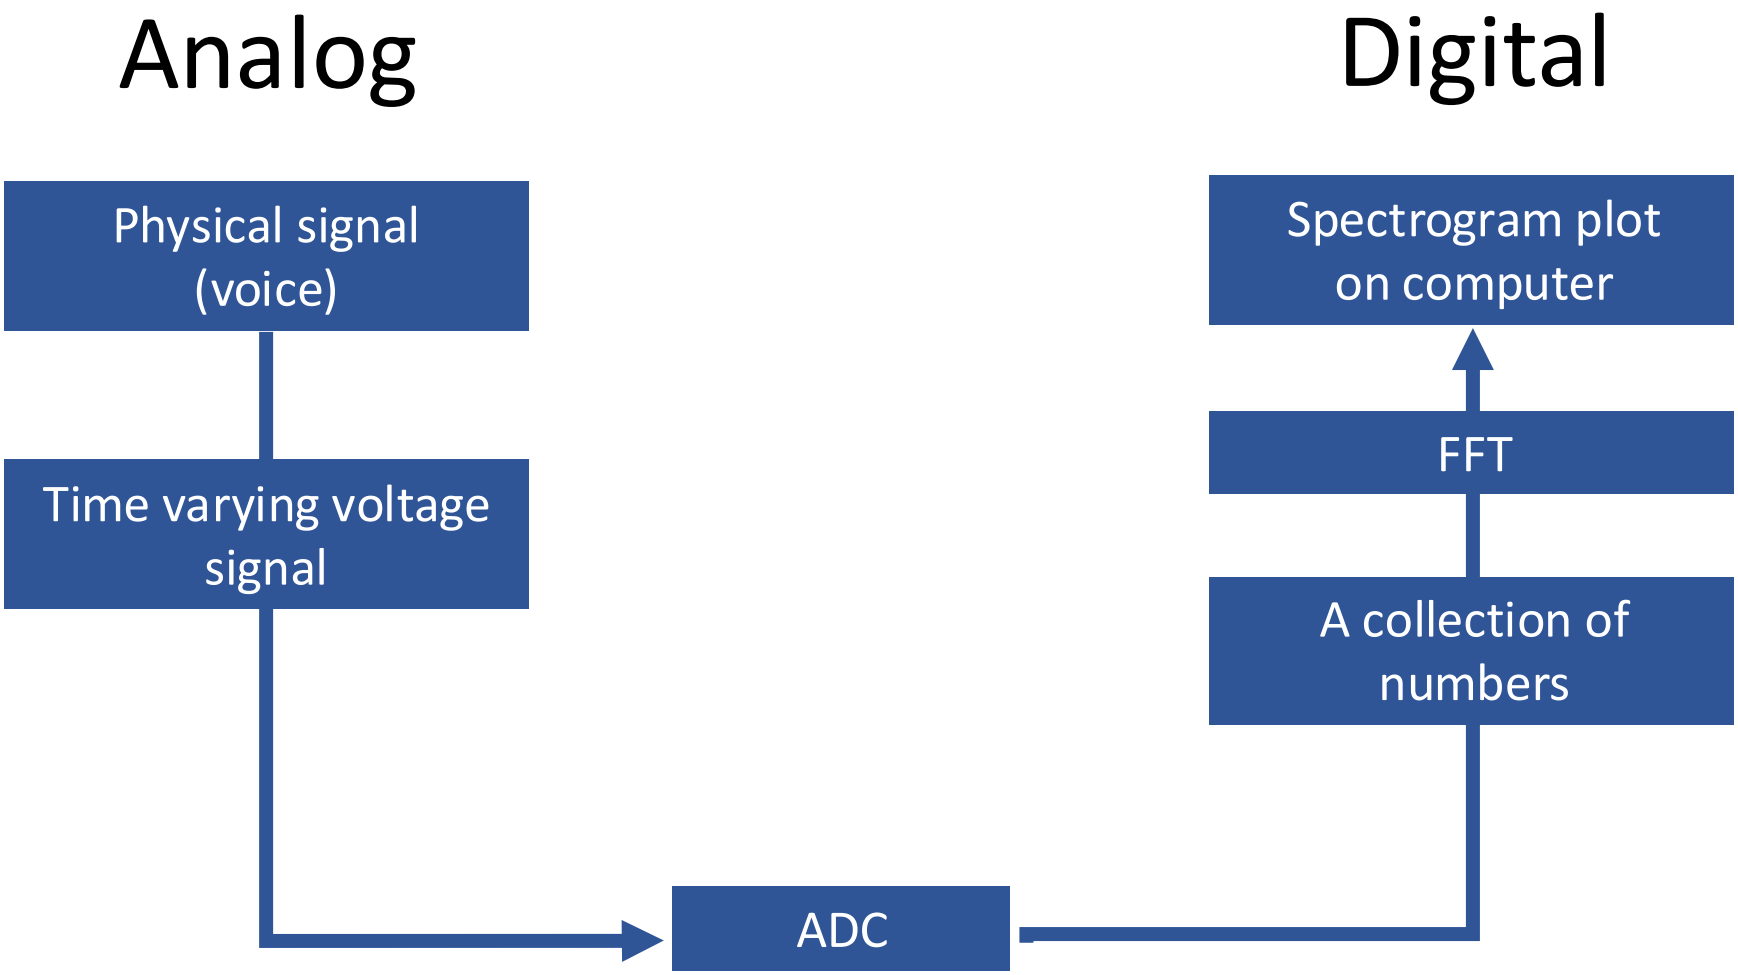

1-dimensional sampling



2-dimensional sampling

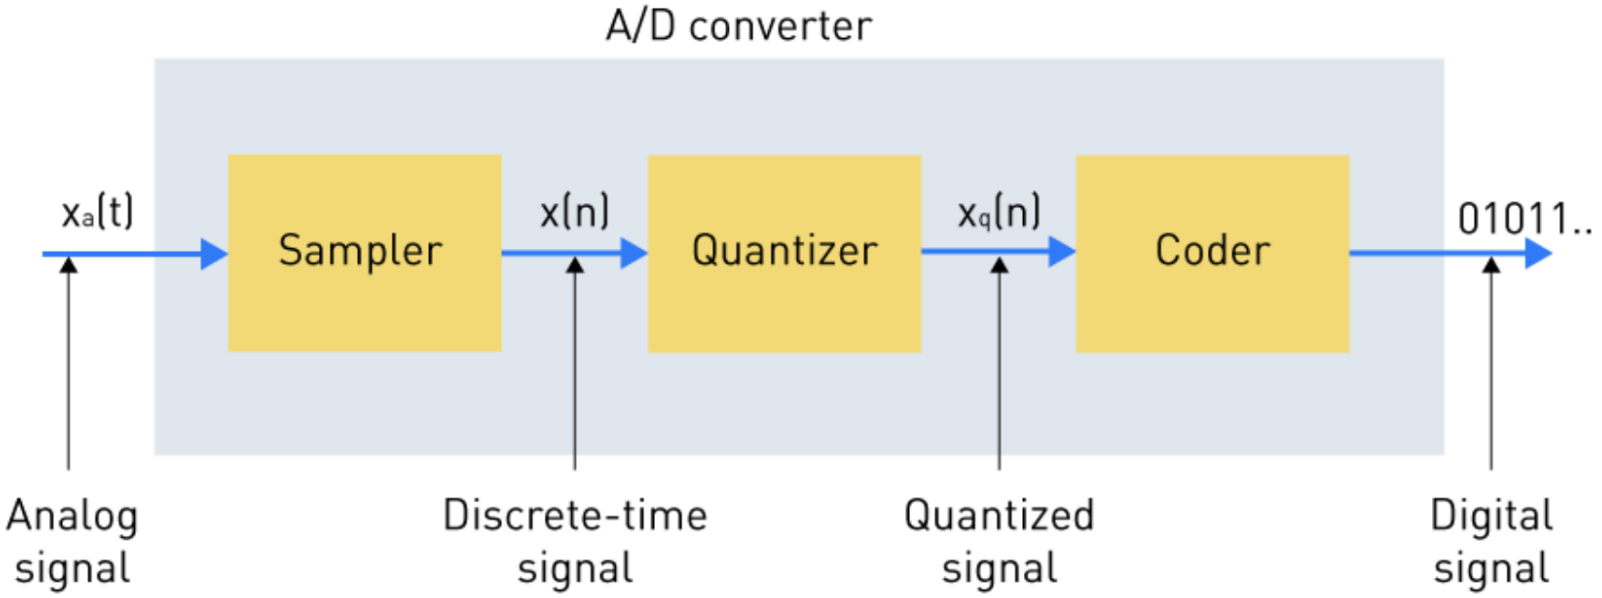


3-dimensional sampling

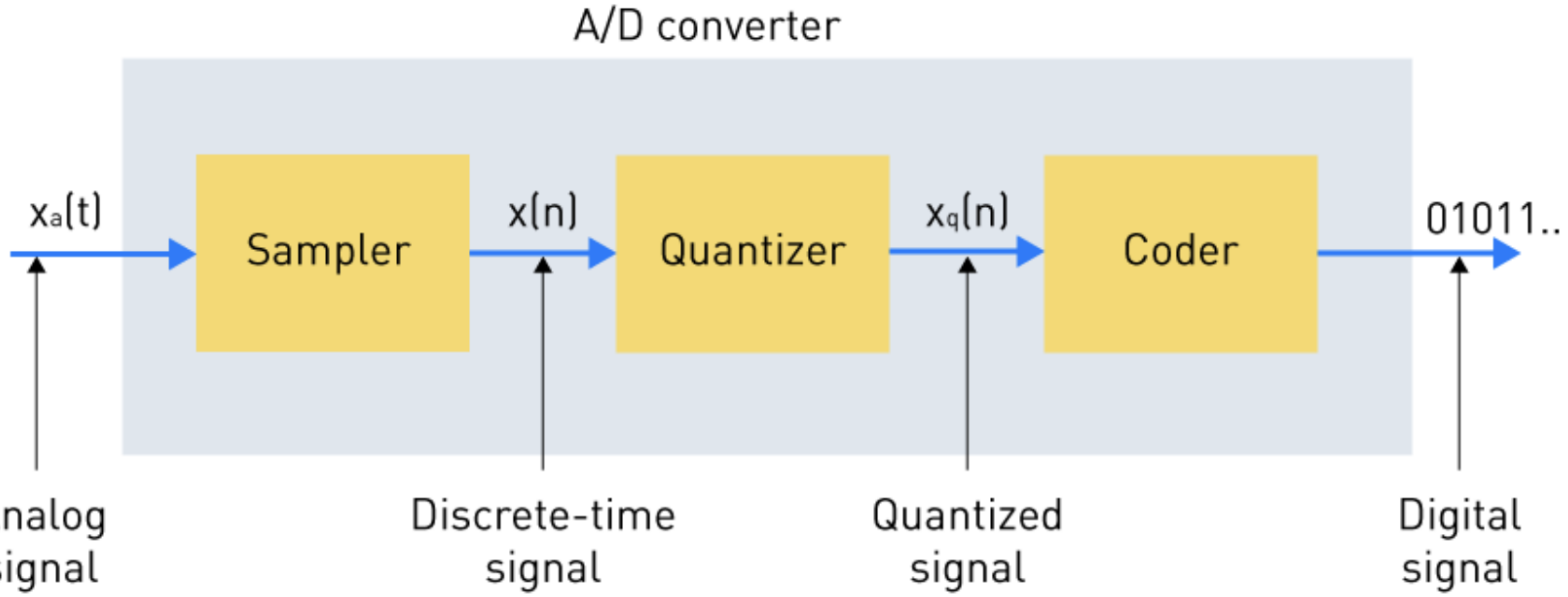


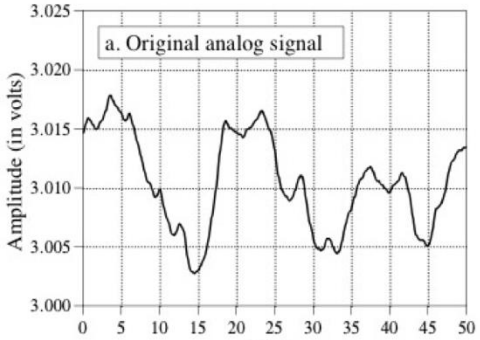

Analog-to-Digital Converter

# Digitization process (What is inside an ADC)

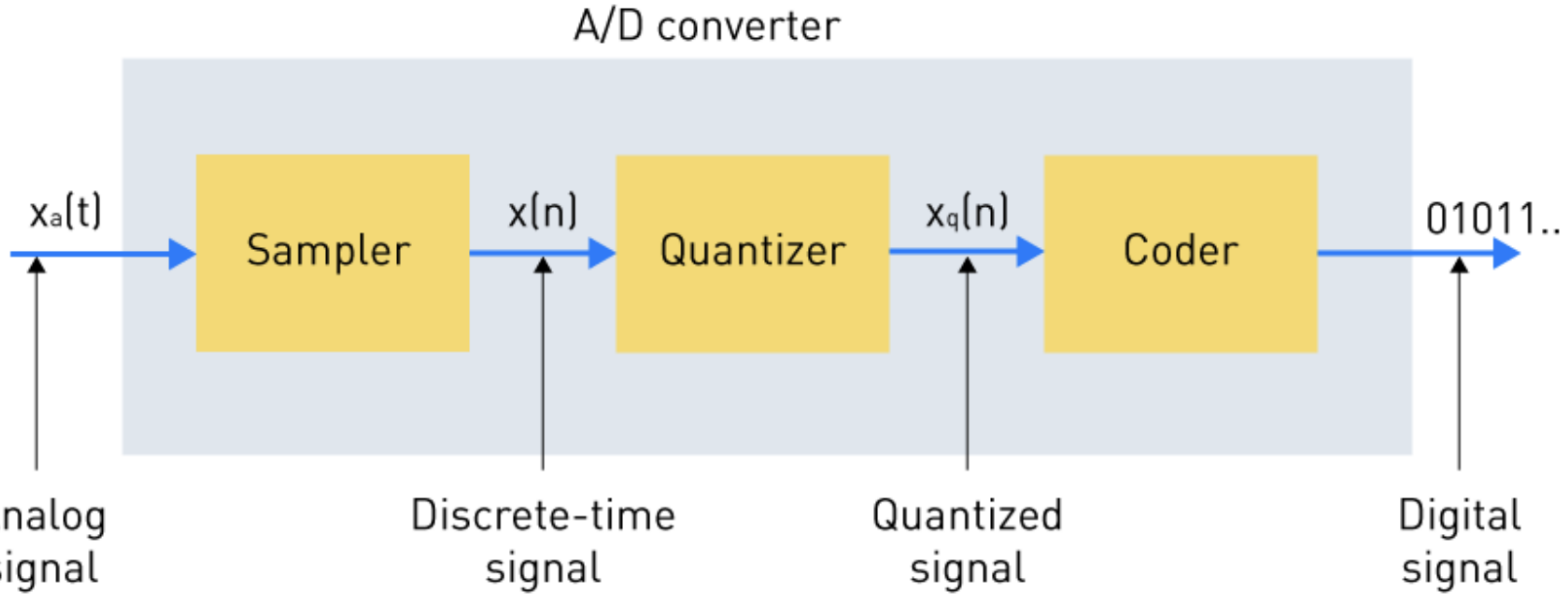


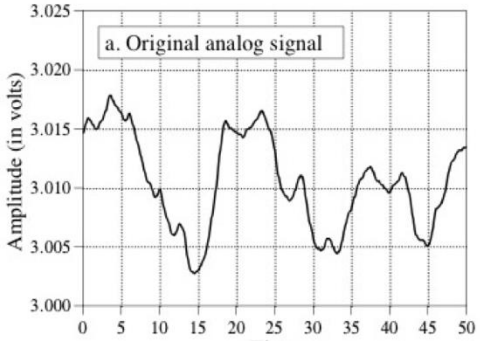
# Digitization process (What is inside an ADC)

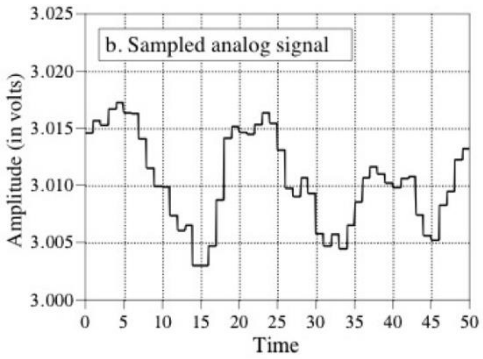


# Digitization process (What is inside an ADC)

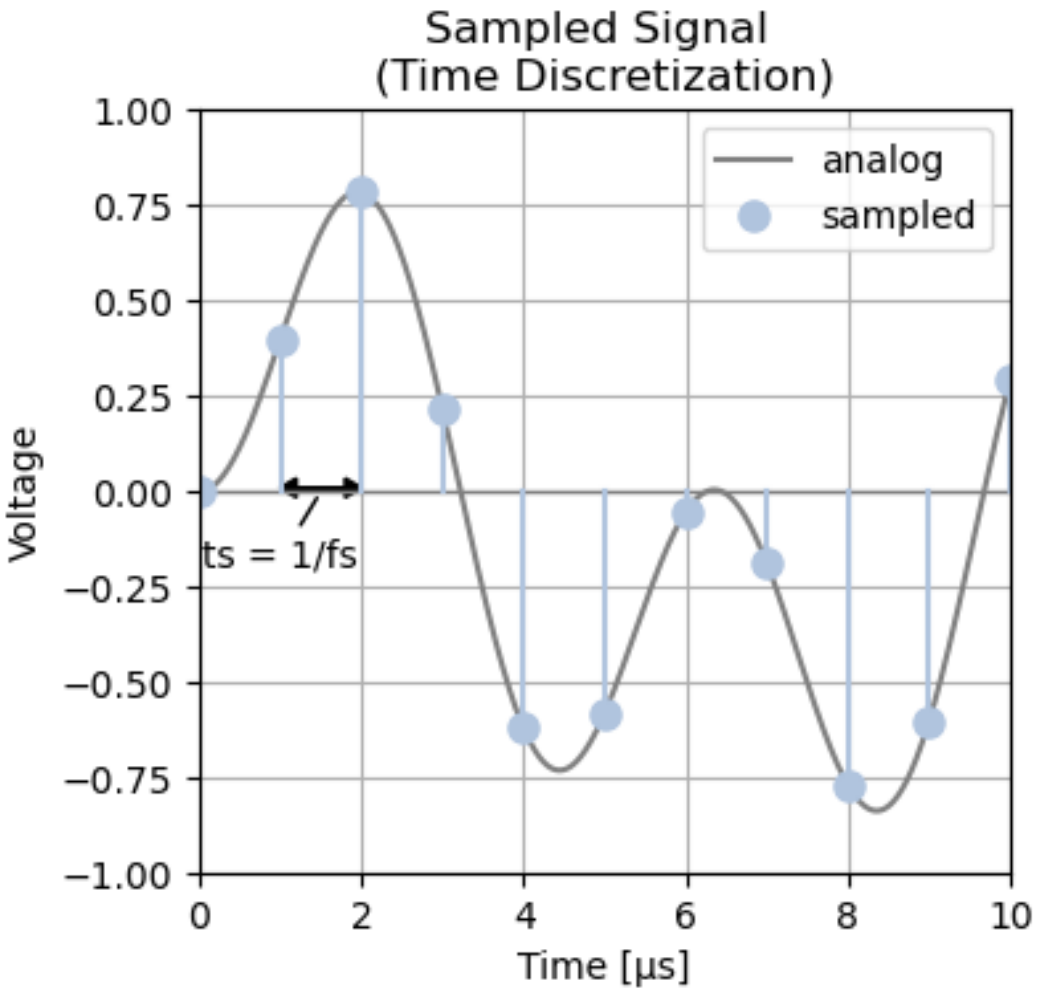


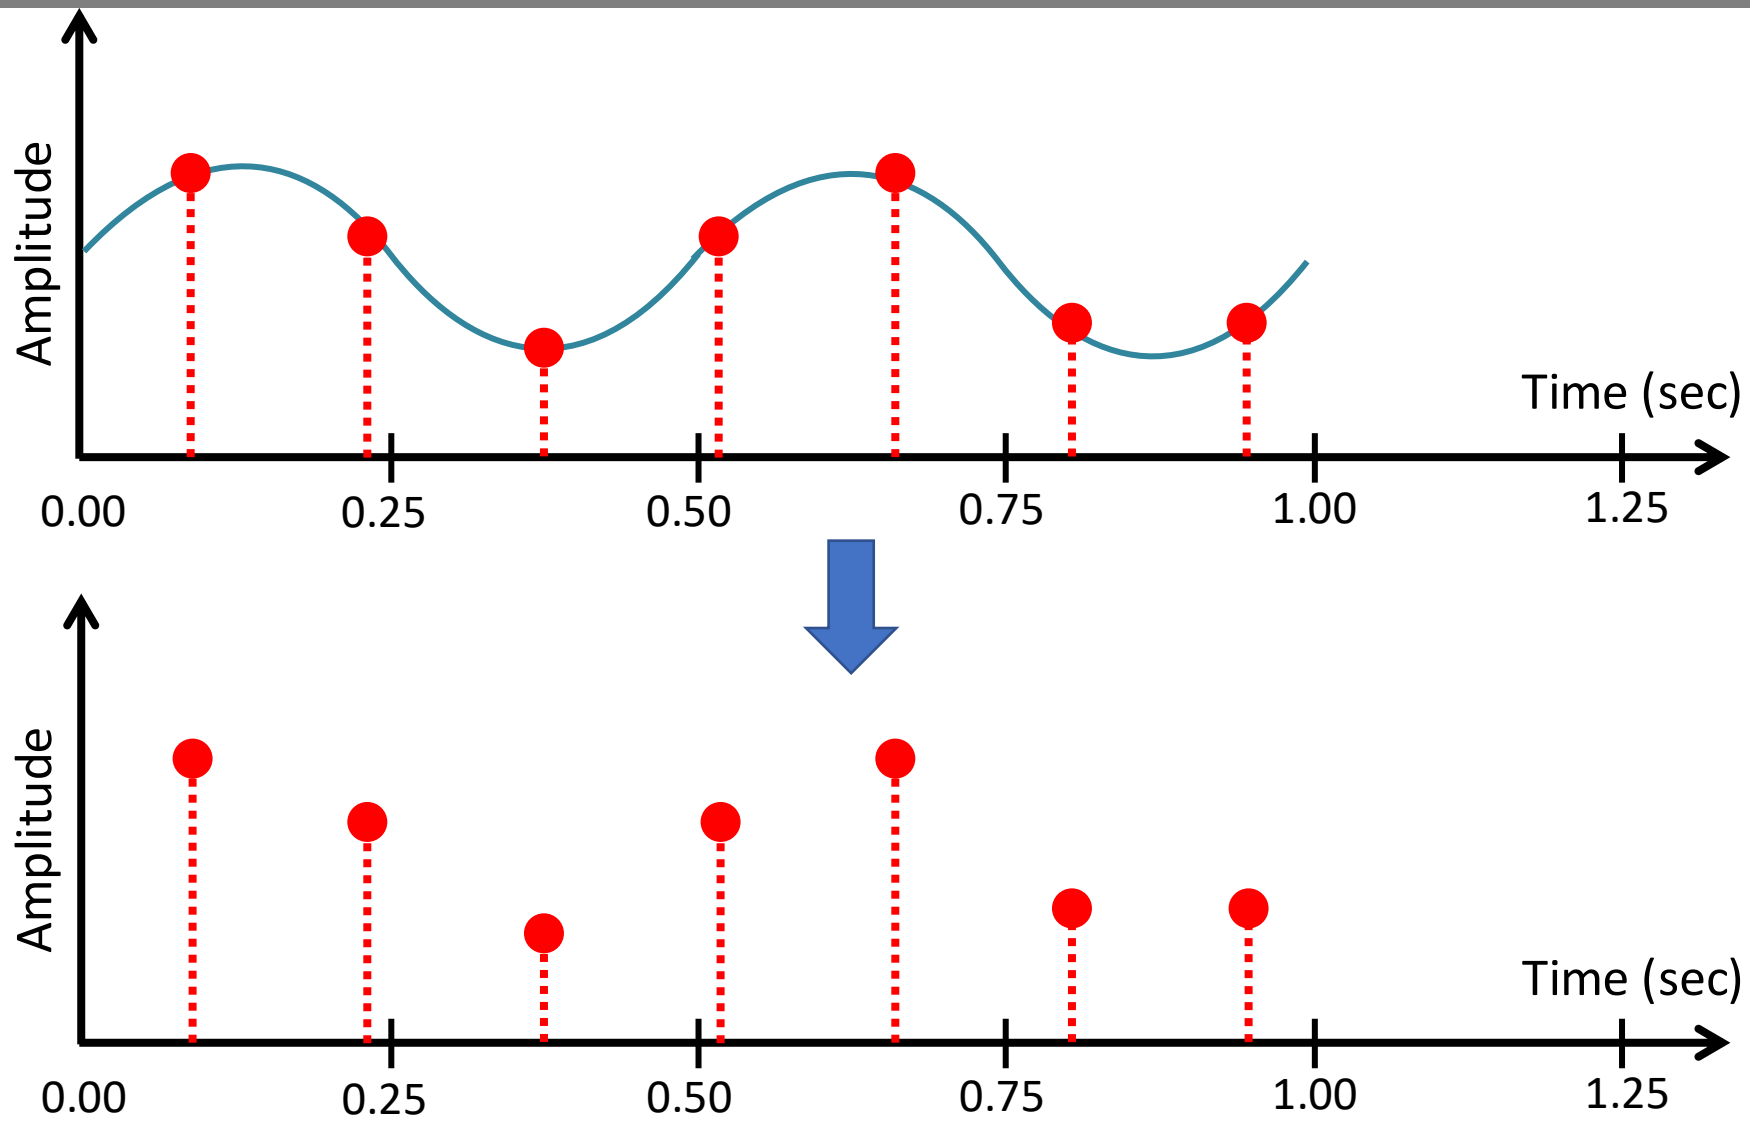




Sample and hold

# Digitization process (What is inside an ADC)

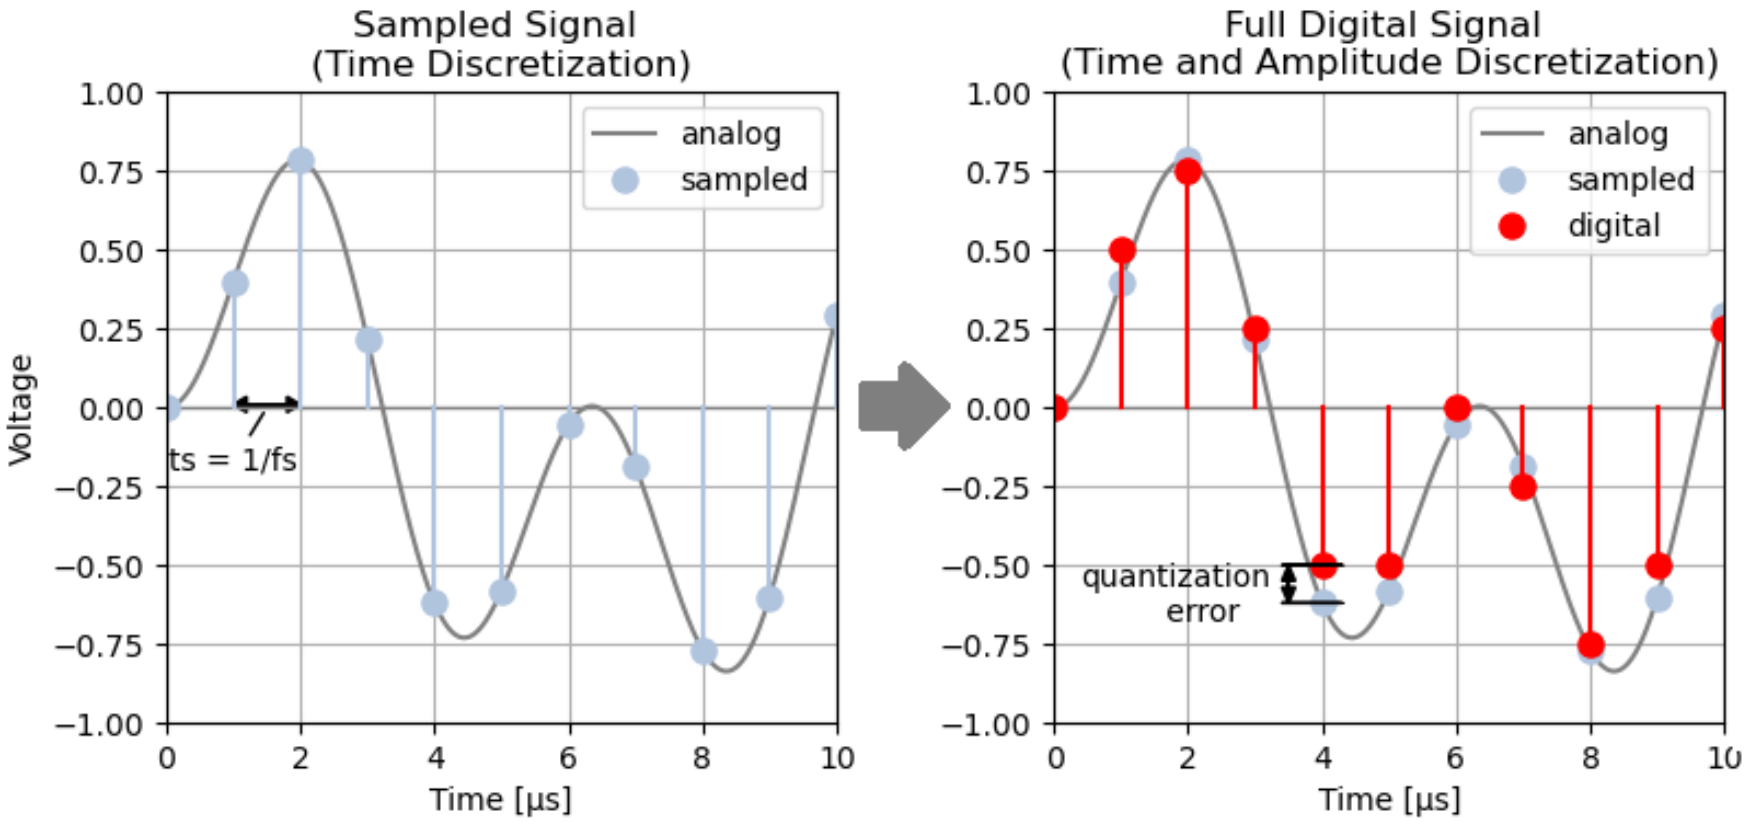


# Discrete in time and in amplitude




[ 0.34 0.22 0.09 0.21 0.30 0.08 0.09 ]

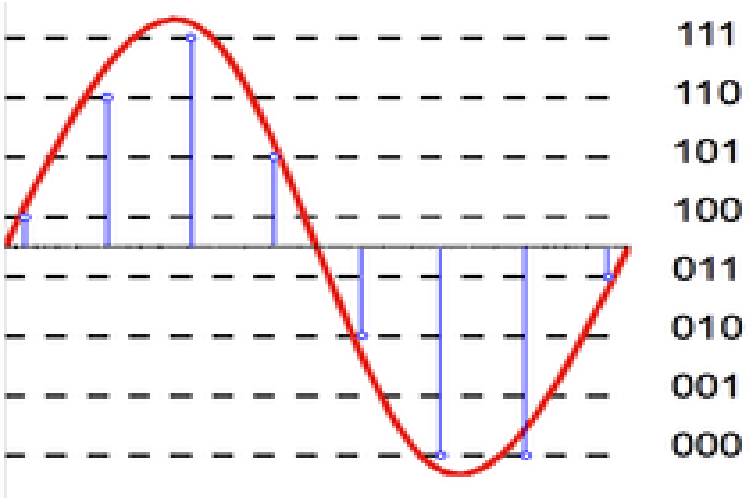
# Digitization process (What is inside an ADC)



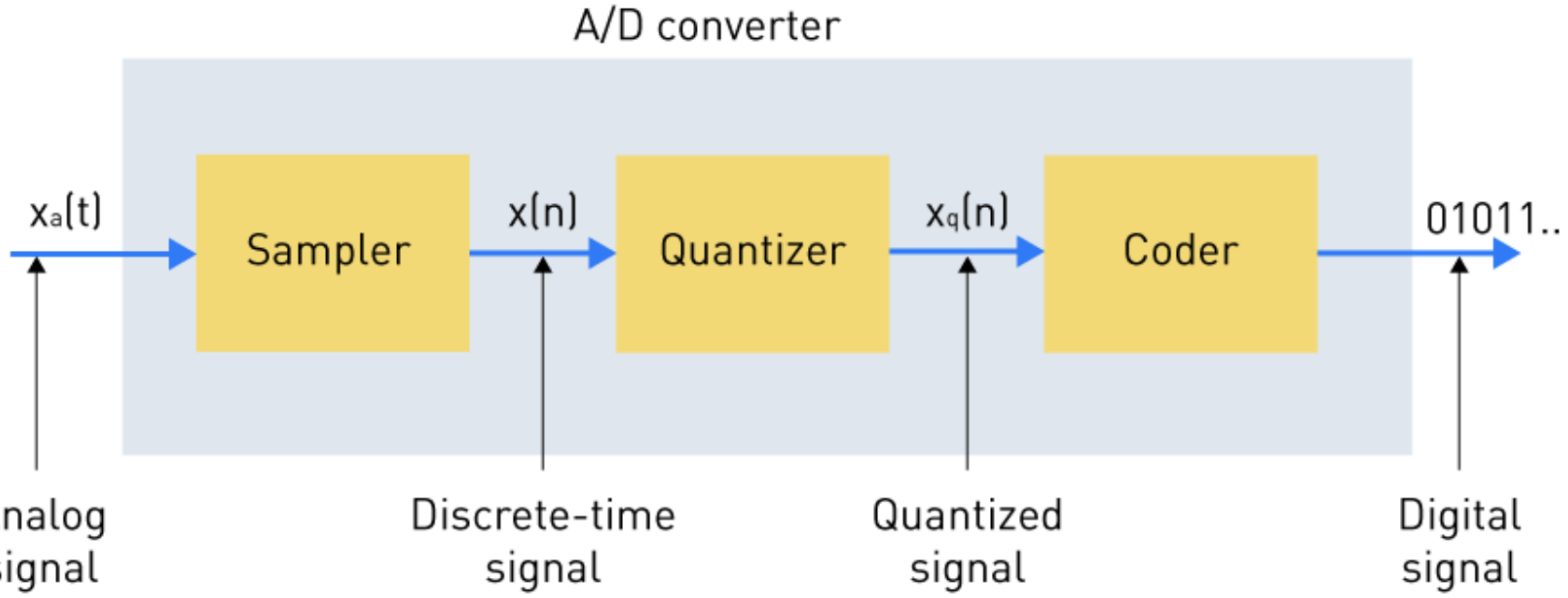
# Quantization level and Bit resolution of an ADC

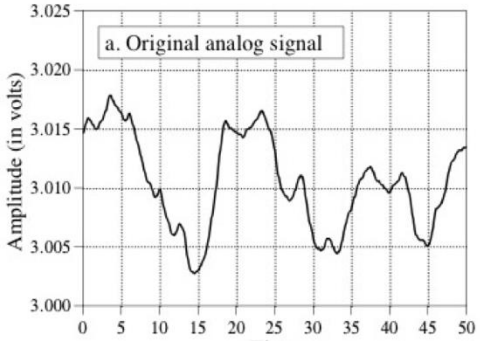


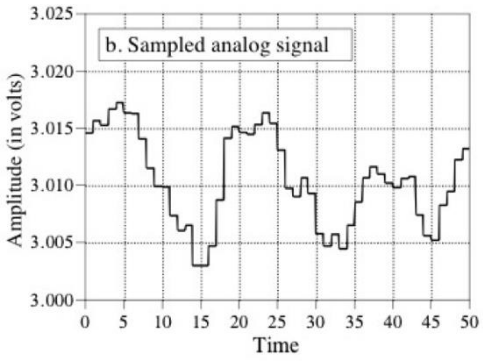
11
10
01
00

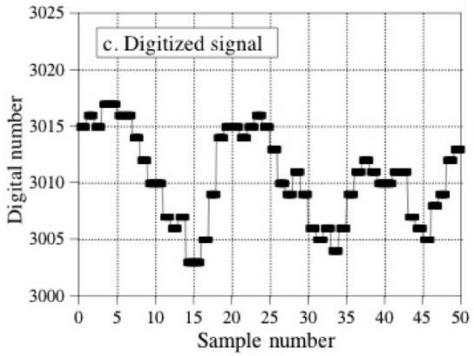


# Digitization process (What is inside an ADC)



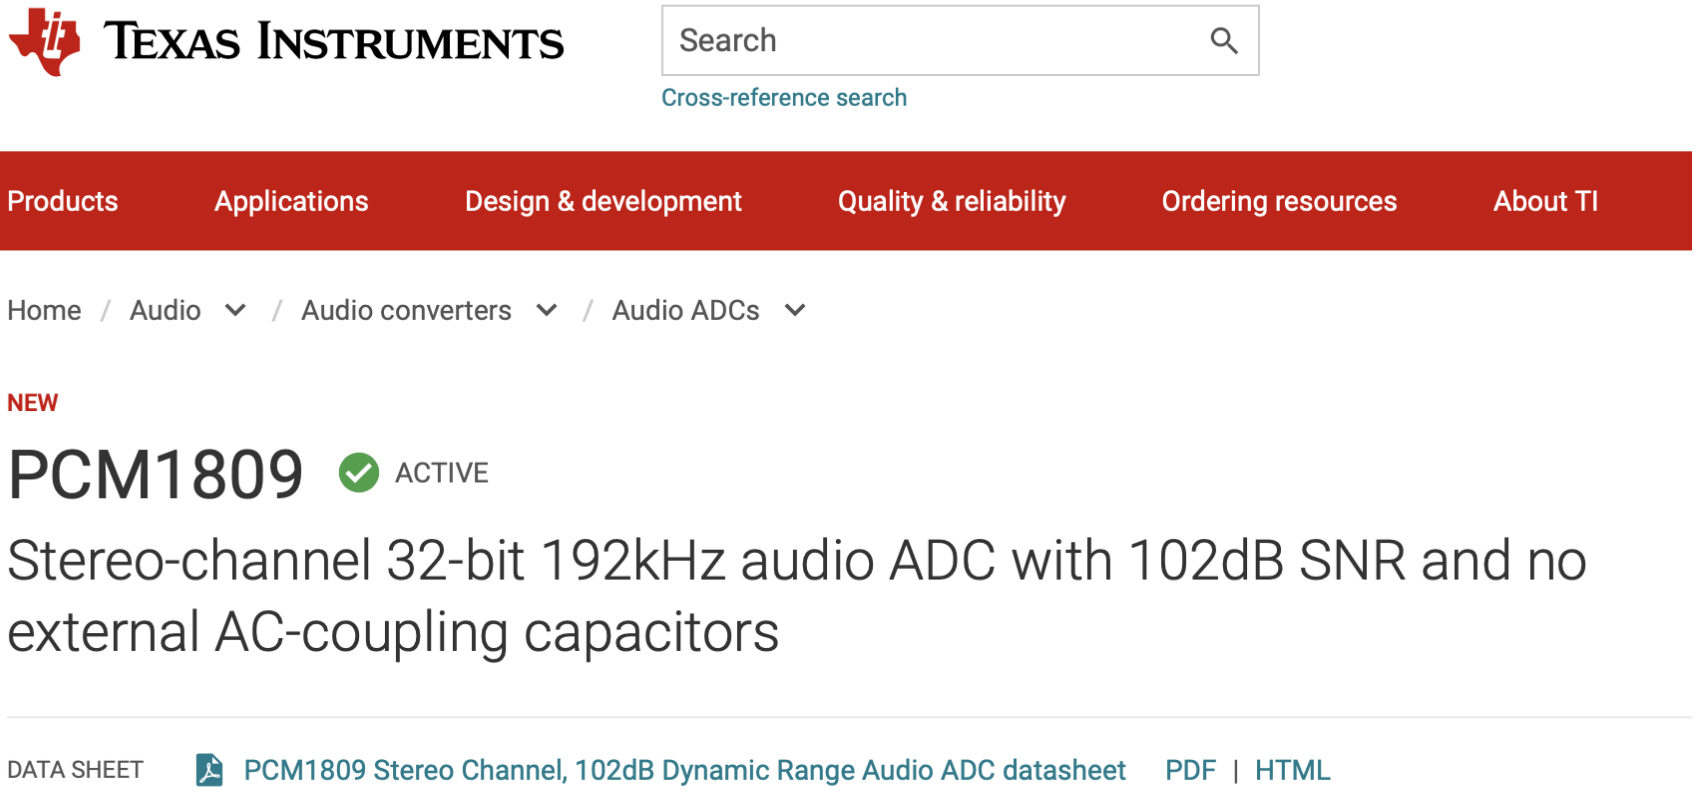




Sample and hold

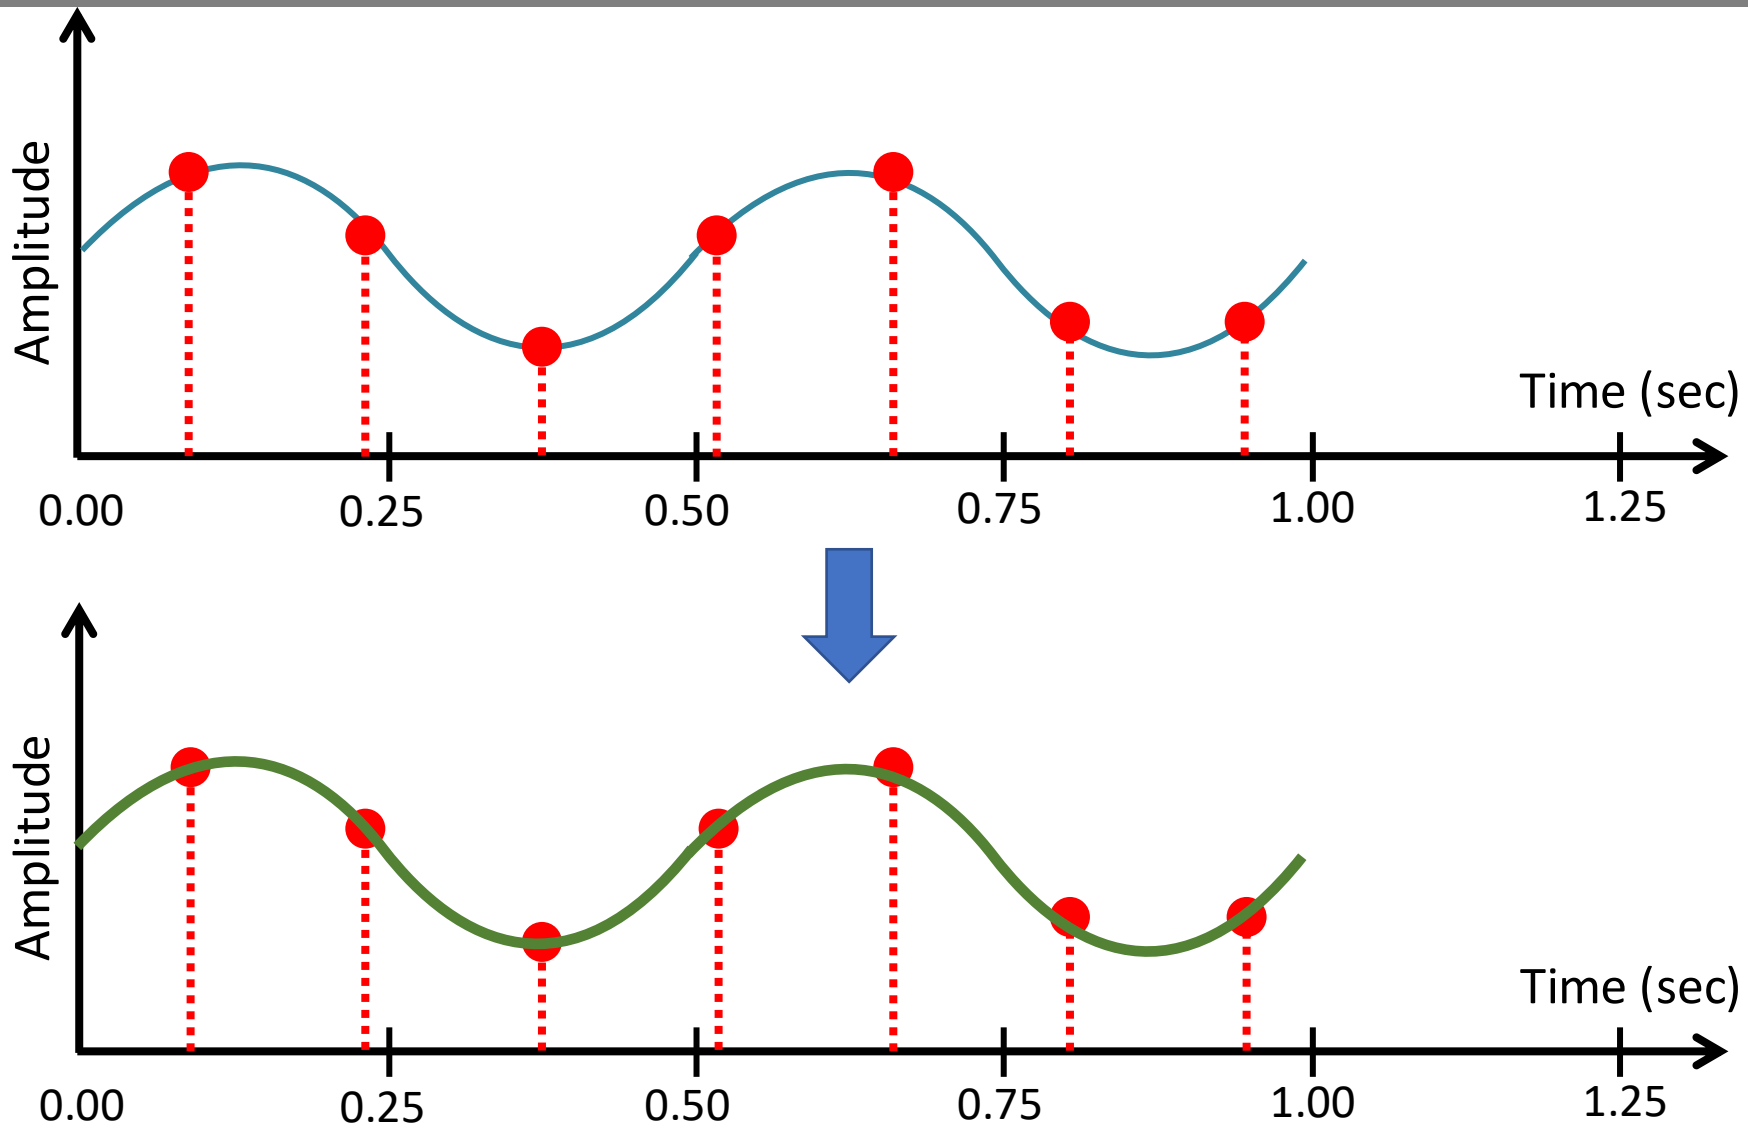
Quantized

# Digitization process (What is inside an ADC)

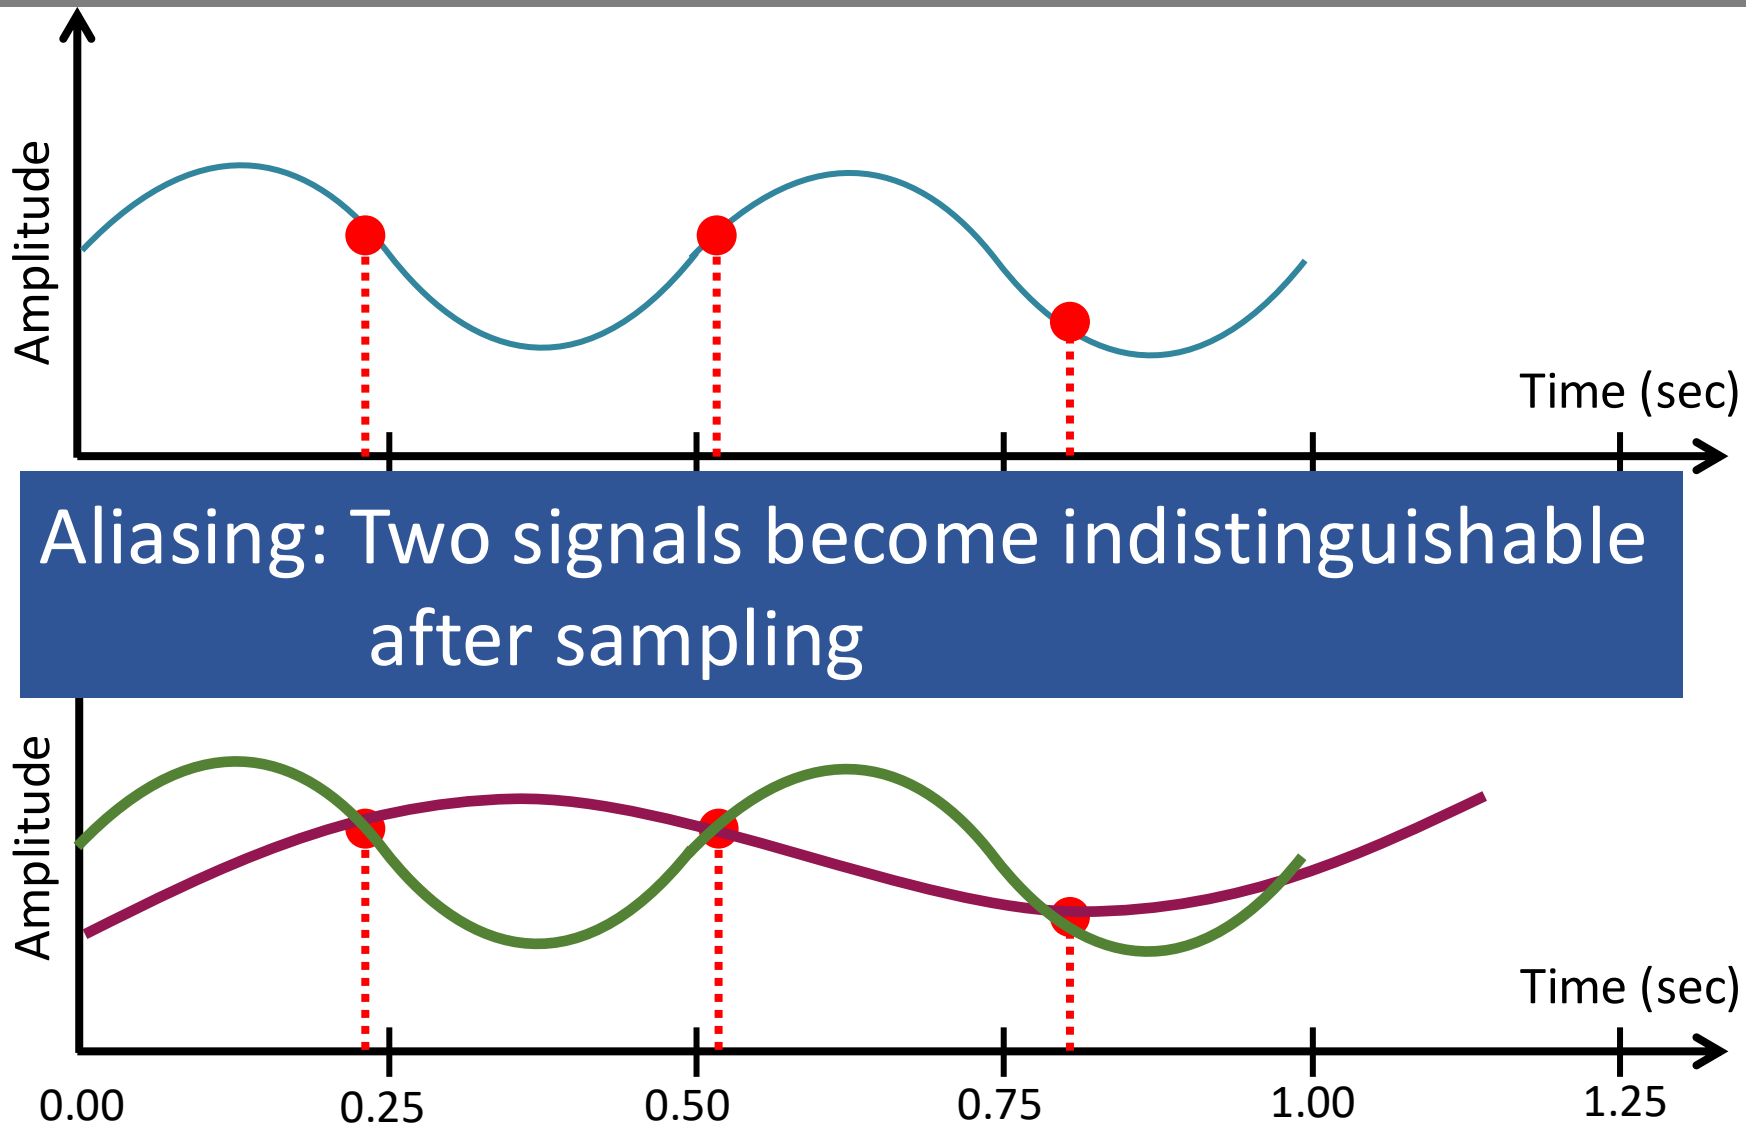


# Gliding over the Sampling Theorem

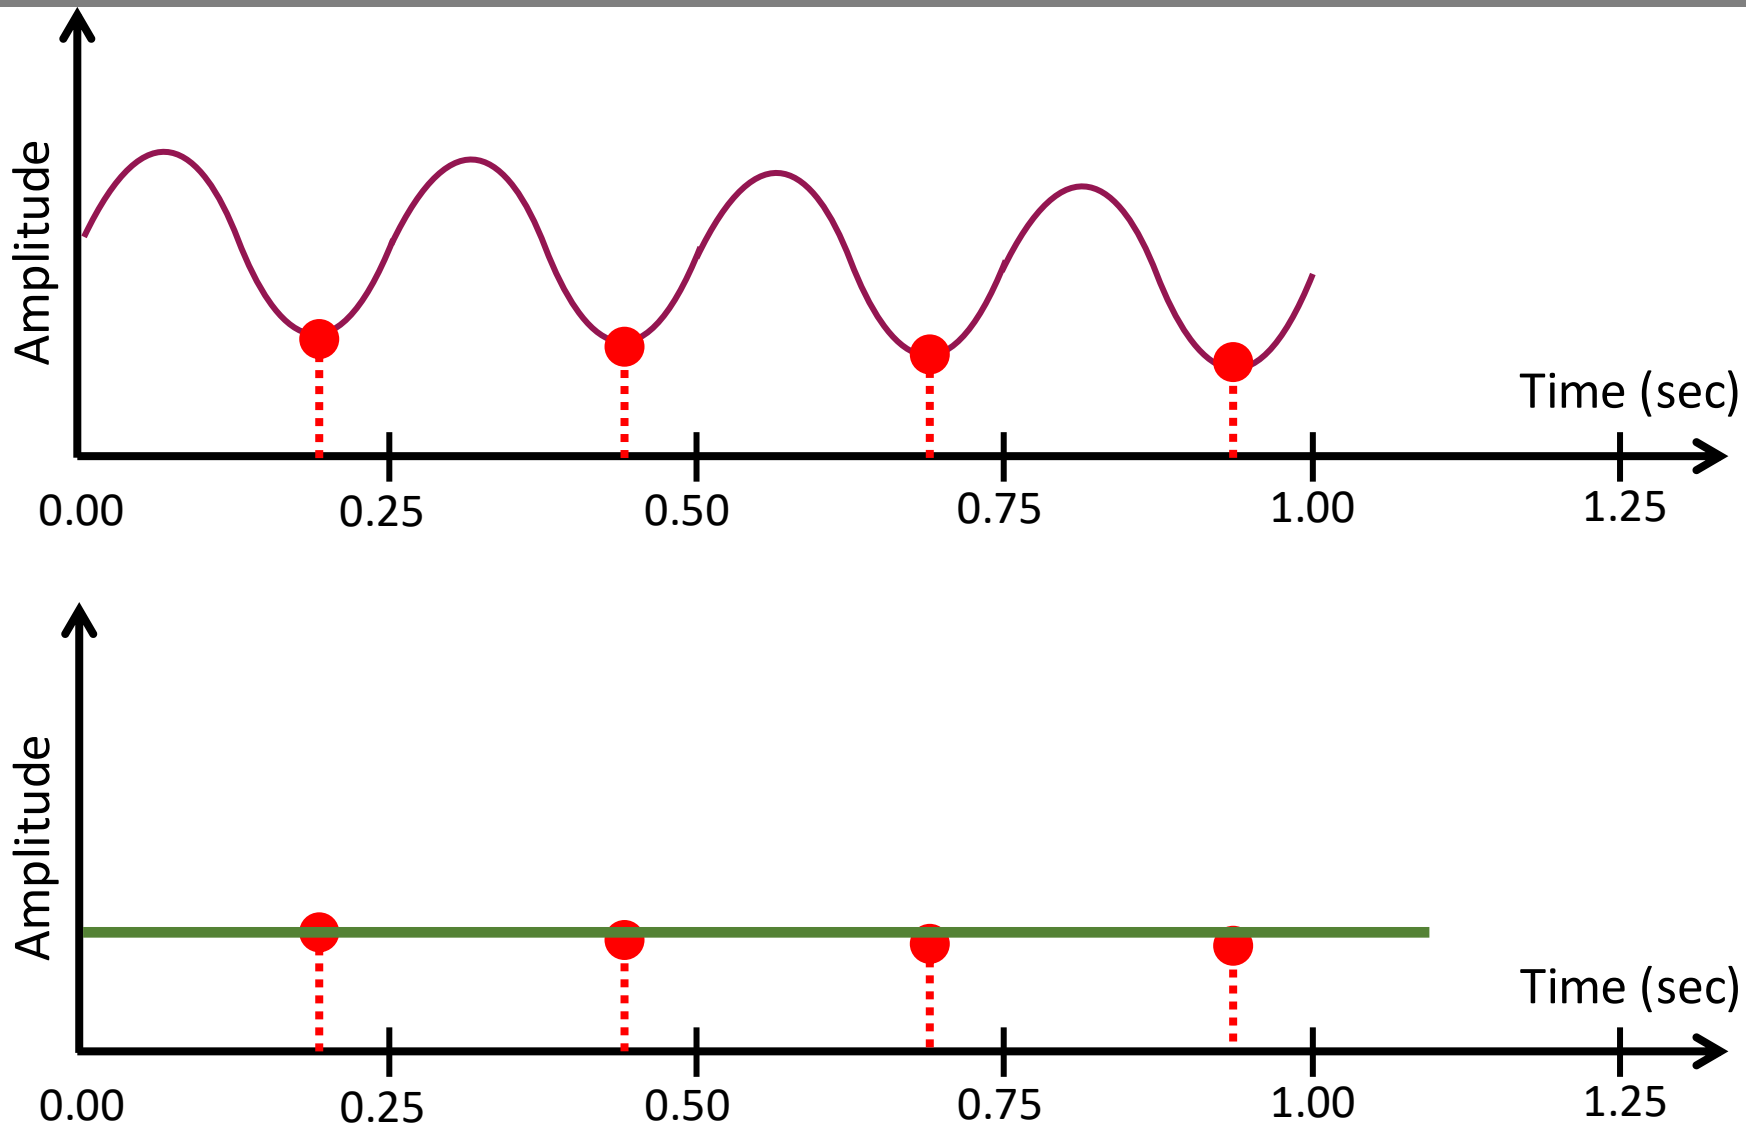
# Sampling theorem



# Sampling theorem



# Aliasing


Amplitude
Time (sec)
0.00
0.25
0.50
0.75
1.00
1.25
Amplitude
Time (sec)
0.00
0.25
0.50
0.75
1.00
1.25

# Aliasing



# Aliasing: Moiré patterns



# Aliasing: Moiré patterns



DEPARTMENT OF
COMPUTER SCIENCE



**Professional Headshot Session**

**Location:** Ground level, near Gannon Auditorium and across from the large screen.

Photos may be used for websites, professional purposes, your LinkedIn profile, etc. For best results, **please avoid fine stripes**, checks, bright colors or busy patterns.

# How to find a good sample rate?

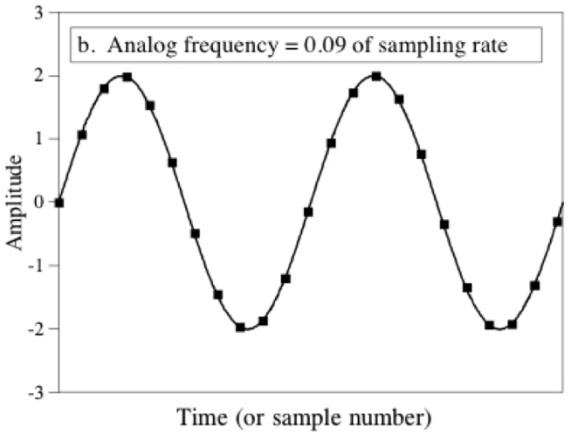
Suppose you sample a signal in some way. If you can **exactly reconstruct** the signal from the samples, then you have done a **proper sampling** and **captured the key** signal **information**

# How to find a good sample rate?

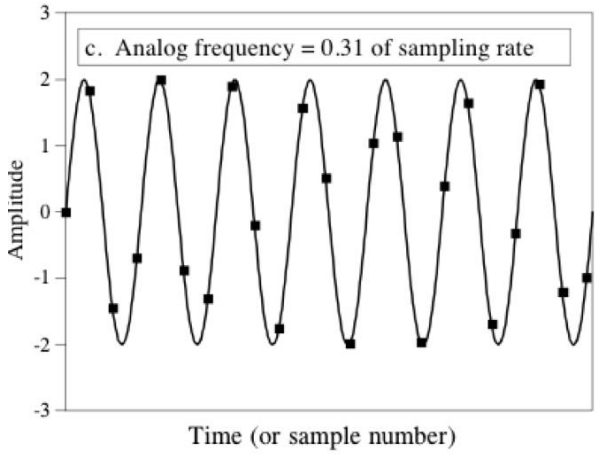
Suppose you sample a signal in some way. If you can **exactly reconstruct** the signal from the samples, then you have done a **proper sampling** and **captured the key** signal **information**

*The* **sampling frequency** $f_s$ , is the number of samples per second. This is to be compared with the signal cyclic frequencies

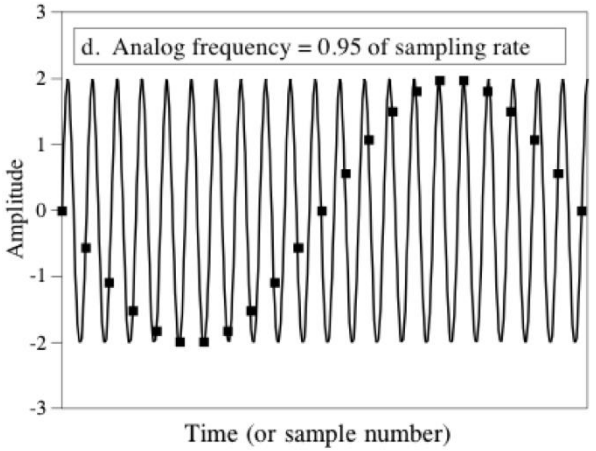
# How to find a good sample rate?




11.1 samples/ cycle




3.2 samples/ cycle



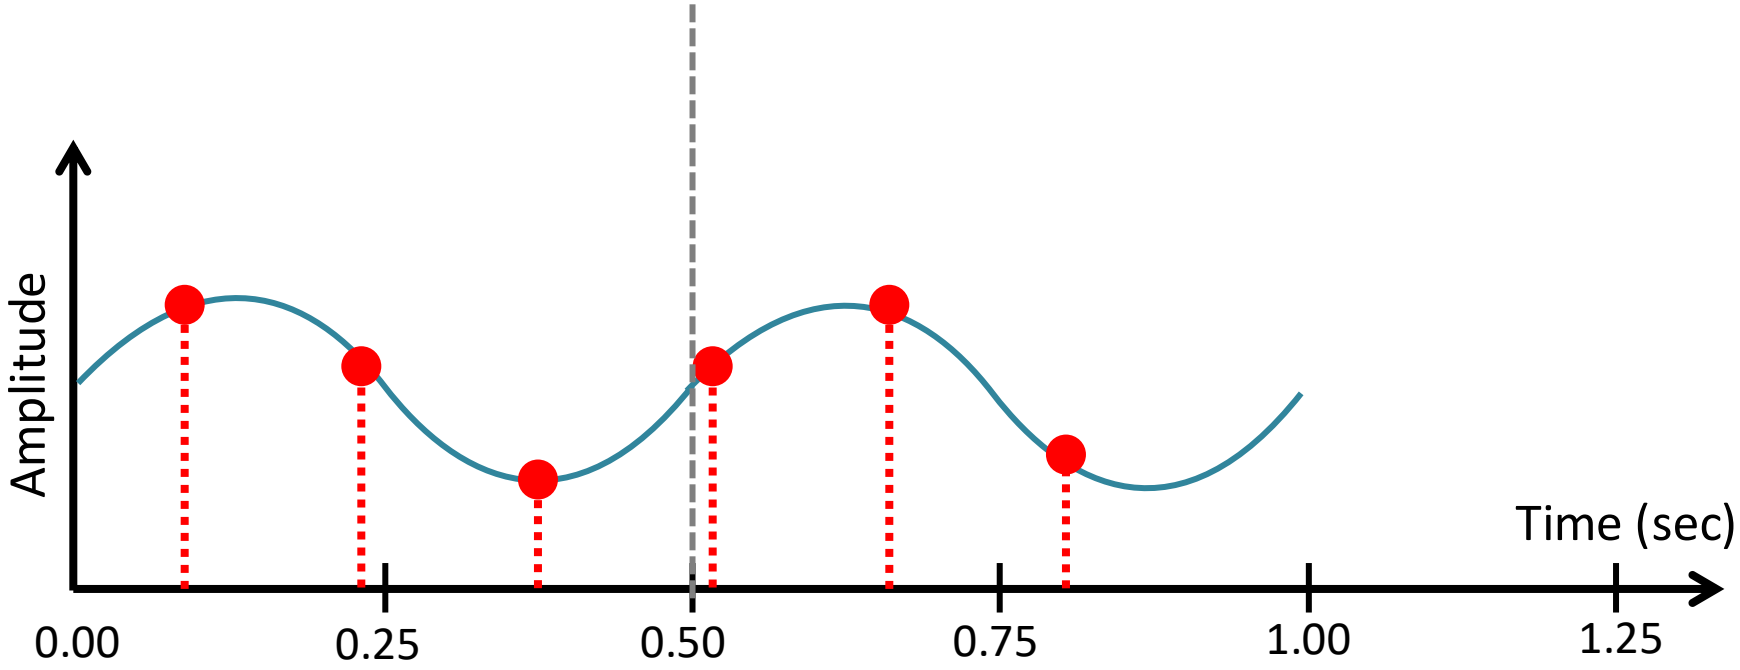

1.05 samples/ cycle

# How to find a good sample rate?

Nyquist-Shannon sampling theorem:
In order to uniquely represent a signal F(t) by a set of samples, the sampling rate must be more than twice the highest frequency component present in F(t).

If sample rate is $f_s$ and maximum frequency we want to record is $f_{max}$, then

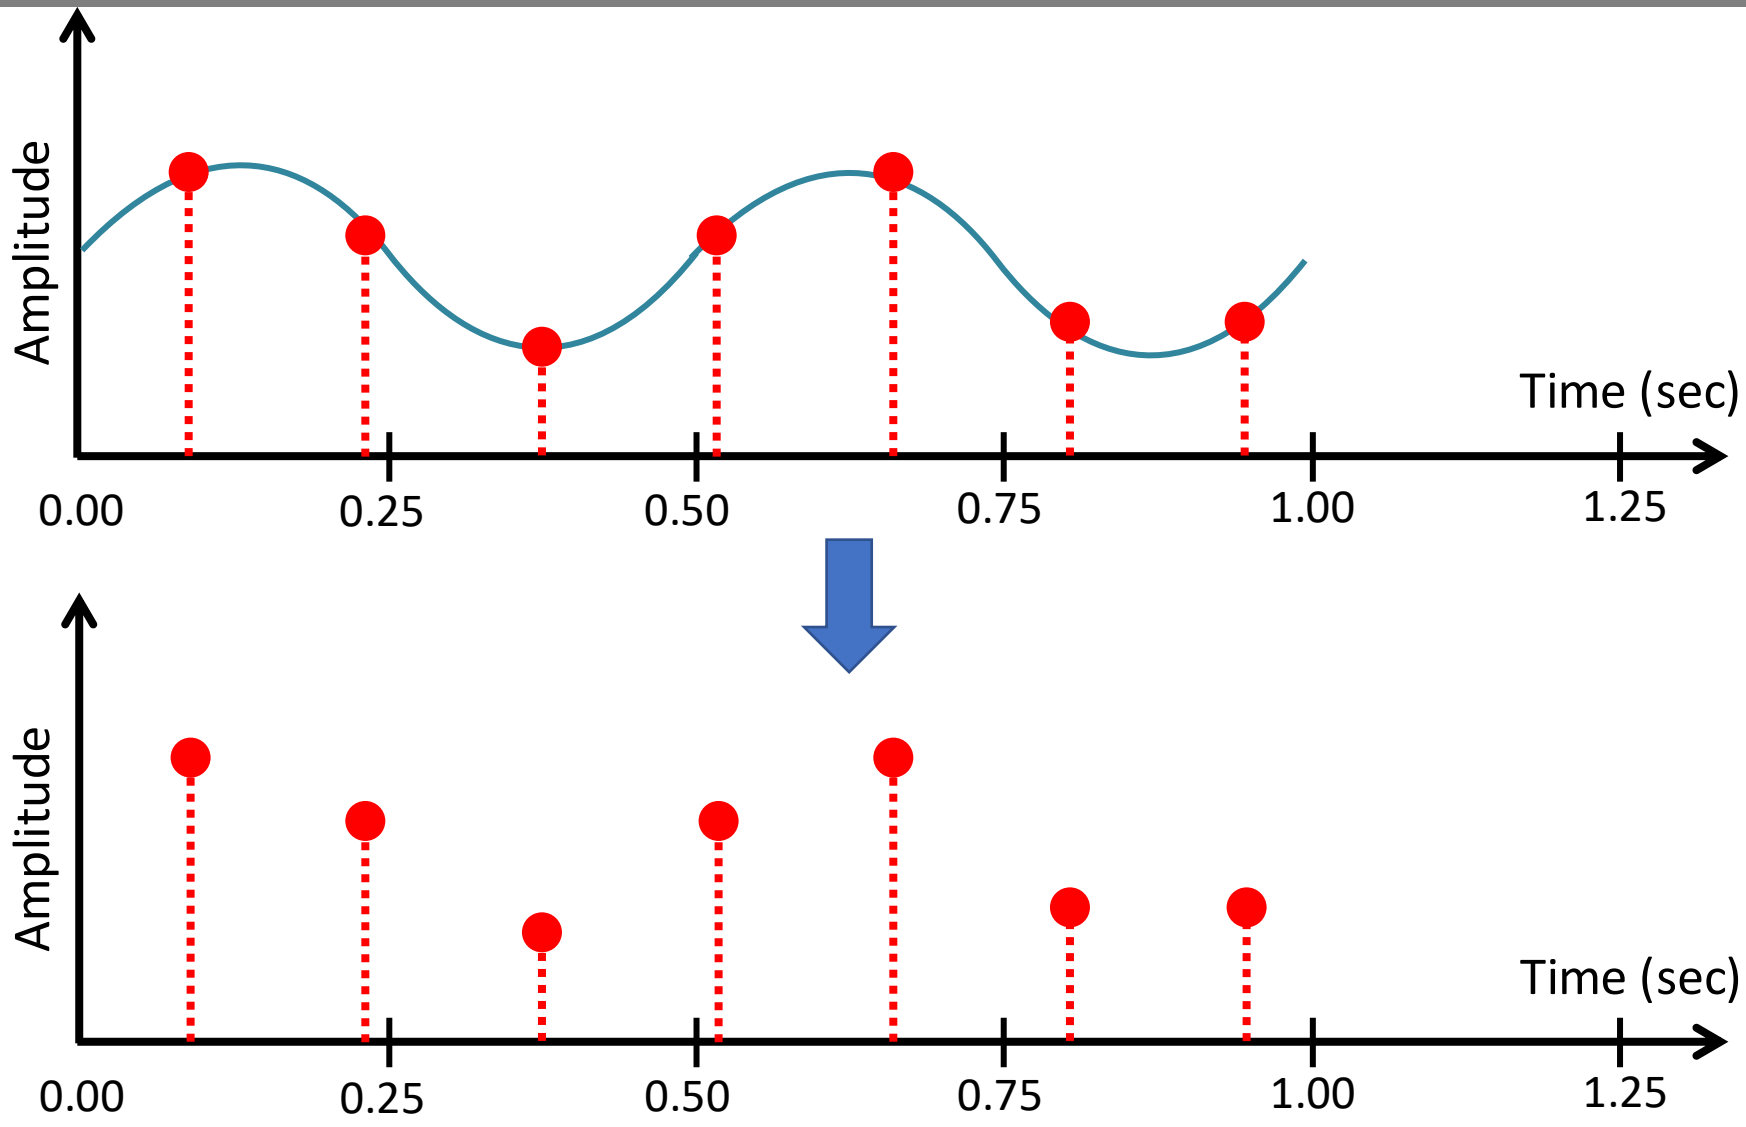
$$f_s > 2f_{max}$$




Nyquist frequency = Maximum alias-free frequency for a given sample rate. ($f_s$ / 2)

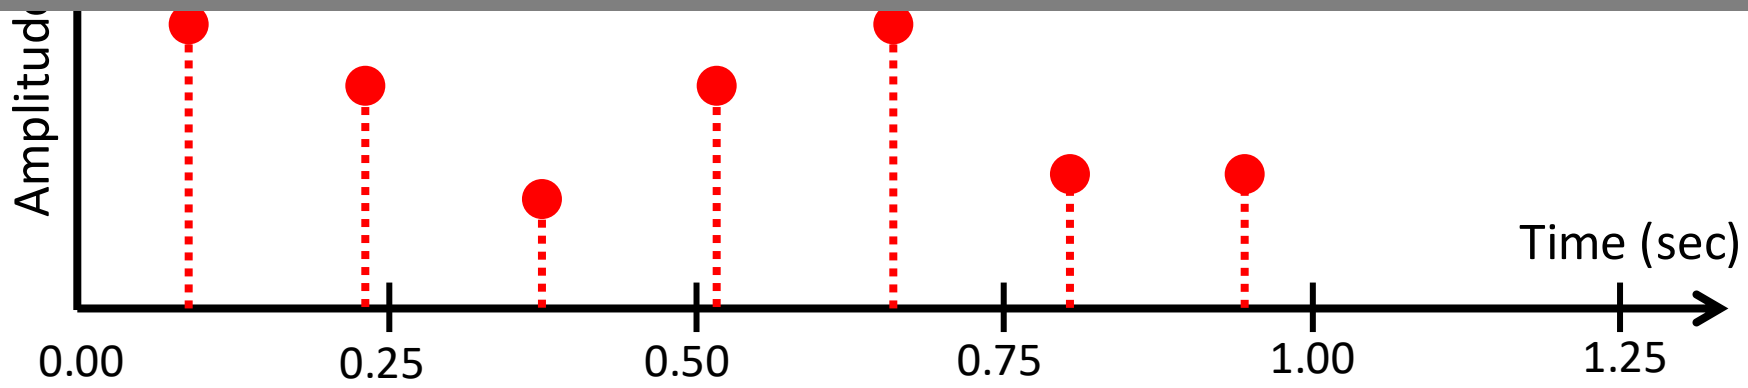
Nyquist rate = Lower bound of sample rate for a signal. ($2f_{max}$)

# Samples of the physical phenomena




[ 0.34 0.22 0.09 0.21 0.30 0.08 0.09 ]

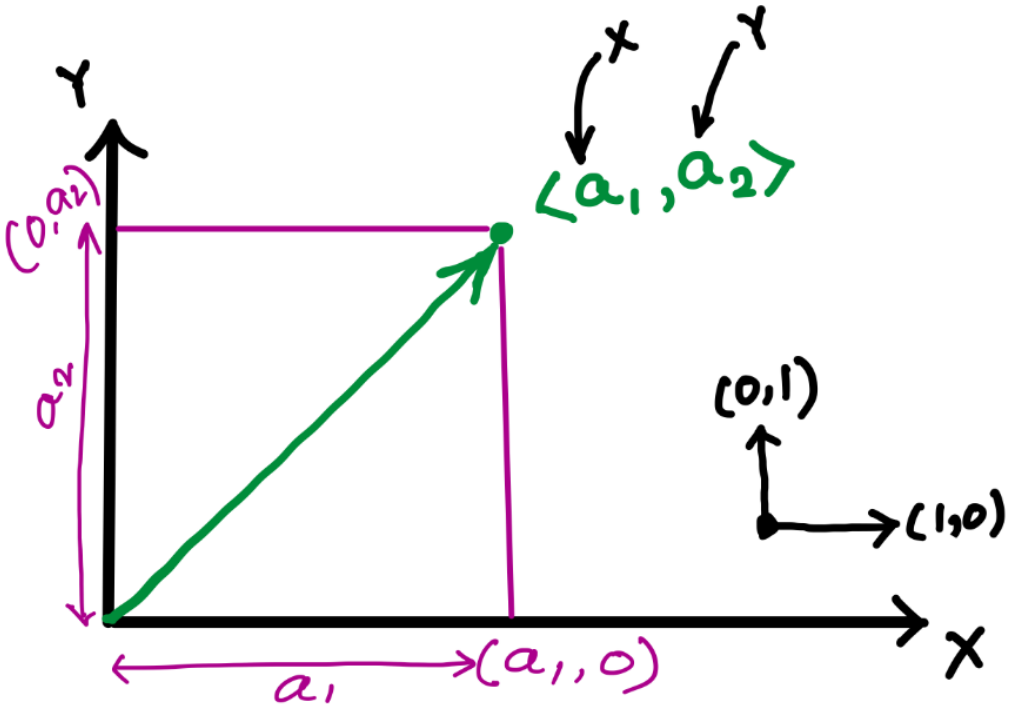
# Samples of the physical phenomena



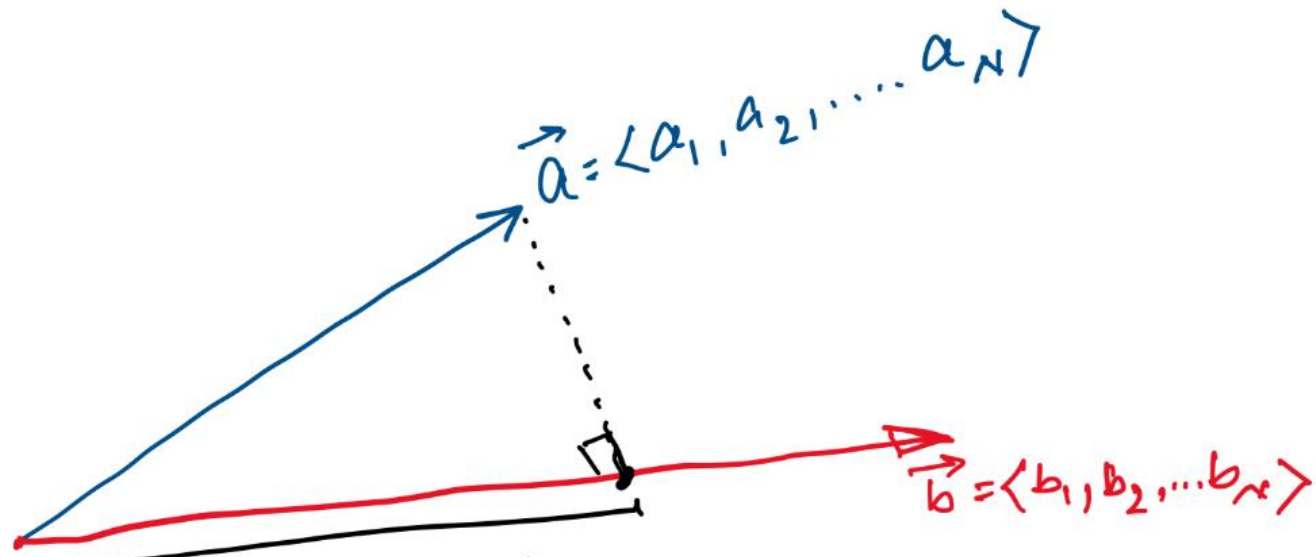

[ 0.34 0.22 0.09 0.21 0.30 0.08 0.09 ]

A point in the N-dimensional space

# Signal as a point in a N-dimensional space



# Projection and Dot product of vectors

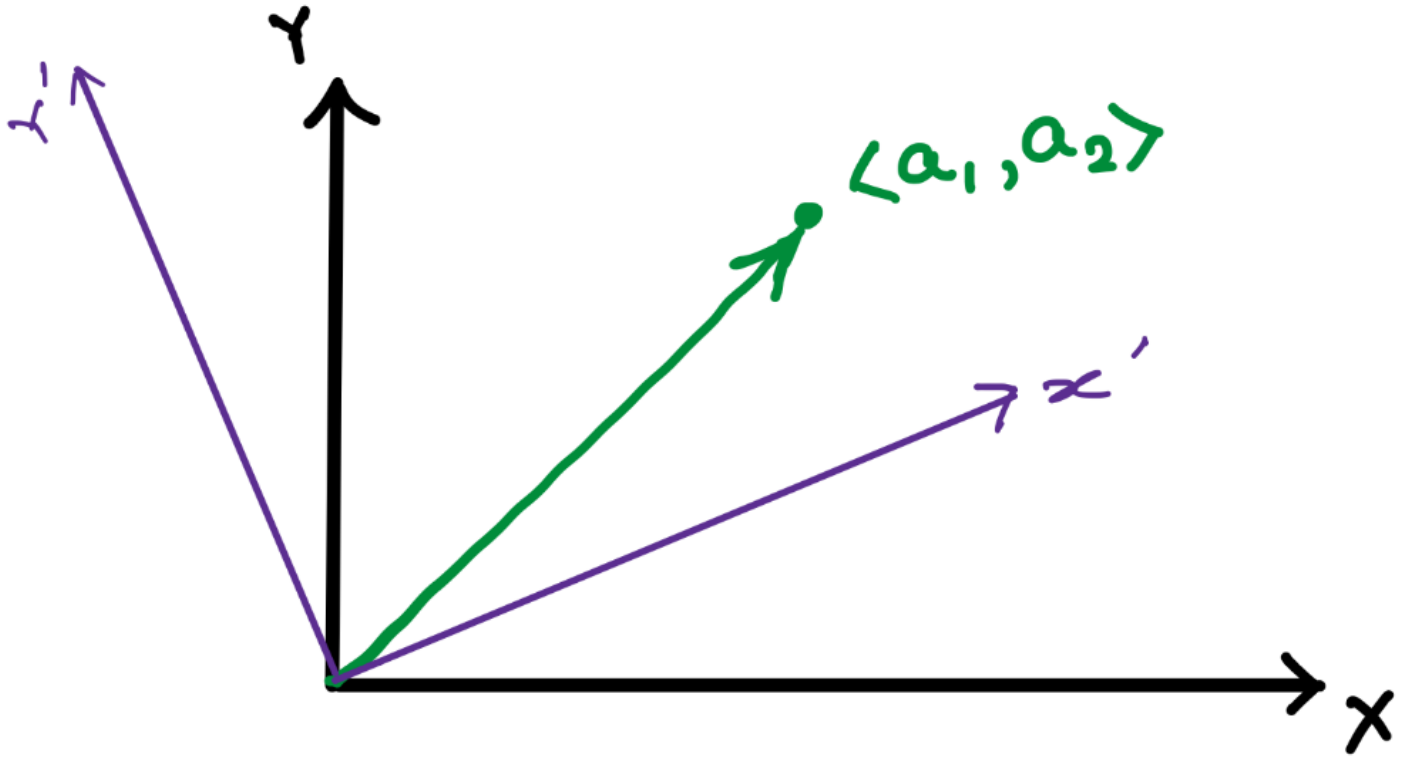



Scalar projection $= \dfrac{a \cdot b}{\|b\|}$

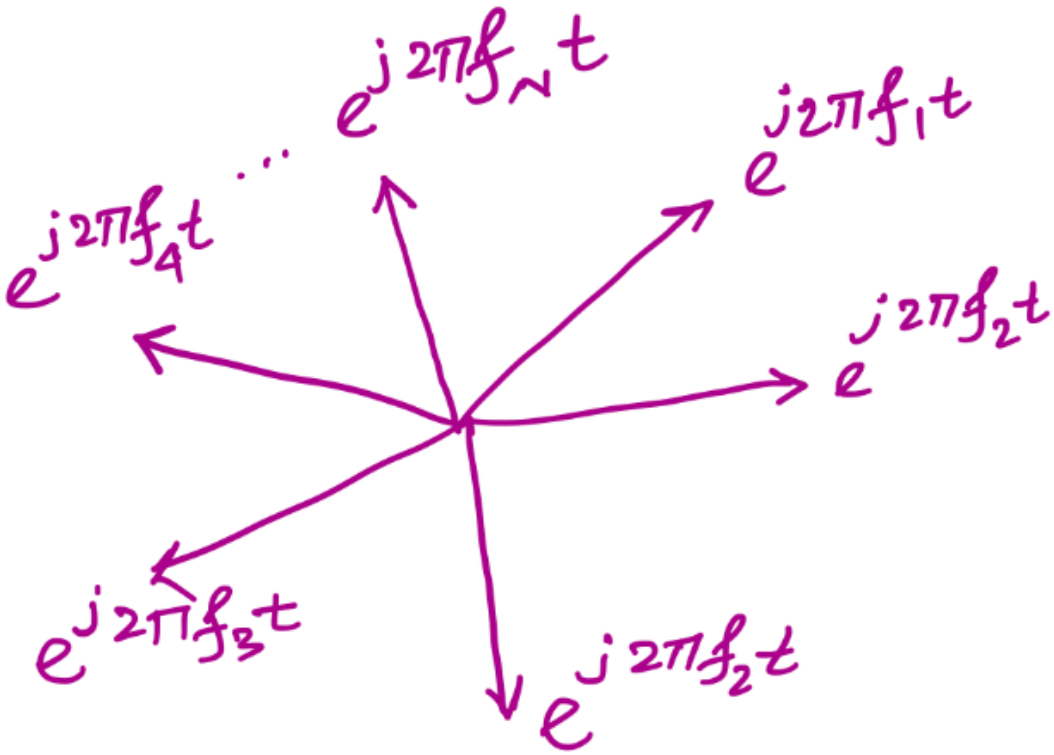
Dot product $= \sum_{i=1}^{N} a_i b_i$

$$\langle x, y \rangle = \int_{-\infty}^{\infty} x(t)\,\overline{y(t)}\,dt$$

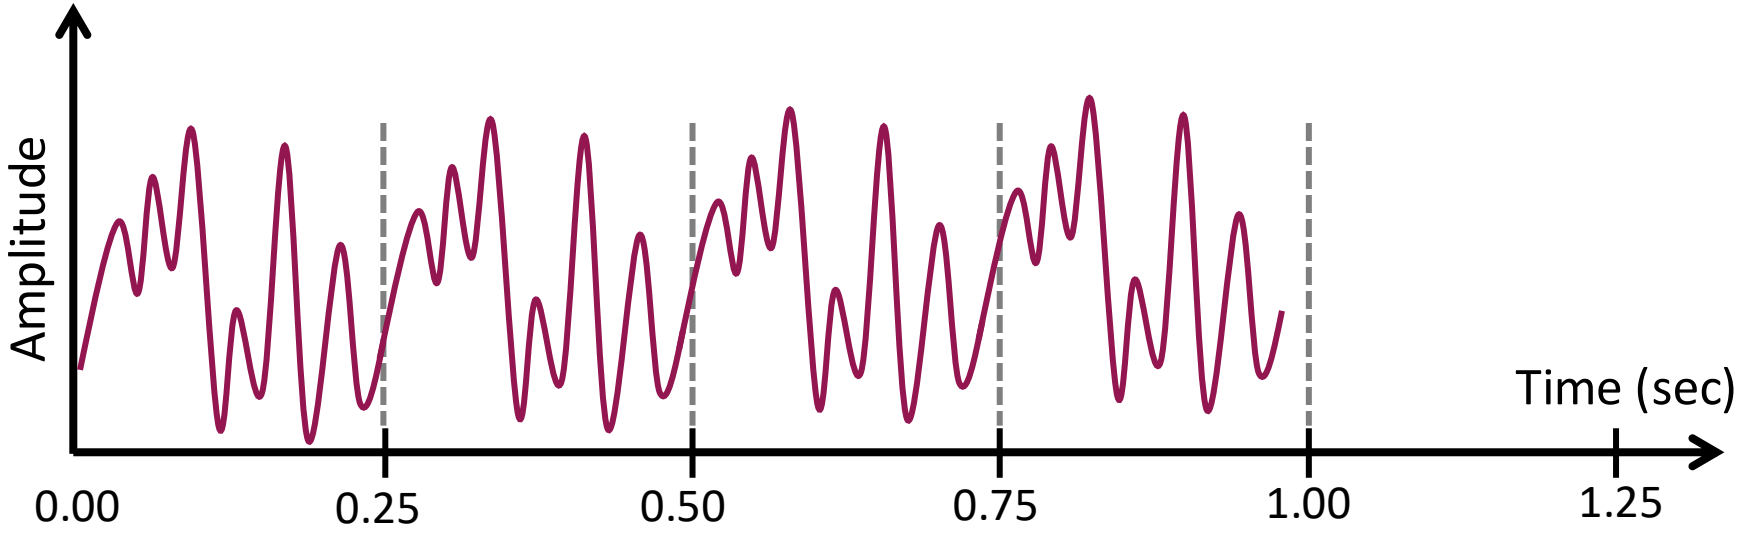
# Change of basis



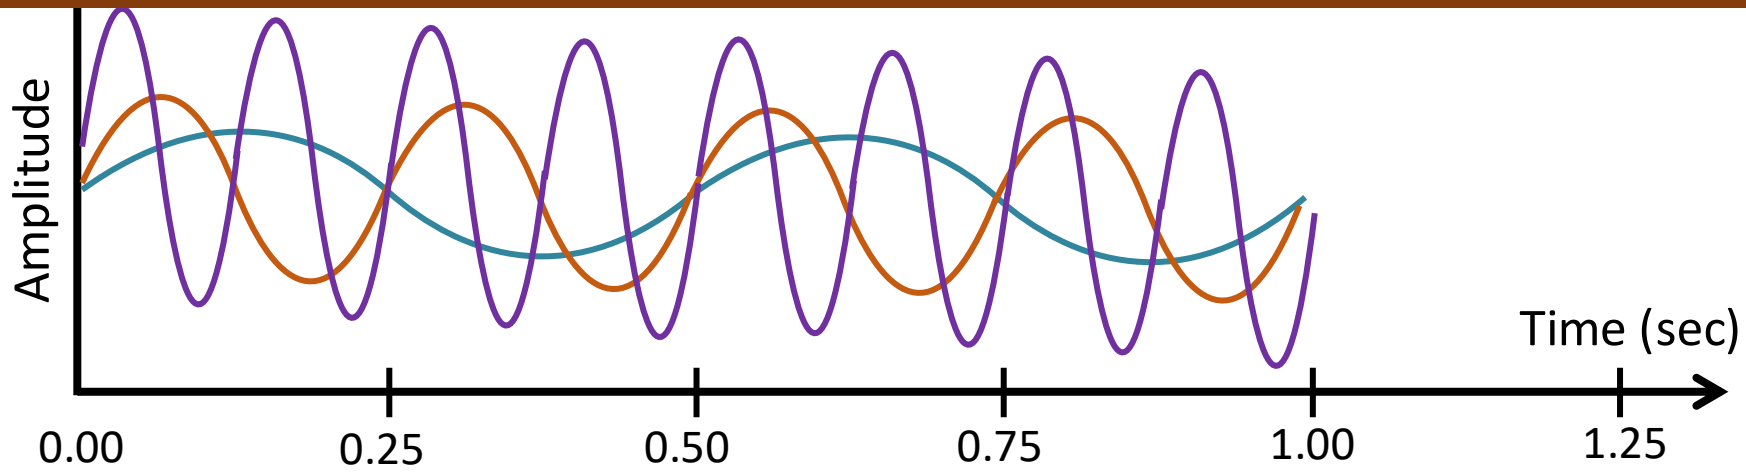
# “Change of basis” to Frequency domain



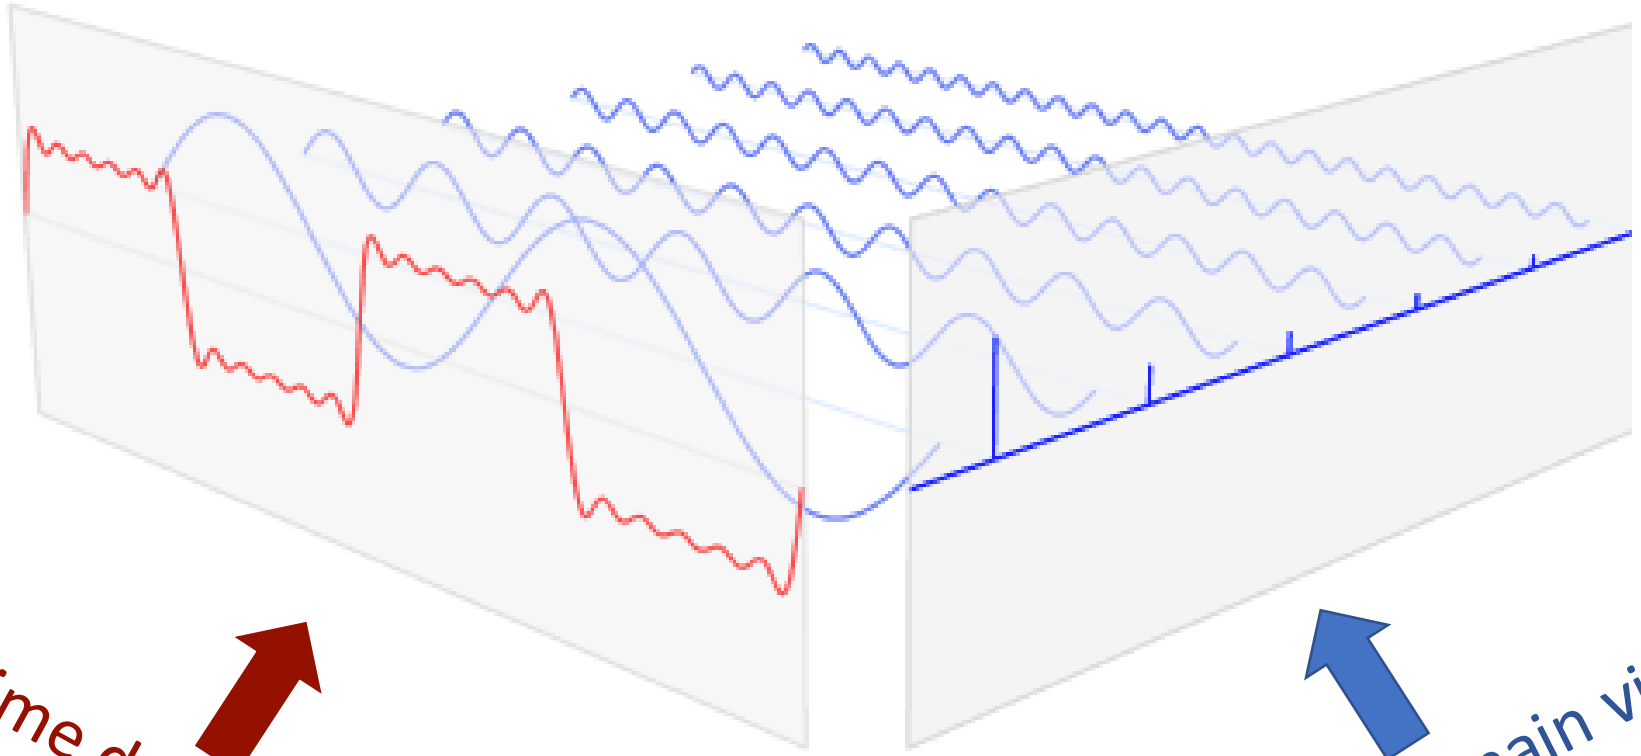
# Recall from past…




The key idea: Any function can be written as the sum of sinusoidal functions.



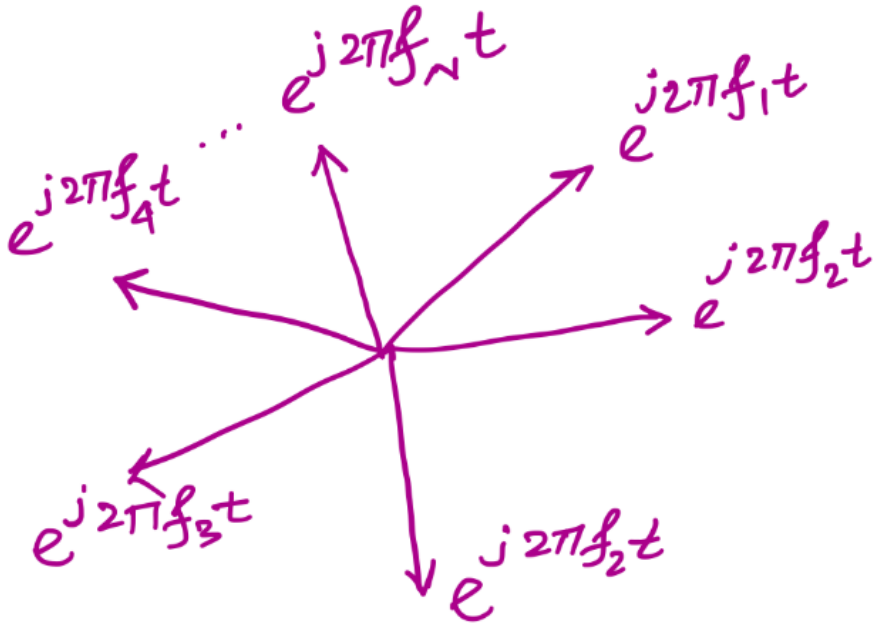
# Time Domain and Frequency Domain



Time domain view

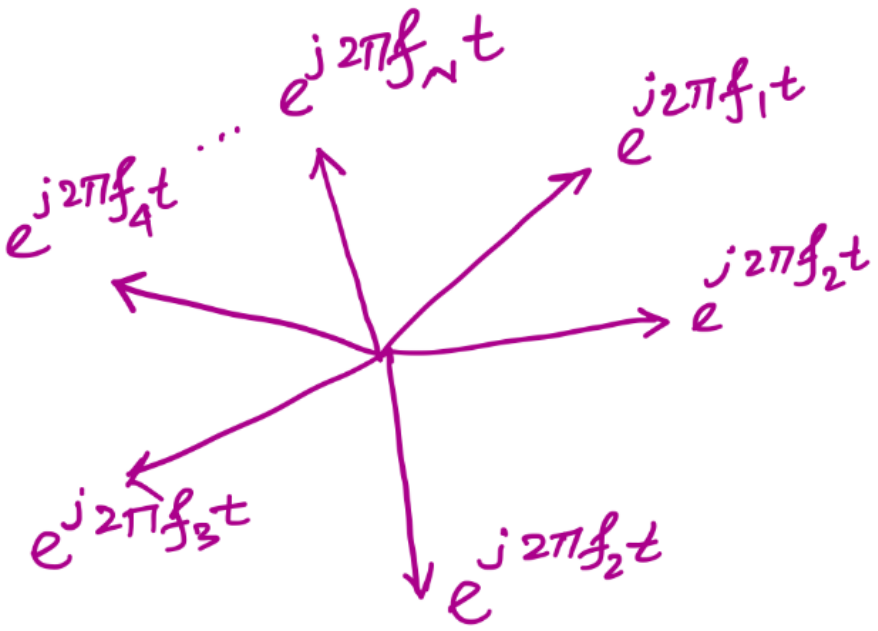
Frequency domain view

# “Change of basis” to Frequency domain




$$X(f) = \int_{-\infty}^{\infty} x(t)\, e^{-j2\pi ft}\, dt = \langle x(t), e^{j2\pi ft} \rangle$$

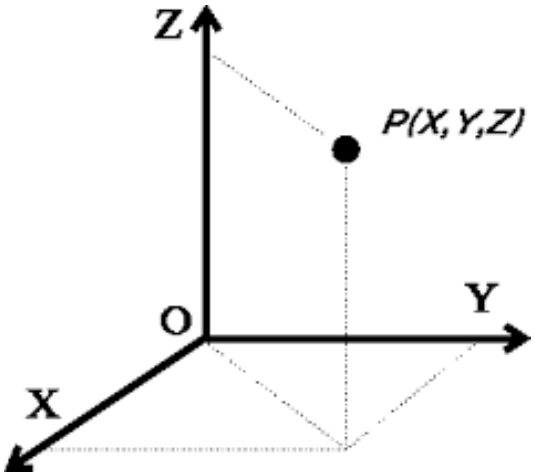
# “Change of basis” to Frequency domain




Fourier Basis

$$X(f) = \int_{-\infty}^{\infty} x(t)\, e^{-j2\pi ft}\, dt = \langle x(t), e^{j2\pi ft} \rangle$$

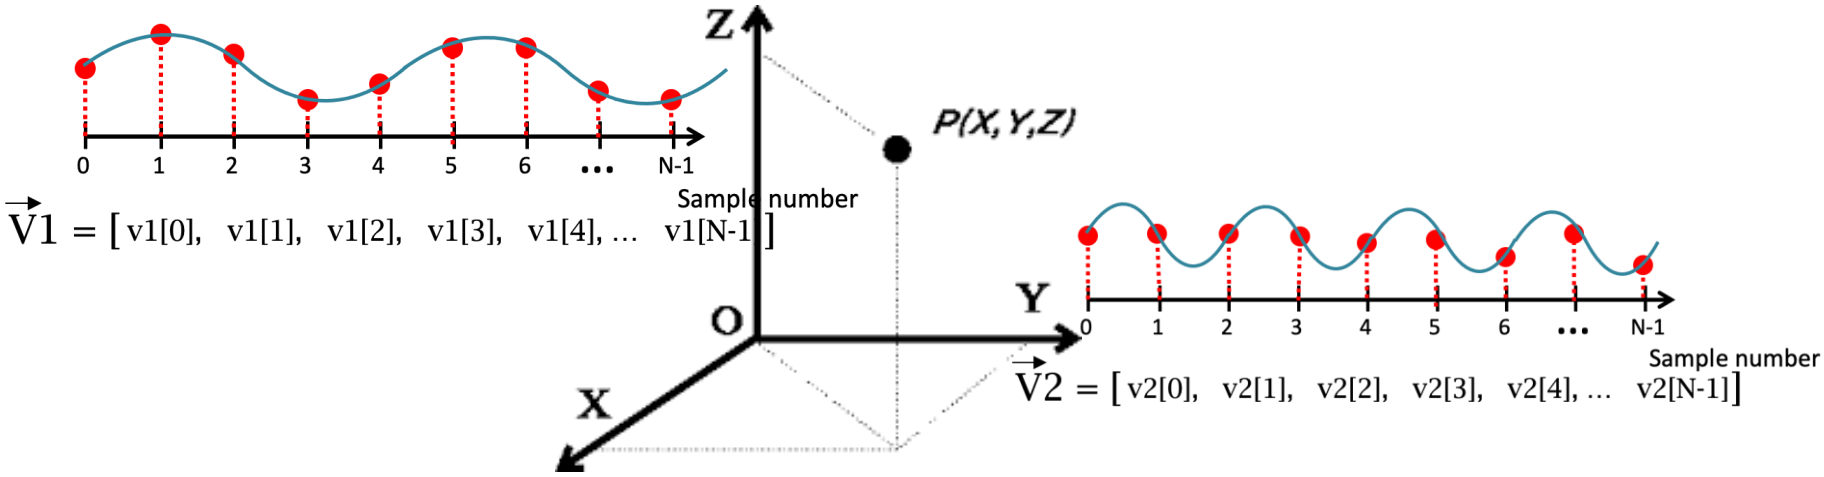
# Finding good basis: Orthogonality




Two vectors $\vec{u}, \vec{v}$ in an inner product space are **orthogonal** if their dot product is zero:

$$\vec{u} \cdot \vec{v} = 0.$$

# Finding good basis: Orthogonality




**What makes two signals orthogonal?**

The formal definition for orthogonal signals is that their inner product (also called the dot product) is zero:

$$\int_{-\infty}^{\infty} x(t)y(t)dt = 0$$

For discrete signals, this becomes:

$$\sum_{n=1}^{N} x[n]y[n] = 0$$