Cyber-Physical System

Cognitive Cyber-Physical system

- The future of intelligent Cyber-physical Systems is oriented towards the implementation, either in software or hardware, of cognitive mechanisms
- However, the predicted future is not something will happen decades from now
  - ✓ Adaptation is a basic form of cognition associated with elementary automatic reactions
  - ✓ Many embedded solutions present in the market introduce adaptation mechanisms at various levels
  - ✓ Technological advances in the hardware permit the execution of sophisticated embedded solutions

The future of embedded systems has started already…

What is Cognitive Cyber-Physical System?

A cognitive Cyber-physical system is a Cyber-physical system that takes advantage of cognitive processes to propose intelligent solutions

Cognitive Processes in Human Brain

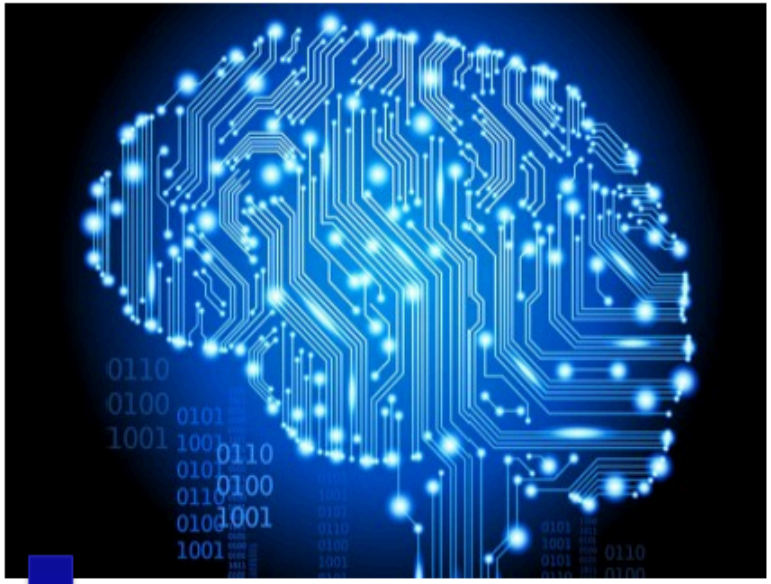

- Providing an engineering-oriented perspective of brain functions
- Functional description of some basic processes of the human brain, with a special focus on emotional and cognitive processing

Apply "lessons" from the brain to embedded systems
(and not to duplicate the way the brain works)

Many mechanisms

- adaptation mechanisms
- energy efficiency
- learning in non-stationary environments
- cognitive fault diagnosis can be immediately cast within a cognitive processing of human brain

Emotional Cognitive Structure

- Brain phenomena can be modeled as a **hierarchy of subsystems** differentiating in time activation and accuracy levels that, possibly, rely on a **memory** containing knowledge suitably stored for decision making

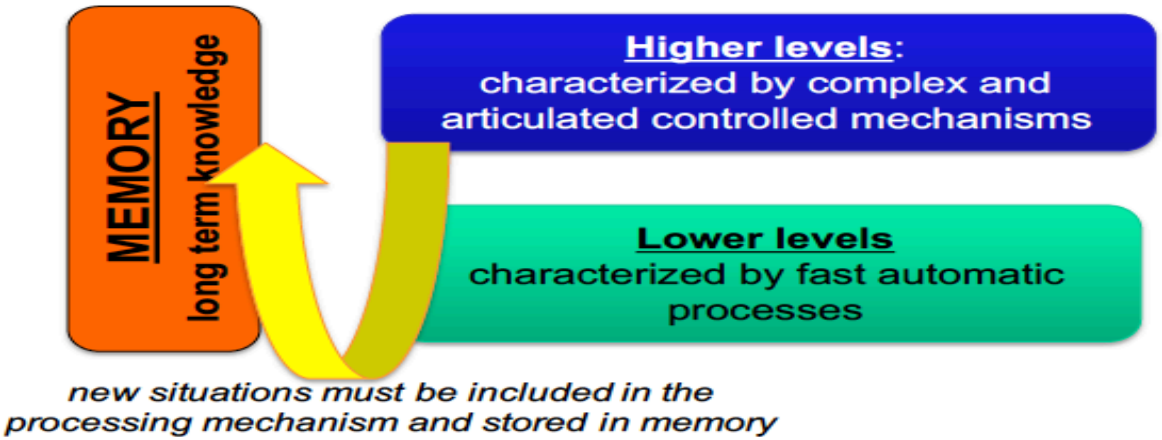

Emotional cognitive structure: lower and upper levels

- **Lower levels (fast automatic processes)**
  - ✓ process stimuli to quickly take a decision and/or provide a prompt reaction
  - ✓ activate higher levels of knowledge processing
- **Higher level (complex controlled mechanisms)**
  - ✓ assess ongoing performance and abort/modify/complete actions and decisions made by lower levels
  - ✓ introduce feedback mechanisms by providing information to improve the learning

The joint activity of automatic and controlled processes allow us for modeling the emotional responses in humans

Automated and Controlled Processes

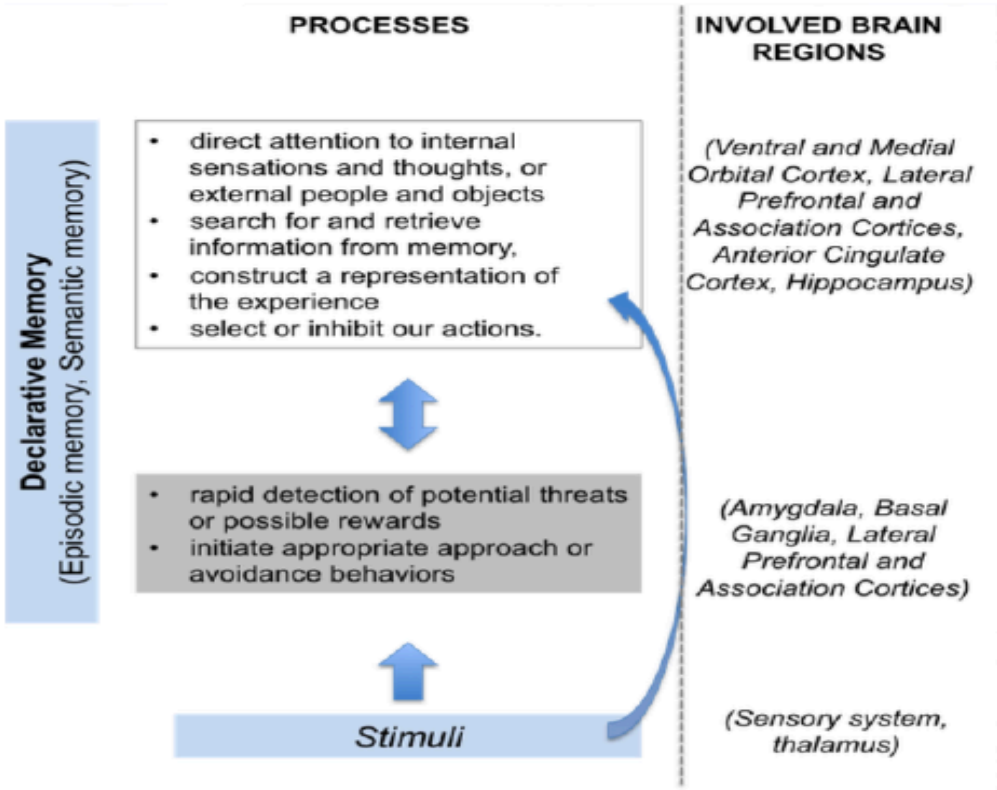

Automated Processes

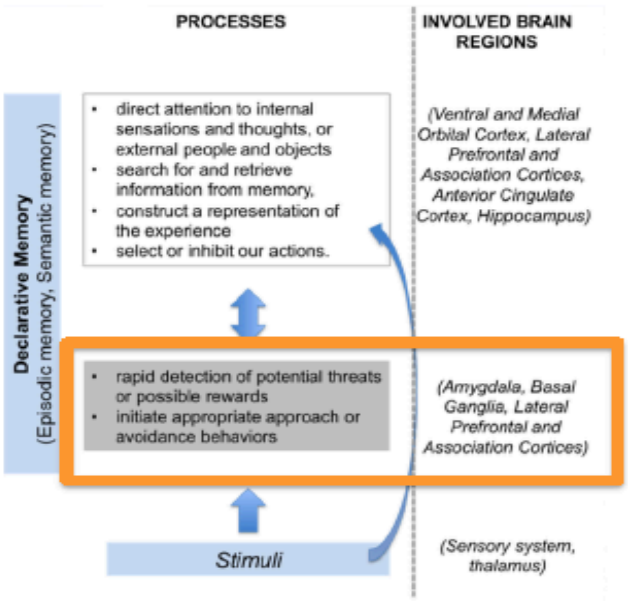

- **Lowest levels of the hierarchical cognitive system**
- Characterized by a **detection-reaction mechanism** designed to
  - ✓ quickly identify potential dangers
  - ✓ plan and schedule actions to get a reward
  - ✓ recall of previously acquired information related the situation/event/emotion
  - ✓ activation of a proper action/reaction, e.g., by increasing the heartbeat or the respiration rate and releasing stress-hormones following a perceived threat

Reduced latency

- **Reduced latency** is more important than keeping under control the false positive rate following a **conservative primordial principle**:
  - *it is much better to react with an unnecessary action following a perceived -possibly new- threat than being insensitive to it*
- Meant to be **quick and effortless in generating an emotional response** to external stimuli such as presentation of faces, objects or events
- **Emotional response** (possibly together with additional environmental information) is then processed to **become part of the knowledge** (i.e., stored in the memory) to be recalled whenever necessary

Controlled Processes

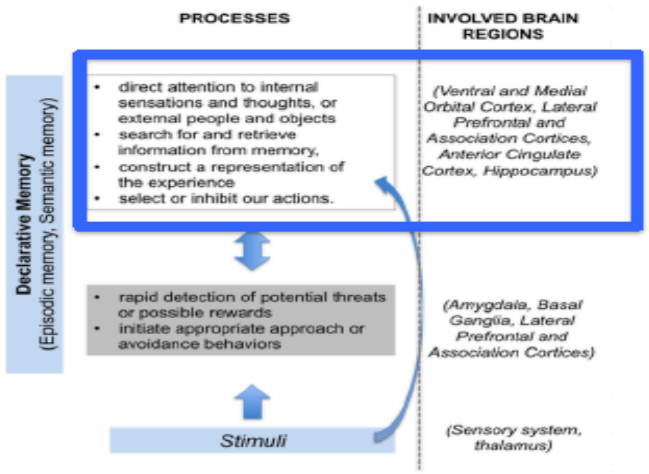

- **Consciously, we also direct attention to** our personal sensations, construct our emotional background, and select or inhibit actions depending on a lifetime experience
- After a preliminary automatic response, **cognitive processes are consciously activated** by higher levels of the cognitive hierarchical system

By deliberately monitoring, activating and processing emotions <u>we can re-interpret and alter their meaning</u> (learning mechanism), change the current personal experience and perception of the world as well as the way we interpret emotions and respond to stimuli

Controlled Processes

- The use of computationally demanding processes in the generation and regulation of emotions is named *controlled emotion processing*
- The **main tasks carried out by controlled processes** can be summarized as
  - ✓ Select or inhibit the actions activated by automatic processes
  - ✓ Construct a representation of the emotional experience over time and perfect it according to the received external stimuli and the final situation outcome
  - ✓ Direct attention to internal sensations and thoughts
  - ✓ Search and retrieve information from the declarative memory

Processes and brain regions

Brain Regions	Function	Operations	Type of Process
Amygdala	Detecting, Learning about Stimuli	Detects potentially threatening stimuli and associates them with appropriate actions	Automatic
Basal Ganglia	Registering Rewards, Acquiring Habits, Selective gating of behavior	Mediates selection and initiation of actions, Automatizes sequences of behavior and reinforced thoughts	Automatic or controlled
Lateral prefrontal/ association cortices	Retrieving and Storing Semantic Emotion Knowledge	Identifies stimuli, differentiates feeling states; attributes emotional qualities to stimuli; repository of regulatory strategies, lay emotion knowledge	Retrieval can be automatic or controlled
Anterior Cingulate Cortex	Conflict Monitoring	Monitors on-going behavior and determines whether a change is necessary or not	Conflicts detected automatically, but making changes takes control
Ventral/ Medial Orbital Frontal Cortex	Context- dependent action selection	Inhibits on-going emotional responses based on analyses of context	Controlled
Hippocampus	Long-term strategies	Understanding spatial relations within the environment	Controlled

Amygdala

- **Complex and differentiated roles**
  - ✓ the detection of a threat by inspecting stimuli patterns
  - ✓ activating a reaction
  - ✓ modulating the long-term memory consolidation of stimuli-stimuli and stimuli-response association in the emotional declarative memory

How to cast it within the cognitive embedded systems framework?

The amygdala may be considered as a module processing stimuli patterns to detect threats. Similar to Change Detection Tests (CDTs)

Basal Ganglia

- **Role**:
  - voluntary motor control (e.g., eyes control)
  - emotional processing leading to learn routine behaviors to attain rewards

Modeling time dependent events by means of dynamic systems or machine learning techniques (e.g., Markov processes) for forecasting purposes and define sequences of actions necessary to achieve a long-term goal by means of planning

Lateral prefrontal and association cortices (LPAC)

- Provides mechanisms enabling **the storage and retrieval of semantic emotion knowledge** as well as using the memory content to asses the relevance of stimuli and events
- **Access to this emotional database** is **automatic** during the generation of an emotional state, or when we **consciously** represent or label emotional states to draw inferences about those emotions we are experiencing

Knowledge management in recurrent adaptive classifiers. This knowledge is organized into concepts, each of which represents a memory of the state. Similar concepts can be fused together to improve/integrate the knowledge over time.

Anterior cingulate cortex (ACC)

- **Role**:
  - ✓ Assess the "congruence" of emotions and feelings that have been generated in response to external stimuli
  - ✓ Forecasting whether external stimuli would induce threats or pain in the future or not
- This capability, which is conscious, is **crucial in the activity planning** to achieve long-term goals.

The ACC activity is conceptually very close to the validation procedure of **hierarchical change detection tests**: in response to an event, signals are jointly evaluated to determine whether what perceived is a false alarm or represents a true change

Orbital and ventral-medial prefrontal cortices (OFC and VM-PFC)

- The role is closely related to the active approach modeling the human behavior:
  - **Emotions, stimuli and memory patterns** automatically generated by lower cognitive levels are **integrated and linked to the long-term memory** to define a "cognitive" high-level response
  - The **"cognitive" ability** resides in the capacity to **connect the memory systems with emotional systems** (amygdala) to evaluate the taken actions and recall associated somatic states

This approach is very close to cognitive analysis in distributed fault diagnosis systems where low level information is integrated and analysed by taking into account also the network topology (to distinguish among faults, change in the env. or model bias)

Hippocampus

- An **old structure of cerebral cortex** that takes part in many declarative memory functions which refer to the memory of facts and events:
  - *Encoding and recalling information from the memory are the two main tasks of this fundamental subsystem.*
- **Interacts with the amygdala** in the formation of short-term memory, a preliminary step for the storage of long-term information

In intelligent embedded systems, the role of the hippocampus can be associated with the ability to recall previously acquired concepts whenever necessary (recurring concepts)

Decision-making involves the orchestration of multiple neural structures and cognitive subsystems,

e.g., VM-PFC, amygdala, LPAC, and hippocampus

Decision Making

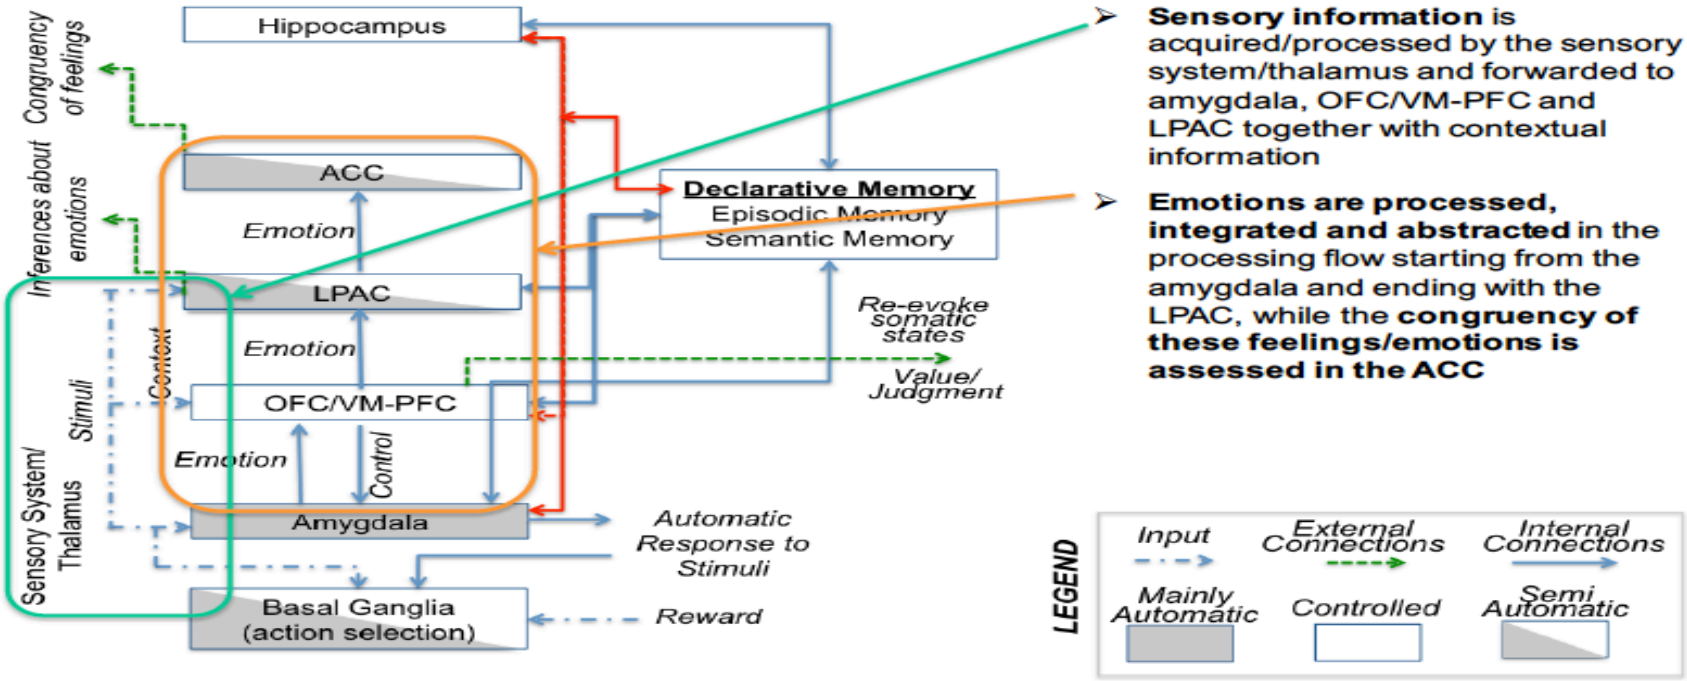

Decision Making (2)

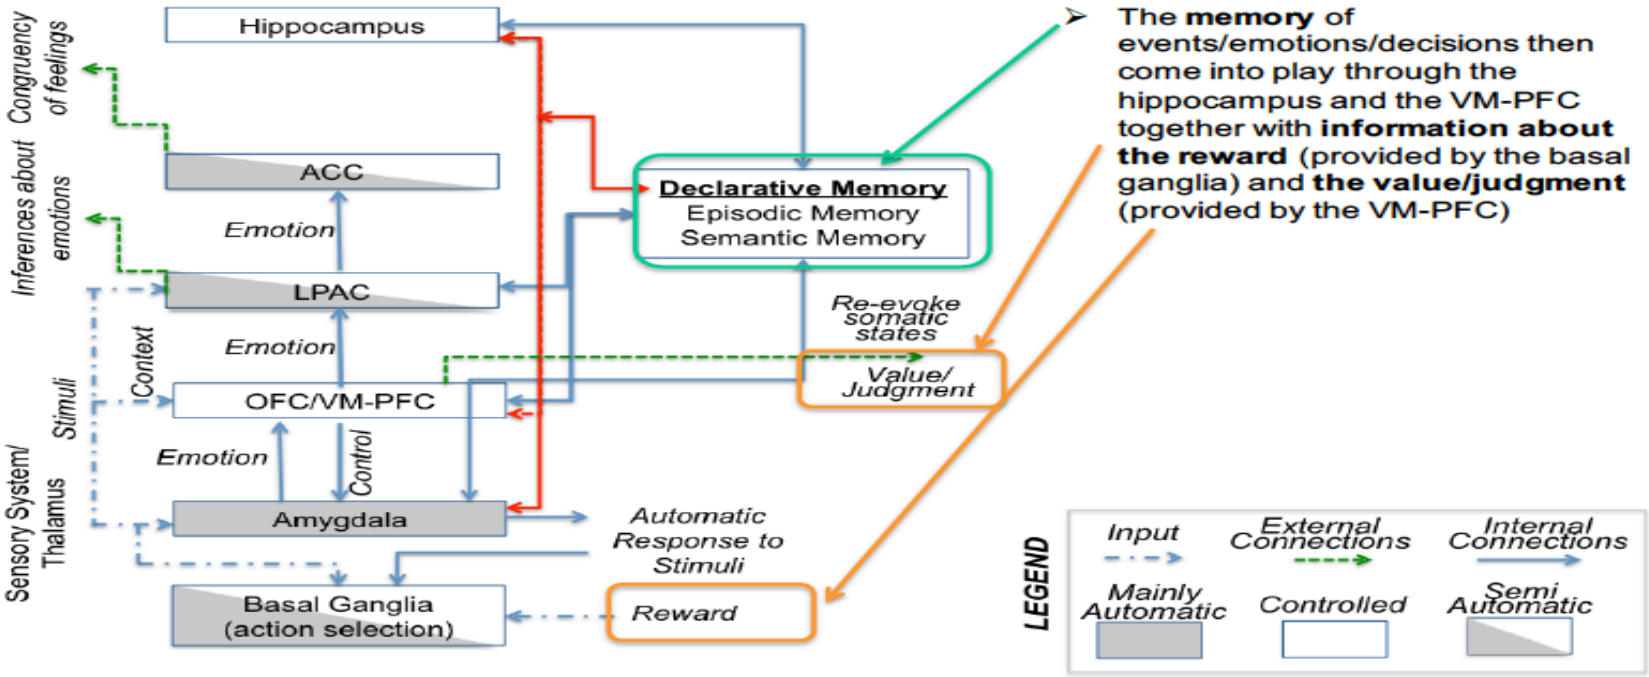

Emotion and decision-making and adaptive classifiers

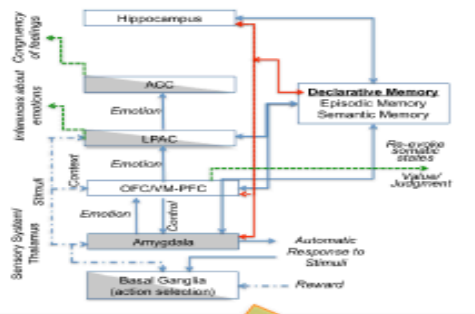

Interestingly, this complex process is very close to the processing in **adaptive classifiers** where an initial decision is initially taken by considering external stimuli and previously acquired information.

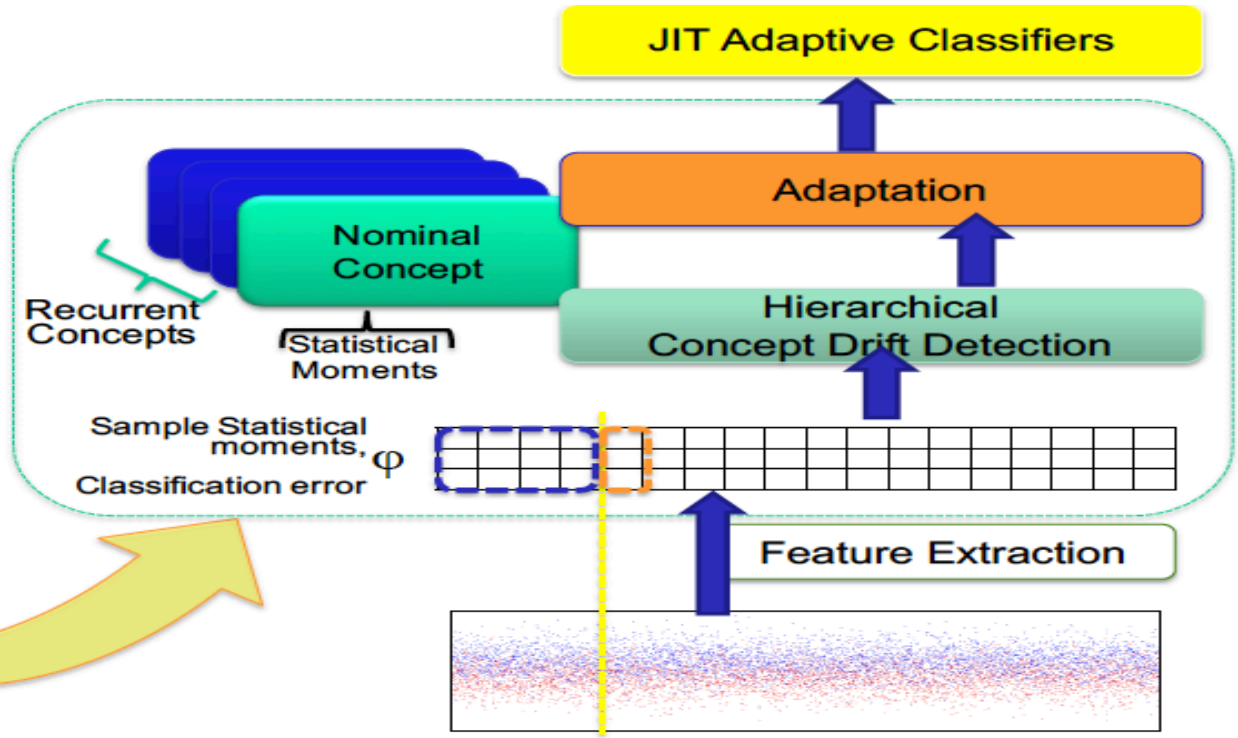

Adaptation in CPS

Need for adaptation

- Many **cyber-physical systems require intelligent mechanisms** to deal with those situations where either evolution or time variance requests a reaction to grant a performance level
- Adaptation is the **basic form** of intelligence to be considered every time the **cyber-physical system has to react quickly**
  - minimize both time-to-reaction and energy consumption
  - adaptive mechanisms are the result of sophisticated techniques

Adaptation: form of intelligence associated with the execution of automatic cognitive processes

Application mechanisms in CPS

	Automatic Processes	Controlled Processes
Single Unit Level	• reducing the energy consumption of the device • maximizing the efficiency of the energy harvesting process	• the implementation of sophisticated mechanisms for adaptive sensing • reprogram the embedded device (when and how)
Group Level	• keeping the units clocks synchronized	• taking advantage of group information to improve accuracy as in distributed clock synchronization

Energy Consumption

- Even though we are good in energy harvesting, **energy consumption is an issue** and should be minimized
- The **Murphy's law** appears to be always there: once you need energy you cannot harvest it…
- **Some strategies to manage the available energy**
  - Throw away what is not needed (keep it simple)
  - Consider incremental applications
  - Duty cycling: The more you sleep the less energy you consume
  - Forecast and react methods
  - Sample what needed
  - Act directly on the Hw

Adaptive Sensing

- **A possible data acquisition procedure:**
  1. The interrupt wakes up the microprocessor, possibly from a deep sleep modality;
  2. The interrupt routine sends the warm up directives to the sensors (if needed) and waits (or the task is inserted in the task wait list) for available data;
  3. If sensors are ready, the sampling procedure is activated, acquiring data and storing them in the memory;
  4. Once the acquisition task is completed the routine terminates.

Measurement chain : Sensor level

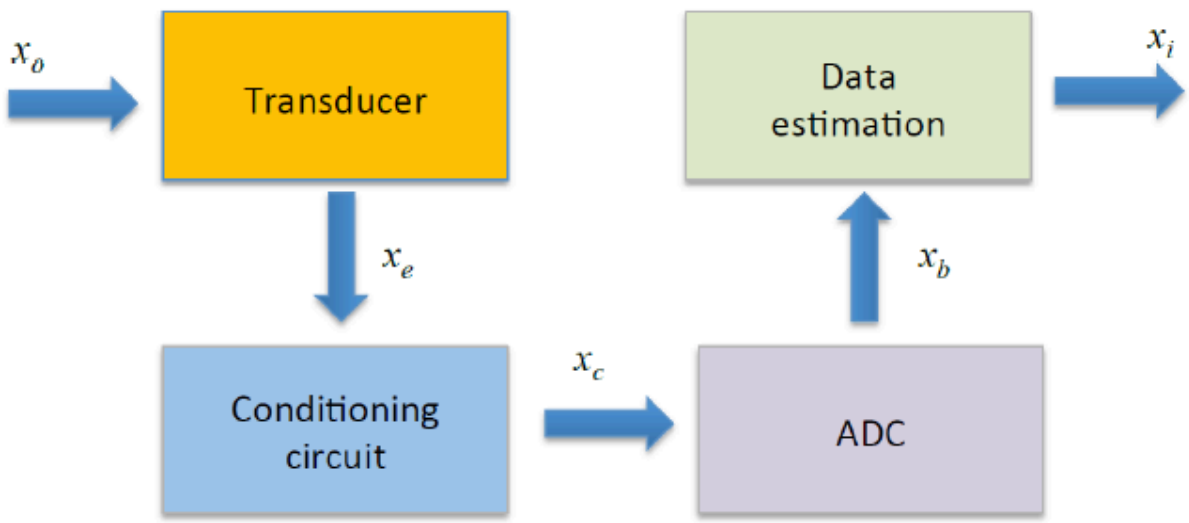

Adaptive Sensing

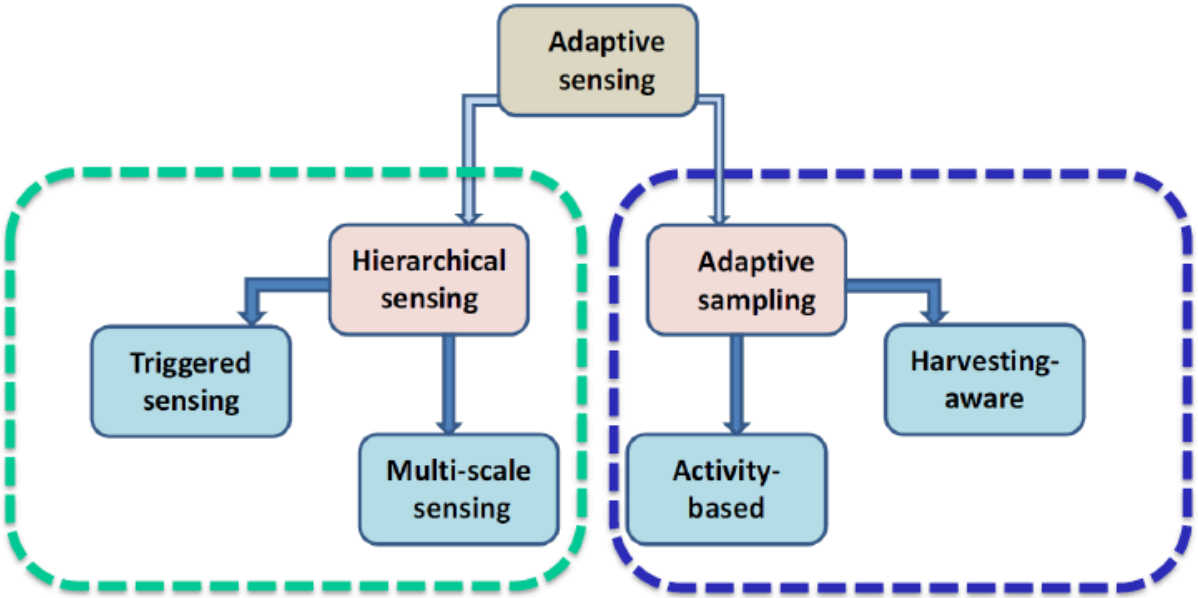

Hierarchical Sensing

- Data acquired with low resolution/low power consumption sensors are processed by simple algorithms. When a potential alarm is detected, high resolution sensors generally characterized by a higher power consumption are activated.

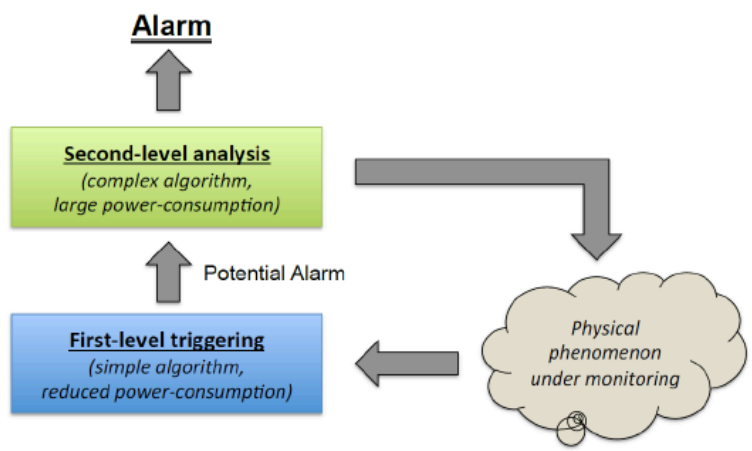

Triggered Sensing

- In triggered sensing **features are used to decide whether to activate an alarm or not**.
- An example is based on **different resolution** (different energy consumption) **case**
- Another example is based on a **CMOS camera, reconfigurable in terms of spatial resolution**:
  - Low resolution scenes of the environment are quickly processed for target detection.
  - If targets are detected, some digital cameras undergo an adaptation phase that, after reconfiguration, provides higher quality images for target detection validation.

Multi-scale Sensing

- In multi-scale sensing, we identify **areas** within the monitoring field that **require a more accurate inspection**
- *The method envisages a lower resolution when the receptive field is less relevant, and a higher resolution when high precision acquisitions are requested.*
- **Example**: An emergency management scenario.
  - The field to be monitored is instrumented with static, low resolution temperature sensors.
  - When an event is detected, a mobile sensor unit (mule) is sent to the area to collect additional high precision measurements.

Adaptive Sampling

- **Adaptive sampling methods modify the sampling rate** based on the information content carried by sensed data, the available residual energy present in the batteries and the incoming or estimated harvested power.

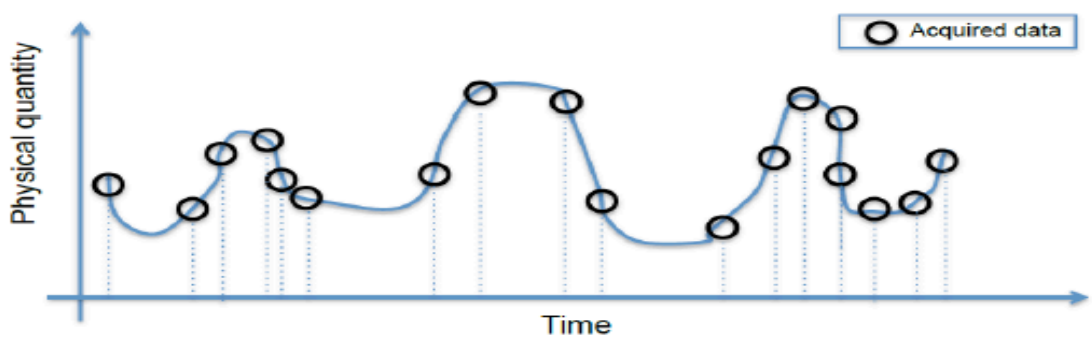

Activity based Sampling

- In model-based sampling temporal and spatial locality redundancy are used to describe incoming data. If a discrepancy exists a new model must be generated

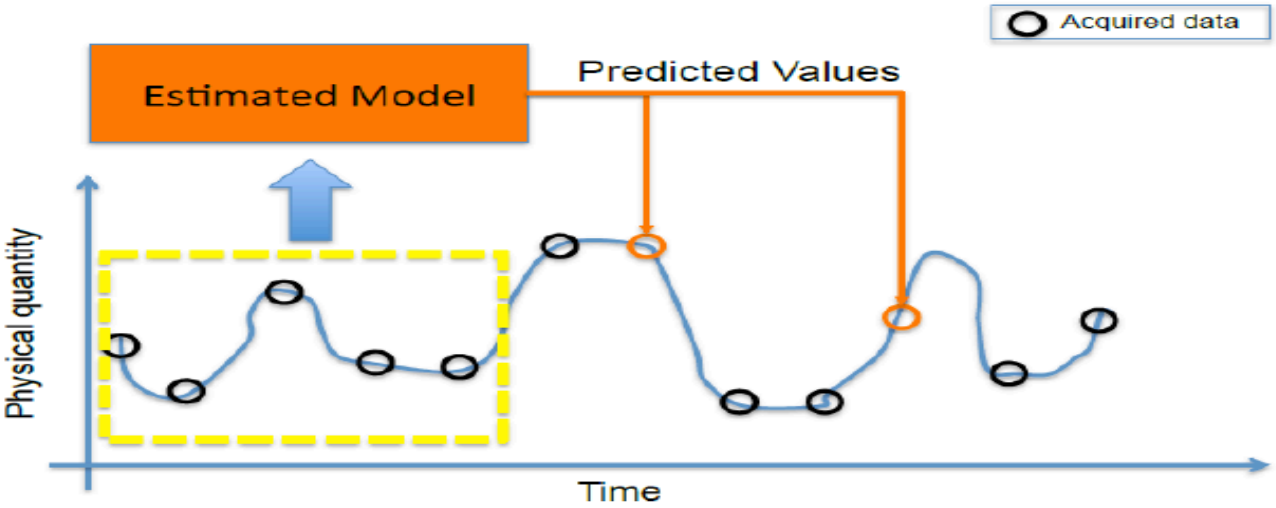

Adaptive Sampling: Harvesting Policy

- In harvesting aware policies, **the sampling frequency is increased or decreased based on information related to the residual energy**.
- **In extreme cases sensors can be switched off**, with strategies depending on
  - The information novelty content expected to be provided by the datastream in the near future (either predicted or estimated on past data)
  - The power consumption of the sensor, also balanced with its accuracy
  - The impact of the sensor information on the application (sensitivity analysis)

Compressive Sensing (CS)

“Measure What Should Be Measured”

Sensing by Sampling

$x \rightarrow$ Sample $\xrightarrow{N}$

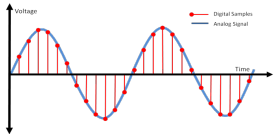

Sensing by Sampling (2)

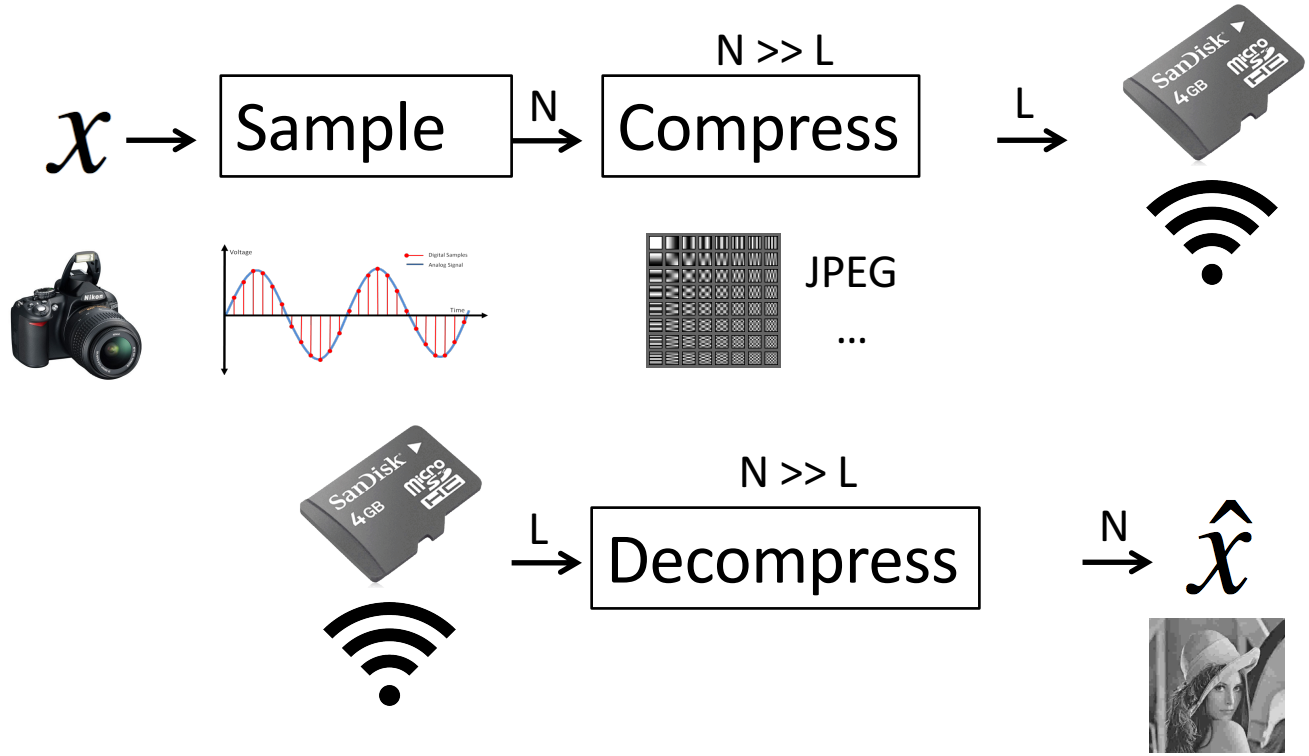

Compression: Toy Example

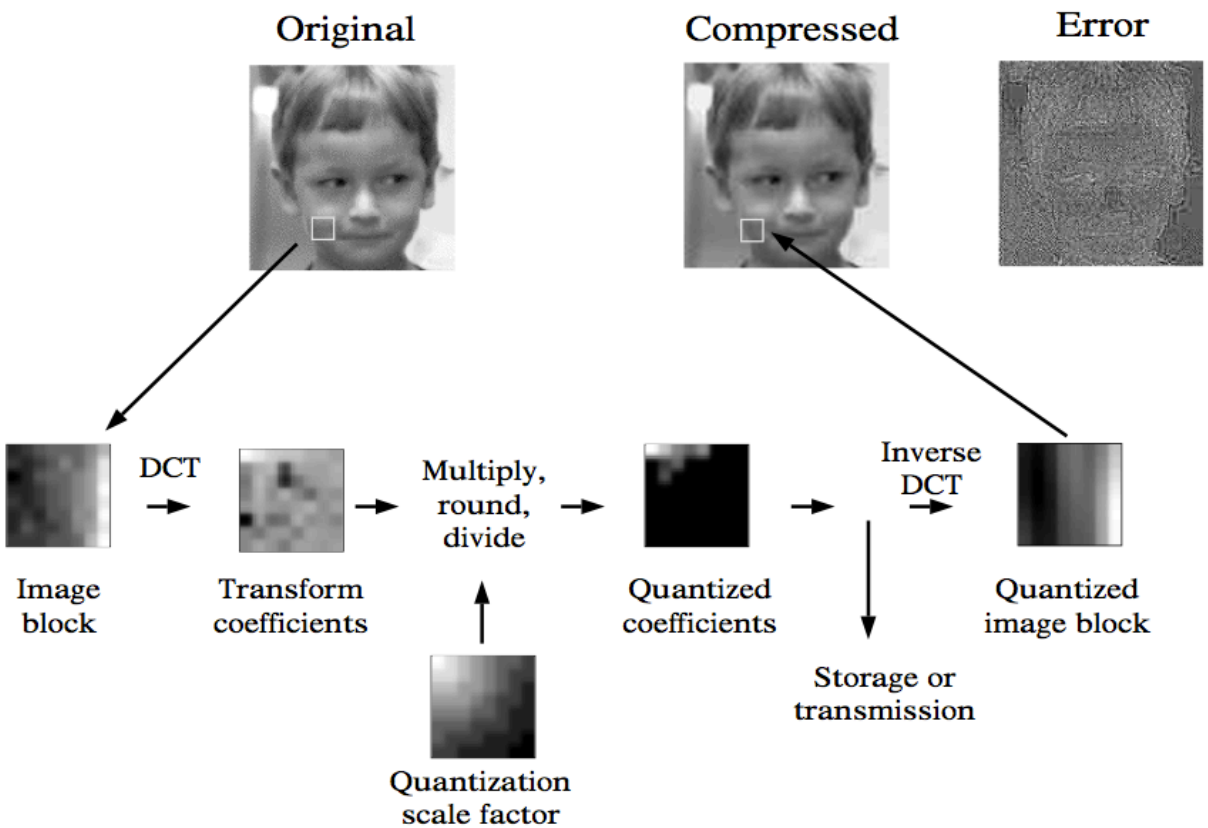

Discrete Cosine Transformation

PICTURE MATRIX

$$\begin{bmatrix} 40 & 24 & 15 & 19 & 28 & 24 & 19 & 15 \\ 38 & 34 & 35 & 35 & 31 & 28 & 27 & 29 \\ 40 & 47 & 49 & 40 & 33 & 29 & 32 & 43 \\ 42 & 49 & 50 & 39 & 34 & 30 & 32 & 46 \\ 40 & 47 & 46 & 35 & 31 & 32 & 35 & 43 \\ 38 & 43 & 42 & 31 & 27 & 27 & 28 & 33 \\ 39 & 33 & 25 & 17 & 14 & 15 & 19 & 26 \\ 29 & 16 & 6 & 1 & -4 & 0 & 7 & 18 \end{bmatrix}$$

→ Transformation →

DCT COEFFICENTS

$$\begin{bmatrix} 239 & 32 & 27 & -12 & 3 & -5 & 3 & 1 \\ 34 & -3 & -19 & 6 & 3 & 0 & -1 & 1 \\ -70 & 2 & 8 & 23 & 9 & 6 & -1 & -1 \\ 5 & 0 & -6 & 11 & -2 & 0 & -1 & 1 \\ -17 & -3 & 6 & 6 & 3 & -1 & 0 & 0 \\ 2 & 4 & 2 & 2 & 1 & -2 & 0 & 1 \\ -3 & 0 & 0 & -1 & -1 & -1 & 0 & 0 \\ 1 & -1 & 3 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Motivation

- Why go to so much effort to acquire all the data when most of the what we get will be thrown away?
- Cant we just directly measure the part that wont end up being thrown away?

Shannon’s sampling theorem

- A band-limited signal with maximum frequency **B** can be **accurately** reconstructed from its uniformly spaced digital samples if the rate of sampling exceeds **2B** (called **Nyquist rate**).
- Independently discovered by Shannon, Whitaker, Kotelnikov and Nyquist.

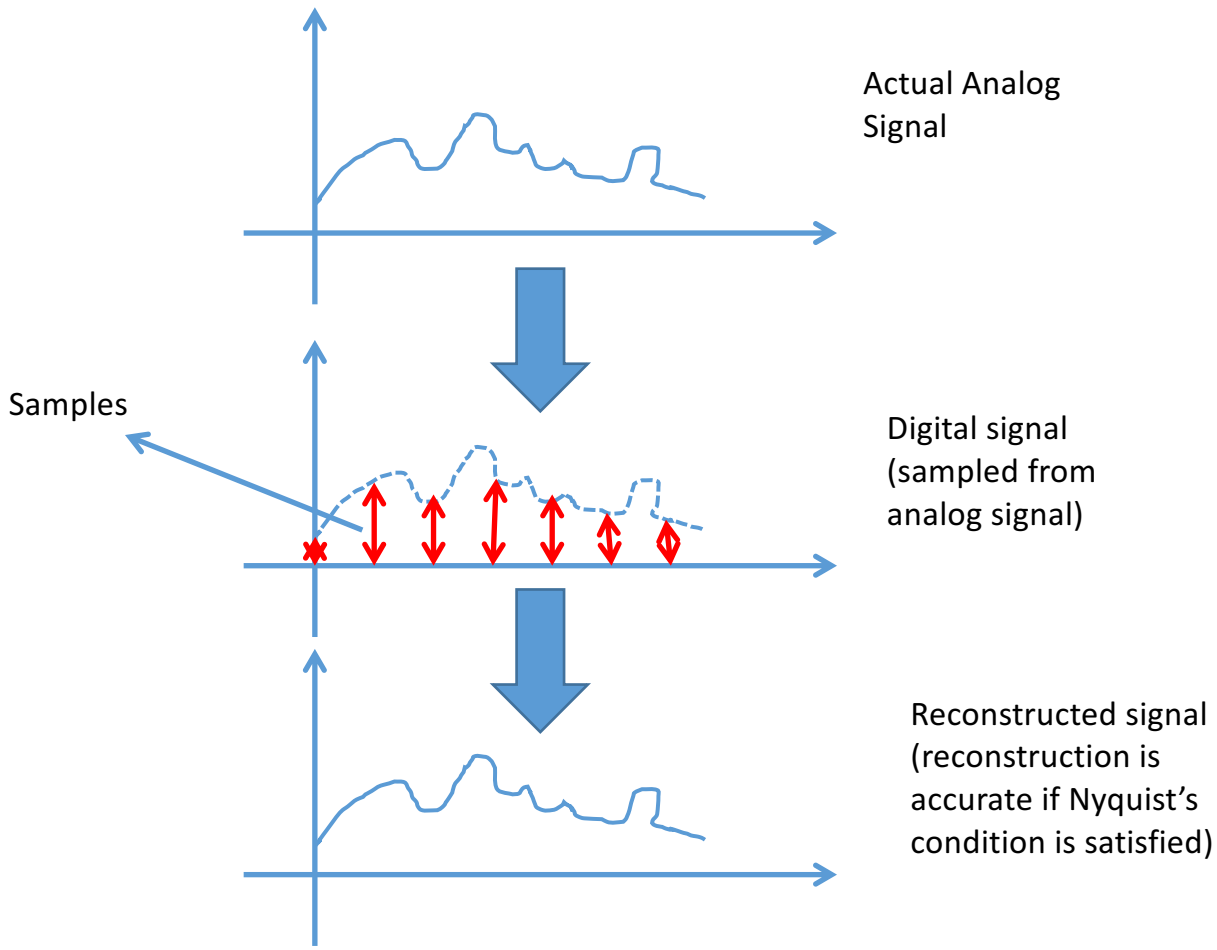
Actual Analog Signal
Samples
Digital signal (sampled from analog signal)
Reconstructed signal (reconstruction is accurate if Nyquist's condition is satisfied)

Band-limited signal

- A truly band-limited signal is an ideal concept and would take infinite time to transmit or infinite memory to store.
- Because a **band-limited signal can never be time-limited** (or space-limited).
- But many naturally occurring signals when measured by a sensing device are approximately band-limited (example: camera blurring function reduces high frequencies).

Aliasing

- Sampling an analog signal with maximum frequency B at a rate **less than or equal to 2B** causes an artifact called **aliasing**.

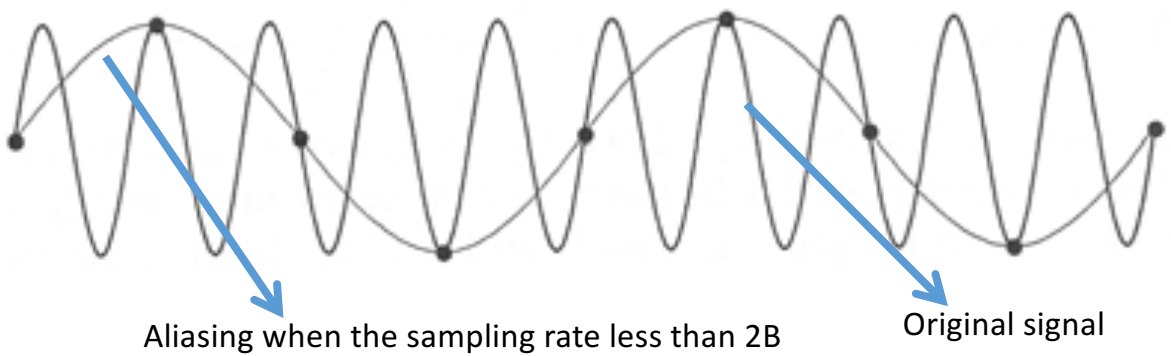

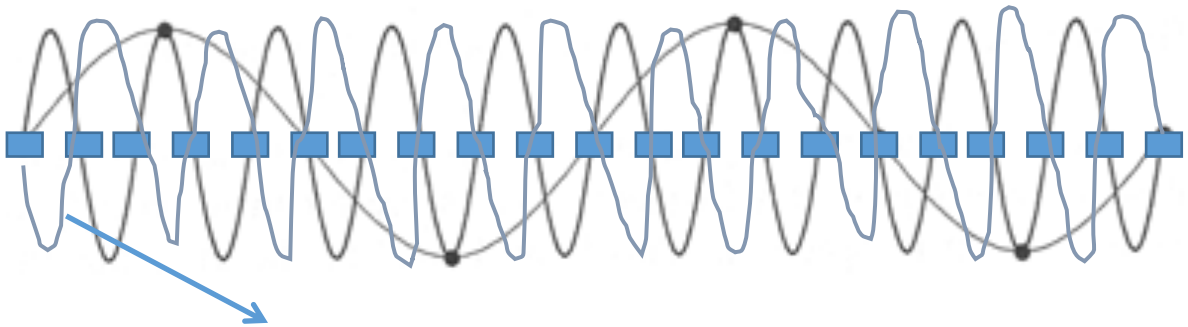

Aliasing when the sampling rate is equal to 2B

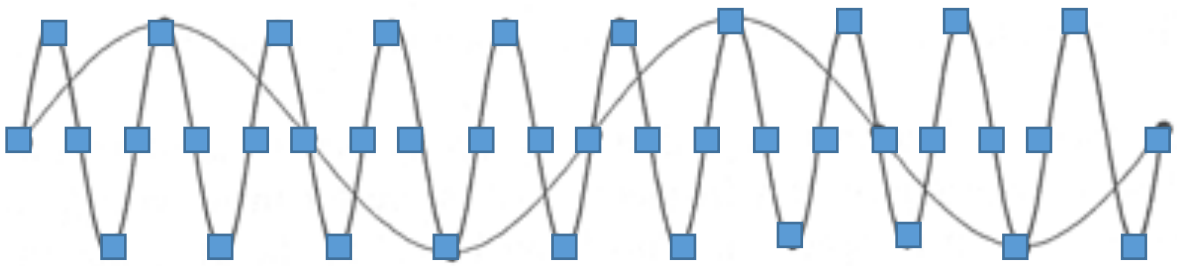

No aliasing when the sampling rate is more than 2B

Whittaker-Shannon Interpolation Formula

- Reconstruction of the analog signal from its digital samples can be done by **polynomial** interpolation (piecewise linear, quadratic, cubic etc.)
- However polynomial interpolation has associated errors which tend to zero only **in the limit of infinite samples**.
- But band-limited signals can be perfectly reconstructed from a **finite number of digital samples** (sampled above Nyquist rate).

Whittaker-Shannon Interpolation Formula

- The optimal reconstruction for band-limited signals from their digital samples proceeds using the **sinc interpolant**. It yields a time-domain (or space-domain signal of infinite extent with values that are very small outside a certain window).
- The associated formula is called Whittaker-Shannon interpolation formula.

Sampling period =1/Sampling rate

$$f(t) = \sum_{n=-\infty}^{+\infty} f(n) \frac{\sin\left(\frac{t-nT}{T}\right)}{\left(\frac{t-nT}{T}\right)} = \sum_{n=-\infty}^{+\infty} f(n)\,\mathrm{sinc}\left(\frac{t-nT}{T}\right) = \left(\sum_{n=-\infty}^{+\infty} f(n)\delta(t-nT)\right) * \mathrm{sinc}\left(\frac{t}{T}\right)$$

Magnitude of the n-th sample

The sinc function (it is the Fourier transform of the rect function)

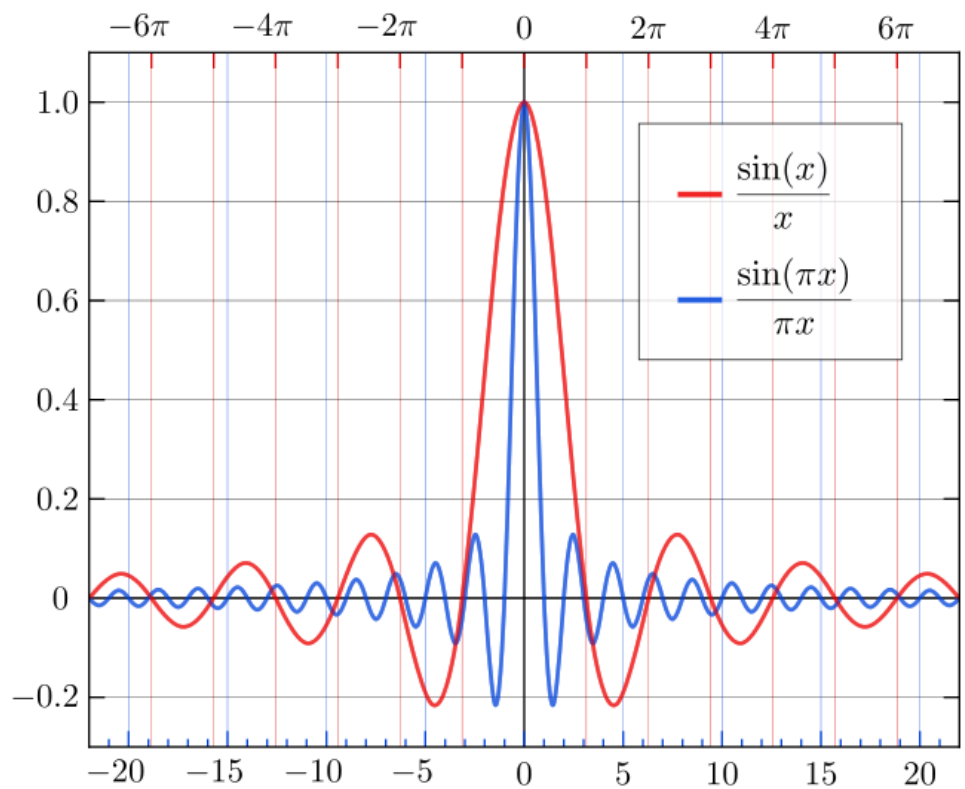

Some limitations of Shannon's theorem

- The samples need to be **uniformly spaced** (there are extensions for non-uniformly spaced samples, with the equivalent of Nyquist rate being an average sampling rate).
- The sampling rate needs to be **very high** if the original signal contains higher frequencies (to avoid aliasing).
- Does not account for several nice properties of naturally occurring signals (except for **band-limitedness**).

What is Compressive Sensing?

- Signal/image compression algorithms like MPEG, JPEG, JPEG-2000 etc. typically measure large amounts of data.
- The data is then converted to a transform domain where a majority of the transform coefficients turn out to have near-zero magnitude and are discarded.
- Example: a digital camera has a detector array where several thousand pixel values are first stored. This 2D array is fed to the JPEG algorithm which computes DCT coefficients of each block in the image, discarding smaller-valued DCT coefficients.

What is Compressive Sensing?

- This “measure and throw” paradigm is wasteful!
- Especially for time-consuming acquisitions like MRI, CT, ultrasound, range data etc.
- Compressive sensing is a new technology where the data are acquired in a compressed format!

What is Compressive Sensing?

- These compressed measurements are then fed to some non-linear optimization algorithm (also called inversion algorithm) to produce the complete signal.
- This part is implemented in software.
- Under suitable conditions that the original signal and the measurement system must fulfill, the reconstruction turns out to have very little or even zero error!

Why compressive sensing?

- Much (as in more than 1000 times) fewer measurements need to be made.
- May dramatically improve acquisition speed for MRI, range images, hyper-spectral data.
- Potential to dramatically improve video-camera frame rates without sacrificing spatial resolution.

CS: Theory

- Let a occurring naturally signal be denoted **x**.
- Let us denote the measurement of this signal (using some device) as **y**.
- Typically, the mathematical relationship between **x** and **y** can be expressed as a linear equation of the form $\mathbf{y} = \Phi\mathbf{x}$.
- Φ is called the measurement matrix (or the "sensing matrix" or the "forward model" of the measuring device).
- For standard digital camera, Φ is represented by a Gaussian blur.

CS: Theory

- Relies on two main principles/requirements:

1. Natural signals/images have a sparse representation in well-known orthonormal bases such as Fourier, DCT, Haar wavelet etc.
2. The measurement matrix Φ is incoherent (poorly correlated) with the signal basis matrix Ψ, i.e. the sensing waveforms have a very dense representation in the signal basis.

Signal Sparsity

- Many signals have sparse representations in standard orthonormal bases (denoted here as Ψ).
- Example:

$$f = \Psi\theta = \sum_{k=1}^{n} \Psi_k \theta_k,$$

$$f \in R^n, \theta \in R^n, \Psi \in R^{n\times n}, \Psi^T\Psi = I,$$

$$\|\theta\|_0 << n$$

Signal Sparsity

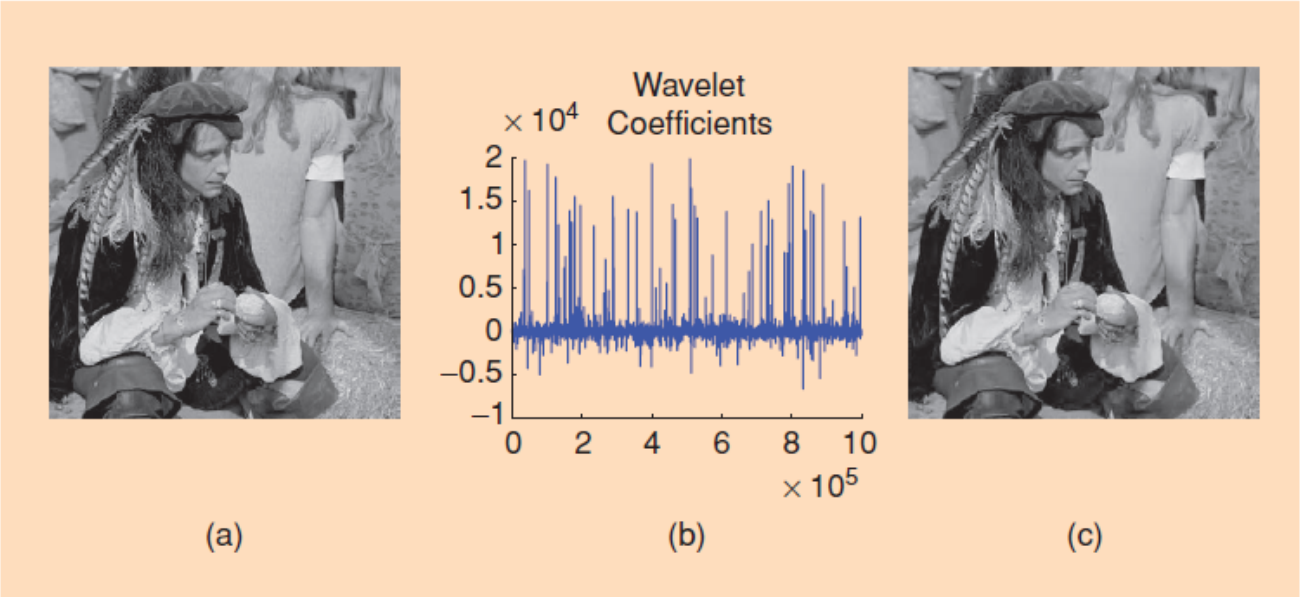

[FIG1] (a) Original megapixel image with pixel values in the range [0,255] and (b) its wavelet transform coefficients (arranged in random order for enhanced visibility). Relatively few wavelet coefficients capture most of the signal energy; many such images are highly compressible. (c) The reconstruction obtained by zeroing out all the coefficients in the wavelet expansion but the 25,000 largest (pixel values are thresholded to the range [0,255]). The difference with the original picture is hardly noticeable. As we describe in "Undersampling and Sparse Signal Recovery," this image can be perfectly recovered from just 96,000 incoherent measurements.

Incoherent Sampling

- Consider a sensing matrix Φ of size m by n, where m is much less than n (if $m >= n$, there is no compression in the measurements!).
- We will assume (for convenience) that each row of Φ is unit-normalized.
- CS theory states that Φ and Ψ should be incoherent with each other.

Incoherent Sampling

- The coherence between $\mathbf{\Phi}$ and $\mathbf{\Psi}$ is defined as:

$$\mu(\Psi,\Phi) = \sqrt{n}\,\max_{\substack{1\le j\le m,\\ 1\le i\le n}} \left| < \Phi_j^T, \Psi_i > \right|$$

- We want this quantity μ to be as small as possible.

'j'-th row of Φ, i-th column of Ψ

- Its value always lies in the range $(1, \sqrt{n})$.

Incoherent sampling: Example

- Let the signal basis Ψ be the Fourier basis. A sampling basis Φ that is incoherent with it is the standard spike (Dirac) basis:

  $$\Phi_k(t) = \delta(t-k)$$

  which corresponds to the simplest sampling basis in space or time.
- The associated coherence value is 1 (why?).
- This has some cool implications for signal reconstruction. We will soon see.

Randomness is cool! ☺

- Sensing matrices whose entries are i.i.d. random draws from Gaussian or Bernoulli (+1/-1) distributions are incoherent with **any** fixed basis Ψ with a very high probability.
- Implication: we want our sensing matrices to be highly correlated with entities like white noise! ☺

Signal Reconstruction

- Let the measured data be given as:

$$y = \Phi f = \Phi \Psi \theta,$$

$$y \in R^m, \Phi \in R^{m \times n}, f \in R^n, m << n$$

- The coefficients of the signal ***f*** in **Φ** – denoted as ***θ*** – and hence the signal ***f*** itself - can be recovered by solving the following constrained minimization problem:

$$\min \|\theta\|_0 \text{ such that } y = \Phi \Psi \theta$$

Signal Reconstruction

- This is a very difficult optimization problem to solve. Hence, a softer version (known as Basis Pursuit) is solved:

$$\min \|\theta\|_1 \text{ such that } y = \Phi\Psi\theta$$

- This is a linear programming problem and can be solved with any LP solver (in Matlab, for example) or with packages like L1-magic.

Signal Reconstruction

- The first use of L1-norm for signal reconstruction goes back to a PhD thesis in 1965 – by Logan.
- L1-norm was used in geophysics way back in 1973 (ref: Claerbout and Muir, “Robust modeling with erratic data”, Geophysics, 1973).
- L1-norm is the most popular sparsity-promoting prior in contemporary CS, but not the only one.
- Other examples: Greedy algorithms like orthogonal matching pursuit, Bayesian methods.

Ref: Candes, Romberg and Tao, "Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information", IEEE Transactions on Information Theory, Feb 2006.

Ref: Donoho, "Compressed Sensing", IEEE Transactions on Information Theory, April 2006.

Theorem 1

- Consider a signal vector ***f*** (having length *n*) with a sparse representation in basis Ψ, i.e. **f= $\Psi\theta$,** where $|\theta|_0 << n$. Suppose ***f*** is measured through the measurement matrix Φ yielding a vector ***y*** of only $m << n$ measurements.

  If $m \geq C\log(n/\delta)\|\theta\|_0 \mu^2(\Psi,\Phi)$ (*C* is some constant) then the solution to the following is **exact** with probability $1-\delta$:

$$\min\|\theta\|_1 \text{ such that } y = \Phi\Psi\theta$$

$$y = \Phi f = \Phi\Psi\theta,$$

$$y \in R^m, \Phi \in R^{m\times n}, f \in R^n, m << n$$

Comments on Theorem 1

- Look at the lower bound on the number of measurements *m*:

$$m \geq C \log(n/\delta) \|\theta\|_0 \mu^2(\Psi, \Phi)$$

- The **sparser** the signal **f** in the chosen basis Ψ, the <u>fewer</u> the number of samples required for exact reconstruction.
- The lower the coherence (**greater the incoherence**) between Ψ and Φ, the <u>fewer</u> the number of samples required for exact reconstruction.
- When δ is smaller, the probability of exact reconstruction is higher. The number of samples required is higher, but the increase in number of samples **is only an additive factor proportional to <u>log δ</u>.**

Comments on Theorem 1

- The algorithm to reconstruct $\boldsymbol{f}$ given $\boldsymbol{y}$ makes **<u>NO ASSUMPTIONS</u>** about the **sparsity** of $\boldsymbol{f}$ in Ψ, or the **location/magnitude** of the non-zero coordinates of $\boldsymbol{\theta}$.
- The algorithm does not have any notion/consideration for "uniformly spaced samples" – unlike Nyquist's theorem.
- Even though we are penalizing the L1 norm (and not L0 norm), we are still getting a reconstruction that is the sparsest (besides being accurate).
- Although L1-norm has been used in sparse recovery for a long time, the optimality proof was established only in 2005 by Donoho, and by Candes, Romberg and Tao.

Comments on Theorem 1: A New Sampling Theorem?

- Let the signal basis Ψ be the Fourier basis. A sampling basis Φ that is incoherent with it is the standard spike (Dirac) basis: $\Phi_k(t) = \delta(t-k)$ which corresponds to the simplest sampling basis in space or time.
- Let the signal ***f*** be sparse in this Fourier basis.
- Then given ANY (**not necessarily uniformly spaced**) *m* Fourier coefficients of ***f***, where

$$m \geq C \log(n/\delta) \|\theta\|_0 \mu^2(\Psi, \Phi), \text{i.e.} \; m \geq C \log(n/\delta) \|\theta\|_0$$

we can get an exact reconstruction of ***f*** with high probability. This acts like a new sampling theorem that does not require uniformly spaced samples or bandlimited signals.

Restricted Isometry Property

- We say the matrix **A=ΦΨ** of size m by n obeys the restricted isometry property of order *S* (*S* <= *m*) if ***A*** approximately preserves the squared magnitude of any *S*-sparse vector ***θ***, i.e.

$$\forall \theta \text{ for which } \|\theta\|_0 \le S(\text{note}: S \le m), \text{we have}$$

$$(1-\delta_S)\|\theta\|^2 \le \|A\theta\|^2 \le (1+\delta_s)\|\theta\|^2,$$

$$\text{where } 0 < \delta_S < 1.$$

- If ***A*** obeys RIP of order S, no *S*-sparse signal can lie in the null-space of ***A*** (if it did, then we would have $\|A\theta\|^2 = 0$ and that obviously does not preserve the squared magnitude of the vector θ).

Restricted Isometry Property (RIP)

- If ***A*** obeys RIP of order *S*, any subset of *S* <= m columns from the matrix ***A*** consists of vectors that are approximately pair-wise orthogonal.
- Note that the difference between two *S*-sparse vectors is 2*S*-sparse. Then if $\delta_{2S} << 1$, we have the following:

$$(1-\delta_{2S})\|\theta_1-\theta_2\|^2 \le \|A(\theta_1-\theta_2)\|^2 \le (1+\delta_{2S})\|\theta_1-\theta_2\|^2,$$

$$\forall\, \theta_1, \theta_2, \text{for which } \|\theta_1\|_0 \le S, \|\theta_2\|_0 \le S$$

- Thus if ***A*** obeys RIP of order 2*S*, it **approximately preserves the squared differences between any two *S*-sparse vectors.**
- The higher the order of the RIP property a matrix obeys, the "better" that matrix.

Theorem 2

- Suppose the matrix **A=ΦΨ** of size *m* by *n* (where sensing matrix **Φ** has size *m* by *n*, and basis matrix **Ψ** has size *n* by *n*) has RIP property of order 2*S* where $\delta_{2S} < 0.41$. Let the solution of the following be denoted as **θ***, (for signal **f = Ψθ**, measurement vector **y=ΦΨθ**):

Then we have:

$$\min \|\theta\|_1 \text{ such that } y = \Phi\Psi\theta$$

$$\|\theta^* - \theta\|_2 \le \frac{C_0}{\sqrt{S}} \|\theta - \theta_S\|_1,$$

$$\|\theta^* - \theta\|_1 \le C_0 \|\theta - \theta_S\|_1$$

θ_S is created by retaining the *S* largest magnitude elements of θ, and setting the rest to 0.

Comments on Theorem 2

- This theorem says that the reconstruction for *S*-sparse signals is **always** exact.
- Whereas Theorem 1 spoke of exact recovery **with high probability**. Hence Theorem 2 is more powerful than Theorem 1.
- For signals that are not *S*-sparse, the reconstruction error is almost as good as what could be provided to us by an **oracle** which knew the *S* largest coefficients of the signal $\boldsymbol{f}$.

Comments on Theorem 2

- We know that most signals are never exactly sparse in standard ortho-normal bases. But they are compressible, i.e. there is a sharp decay in the coefficient magnitude with many coefficients being nearly zero!
- Theorem 2 handles such compressible signals robustly. Theorem 1 did not apply to compressible signals! Here again, Theorem 2 is more powerful than Theorem 1.

Compressive Sensing under Noise

- So far we assumed that our measurements were exact, i.e.
  .
  $$y = \Phi f = \Phi \Psi \theta$$
- But practical measurements are always noisy so that
  .
  $$y = \Phi f + \eta = \Phi \Psi \theta + \eta$$
- Under the same assumption as before, we can estimate $\boldsymbol{\theta}$ by solving the following problem:

$$\min \|\theta\|_1 \text{ such that } \|y - \Phi \Psi \theta\|_2^2 \le \varepsilon$$

Convex problem, called as second-order cone program. Can be solved efficiently with standard packages including MATLAB and L1-MAGIC

Theorem 3

- Suppose the matrix **A=ΦΨ** of size *m* by *n* (where sensing matrix **Φ** has size *m* by *n*, and basis matrix **Ψ** has size *n* by *n*) has RIP property of order 2*S* where $\delta_{2S} < 0.41$. Let the solution of the following be denoted as **θ***, (for signal **f = Ψθ**, measurement vector **y=ΦΨθ**):

Then we have:

$$\min \|\theta\|_1 \text{ such that } \|y - \Phi\Psi\theta\|_2^2 \le \varepsilon$$

θ_S is created by retaining the S largest magnitude elements of θ, and setting the rest to 0.

$$\|\theta^* - \theta\|_2 \le \frac{C_0}{\sqrt{S}} \|\theta - \theta_S\|_1 + C_1 \varepsilon$$

Comments on Theorem 3

- Theorem 3 is a direct extension of Theorem 2 for the case of noisy measurements.
- It states that the solution of the convex program (see previous slides) gives a reconstruction error which is the sum of two terms: (1) the error of an oracle solution where the oracle told us the S largest coefficients of the signal $\boldsymbol{f}$, and (2) a term proportional to the noise variance.
- The constants C_0 and C_1 are very small (less than or equal to 5.5 and 6 respectively for $\delta_{2S} = 0.25$).

Randomness is super-cool!☺

- Consider any fixed orthonormal basis $\mathbf{\Psi}$.
- For a sensing matrix $\mathbf{\Phi}$ constructed in any of the following ways, the matrix **A=** $\mathbf{\Phi\Psi}$, will obey the RIP of order *S* with overwhelming probability provided that the number of samples $m \geq CS\log(n/S)$:

1. $\mathbf{\Phi}$ contains entries from zero-mean Gaussian with variance *1/m*.
2. $\mathbf{\Phi}$ contains entries with values $\pm 1/\sqrt{m}$ with probability 0.5 (symmetric Bernoulli) .
3. Columns of $\mathbf{\Phi}$ are sampled uniformly randomly from a unit sphere in *m*-dimensional space.

Randomness is super-super-cool!☺

- These properties of random matrices hold true for <u>any</u> orthonormal basis Ψ.
- As such, we do not need to tune Φ for a given Ψ.
- But there are methods that can give further improved results if you do such a tuning. We may study those techniques later on during the course.

L_1 norm and L_0 norm

- There is a special relationship between the following two problems (which we will study in compressive sensing later on):

$$y = A\theta$$
$$y \in R^n, A \in R^{m \times n}$$
$$\theta \in R^n (m << n)$$
$$\theta^* = \min \|\theta\|_0 \text{ subject to } y = A\theta$$

$$y = A\theta$$
$$y \in R^n, A \in R^{m \times n}$$
$$\theta \in R^n (m << n)$$
$$\theta^* = \min \|\theta\|_1 \text{ subject to } y = A\theta$$

The L_1 norm is a "softer" version of the L_0 norm. Other Lp-norms where $0 < p < 1$ are possible and impose a stronger form of sparsity, but they lead to non-convex problems. Hence L_1 is preferred.

$$\|x\|_p = \left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p}$$

L_1 is good, L_2 is bad. Why?

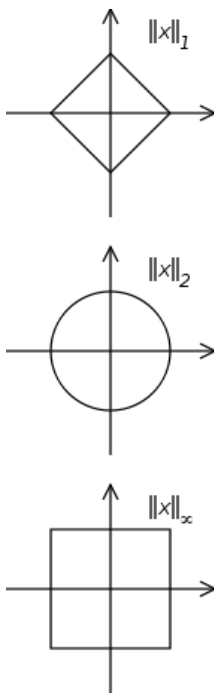

<u>Unit Lp-circle</u> (defined below) for different values of p. **Left side**, top to bottom, p = 1 (square with diagonals coinciding with Cartesian axes), p = 2 (circle in the conventional sense), p = infinity (square with sides parallel to the axes); **right side**, p = 2/3 (astroid).

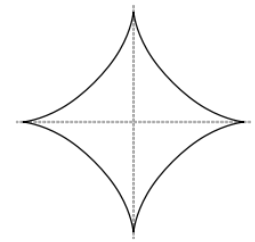

$$(|x|^p + |y|^p)^{\frac{1}{p}} = 1$$

In 2-D

$$\left(\sum_{i=1}^{d} |x_i|^p\right)^{\frac{1}{p}} = 1$$

In d-dimensions

L_1 is good, L_2 is bad. Why?

- Look at the basic problem we are trying to solve:

$$\min \|\alpha\|_p \text{ such that } y = \Phi\alpha$$

- In 2D (i.e. $\boldsymbol{\alpha}$, $\boldsymbol{y}$ being 2D vectors and $\mathbf{y} = \boldsymbol{\Phi\alpha}$ being a line in 2D space), the problem can be represented as the following diagram:

Ref: Romberg, "Imaging via compressive sampling", IEEE Signal Processing Magazine, March 2008.

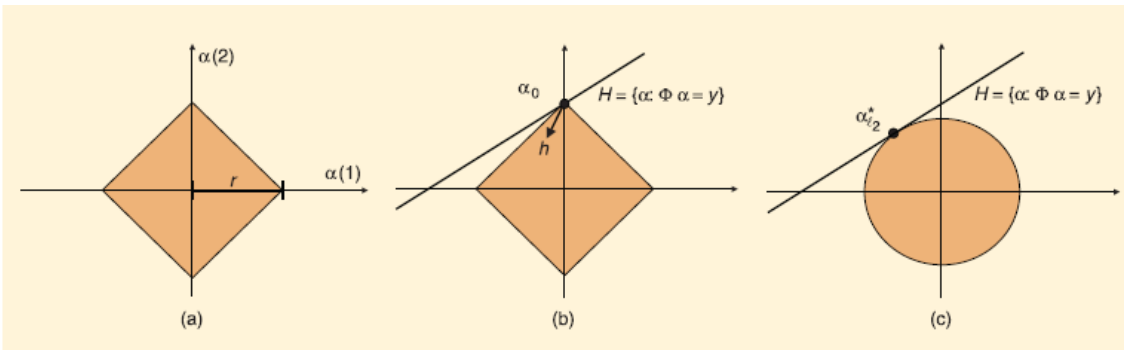

[FIG3] Geometry of ℓ_1 recovery. (a) ℓ_1 ball of radius r; the orange region contains all $\alpha \in \mathbb{R}^2$ such that $|\alpha(1)| + |\alpha(2)| \le r$. (b) Solving the $\min -\ell_1$ problem (8) allows us to recover a sparse α_0 from $y = \Phi\alpha_0$, as the anisotropy of the ℓ_1 ball favors sparse vectors. Note that the descent vectors h pointing into the ℓ_1 ball from α_0 will be concentrated on the support of α_0. (c) Minimizing the ℓ_2 norm does not recover α_0. Since the ℓ_2 ball is isotropic, the $\min -\ell_2$ solution $\alpha^*_{\ell_2}$ will in general not be sparse at all.

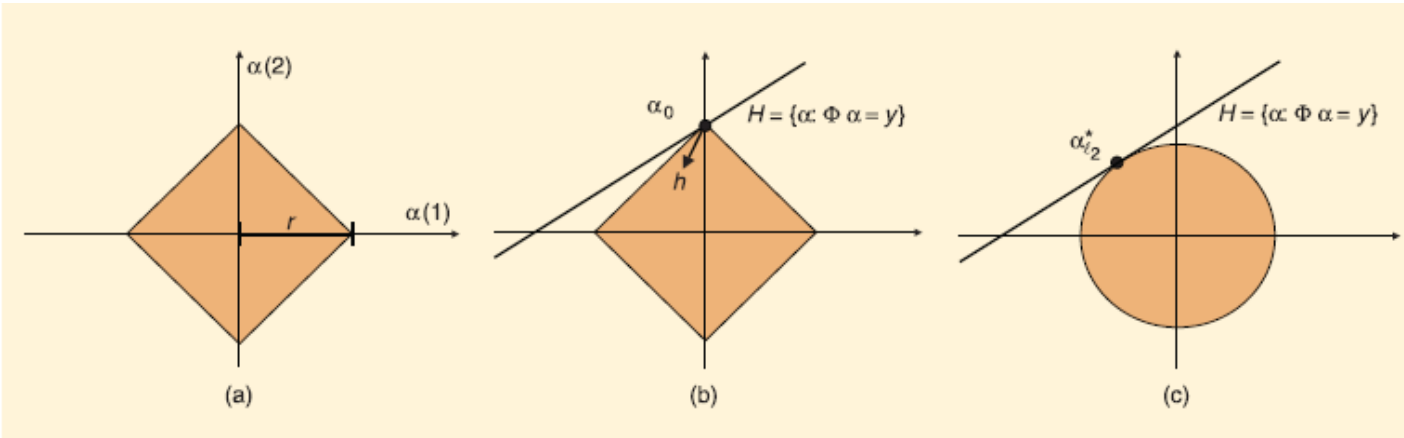

[FIG3] Geometry of ℓ_1 recovery. (a) ℓ_1 ball of radius r; the orange region contains all $\alpha \in \mathbb{R}^2$ such that $|\alpha(1)| + |\alpha(2)| \leq r$. (b) Solving the $\min -\ell_1$ problem (8) allows us to recover a sparse α_0 from $y = \Phi\alpha_0$, as the anisotropy of the ℓ_1 ball favors sparse vectors. Note that the descent vectors h pointing into the ℓ_1 ball from α_0 will be concentrated on the support of α_0. (c) Minimizing the ℓ_2 norm does not recover α_0. Since the ℓ_2 ball is isotropic, the $\min -\ell_2$ solution $\alpha^*_{\ell_2}$ will in general not be sparse at all.

$$\min \|\alpha\|_p \text{ such that } y = \Phi\alpha$$

• Note that the line **y = Φα** represents a HARD constraint, i.e. any vector α must satisfy this constraint.

•We want to find the vector α with minimum Lp-norm which satisfies this constraint.

•So we grow circles of increasing radius (denoted as *r* and defined in the sense of the Lp-norm) starting from the origin, until the circle touches the line **y = Φα.**

• The very first time the circle touches the line, gives us the solution (the second time, the circle will have a greater radius and hence the solution vector α will have higher Lp-norm).

• For the L1-norm, the solution vector α will be sparse (see middle figure). For the L2-norm, the solution vector will not be sparse (rightmost figure).

L_1 is good, L_2 is bad. Why?

- The L1-norm minimization (subject to the chosen constraint) is guaranteed to give the sparsest possible solutions provided the conditions of theorems 1, 2 and 3 are met.

- Note that we get the sparsest possible solutions even though the L1-norm is a softened version of the L0-norm!

Compressive Sensing: Toy Example with images

- We will compare image reconstruction by two different methods.
- We won't do true compressive sensing. Instead we'll apply different types of sensing matrices Φ synthetically (in software) to an image loaded into memory, merely for example sake.
- In the first method, we compute the DCT coefficients of an image (size N = 512 x 512), and reconstruct the image using the m largest absolute value DCT coefficients (and a simple matrix transpose). m is varied from 1000 to 30000.
- Note that this is equivalent to using a sensing matrix Φ of size m by N.

Compressive Sensing: Toy Example with images

- In the second method, we reconstruct the image using the first 1000 largest DCT coefficients and a random selection of the remaining ones, yielding a total of *m* coefficients, where *m* is again varied from 1000 to 30000.
- The sensing matrix Φ is appropriately assembled.
- In this case, the reconstruction proceeds by solving the following optimization problem:

$$\min \|\theta\|_1 \text{ such that } y = \Phi\Psi\theta$$

Compressive Sensing: Toy Example with images

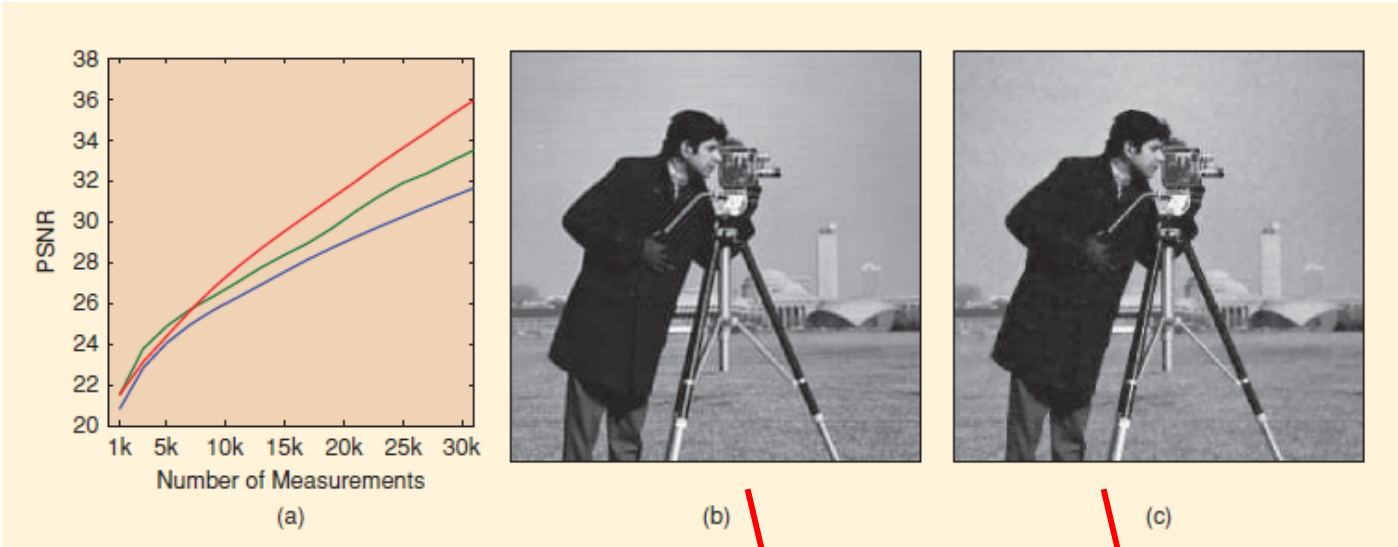

[FIG2] Coded imaging simulation. (a) Recovery error versus number of measurements for linear DCT acquisition (blue), compressive imaging (red), and DCT imaging augmented with total-variation minimization (green). The error is measured using the standard definition of peak signal-to-noise ratio: $\mathrm{PSNR} = 20\log_{10}(255 \cdot 256/\|X - \tilde{X}\|_2)$. (b) Image recovered using linear DCT acquisition with 21,000 measurements. (c) Image recovered using compressive imaging from 1,000 DCT and 20,000 noiselet measurements.

Better edge preservation with the randomly assembled Φ – which actually uses some high frequency information (unlike conventional DCT which discards higher frequencies)

Compressive Sensing: Toy Example with images

- In the second method, the reconstruction results are of better quality and higher PSNR (**red** curve) as compared to the first method (**blue** curve).
- The **green** curve indicates the reconstruction results with the following minimization problem, but using the same set of DCT coefficients as in the first case:

$$\min \|\theta\|_1 \text{ such that } y = \Phi\Psi\theta$$

- The results of this third method (**green** curve) are intermediate between the first two. It indicates the **superiority of the L1-method (green) over simple pseudo-inverse or matrix transpose (blue)**. It also indicates that **choosing a random set of DCT coefficients (red) is better than taking only the low frequency DCT coefficients (green)**.

Compressed Sensing for Piecewise Constant Signals

- In the case of piecewise constant signals/images, the gradients are sparse (zero everywhere except along edges).
- For such images, we can solve the following problem (based on total variation or TV-norm):

$$\min TV(f) \text{ such that } y = \Phi\Psi^T f,$$

$$TV(f) \leftarrow \iint \sqrt{f_x^2 + f_y^2}\, dxdy$$

$$= \text{L1} - \text{norm of the 2D gradient vector}$$

This optimization problem is called a second order cone program. It can be solved efficiently using packages like L1-magic. We will not go into more details, but the documentation of L1-magic is a good starting point for the curious readers.

Ref: Candes, Romberg and Tao, "Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information", IEEE Transactions on Information Theory, Feb 2006.

Theorem 4 (similar to Theorem 1)

- Consider a piecewise constant signal ***f*** (having length *n*). Suppose coefficients of ***f*** are measured through the measurement matrix **Φ** yielding a vector **y** of only *m* << *n* measurements.

  If (*C* is some constant) then the solution to the following is **exact** with probability $1-\delta$:

$$m \geq C \log(n/\delta) \|\theta\|_0 \mu^2(\Psi, \Phi)$$

$$\min_f TV(f) \text{ such that } y = \Phi\Psi^T f$$

$$y \in R^m, \Phi \in R^{m \times n}, f \in R^n, m << n$$

Experiment by Candes, Romberg and Tao

Ref: Candes, Romberg and Tao, “Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information”, IEEE Transactions on Information Theory, Feb 2006.

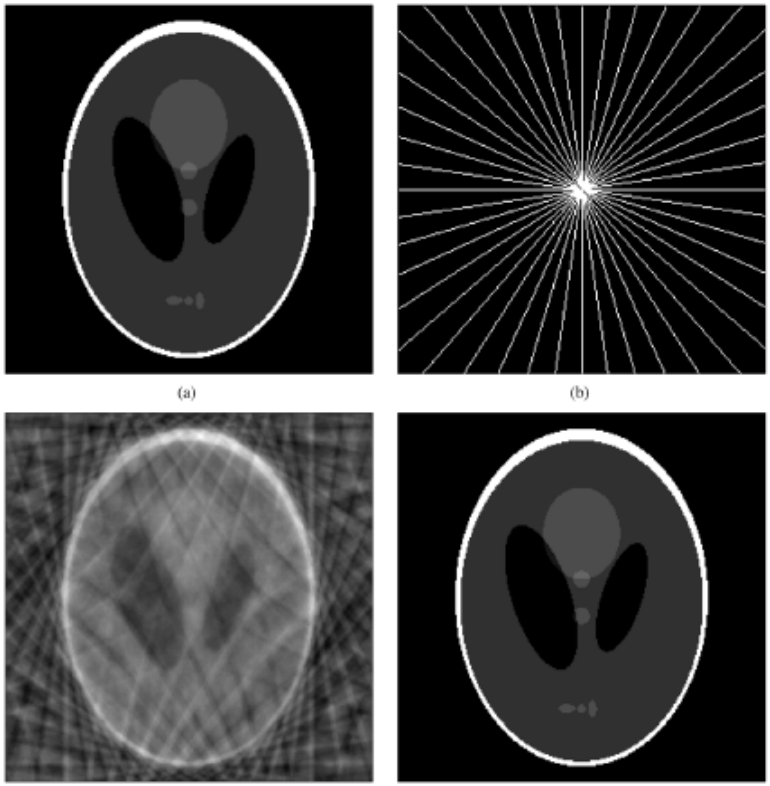

Fig. 1. Example of a simple recovery problem. (a) The Logan–Shepp phantom test image. (b) Sampling domain Ω in the frequency plane; Fourier coefficients are sampled along 22 approximately radial lines. (c) Minimum energy reconstruction obtained by setting unobserved Fourier coefficients to zero. (d) Reconstruction obtained by minimizing the total variation, as in (1.1). The reconstruction is an exact replica of the image in (a).

Experiment by Candes, Romberg and Tao

- The top-left image (previous slide) is a standard phantom used in medical imaging - called as Logan-Shepp phantom.
- The top-right image shows 22 radial directions with 512 samples along each, representing those Fourier frequencies (remember we are dealing with 2D frequencies!) which were measured.
- Bottom left: reconstruction obtained using inverse Fourier transform, assuming the rest of the Fourier coefficients were zero.

Experiment by Candes, Romberg and Tao

- Bottom right (two slides back): reconstruction obtained by solving the following problem **(associated reconstruction error is ZERO)**:

$$f_{true} = \min_f TV(f) \text{ such that } \forall \omega \in \Omega, \hat{f}_{measured}(\omega) = \hat{f}(\omega).$$

Ω is the set of frequencies for which Fourier coefficients were measured.

Equivalent!

$$f_{true} = \min_f TV(f) \text{ such that } \hat{f}_{measured} = I_\Omega \Psi^T f,$$

$\Psi =$ Fourier Basis,

I_Ω : created by choosing only a subset of rows of identity matrix, selects just those Fourier coefficients belonging to Ω.

$$f_{true} = \min_f TV(f) \text{ such that } \hat{f}_{measured} = I_\Omega \Psi^T f,$$

$\Psi = \text{Fourier Basis},$

I_Ω : created by choosing only a subset of rows of identity matrix, selects just those Fourier coefficients belonging to Ω.

Equivalent!

$$f_{true} = \min_f TV(f) \text{ such that } \nabla\hat{f}_{measured} = I_\Omega \Psi^T (\nabla f),$$

L1-norm of vector of gradient magnitudes

Remember: a piecewise constant image has sparse gradients!

- Define

$$d(x,y) \leftarrow (g(x+1,y) - g(x,y)) + i(g(x,y+1) - g(x,y))$$

$$F(d(x,y)) = (1 - i\omega_x)F(g(x,y)) + i(1 - i\omega_y)F(g(x,y))$$

$$= (1 + i - i\omega_x + \omega_y)F(g(x,y))$$

- As Fourier transform of a derivative of a function is an appropriately scaled version of the original function, we can assume that we have with us the measurements of the Fourier transform of the derivative of the image.

Experiment by Candes, Romberg and Tao

- Notice that the **gradient magnitudes** of the phantom image are **sparse**.
- Our measurement matrix in this case is a partial Fourier matrix. It is known to be highly incoherent with the spatial domain!
- The L1-norm we are penalizing in this case turns out to be the total variation (**sum total of gradient magnitudes** over the entire image)!
- It was this experiment by Candes, Romberg and Tao (done around 2005) that inspired modern compressive sensing theory!

Experiment by Candes, Romberg and Tao

- Compare the following two problems:

$$\min \|\theta\|_1 \text{ such that } y = \Phi f = \Phi \Psi \theta \qquad \min_f TV(f) \text{ such that } y = \Phi \Psi^T f$$

- In the former case, we are measuring the image/signal (**spatial domain**) through Φ.
- In the second case, we are making measurements in the **frequency domain** (Fourier coefficient magnitudes) through Φ.

Ref: Logan, “Properties of high-pass signals”, PhD Thesis, Columbia University, 1965

Observations by Logan (1965)

- Consider a band-limited signal **f** (of maximum frequency *B*) corrupted by **impulse noise** (i.e. spikes) of arbitrary magnitude, yielding the corrupted signal:

$$y(t) = f(t) + \eta(t)$$

- Let the support of the impulse noise be sparse in time (valid assumption), having length *T*.
- Consider the following reconstruction criterion:

$$\min_f \|y - f\|_1 \text{ subject to } \forall \omega > B, \hat{f}(\omega) = 0$$

Observations by Logan (1965)

- It turns out that this reconstruction yields **PERFECT** results **regardless of the noise magnitude**, provided that the following condition holds:

$$TB \leq \frac{\pi}{2}$$

- This result exploits the band-limitedness of the signal **f**, but does not impose constraints on the sampling rate.
- This was a very early application of L1-norm optimization for signal reconstruction.

Need for blind compressive sensing

- Take another look at Theorem 2.
- The reconstruction error is upper bounded by the quantity $\|\theta - \theta_S\|_1$. This error depends on how compressible **f** is in the chosen basis $\boldsymbol{\Psi}$ (Note: $\mathbf{f}=\boldsymbol{\Psi}\boldsymbol{\theta}$).
- It is in our hands to choose a "good" basis $\boldsymbol{\Psi}$, for which the error will be low. $\|\theta - \theta_S\|_1$
- This basis can be obtained by learning the basis $\boldsymbol{\Psi}$ from the compressed data on the fly, or based on some offline training. The former is called as **blind compressive sensing**, and can be used to obtain better reconstruction results.

Other extensions of the theory

- Theorems 1, 2, 3 apply for orthonormal bases. The theory has been recently extended to handle **over-complete bases** as well (**Ref:** Candes et al, "Compressed sensing with coherent and redundant dictionaries", *Applied and Computational Harmonic Analysis,* 2010).
- Theorems 1, 2, 3 assume that the measurement matrix Φ is known to us accurately. This may not happen in practice. In such a case, we have:

$$y = (\Phi + \Delta\Phi)\Psi\theta + \eta = \Phi\Psi\theta + \boxed{(\Delta\Phi\Psi\theta + \eta)}$$

Signal dependent noise

CS theory has been extended to handle such cases, ref: Yang, Zhang and Xie, "Robustly Stable Signal Recovery in Compressed

Sensing with Structured Matrix Perturbation".

THANK YOU