„Signal Processing“

→ Compressed Sensing and related tales

Part 1: Compressed Sensing Primer

Dr.-Ing. Florian Römer
Florian.Roemer@izfp.fraunhofer.de

Fraunhofer IZFP

Part 1: Key Take-Aways

What's behind this?

- „Breaking" Shannon-Nyquist: how?
- What is the role of sparsity?

How do we measure?

- What is the role of incoherence?
- What is the role of randomness?

How do we reconstruct?

- How to find the solution when there are infinitely many?
- How to solve the reconstruction problem efficiently?

Recap: Shannon-Nyquist sampling

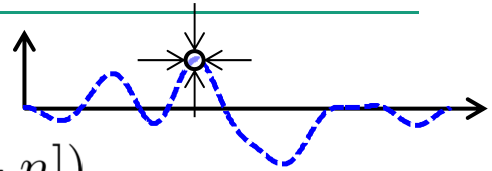

- Shannon-Nyquist-Whittaker theorem (SNWT)

$$\forall x(t) \in \mathcal{B}_B : \quad x(t) = \sum_n x(n/B) \cdot \mathrm{sinc}(\pi[tB - n])$$

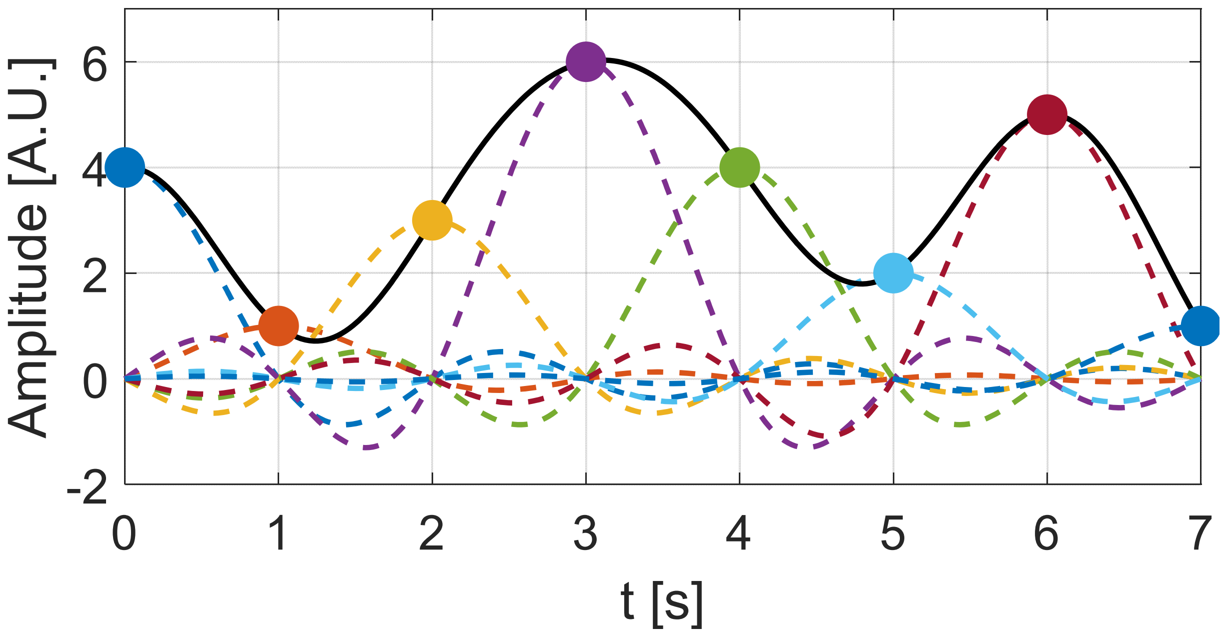

Compressed Sensing: Measure signals at sub-Nyquist rates w/o loss of information?

Compressed Sensing: Intuition

- SNWT: *sufficient* condition, but not necessary. We can violate it if
  - We have prior knowledge
  - We change the way we sample

Example

- Bandlimited Signal (< 500 Hz)
  - Shannon-Nyquist: $N = 1000$ samples per second sufficient for *any* possible signal
  - What if we knew the signal is perfectly sinusoidal?
    - $N = 3$ samples sufficient (conceptually), if placed properly
    - Example (jupyter)

Compressed Sensing: Intuition

- SNWT: **sufficient** condition, but not **necessary**. We can violate it if
  - We have prior knowledge
  - We change the way we sample

- Nyquist-Samples often have **redundancies**
  - Signals have **structure**
  - Only a part of the signals is **relevant**
  - Often, **prior knowledge** is available

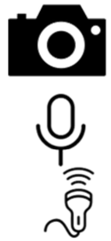

RAW => JPEG: ca. 1/10-1/50*

PCM => MP3: ca. 1/6*

Ultrasound: ca. 1/10-1/500

* „lossless“: ca. ½

- So, we *can* do it, *should* we?
  - Yes: Sub-Nyquist-Sampling has many practical **advantages**
  - Saving: time, energy, data rate/size
  - Very application-specific

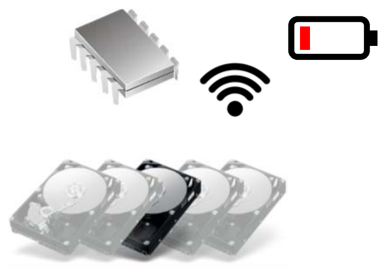

David Brady: „One can regard the possibility of digital compression as a failure of sensor design. If it is possible to compress measured data, one might argue that too many measurements were taken.“

What do we gain?

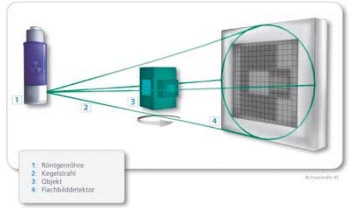

- Potentially: reduction of measurement time
  - Depending on application, e.g. for MRT/CT, imaging

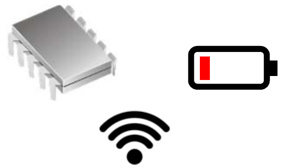

- Potentially: reduction of energy consumption
  - (a) Data acquisition itself (e.g. slower system clock)
  - (b) Data transmission, esp. for autonomous sensors

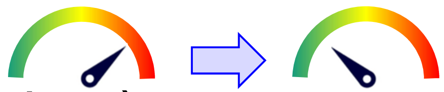

- Potentially: reduction of data rate
  - Interesting for high rate applications („fast pre-sorting“)

- Reduction of total amount of data (yes, but)
  - Better: do not record irrelevant data in the first place

Step 1: Prior knowledge => sparsity

- Basic idea: Set of possible (relevant) **signal** parts = **known**.
  Observed signal = combination of a *few* of these.
  For example…

Harmonics Spikes Both! etc

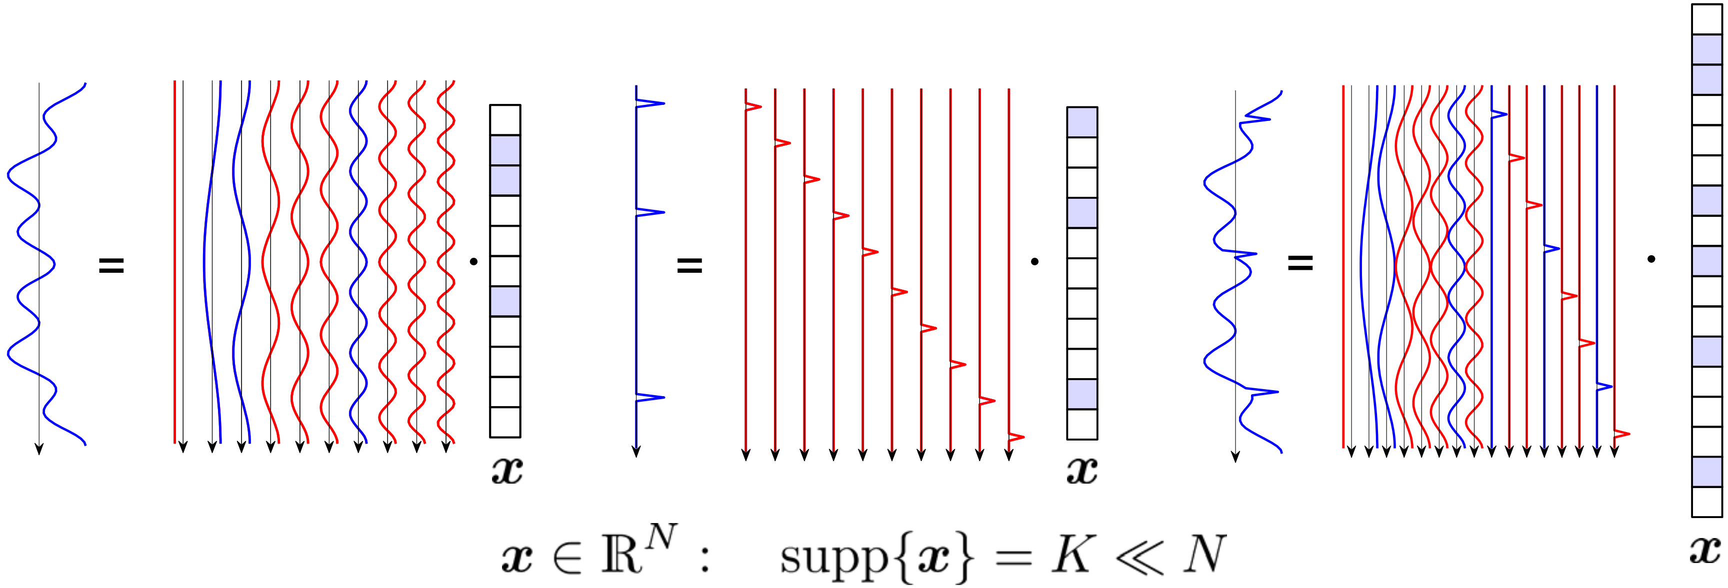

$$\boldsymbol{x} \in \mathbb{R}^N : \quad \mathrm{supp}\{\boldsymbol{x}\} = K \ll N$$

Step 2: Measurement => Incoherence

- Generalized sampling (linear)

$$y[n] = \int p_n(t) y(t) \mathrm{d}t$$

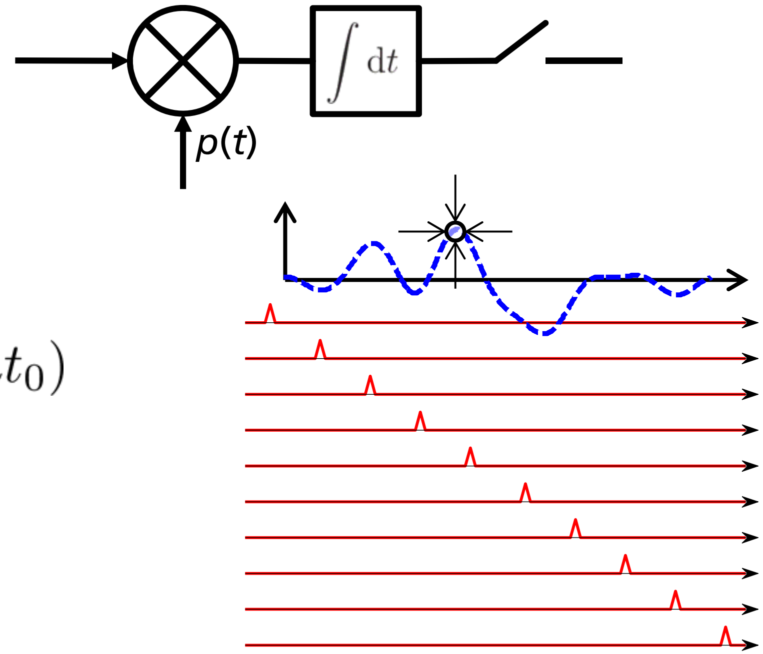

- Special case ideal sample&hold

$$p_n(t) = \delta(t - nt_0) \quad \Rightarrow y[n] = y(nt_0)$$

- How to choose $p_n(t)$?
  - Example (jupyter)

    => Incoherence is the key!

    => Randomness as a natural source of incoherence

Random Matrices as a Measuring Kernels

- How to build incoherent bases without knowing the basis?
- “Randomness” is incoherent with any “structure”!

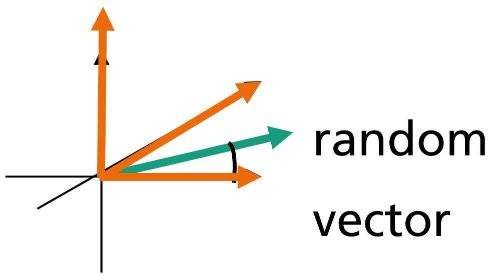

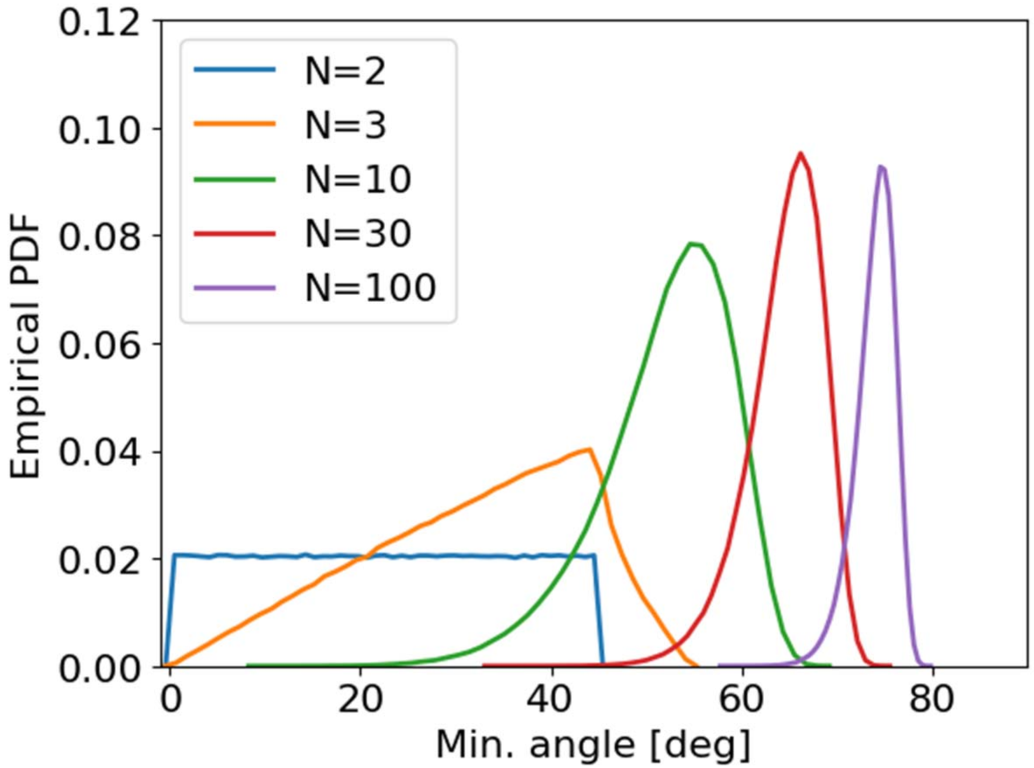

Step 3: Reconstruction

- Linearity: sparse model ✓ measurements ✓
  - Mapping from coefficients to measurements is linear. Solve LSE!

$$\boldsymbol{d} = \boldsymbol{\Psi} \cdot \boldsymbol{x}$$

(M,1) (M,N) (N,1)

Sub-Nyquist: M < N => underdetermined

But: $\mathrm{supp}\{\boldsymbol{x}\} = K < M < N$

=> Seek out *sparse* solution!

$$\boxed{\begin{aligned} \boldsymbol{x}^\star &= \arg\min_{\boldsymbol{x}} \mathrm{supp}\{\boldsymbol{x}\} \\ \text{s.t.} \quad & \boldsymbol{d} = \boldsymbol{\Psi} \cdot \boldsymbol{x} \end{aligned}}$$

$$\|\boldsymbol{x}\|_{\ell_p} = \sqrt[p]{\sum_{k=1}^{N} |x_k|^p}$$

However: This problem is NP hard.

Step 3: Reconstruction

- Linearity: sparse model ✓ measurements ✓
  - Mapping from coefficients to measurements is linear. Solve LSE!

$$\boldsymbol{d} = \boldsymbol{\Psi} \cdot \boldsymbol{x}$$

(M,1) (M,N) (N,1)

Sub-Nyquist: $M < N$ => underdetermined

But: $\mathrm{supp}\{\boldsymbol{x}\} = K < M < N$

=> Seek out *sparse* solution!

$$\boxed{\begin{aligned} \boldsymbol{x}^\star &= \arg\min_{\boldsymbol{x}} \|\boldsymbol{x}\|_{\ell_1} \\ \text{s.t.} \quad & \boldsymbol{d} = \boldsymbol{\Psi} \cdot \boldsymbol{x} \end{aligned}}$$

$$\|\boldsymbol{x}\|_{\ell_p} = \sqrt[p]{\sum_{k=1}^{N} |x_k|^p}$$

This problem has efficient solutions.
Under certain conditions, it gives the same solution as the l0 problem.

Step 3: Geometric intuition

$$\boldsymbol{x}^\star = \boxed{\arg\min_{\boldsymbol{x}} \|\boldsymbol{x}\|_{\ell p}}$$

$$\text{s.t.} \quad \boxed{\boldsymbol{d} = \boldsymbol{\Psi} \cdot \boldsymbol{x}}$$

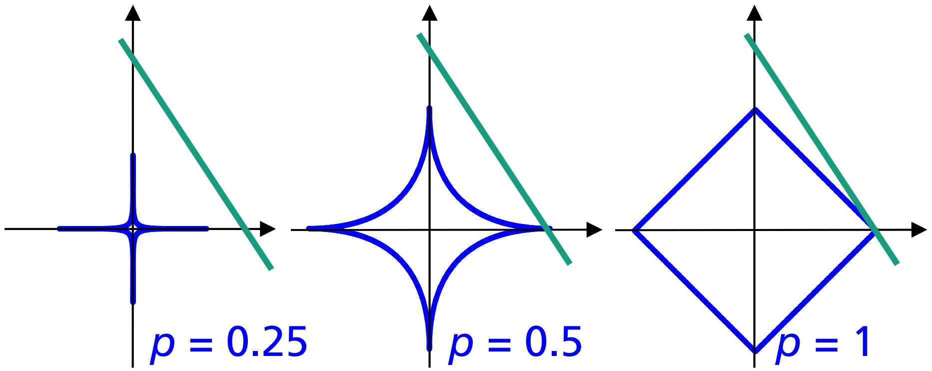

- Sparse solutions are "on the axes".
  - For $p \leq 1$: "arms" of the p-norm balls find them.
  - For $p \geq 1$: p-norm ball is convex.
  - => p = 1 is the best compromise.

$$\|\boldsymbol{x}\|_{\ell p} = \sqrt[p]{\sum_{k=1}^{N} |x_k|^p}$$

Fraunhofer IZFP

Escape Velocities for the L1-ball

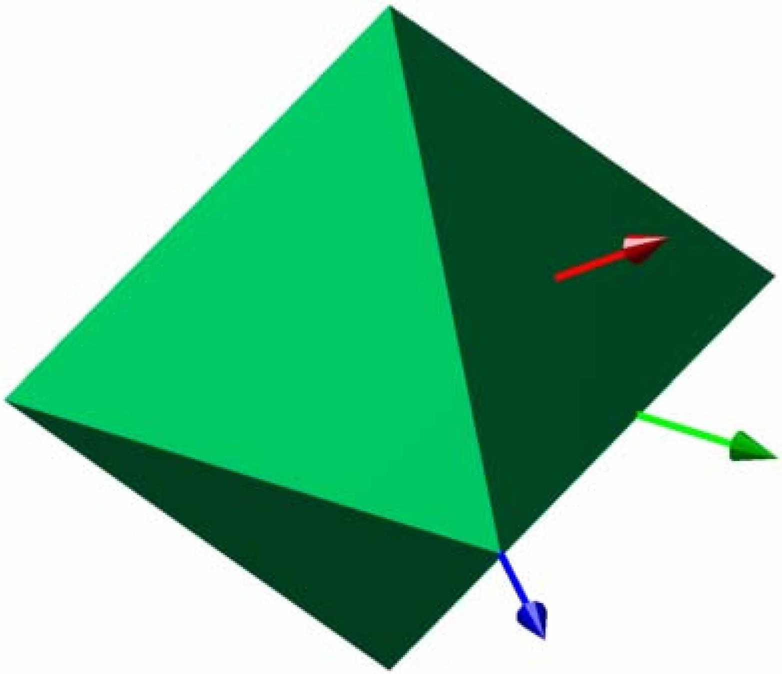

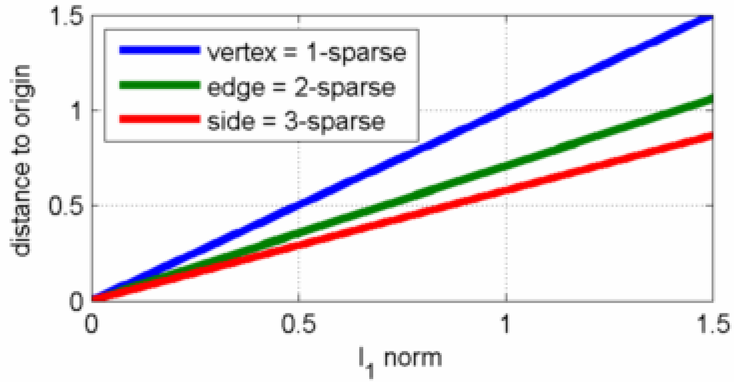

- Vertices move faster than edges, which move faster than sides
- They correspond to 1-sparse, 2-sparse, and 3-sparse respectively

Identifiability: L0 vs. L1

- In general, for l0 we have $K < \frac{1}{2}(\underbrace{\text{k-rank}(\mathbf{\Psi})}_{\text{Kruskal-rank} \geq r \text{ if all sets of } r \text{ columns are linearly independent}} + 1) \leq \frac{M+1}{2}$

Kruskal-rank ≥ *r* if **all** sets of *r* columns are linearly independent

$$\boxed{M \geq 2K + 1}$$

NB: there are 2*K* degrees of freedom!
(support indices + amplitudes)

- whereas for l1 the bounds are of the form

$$\boxed{M \geq C \cdot K \cdot \underbrace{\log(N/K)}_{\text{l1 “penalty”}}}$$

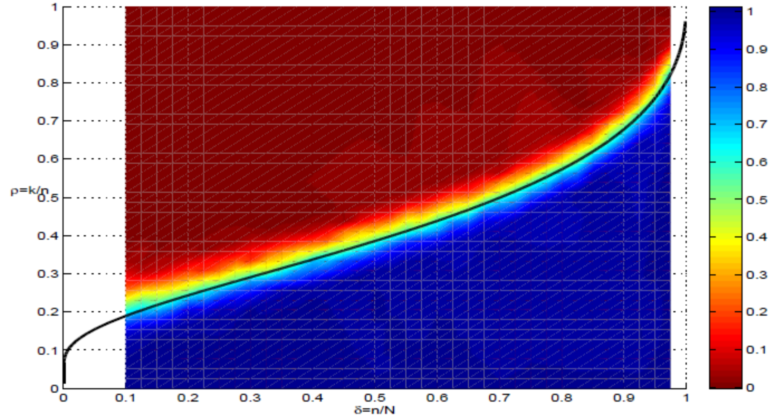

Practical solvers

- Two equivalent formulations
  - BPDN[1] $$\min \|\mathbf{x}\|_1 \quad \text{s.t.} \quad \|\mathbf{d} - \mathbf{\Psi x}\|_2 \leq \epsilon$$
  - LASSO[2] $$\min \|\mathbf{d} - \mathbf{\Psi x}\|_2 + \lambda \|\mathbf{x}\|_1$$
- The zoo of algorithms is gigantic. Two commonly found examples
  - OMP[3]: to fit $\|\mathbf{d} - \mathbf{\Psi x}\|_2$ iteratively, select matching columns in a greedy manner, project residual data on the complement subspace

[1]BPDN = Basis Pursuit DeNoising
[2]LASSO = Least Absolute Shrinkage and Selection Operator
[3]OMP = Orthogonal Matching Pursuit

OMP example for a three-sparse problem

$$\mathbf{d} = \mathbf{\Psi x}$$

$$n_1 = \arg\max_n |\mathbf{d}^{\mathrm{T}} \cdot \psi_n|$$

$$\mathbf{d}_1 = \mathbf{d} - \mathbf{\Psi}_{n_1} \mathbf{\Psi}_{n_1}^{+} \mathbf{d}$$

$$n_2 = \arg\max_n |\mathbf{d}_1^{\mathrm{T}} \psi_n|$$

$$\mathbf{d}_2 = \mathbf{d} - \mathbf{\Psi}_{[n_1,n_2]} \mathbf{\Psi}_{[n_1,n_2]}^{+} \mathbf{d}$$

$$n_3 = \arg\max_n |\mathbf{d}_2^{\mathrm{T}} \psi_n|$$

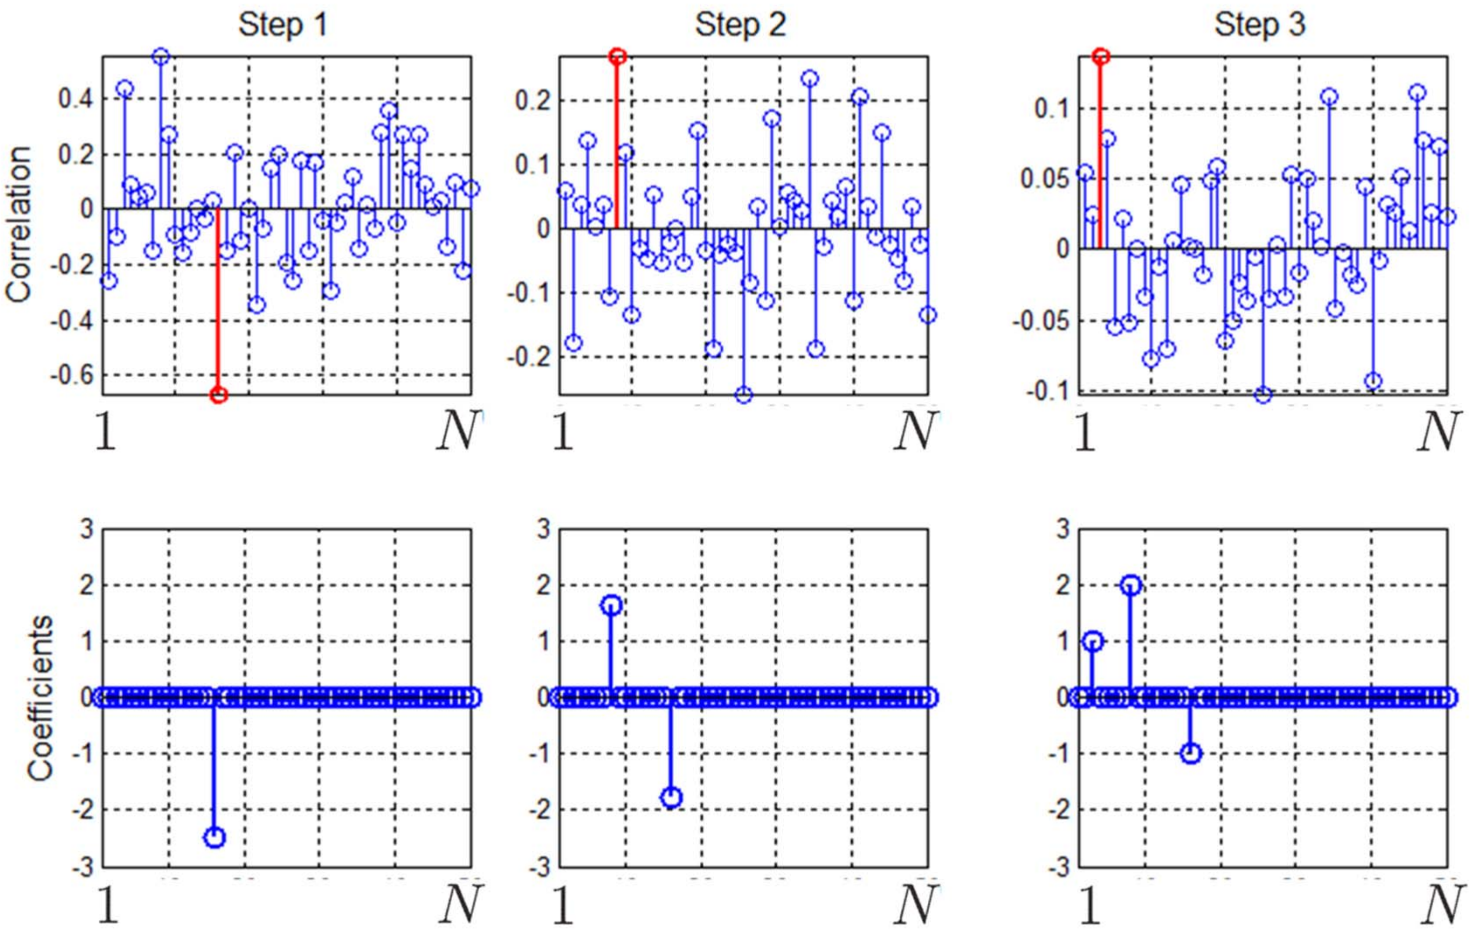

Practical solvers

- Two equivalent formulations
  - BPDN[1] $\min\|\mathbf{x}\|_1 \quad \text{s.t.} \quad \|\mathbf{d}-\mathbf{\Psi x}\|_2 \leq \epsilon$
  - LASSO[2] $\min\|\mathbf{d}-\mathbf{\Psi x}\|_2 + \lambda\|\mathbf{x}\|_1$
- The zoo of algorithms is gigantic. Two commonly found examples
  - OMP[3]: to fit $\|\mathbf{d}-\mathbf{\Psi x}\|_2$ iteratively, select matching columns in a greedy manner, project residual data on the complement subspace
  - (F)ISTA[4]: to minimize $\|\mathbf{d}-\mathbf{\Psi x}\|_2$, perform gradient steps, regularize by (soft) thresholding the coefficients.

[1]BPDN = Basis Pursuit DeNoising
[2]LASSO = Least Absolute Shrinkage and Selection Operator
[3]OMP = Orthogonal Matching Pursuit
[4]FISTA = Fast Iterative Shrinkage and Thresholding Algorithm

Iterative Shrinkage and Thresholding

- To minimize $\|\mathbf{d}-\mathbf{\Psi x}\|_2^2$ consider its gradient

$$\|\mathbf{d}-\mathbf{\Psi x}\|_2^2 = \mathbf{d}^{\mathrm{T}}\mathbf{d}-\mathbf{d}^{\mathrm{T}}\mathbf{\Psi x}-\mathbf{x}^{\mathrm{T}}\mathbf{\Psi}^{\mathrm{T}}\mathbf{d}-\mathbf{x}^{\mathrm{T}}\mathbf{\Psi}^{\mathrm{T}}\mathbf{\Psi x}$$

$$\Rightarrow \nabla_{\mathbf{x}} = 2\mathbf{\Psi}^{\mathrm{T}}\mathbf{d}-\mathbf{\Psi}^{\mathrm{T}}\mathbf{\Psi x} = 2\mathbf{\Psi}^{\mathrm{T}}(\mathbf{d}-\mathbf{\Psi x})$$

- Gradient step $\quad \mathbf{x} \leftarrow \mathbf{x}+\mu\mathbf{\Psi}^{\mathrm{T}}(\mathbf{d}-\mathbf{\Psi x})$
- But then $\boldsymbol{x}$ is not sparse. Regularizer? Soft thresholding:

$$\mathbf{x} \leftarrow \tau(\mathbf{x})$$

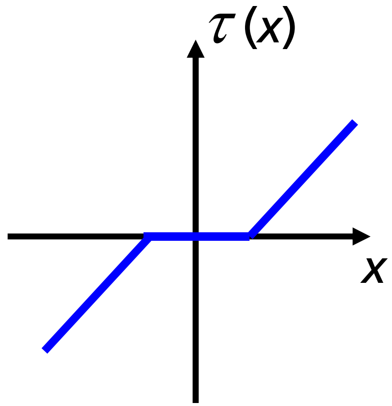

- Fast?
  - Apply momentum methods
  - Smart step size adaptation

But how do I measure incoherently?

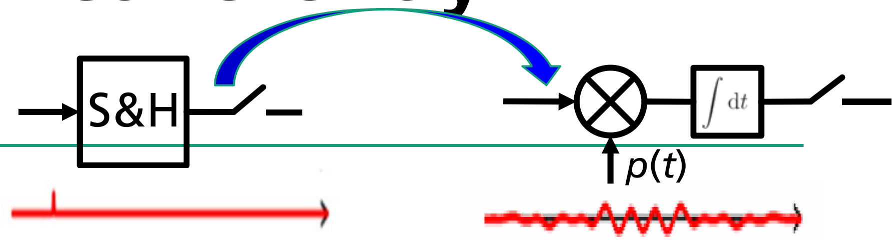

- That depends…
  - Localized signals (spikes) => extended sampling functions (MLBS, harmonics)
    - Random Demodulator, MWC, Fourier sampling
  - Highly correlated signals => subsampling
    - Spatial / angular subsampling, irregular arrays
  - Integral sensors => „Scrambling“
    - Single pixel camera, shattered lens imaging, coded aperture

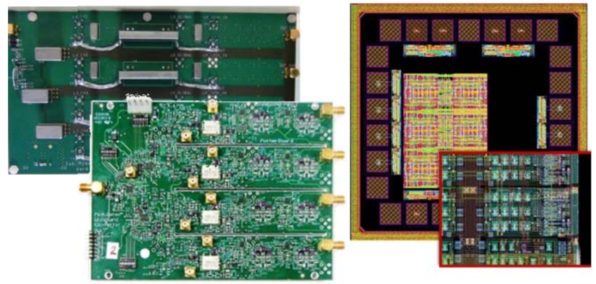

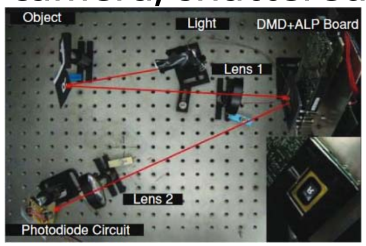

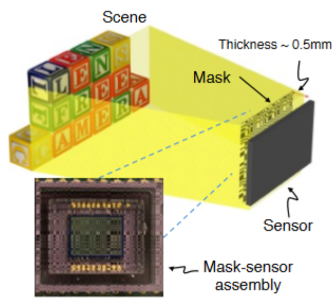

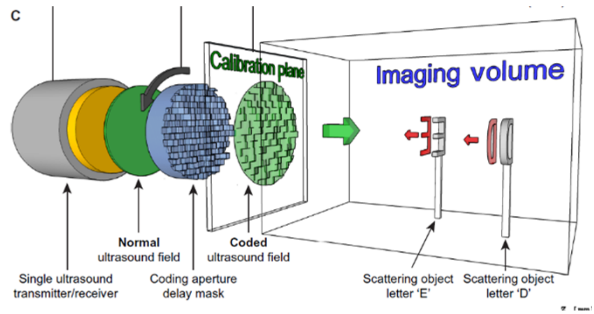

Bigger picture: latent space representation

- Behind the scenes: existence of efficient latent space representation
  - In CS obtained through linear projections.
  - NB: also a common component in ML/DL (self-attention in transformers, latent representation in autoencoders)
    - Their advantage: not limited to linear compression

- Autoencoder as a CS sensor?
  - Sure, why not.
  - Sensible especially if you can implement the encoder part (after freezing the weights) in hardware (FPGA?).

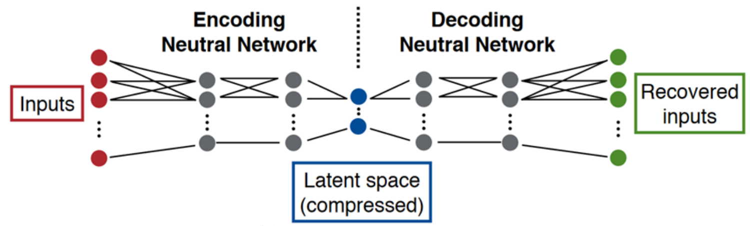

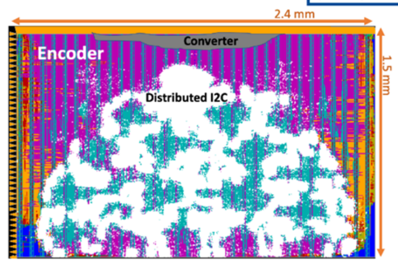

Neural Network (NN) autoencoder for front-end data compression on an ASIC, based on the CMS High-Granularity Endcap Calorimeter (HGCal) at HL-LHC

https://lss.fnal.gov/archive/2020/slides/fermilab-slides-20-121-e.pdf

Applications (Selection)

MRI

X-ray CT

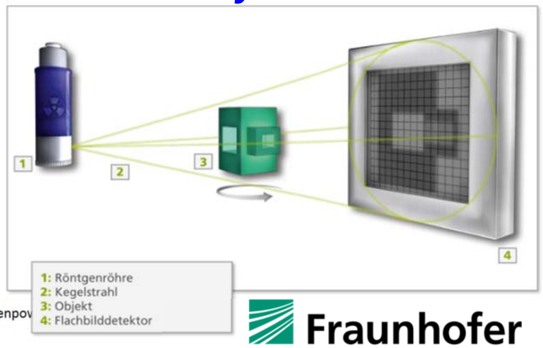

Fraunhofer IIS

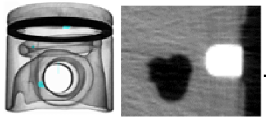

firefly [flexpress]

Ultrasound

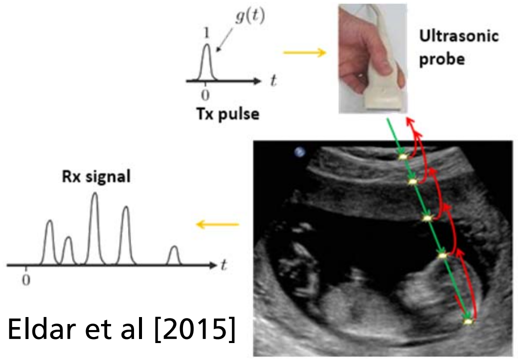

Eldar et al [2015]

Hyper-spectral imaging

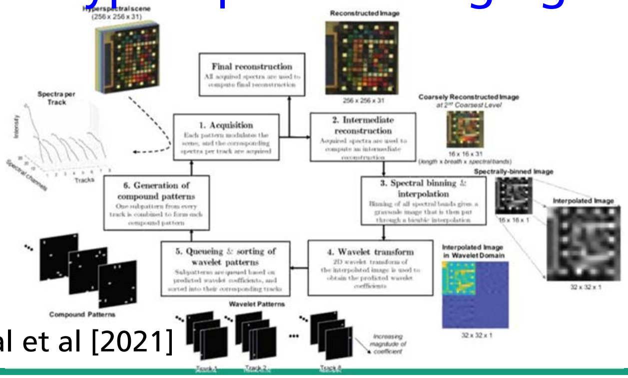

Kubal et al [2021]

Radar

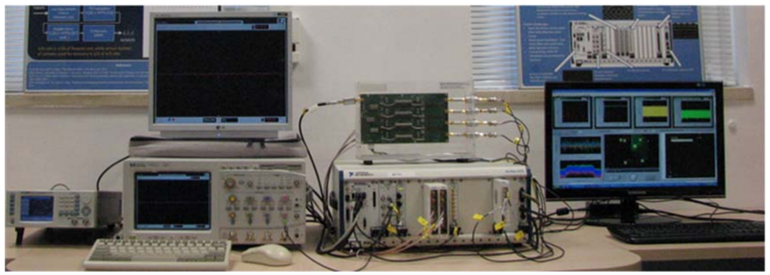

Eldar et al [2017]

etc.

Part 1: Key Take-Aways

What's behind this?

- Shannon-Nyquist is sufficient, but not necessary. With *prior knowledge*, perfect reconstruction from fewer samples is possible.
- We formalize „prior knowledge" through sparsity, as a measure for signal complexity.

How to measure?

- We need incoherent measurement functions.
- Pseudo-randomness can be a source of incoherence.

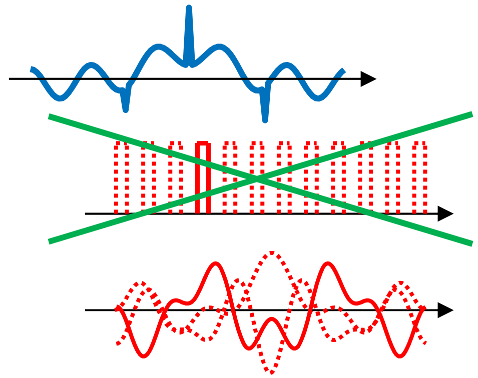

How do we reconstruct?

- Look for the one signal that maximizes sparsity.
- Can be solved efficiently through relaxation.

$$\min \|\mathbf{x}\|_0 \text{ s.t. } \mathbf{y} = \mathbf{A}\,\mathbf{x}$$

$$\Downarrow$$

$$\min \|\mathbf{x}\|_1 \text{ s.t. } \mathbf{y} = \mathbf{A}\,\mathbf{x}$$

„Signal Processing“

➔ Compressed Sensing
and related tales

Part 2: Compressed Sensing in Ultrasound

Dr.-Ing. Florian Römer
Florian.Roemer@izfp.fraunhofer.de

Ultrasound NDT „powered by" Compressed Sensing

Improve US-NDT inspired by innovations in the CS field

Data reduction

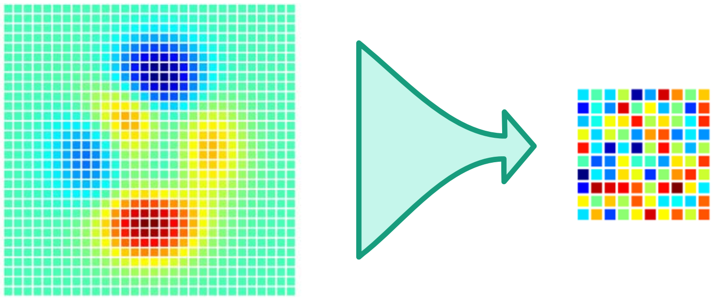

=> Design of Experiment problem: where/how to sample optimally?

- Q1: Temporal sampling: how?
  - Fourier subsampling
- Q2: Spatial sampling: where?
  - Spatial subsampling für synthetic aperture, sparse array Design
- Q3: Spatial sampling: how?
  - Optimal excitation, coded signals

Enhanced reconstruction

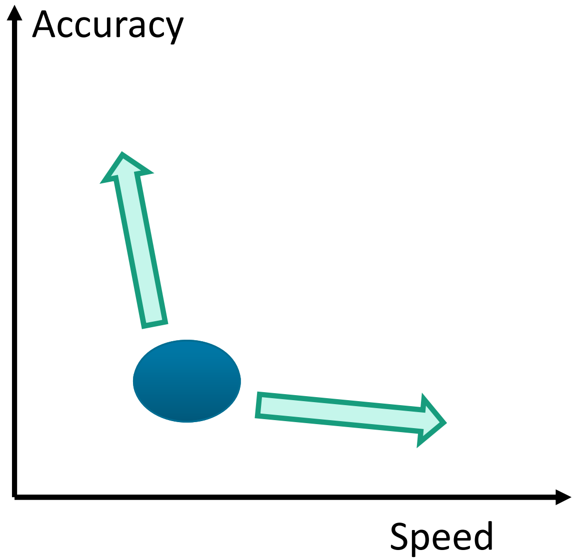

=> Some of these can be applied to fully sampled data…

Compressed Sensing for Ultrasound: Intuition

- In CS we change the way we sample analog signals to avoid capturing redundancy
  ➔ reduce data size, data rate, acquisition time, energy

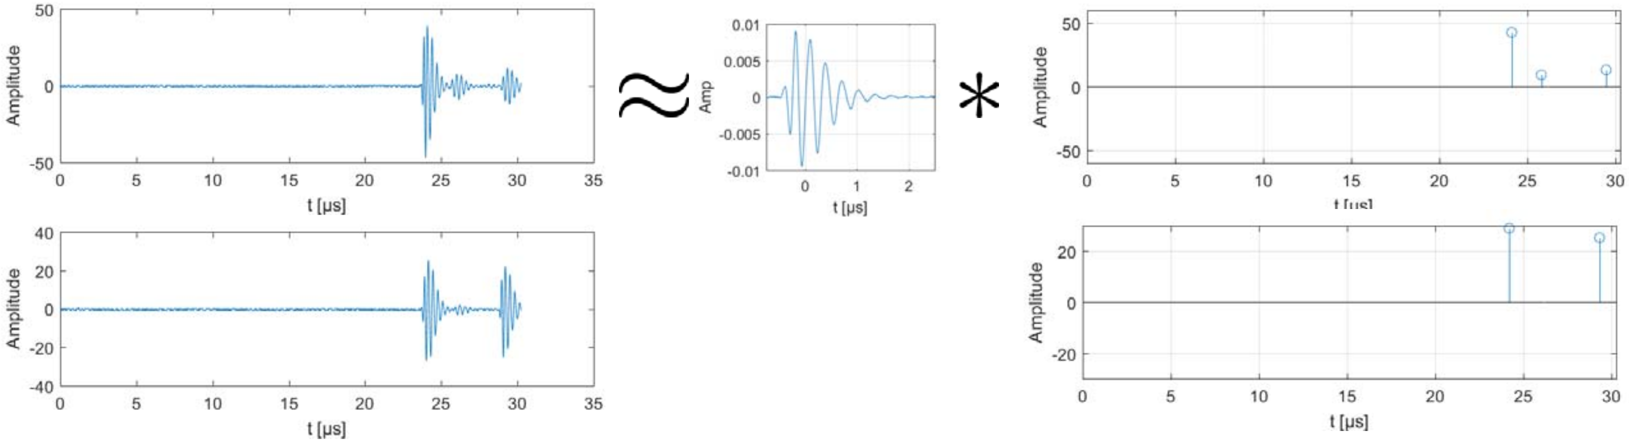

- **Observation 1**: A-Scans are often sparse in time domain. Sampling many zeros is redundant
  ➔ we can measure in frequency domain instead (i.e., "a few" Fourier coefficients)
- **Observation 2**: adjacent A-Scans highly correlated
  ➔ we can apply subsampling (to some degree)

Subsampling in space and frequency

- Note: requires suitable reconstruction method $\min \| \boldsymbol{A}\, \boldsymbol{x} - \boldsymbol{b} \|_2 + \mu \|\boldsymbol{x}\|_1$ (e.g., FISTA)

TECHNISCHE UNIVERSITÄT ILMENAU Fraunhofer IZFP

Compressed Sensing in Time

Why do we talk about Fourier sensing so much? What is the link?

- Consider the generalized linear sampling operator $y_n = \int_T x(t) b_n(t) \mathrm{d}t$

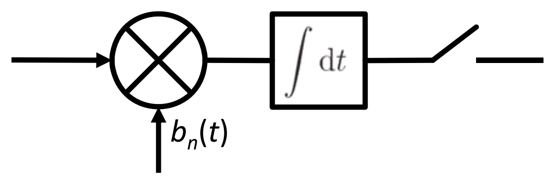

- (aWLOG), assume that $x(t)$ is limited in time $[0,T]$. => WLOG choose $b_n(t)$ to be time-limited as well.

$$x(t) = \sum_\ell X[\ell] \mathrm{e}^{-\jmath 2\pi\ell/T} \qquad b_n(t) = \sum_\ell B_n[\ell] \mathrm{e}^{-\jmath 2\pi\ell/T} \qquad \forall t \in [0,T]$$

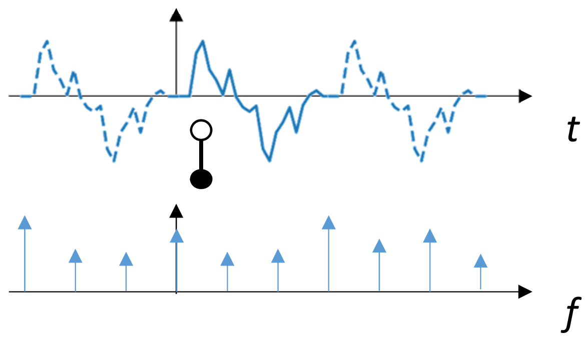

- Then, it is easy to see that

$$y_n = \int_T x(t) p_n(t) \mathrm{d}t = \sum_\ell X[\ell] B_n[\ell]$$

i.e., our samples are linear combinations of the Fourier coefficients of $x(t)$. Our choice of $b_n(t)$ determines the weights.

- Special case $B_n[\ell] = \delta[\ell - m_n] \Rightarrow y_n = X[m_n]$

TECHNISCHE UNIVERSITÄT ILMENAU Fraunhofer IZFP

Fourier Sampling: Architectures

How to build a Fourier Subsampling Sensor?

- Bank of analog narrowband (I/Q) filters

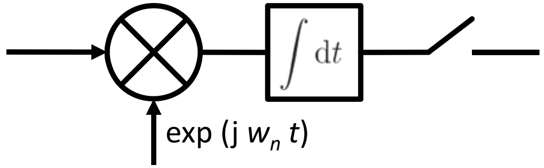

- Modulated Wideband Converter (MWC)
  - Delivers an (invertible) linear combination of FC

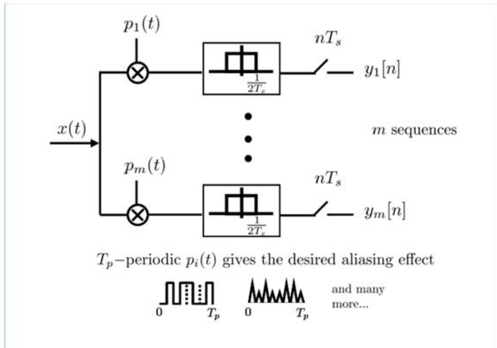

Figure 2. The modulated wideband converter

Figure 3. Circuit implementation. The left board realizes the analog mixing and filtering. A sign waveform alternating at 2 GHz is implemented in the right board.

https://www.weizmann.ac.il/math/yonina/software-hardware/software/sub-nyquist-sampling-modulated-wideband-converter

- Digital FPGA-based compression
  - e.g., FFT+subsel., chirp-Z, Goertzel

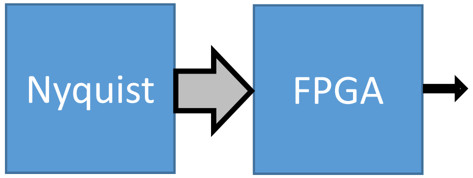

- Analog computing DFTs, …

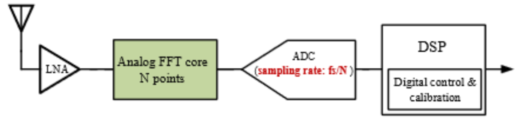

J. Li et. al., PRIME 2022

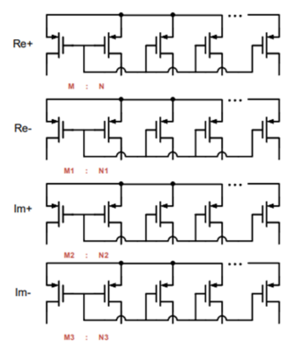

- XG-1 + XF-1 (ASIC)

TECHNISCHE UNIVERSITÄT ILMENAU
PRIME

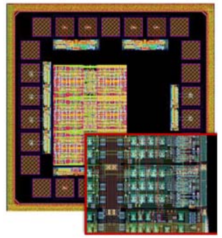

- …

TECHNISCHE UNIVERSITÄT ILMENAU
Fraunhofer IZFP

Fourier Sensing vs. Scrambled Fourier Sensing

- Architectures

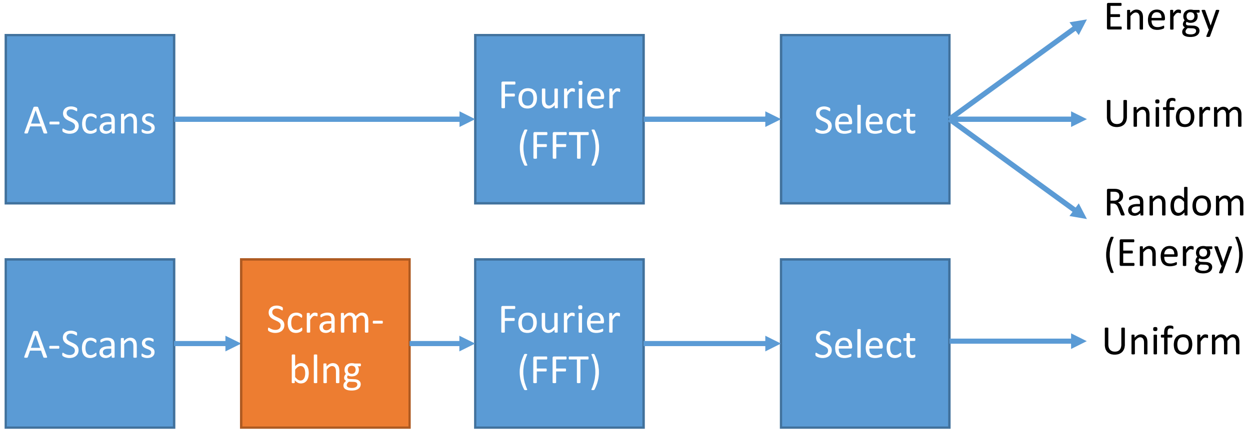

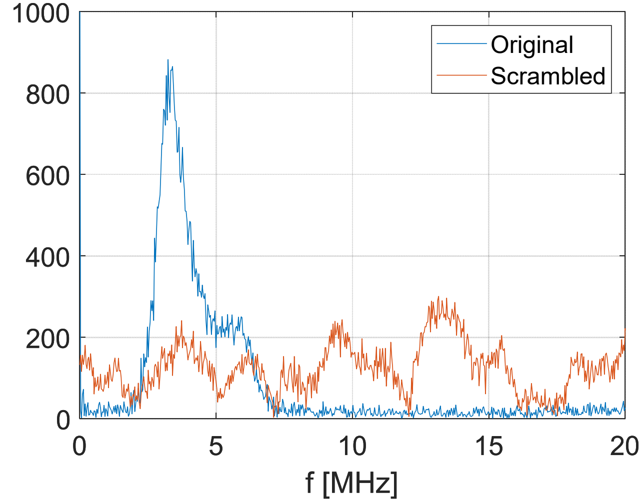

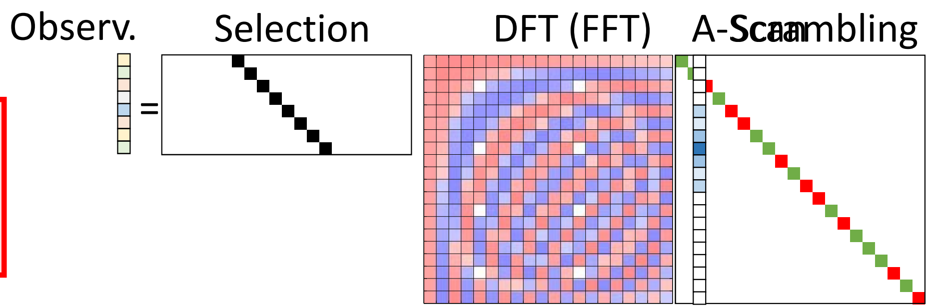

Recovery: for $n_f \geq C s_{\max} \delta^{-2} \log^2(1/\delta) \log N_t \log(s_{\max}) \log^2(s_{\max}\delta)$ scrambled Fourier sensing is guaranteed to recover the A-Scans with probability $\mathrm{e}^{-C_2 \delta}$.

TECHNISCHE UNIVERSITÄT ILMENAU

Fraunhofer IZFP

Other variations of Fourier sensing

MWC, why is it related?

- Build
  - Sequence generator for PN sequences $p_n(t)$
  - analog multiplier + integrator
- What you get is $\boldsymbol{y} = \boldsymbol{B}\,\boldsymbol{x}$, where $\boldsymbol{x}$ are the desired Fourier coefficients and $\boldsymbol{B}$ depends on your sequences

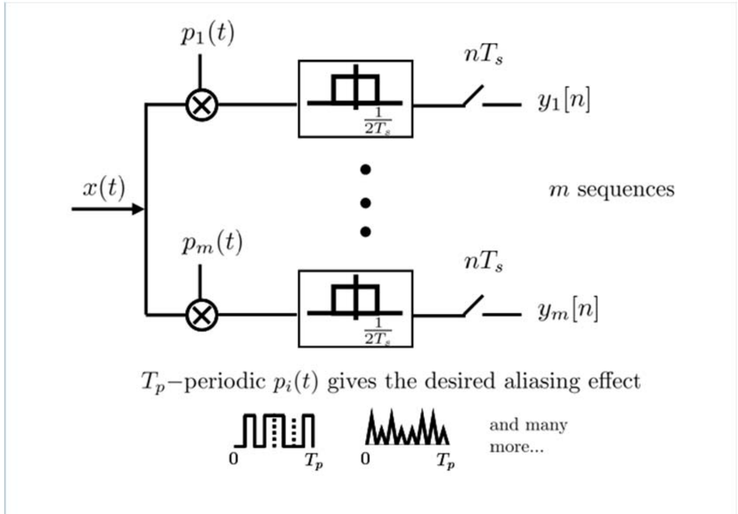

Figure 2. The modulated wideband converter

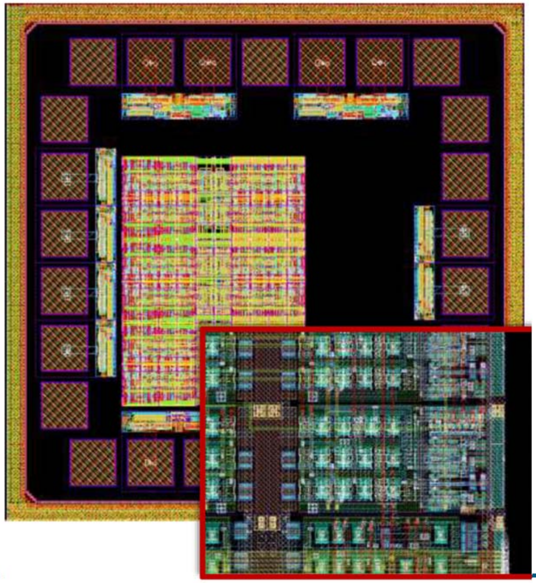

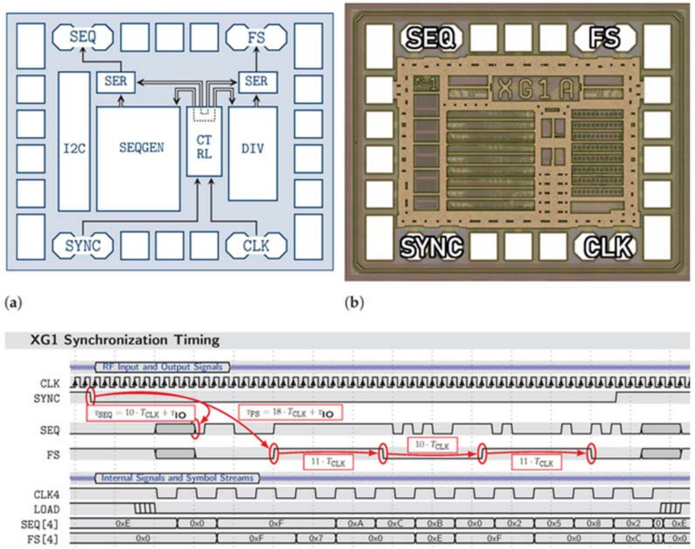

TECHNISCHE UNIVERSITÄT ILMENAU

Fraunhofer IZFP

When to choose what?

- With so many architectures in place, what should we choose?
  - It depends on the application.

- For analog hardware implementations
  - MWC is a tried and tested concept that scales well.
  - For very few Fourier coefficients, a direct Fourier sensing may be even simpler.

- For quasi-digital implementations (FPGA)
  - If we know where the signal energy focusses => Fourier sensing
  - If we don‘t (multiband/wideband) => Scrambled Fourier sensing
  - BTW: Optimal FC selection is an interesting Design of Experiment problem!

Some numerical results: Scrambled Fourier Sensing + FISTA reconstruction

- Specimen: „MUSE“, recorded with an automated scan
  0.5mm grid, 2 MHz, immersion

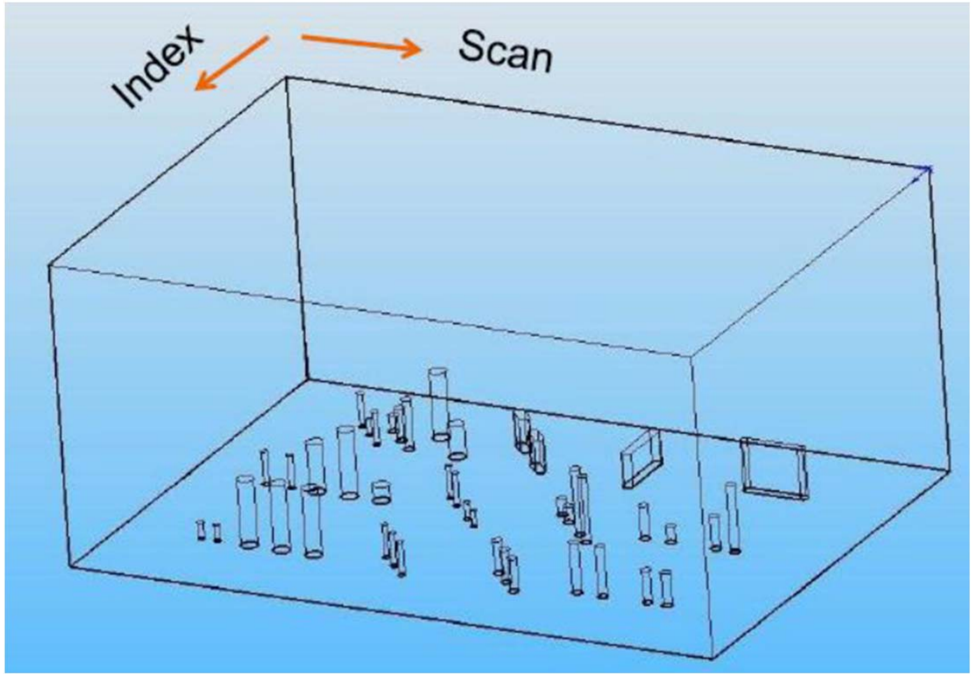

Some numerical results: Scrambled Fourier Sensing + FISTA reconstruction

- Specimen: „MUSE“, recorded with an automated scan 0.5mm grid, 2 MHz, immersion

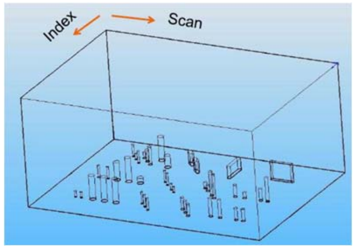

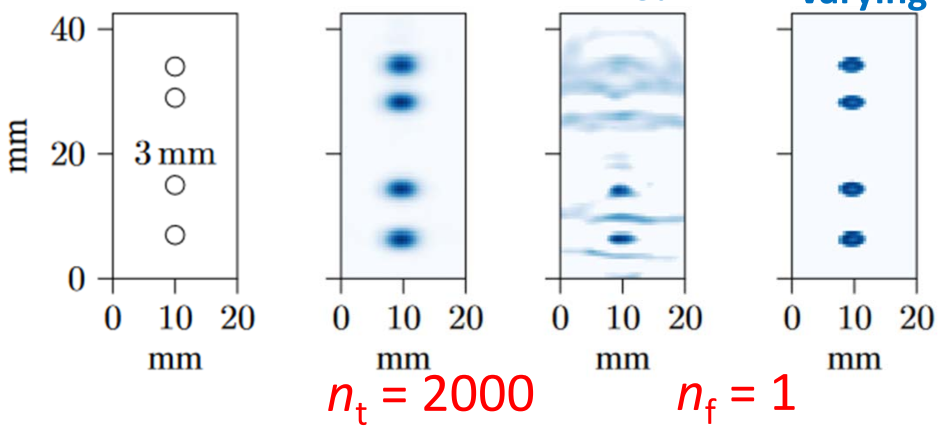

Top view (C image)

Side view (B image)

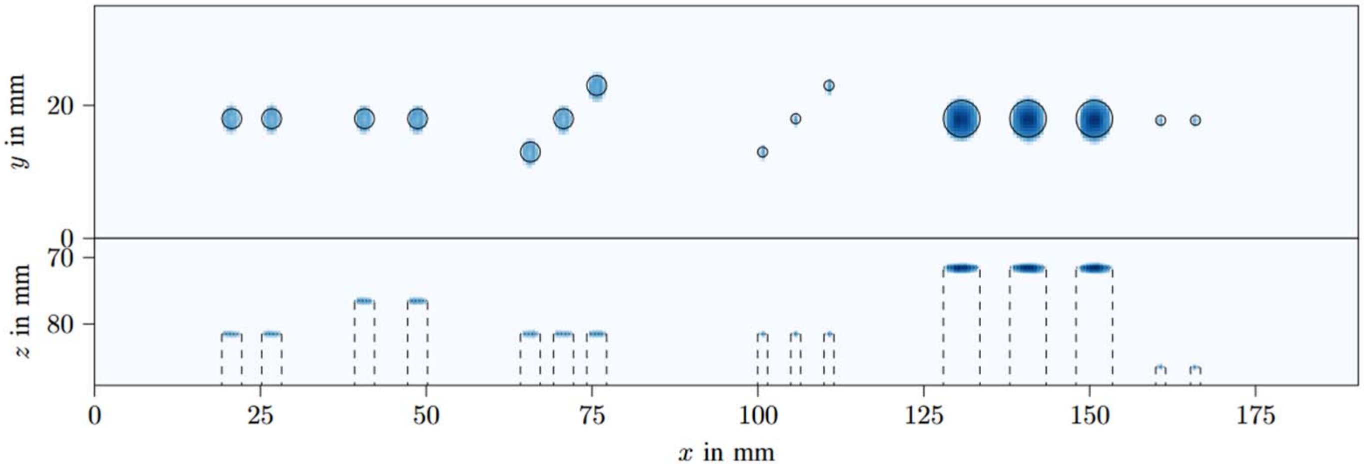

TECHNISCHE UNIVERSITÄT ILMENAU

Fraunhofer IZFP

Ultrasound NDT „powered by“ Compressed Sensing

Improve US-NDT inspired by innovations in the CS field

Data reduction

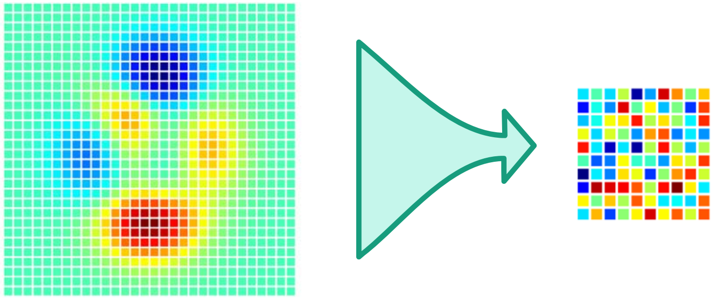

=> Design of Experiment problem: where/how to sample optimally?

- Q1: Temporal sampling: how?
  - Fourier subsampling
- Q2: Spatial sampling: where?
  - Spatial subsampling für synthetic aperture, sparse array Design
- Q3: Spatial sampling: how?
  - Optimal excitation, coded signals

Enhanced reconstruction

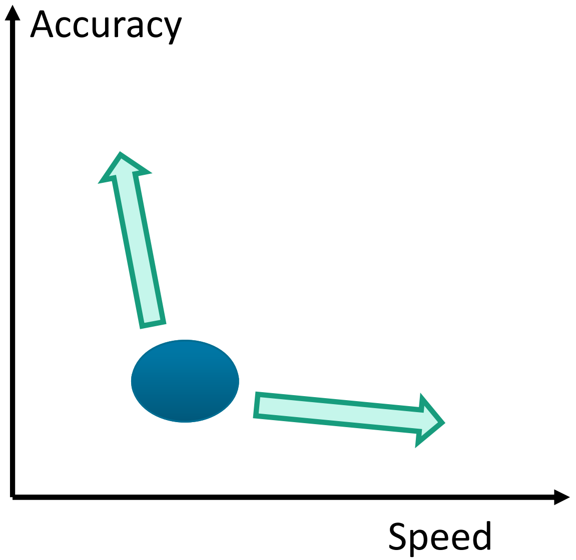

=> Some of these can be applied to fully sampled data…

TECHNISCHE UNIVERSITÄT ILMENAU

Fraunhofer IZFP

A particularly relevant usecase: Full Matrix Capture

- Elements transmit sequentially, receive in parallel
- For *M* elements:
  - *M* measurement cycles
  - M^2 A-Scans (M x M x N_T samples)
- Drawbacks:
  - *M* cycles need time
  - *M* parallel RF chains: hardware complexity
  - Data size $O(M^2)$

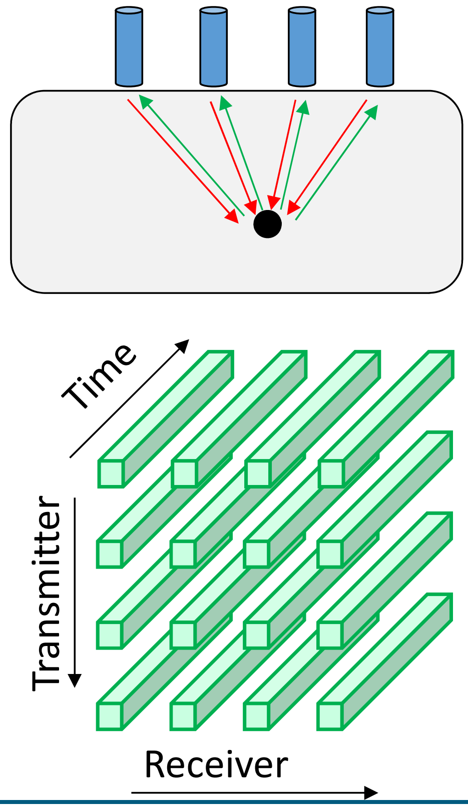

A particularly relevant usecase: Full Matrix Capture

- Elements transmit sequentially, receive in parallel
- For *M* elements:
  - *M* measurement cycles
  - M^2 A-Scans (M x M x N_T samples)
- Drawbacks:
  - *M* cycles need time
  - *M* parallel RF chains: hardware complexity
  - Data size O(M^2)
- Spatiotemporal compression
  - Due to strong spatial correlation: spatial subsampling as a valid CS strategy

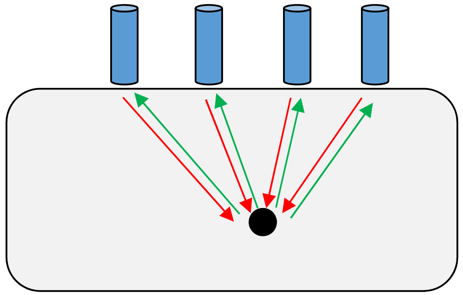

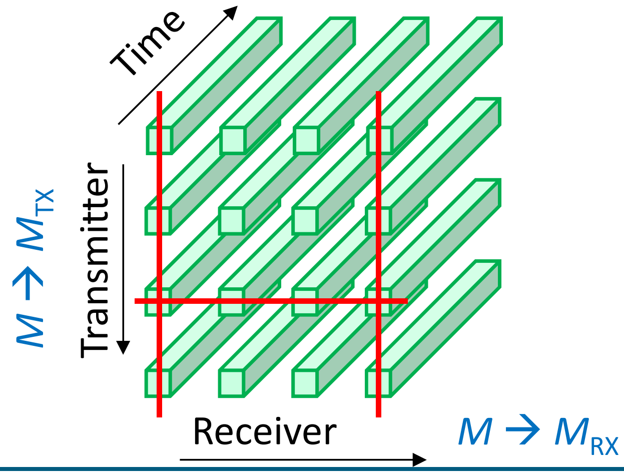

A particularly relevant usecase: Full Matrix Capture

- Elements transmit sequentially, receive in parallel
- For M elements:
  - M measurement cycles
  - M^2 A-Scans ($M \times M \times N_T$ samples)
- Drawbacks:
  - M cycles need time
  - M parallel RF chains: hardware complexity
  - Data size $O(M^2)$
- Spatiotemporal compression
  - Due to strong spatial correlation: spatial subsampling as a valid CS strategy
  - Sparsity in time => subsampling in frequency
- Overall: (M, M, N_T) => (M_{TX}, M_{RX}, N_F)

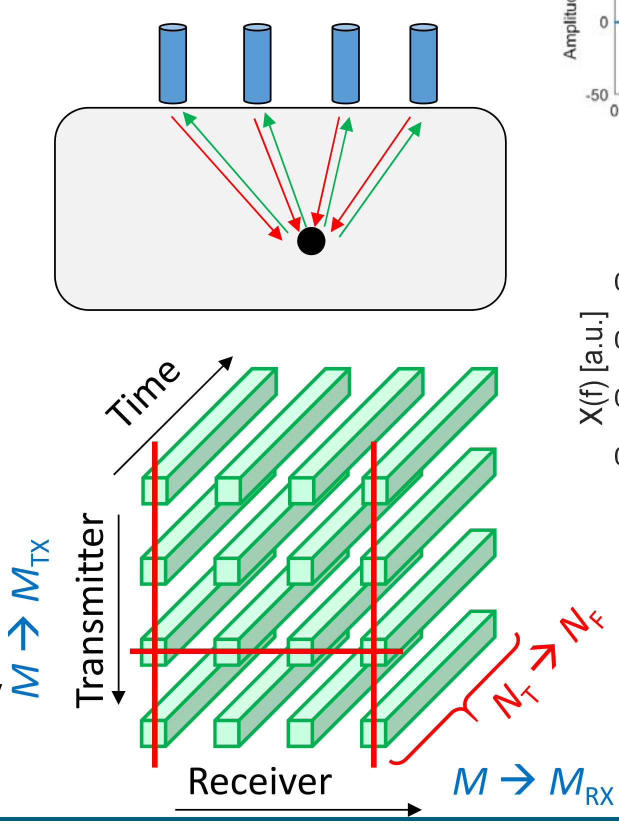

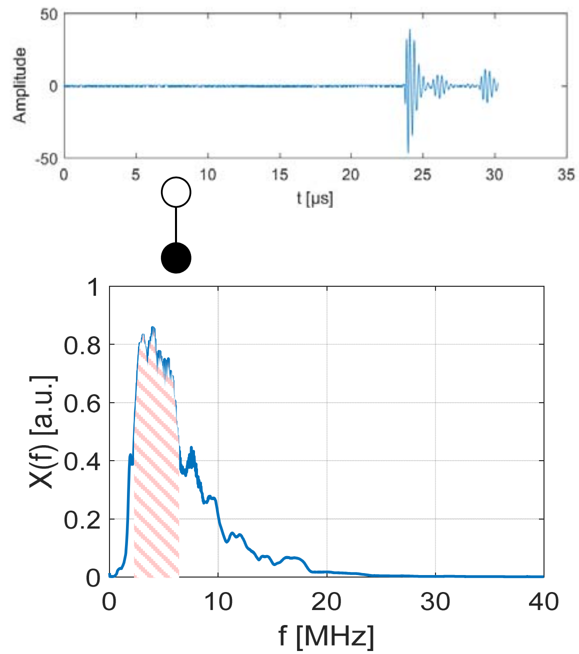

Spatiotemporal compression strategies

- 3-D **separable** design
  - For each selected TX, use the same RX channels and Fourier bins

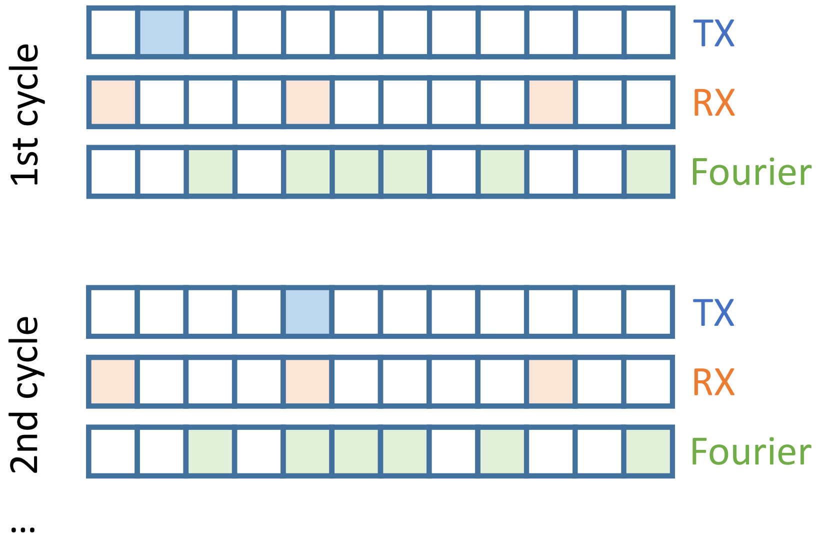

- 3-D **joint** design
  - For each selected TX, allow a new set of RX channels. For each TX and RX, allow new Fourier bins

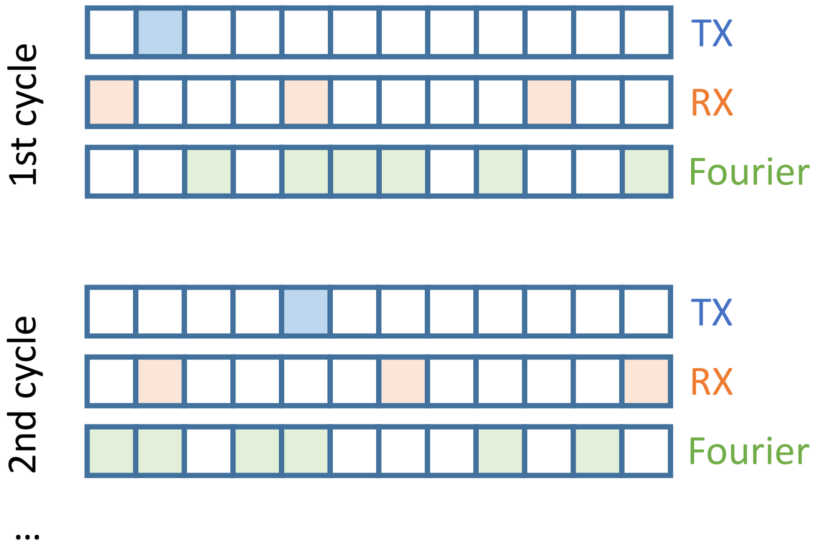

TECHNISCHE UNIVERSITÄT ILMENAU

Fraunhofer IZFP

How to solve the 3-D design of experiment problem?

It depends on the complexity…

Problem size M

Exhaustive search — Greedy search — Genetic Algorithms, Annealing — ML-based: Learn-2-Select, Deep Probabilistic Subsampling

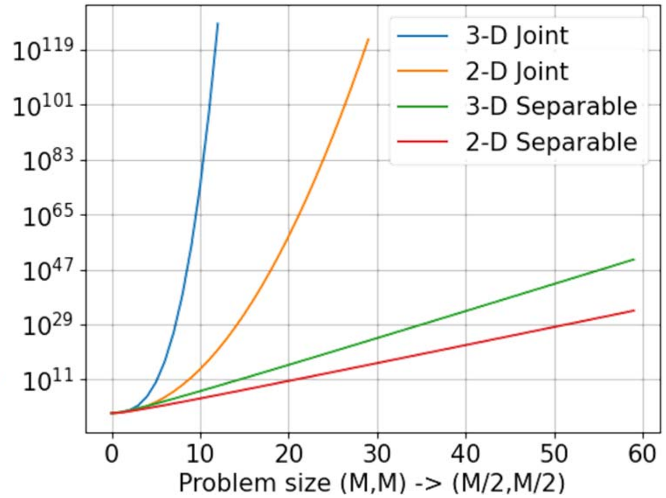

Separable design: $\binom{M}{M_T}\cdot\binom{M}{M_R}\cdot\binom{N_T}{N_F}$

Joint design $\binom{M}{M_T}\cdot\binom{M}{M_R}^{M_T}\cdot\binom{N_T}{N_F}^{M_T M_R}$

- Criterion to choose?
  - Many choices: task-based end2end, image quality, coherence, …
  - Our of our favorites: information theory based criterion: **Cramer-Rao Bound**
    - Quantifies ultimate **uncertainty** (variance) about scattering location
  - For point-wise criteria: Minmax approach => find pattern with best worst-case performance

$$\min_{m_{RX}, m_{TX}, n_T} \max_{\{x,z\}\in \mathrm{ROI}} CRB$$

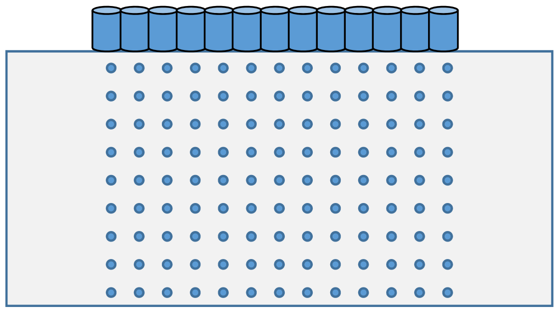

TECHNISCHE UNIVERSITÄT ILMENAU
Fraunhofer IZFP

Design of Experiment via Cramér-Rao Bounds (CRBs)

- Treat the problem as estimation of point scattering locations (x,z): CRB delivers algorithm-independent bound (“information content”)
  - Forward model: $\boldsymbol{y} = f(x_\mathrm{s}, z_\mathrm{s})$
  - Estimator: $\hat{x}_\mathrm{s}, \hat{z}_\mathrm{s} = T(\boldsymbol{y})$
  - Unbiasedness: $$\mathbb{E}\{\hat{x}_\mathrm{s}\} = x_\mathrm{s}$$ $$\mathbb{E}\{\hat{z}_\mathrm{s}\} = z_\mathrm{s}$$
  - Then:

$$\mathrm{COV}\left\{[x_\mathrm{s}, z_\mathrm{s}]^\mathrm{T}\right\} \succeq \boldsymbol{C} = \boldsymbol{J}^{-1}$$

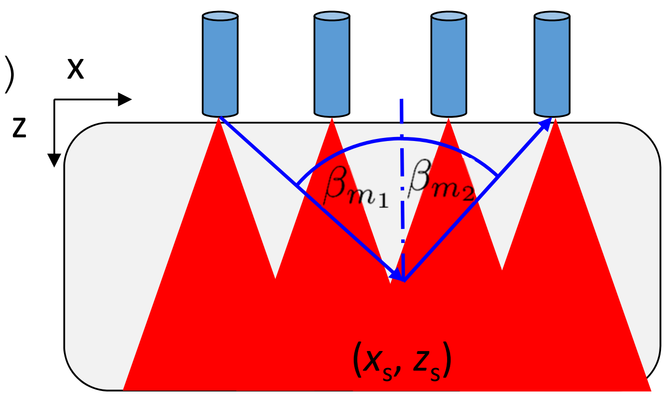

- We can show that

$$\boldsymbol{J} = C \cdot \sum_{m_1, m_2} \begin{bmatrix} (\sin\beta_{m_1} + \sin\beta_{m_2})^2 & (\sin\beta_{m_1} + \sin\beta_{m_2})(\cos\beta_{m_1} + \cos\beta_{m_2}) \\ (\sin\beta_{m_1} + \sin\beta_{m_2})(\cos\beta_{m_1} + \cos\beta_{m_2}) & (\cos\beta_{m_1} + \cos\beta_{m_2})^2 \end{bmatrix} \cdot G^2(m_1, m_2)$$
$$+ C_2 \cdot \sum_{m_1, m_2} \begin{bmatrix} \dot{G}_x(m_1, m_2)^2 & \dot{G}_x(m_1, m_2) \cdot \dot{G}_z(m_1, m_2) \\ \dot{G}_x(m_1, m_2) \cdot \dot{G}_z(m_1, m_2) & \dot{G}_z(m_1, m_2)^2 \end{bmatrix}$$

- CRB can be used to build spatial subselection strategies (exhaustive search / greedy)

Sample result: DOE through CRB minimization

Exhaustive search

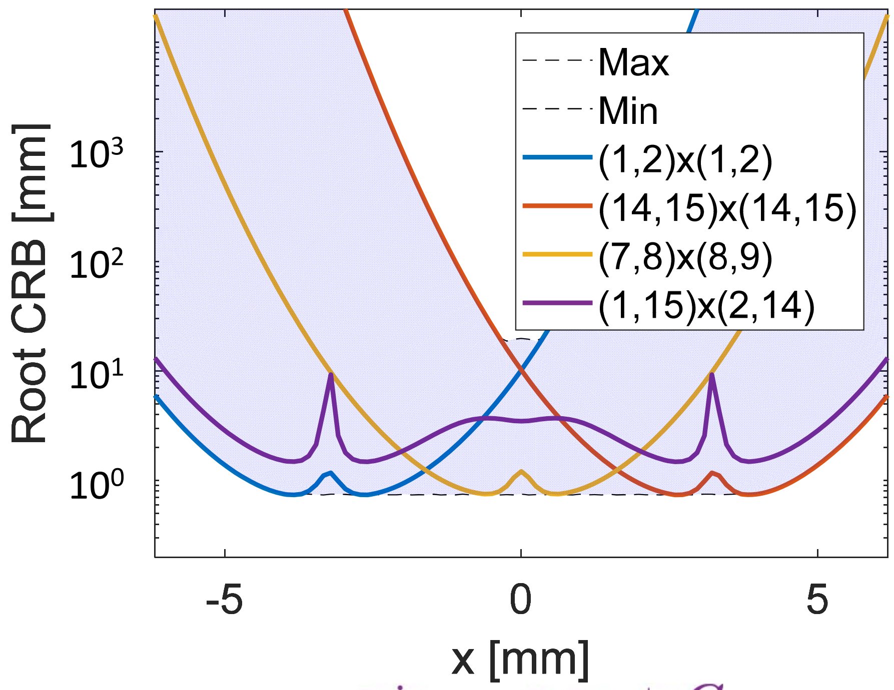

$$\min_{\{m_{\mathrm{TX}}\},\{m_{\mathrm{RX}}\}} \max_{x_{\mathrm{s}},z_{\mathrm{s}}} \operatorname{tr} \boldsymbol{C}$$

- M = 15 elements
  - choose 2 TX, 2 RX (50x compression)
- 1 mm pitch
- c = 5900 m/s
- Gaussian pulse T = 2 µs, P_{N} = 0.1
- Single defect z_{s} = 15 mm
- Gaussian apodization

$$G_m(x,z) = \mathrm{e}^{-\frac{\left((x-x_m)^2+z^2\right)^2}{z^2 \tan^2 \theta}}$$

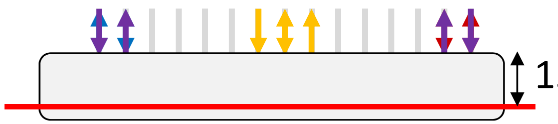

TECHNISCHE UNIVERSITÄT ILMENAU

Fraunhofer IZFP

Sample result: DOE through CRB minimization

Exhaustive search

$$\min_{\{m_{\mathrm{TX}}\},\{m_{\mathrm{RX}}\}} \max_{x_{\mathrm{s}},z_{\mathrm{s}}} \operatorname{tr} \boldsymbol{C}$$

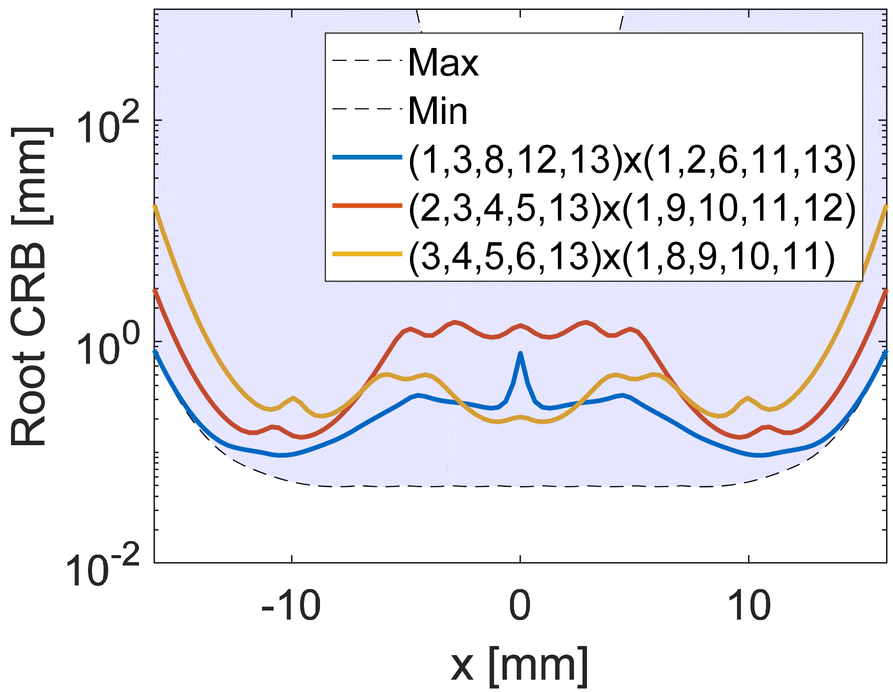

- M = 13 elements
  - choose 5 TX, 5 RX (7x compression)
- 2 mm pitch
- c = 5900 m/s
- Gaussian pulse T = 2 µs, P_{N} = 0.01
- Single defect z_{s} = 15 mm
- Gausssian apodization

$$G_m(x,z) = \mathrm{e}^{-\frac{\left((x-x_m)^2+z^2\right)^2}{z^2\tan^2\theta}}$$

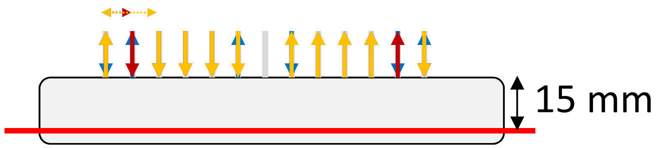

TECHNISCHE UNIVERSITÄT ILMENAU Fraunhofer IZFP

Greedy selection approach in 2-D

- Separable 2-D design

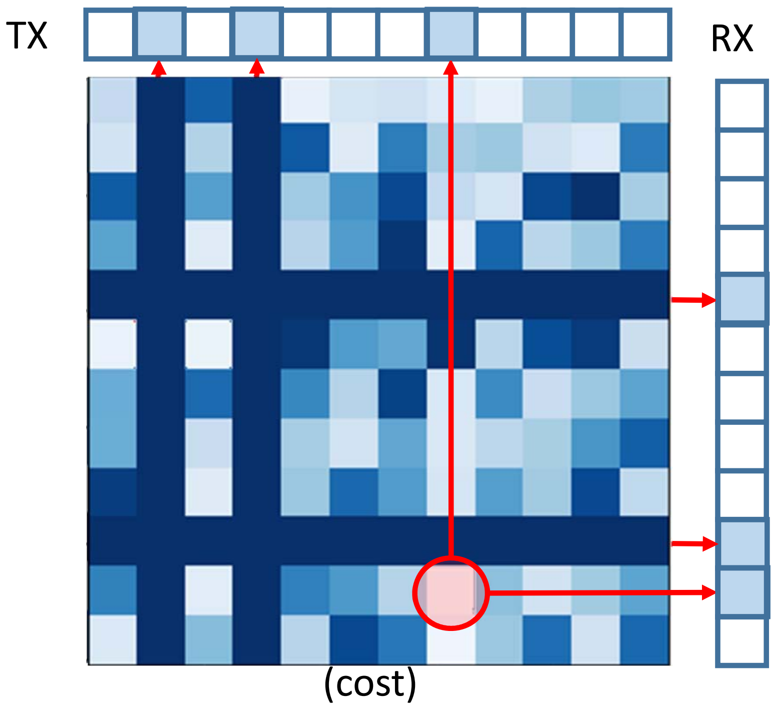

- Joint 2-D design

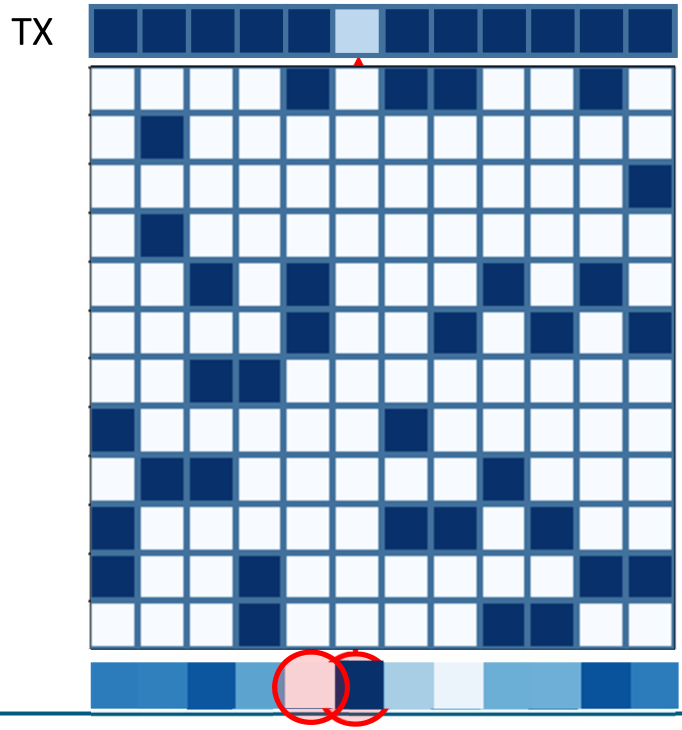

How to evaluate?

Criterion: Point Spread Function

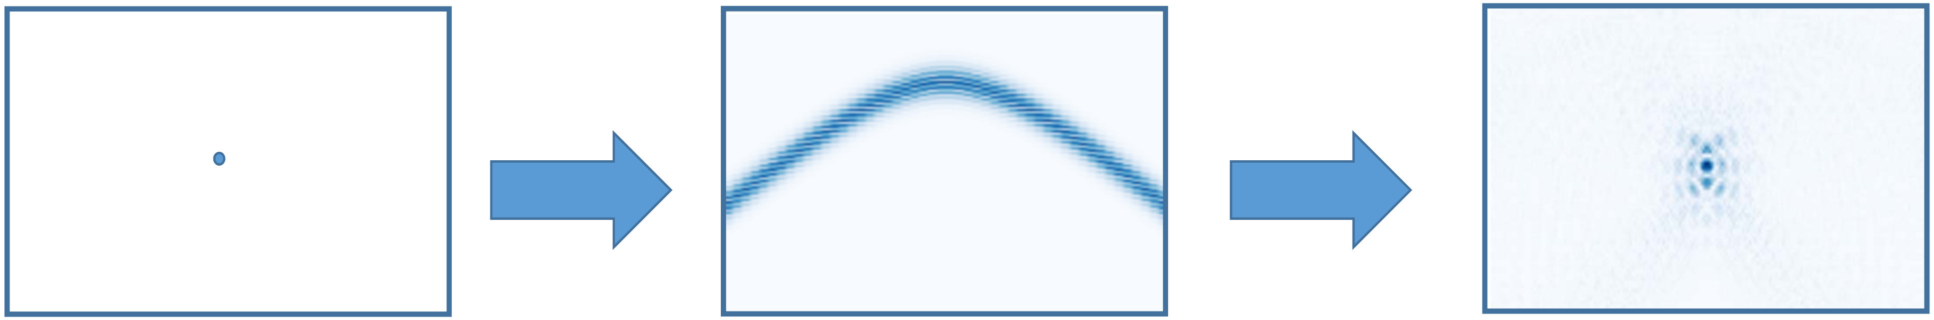

1) Generate ideal point scatterer (x,z) 2) Generate B-Scan (forward model) 3) CS measurement + image reconstruction

→ Width of the PSF provides a measure of achievable resolution (in a linear imaging device)
  → Cross sections to study vertical / horizontal resolution
  → Area under the PSF = „Array Performance Indicator (API)“

$$\mathrm{API} = \frac{N_{\mathrm{thresh}} \cdot \Delta_x \cdot \Delta_z}{\lambda_c^2}$$

TECHNISCHE UNIVERSITÄT ILMENAU
Fraunhofer IZFP

Simulation results

Single defect (to evaluate PSF)

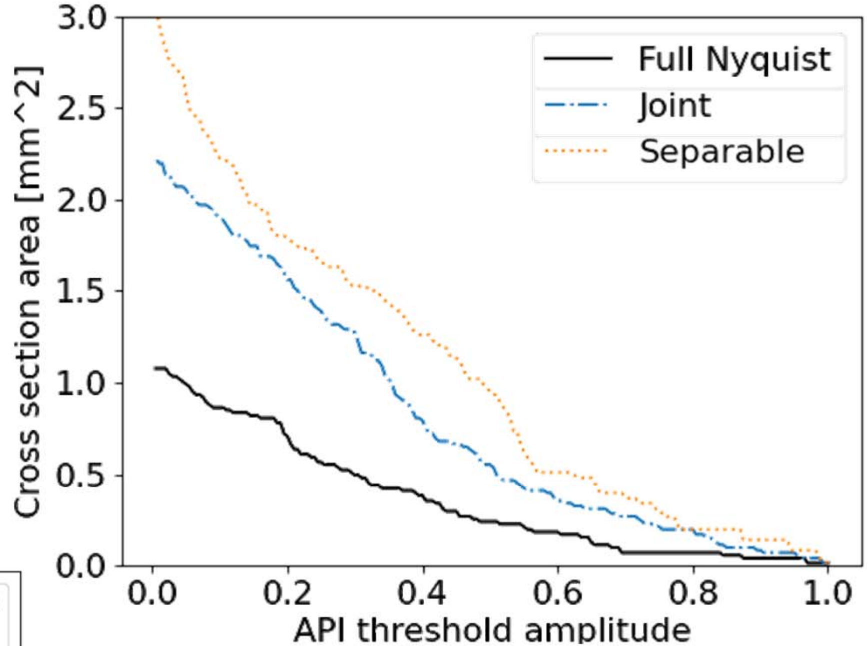

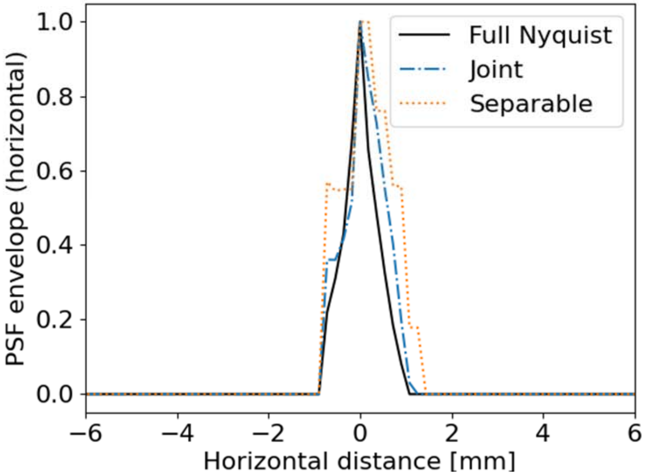

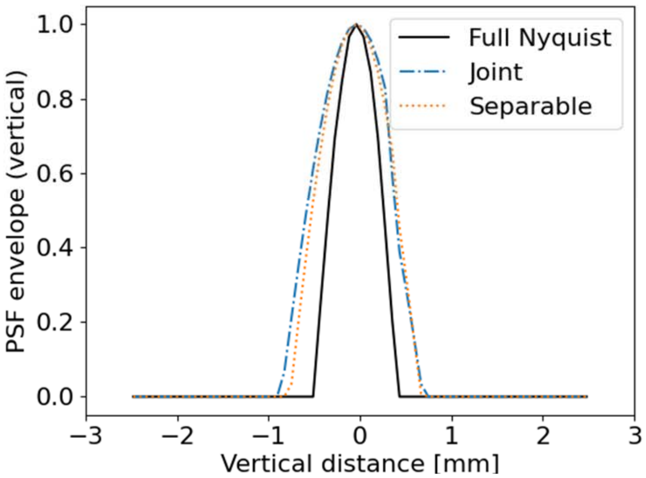

$N_T = 1525 \Rightarrow N_F = 31$ $c = 5900$ m/s

$M = 16 \Rightarrow M_T = 3$ $f = 4.54$ MHz

$M = 16 \Rightarrow M_R = 4$ $\lambda = 1.9$ mm

TECHNISCHE UNIVERSITÄT ILMENAU Fraunhofer IZFP

Measurement results: test specimen

Linear Scale
Log Scale

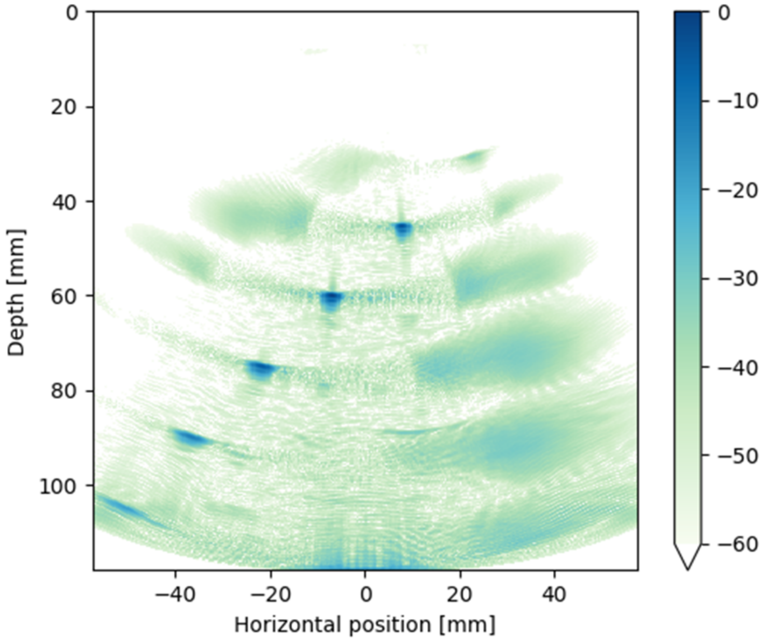

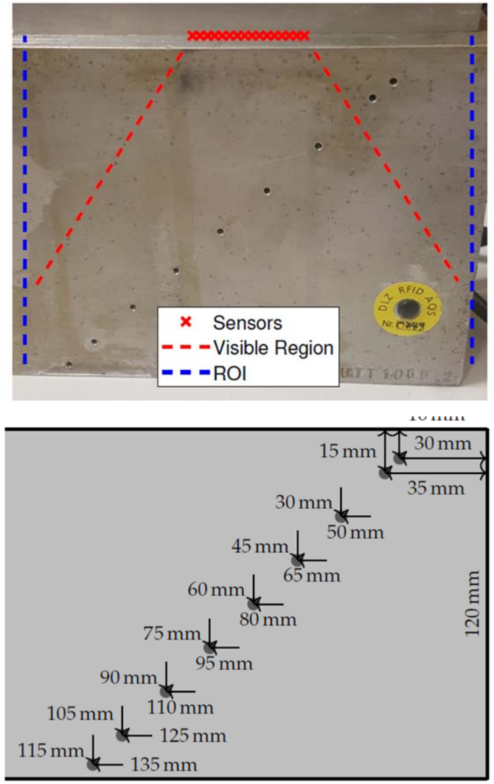

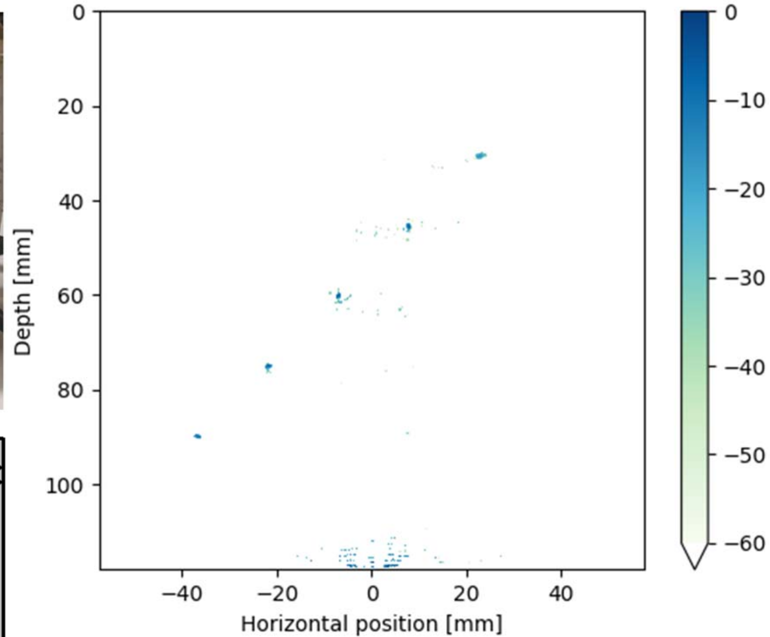

No compression
(100% of the data)
TFM

>12x spatial comp.
>7x temporal comp.
(1% of the data)
FISTA

Outlook: DOE via Machine Learning

- **Realistic** problem **size**s (M = 64, 128, …): exhaustive search fails (exponential complexity), greedy methods fail (local minima)
- Idea: **leverage** power of **neural networks** to find good solutions to the non-convex problem
- Problem: selection problem is combinatorial, **non-differentiable**
- Solution: **Gumbel soft-max reparametrization trick**
  - Rows of Phi: samples of the categorical distribution
- **Cost** function: final image **reconstruction** error!
- Estimation method: **Learned ISTA**

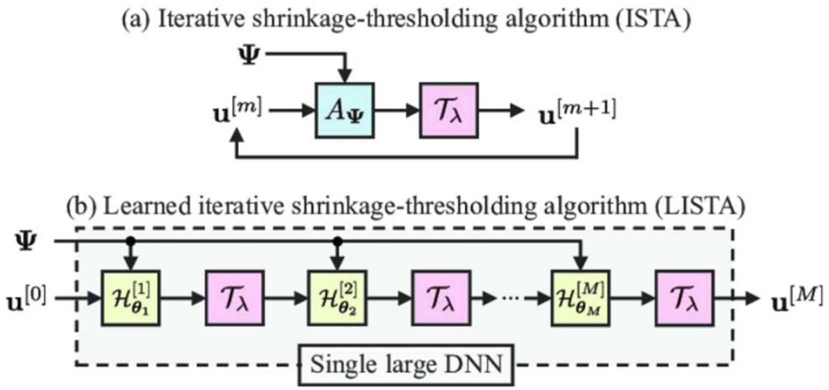

Figure reproduced from Masuyaka et. al. 2020, LISTA originally due to LeCun et al 2010.

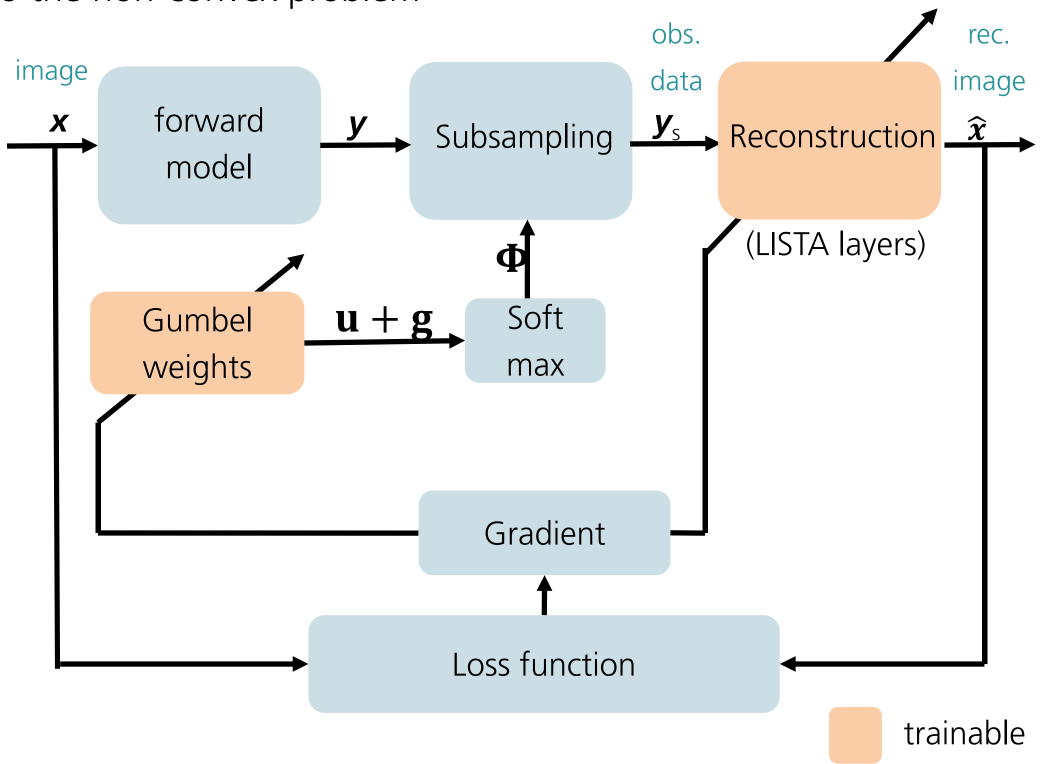

TECHNISCHE UNIVERSITÄT ILMENAU

Fraunhofer IZFP

Outlook: DOE via Machine Learning

Comparison with other designs

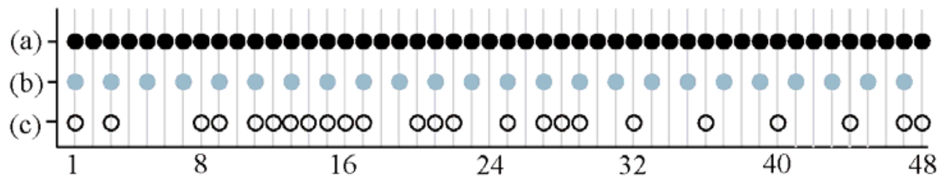

Note: 48 choose 24 = 32.2 trillion

- Comparing reconstructed images performance

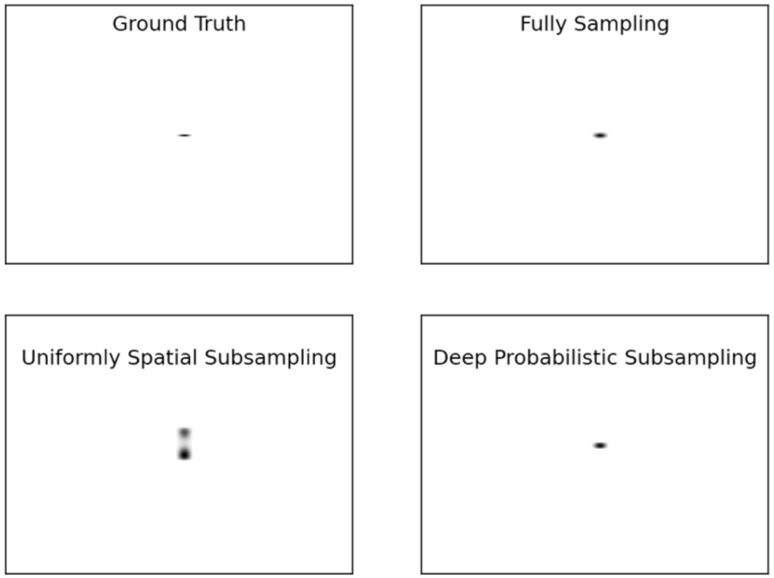

Fig. 2. Comparison of reconstructed images

- Evaluating CNR of 1000 images

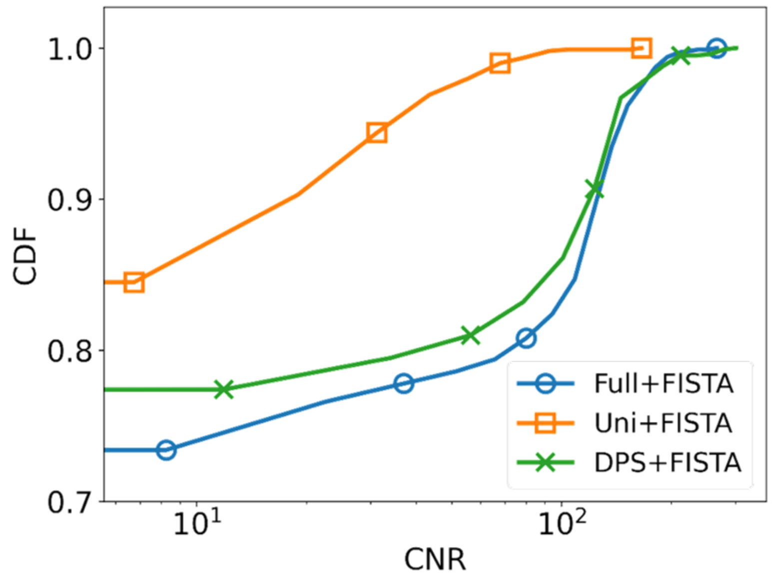

Fig. 3. Comparison w.r.t. CNR

Summary

- Compressed Sensing for NDT: very **promising**!
  - Reduction of scan time, data rate, power consumption

- Successfully **demonstrated** in **X-ray CT** (inline inspection)

- **Ultrasound**:
  - data reduction in time and space
  - real-time reconstruction for assistance systems

- Challenges **ahead**, e.g.,
  - Complex **forward models**, e.g., Ultrasound tomography
  - **data**-driven **+ physics**-driven processing
  - Lack of training data: **artificial training data** generation