



Article

Frequency-Constrained QR: Signal and Image Reconstruction

Harrison Garrett  and David G. Long * 

Electrical and Computer Engineering Department, Brigham Young University, Provo, UT 84602, USA; hcgarret@byu.edu

* Correspondence: long@ee.byu.edu

Abstract: Because a finite set of measurements is limited in the amount of spectral content it can represent, the reconstruction process from discrete samples is inherently band-limited. In the case of 1D sampling using ideal measurements, the maximum bandwidth of regular and irregular sampling is well known using Nyquist and Gröchenig sampling theorems and lemmas, respectively. However, determining the appropriate reconstruction bandwidth becomes difficult when considering 2D sampling geometries, samples with variable apertures, or signal to noise ratio limitations. Instead of determining the maximum bandwidth a priori, we derive an inverse method to simultaneously reconstruct a signal and determine its effective bandwidth. This inverse method is equivalent to incrementally computing a band-limited inverse using a frequency-constrained QR decomposition (FQR). Comparisons between reconstruction results using FQR and QR decompositions illustrate how FQR is less sensitive to noisy measurement errors, but it is more sensitive to high-frequency components. These methods are particularly useful in the reconstruction of remote sensing images from such as microwave radiometers and scatterometers.

Keywords: pseudo-inverse; band-limited; resolution enhancement; variable apertures

1. Introduction

Signal reconstruction from a finite set of measurements is a common problem found in signal and image processing. When measurements are modeled as idealized delta-like samples, the Nyquist sampling theorem [1] and Gröchenig's sampling lemma [2] can be used to determine a reconstruction's maximum spectral band limit in the 1D sampling case. However, when dealing with measurements with variable apertures, SNR limitations, or 2D sampling considerations, finding the reconstructable bandwidth can be difficult [3].

To address these reconstruction issues, we review the concept of a band-limited sampling matrix, its inverse, and an iterative method to determine a reconstruction's reconstructable bandwidth as treated in [4]. In this paper, we present an inverse method for simultaneously reconstructing a signal and determining its effective bandwidth, i.e., the maximum amount of the recoverable signal bandwidth without significant noise amplification. This work is motivated by the need for finer resolution from microwave remote sensors such as radiometers and scatterometers, which have irregular sampling with variable apertures [5–8].

We first consider the need for a band-limited inverse by reviewing how a discrete sampling matrix is inherently band-limited in order to find a unique, nonaliased solution. We also review how noise is amplified and how model aliasing occurs in the computation of a sampling matrix inverse. We then show how to incrementally invert a sampling matrix's band-limited content and derive a method for computing an incremental band-limited



Academic Editor: Paolo Pampaloni

Received: 21 November 2024

Revised: 15 January 2025

Accepted: 26 January 2025

Published: 29 January 2025

Citation: Garrett, H.; Long, D.G. Frequency-Constrained QR: Signal and Image Reconstruction. *Remote Sens.* **2025**, *17*, 464. <https://doi.org/10.3390/rs17030464>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

inverse using a frequency-constrained QR decomposition (FQR). We then compare the difference in reconstruction using FQR and QR decompositions.

2. Materials and Methods

2.1. Discrete Sampling

To motivate the need for a bandlimited inverse, we consider how finite measuring systems are inherently band-limited. We start by defining an ideal finite measurement as a sample of a continuous linear system,

$$z_i = \int m_i(r) a(r) dr, \quad (1)$$

where z_i is the i th discrete measurement, $m_i(r)$ is the measurement's response function (MRF), $a(r)$ is the truth signal, and r is a positional index. Given a finite number of measurements N , inverting the measurement system in Equation (1) is an ill-posed problem, since the existence, uniqueness, and stability of a solution depend on the configuration, spacing, and number of $m_i(r)$, as well as assumptions about $a(r)$.

To illustrate the difficulty in estimating $a(r)$ from a set of z_i , imagine if each $m_i(r)$ is a delta function centered at position r_i (i.e., $m_i(r) = \delta(r - r_i)$) and then $z_i = a(r_i)$. Although each z_i may accurately represent the value of the truth signal at the point r_i , it is impossible to uniquely infer the entire continuous signal ($a(r)$) without further information, since an infinite number of signals may also cross at the same discrete values. However, if $a(r)$ is assumed to be band-limited, sufficient sampling configurations can uniquely determine $a(r)$. Under band-limited sampling theory, only a periodic band-limited signal $a(r)$ can be perfectly reconstructed from a finite set of samples [3].

We define a band-limited signal $a(r)$ to be a signal which only possesses frequency content below a maximum cutoff frequency, or band limit B (i.e., the Fourier transform of $a(r)$ has a finite region of support). Let $A(\omega)$ be the Fourier transform of $a(r)$, $A(\omega)$ that is band-limited if $|A(\omega)| = 0 \forall \omega > B$. According to Nyquist, a uniformly sampled band-limited signal can be perfectly reconstructed from a set of ideal discrete samples $a[n]$ as

$$a(r) = \sum_n a[n] d(r - r_n), \quad (2)$$

where $a[n] = a(r_n)$, $d(r)$ is a Dirichlet kernel (i.e., band-limited identity function) corresponding to the band limit of $a(r)$, and n represents an arbitrary sample index corresponding to position r_n when the sample distances between r_n is at least twice the highest frequency. Since $d(r)$ is a periodic band-limited identity, $a(r)$ is periodic. Note that Equation (2) is equivalent to bandlimited interpolation. Even though multiple sampling configurations (i.e., the particular spacing and number of sample positions r_n) can satisfy the reconstruction in Equation (2), all such configurations are equivalent under the band-limited assumption [3]. For convenience, we can assume that r_n is spaced on a uniform grid, where the grid spacing is $\Delta r \leq 1/(2f_c)$ with f_c as the highest frequency within the band limit of $a(r)$ (i.e., $1/(2f_c)$ is the Nyquist sampling distance).

The relationship between z_i and $a[n]$ can be found by combining Equations (1) and (2),

$$z_i = \int m_i(r) \left[\sum_n a[n] d(r - r_n) \right] dr,$$

$$z_i = \sum_n \left[\int m_i(r) d(r - r_n) dr \right] a[n], \quad (3)$$

$$z_i = \sum_n m_i[n] a[n], \quad (4)$$

$$m_i[n] = \int m_i(r) d(r - r_n) dr, \quad (5)$$

where $m_i[n]$ is a discrete band-limited sample of $m_i(r)$, and $d(r - r_n)$ is the Dirichlet kernel function centered at r_n . Note that $m_i[n] \neq m_i(r_n)$, unless $m_i(r)$ is already band-limited within $d(r)$'s corresponding bandlimit.

Equation (3) can be expressed in vector and matrix form as

$$\begin{aligned} z_i &= \mathbf{m}_i^T \mathbf{a}, \\ \mathbf{z} &= \mathbf{M} \mathbf{a}, \end{aligned} \quad (6)$$

where \mathbf{m}_i and \mathbf{a} are length 'N' vectors of $m_i[n]$ and $a[n]$, respectively, \mathbf{z} is a length M vector of z_i , and \mathbf{M} is a matrix with rows of \mathbf{m}_i^T and is $[M \times N]$, where M is the total number of z_i , and N is the total number of r_n . From Equation (6), it can be seen that the band-limited reconstruction process is equivalent to finding the inverse of \mathbf{M} ; however, the invertibility of \mathbf{M} depends on the chosen band-limited discretization and spacing of each MRF.

In the case of uniformly spaced samples and constant MRF apertures, i.e., $m_i = m_j$, \mathbf{M} is square and invertible (assuming each $m_i[n]$ is defined over the assumed band limit). When measurements are dependent, i.e., they do not increase the rank of the inverse of \mathbf{M} , they do not improve the resolution of the estimate but do provide redundancy for noise reduction. Irregular sampling and variable apertures tend to have more dependent measurements. A unique band-limited solution only exists if a band-limited inverse of \mathbf{M} exists.

In order to define the band-limited inverse of \mathbf{M} , it is helpful to study the discretization of ideal and band-limited measurements. Under band-limited assumptions, the discrete equivalent to an ideal measurement is found by substituting $m_i(r) = d(r - r_i)$ into Equation (5):

$$m_i[n] = \int d(r - r_i) d(r - r_n) dr = d(r_n - r_i) = d_i[n], \quad (7)$$

where $d_i[n]$ is a Dirichlet kernel centered at r_i that is discretely sampled over the grid defined by r_n .

To express Equation (2) in matrix form, we define a band-limited identity matrix \mathbf{D} constructed with rows $d(r - r_n)$ sampled over the grid defined by r_n . The multiplication of \mathbf{a} by \mathbf{D} is an identity operation, as \mathbf{a} is unchanged by the band limit of \mathbf{D} . This band-limited identity operation, $\mathbf{a} = \mathbf{D} \mathbf{a}$, is the discrete band-limited reconstruction equivalent to Equation (2). Using this result, we can express a band-limited inverse of \mathbf{M} :

$$\begin{aligned} \mathbf{D} &= \mathbf{W} \mathbf{M}, \\ \mathbf{a} &= \mathbf{W} \mathbf{z}, \end{aligned} \quad (8)$$

where \mathbf{W} is a band-limited inverse of \mathbf{M} . The existence of \mathbf{W} depends on \mathbf{M} being full rank over a band-limited subspace. When this criteria is met, \mathbf{W} is the SVD pseudo-inverse of \mathbf{M} . When \mathbf{M} is full row rank, \mathbf{W} is the Moore–Penrose pseudo-inverse, and \mathbf{D} becomes the identity matrix. The existence of \mathbf{W} depends on the sampling positions r_i , the choice

of discrete positions r_n , and the band limit. For a more rigorous study on how sampling conditions affect the existence of this band-limited inverse, see [1,4,9,10].

In summary, an ideal band-limited reconstruction can be accomplished from a finite number of measurements if a band-limited inverse of \mathbf{M} exists, where \mathbf{M} is the discrete band-limited sampling matrix of all measurement response functions.

2.2. Background

2.2.1. Sampling Matrix Inversion

The process of estimating a discrete signal \mathbf{a} from a discrete set of measurements \mathbf{z} is equivalent to estimating a particular inverse of the sampling matrix \mathbf{M} . The estimated signal can be found by using the sampling inverse of \mathbf{M} ,

$$\hat{\mathbf{a}} = \mathbf{M}^\dagger \mathbf{z}, \quad (9)$$

where $\hat{\mathbf{a}}$ is the estimated signal, and \mathbf{M}^\dagger is a unique inverse of the sampling matrix \mathbf{M} . While there may be a plethora of matrix inversion methods that could be used to estimate \mathbf{M}^\dagger , their various constraints, assumptions, and approximations can affect the quality of the reconstructed signal $\hat{\mathbf{a}}$.

2.2.2. Inverse Error Amplification

In order to better understand noise amplification and reconstruction artifacts, we consider the additive noise model:

$$\tilde{\mathbf{z}} = \mathbf{M}\mathbf{a} + \boldsymbol{\nu}, \quad (10)$$

where $\tilde{\mathbf{z}}$ is a set of noisy measurements, and $\boldsymbol{\nu}$ is measurement noise. Assuming \mathbf{M}^\dagger exists, let

$$\mathbf{a}_\nu = \mathbf{M}^\dagger \boldsymbol{\nu}, \quad (11)$$

where \mathbf{a}_ν is the inverse noise. Let

$$\mathbf{a}_\mu = (\mathbf{I} - \mathbf{M}^\dagger \mathbf{M})\mathbf{a}, \quad (12)$$

where \mathbf{a}_μ is the aliasing of inverse model artifacts,

$$\tilde{\mathbf{a}} = \mathbf{a} + \mathbf{a}_\mu + \mathbf{a}_\nu, \quad (13)$$

and $\tilde{\mathbf{a}}$ is the estimated signal corrupted by amplified noise and model artifacts.

In the ideal case, \mathbf{M}^\dagger is a perfect inverse of \mathbf{M} , i.e., $\mathbf{M}^\dagger \mathbf{M} = \mathbf{I}$ and $\mathbf{a}_\mu = \mathbf{0}$, such that there are no model errors or reconstruction artifacts. However, when sampling model errors are introduced into the computation of \mathbf{M}^\dagger , $\mathbf{M}^\dagger \mathbf{M} \neq \mathbf{I}$, the content from \mathbf{a} can be aliased in \mathbf{a}_μ . Additionally, since \mathbf{M}^\dagger is sensitive to small errors, computational errors can be amplified such that features in \mathbf{a}_ν can overpower recoverable signal features.

2.2.3. Deconvolution Example

To illustrate how errors are amplified and aliased in the inversion of a sampling matrix, we examine the case of discrete deconvolution using a shift invariant sampling function \mathbf{m} . The measurement vector is

$$\begin{aligned} \mathbf{z} &= \mathbf{m} \circledast \mathbf{a} \\ &= \mathbf{M}\mathbf{a}, \end{aligned} \quad (14)$$

$$\mathbf{M} = \begin{bmatrix} m_0 & m_1 & \dots & m_{N-1} \\ m_{N-1} & m_0 & \vdots & m_{N-2} \\ m_{N-2} & m_{N-1} & \ddots & m_{N-3} \\ \vdots & \vdots & \ddots & \vdots \\ m_1 & m_2 & \dots & m_0 \end{bmatrix} \quad (15)$$

where \mathbf{z} is the convolution of vectors \mathbf{m} and \mathbf{a} , (\circledast) is the circular convolution operator, and \mathbf{M} is a circulant matrix of \mathbf{m} . Because \mathbf{M} is a circulant matrix, it can be factored using the Discrete Fourier Transform (DFT) matrix \mathbf{F} [11]:

$$\begin{aligned} \mathbf{M} &= \mathbf{F}^* (\mathbf{F}\mathbf{m} \cdot \mathbf{F}) \\ &= \mathbf{F}^* \text{diag}(\mathbf{F}\mathbf{m})\mathbf{F} \end{aligned} \quad (16)$$

where $*$ denotes the conjugate transpose, (\cdot) is an element-wise multiplier, $\text{diag}(\mathbf{x})$ denotes forming a diagonal matrix with the vector \mathbf{x} , and \mathbf{F} is normalized such that $\mathbf{F}^*\mathbf{F} = \mathbf{I}$.

The deconvolution of \mathbf{m} can likewise be represented in matrix form. In the ideal case, the deconvolution of \mathbf{m} is \mathbf{w} , where

$$\mathbf{w} = \mathbf{F}^* (\mathbf{F}\mathbf{m})^{-1}. \quad (17)$$

The deconvolution then is

$$\begin{aligned} \tilde{\mathbf{a}} &= \mathbf{w} \circledast \mathbf{z} \\ &= \mathbf{W}\mathbf{z}, \\ &= \mathbf{W}\mathbf{M}\mathbf{a}, \end{aligned} \quad (18)$$

where \mathbf{W} is a circulant matrix of \mathbf{w} , and \mathbf{W} is also the ideal inverse of \mathbf{M} . Then, we have the following:

$$\begin{aligned} \mathbf{W} &= \mathbf{F}^* \text{diag}(\mathbf{F}\mathbf{m})^{-1} \mathbf{F}, \\ &= \mathbf{M}^{-1}. \end{aligned} \quad (19)$$

However, the deconvolution of \mathbf{m} is sensitive to small spectral values within $\mathbf{F}\mathbf{m}$. The inversion of small spectral values is a common cause of noise amplification. When the deconvolution does not match or equalize the spectrum of $\mathbf{F}\mathbf{m}$, aliasing of spectral features can occur as follows:

$$\mathbf{a}_\mu = \mathbf{F}^* \text{diag}(1 - (\mathbf{F}\mathbf{w} \cdot \mathbf{F}\mathbf{m}))\mathbf{F}\mathbf{a} \quad (20)$$

where \mathbf{a}_μ defines how spectral features are aliased by deconvolution.

To demonstrate noise amplification and model error aliasing, a simple Gaussian MRF was used to blur a 300 sample truth signal with 3 Dirichlet kernels of different degrees. The deconvolution of this synthetically blurred scene was performed twice. The first deconvolution was done after noise ν added to the scene in order to achieve a 30 dB SNR. The second deconvolution was done after adding corruption to the Gaussian MRF to achieve a 10 dB SNR of corruption. Both deconvolutions are shown in Figure 1a. The FFT of

the ideal Gaussian MRF, i.e., $\mathbf{F}_w \cdot \mathbf{F}_m$ alongside its equalization at an expected 30 dB SNR level and equalization using the corrupted MRF, are shown in Figure 1b. Note that even a small amount of error, without any conditioning, significantly distorted the equalization such that $1 - (\mathbf{F}_w \cdot \mathbf{F}_m)$ amplified the noise floor to a visible level.

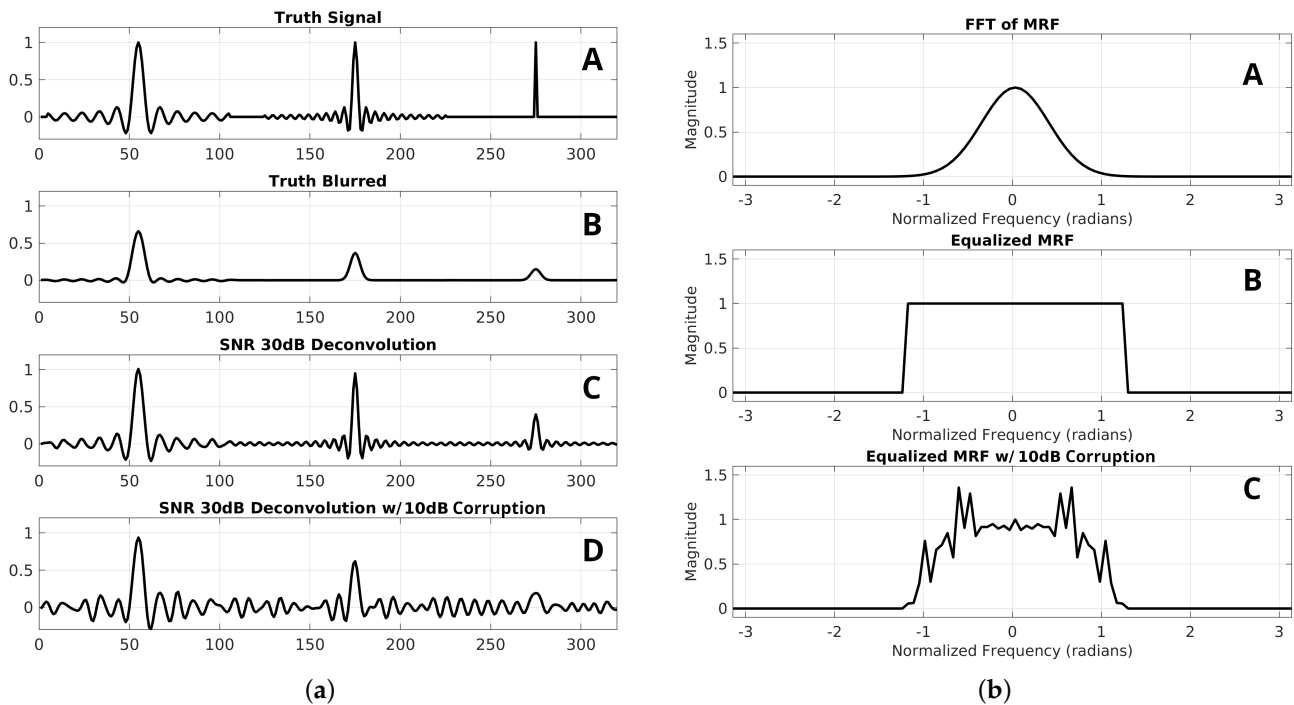


Figure 1. Deconvolution example. (a) An example deconvolution. (A) Truth signal of two Dirichlet kernels and one Kronecker delta function. (B) Truth signal blurred by a Gaussian MRF function. (C) Deconvolution result after equalization with a band-limited deconvolution on a scene with 30 dB SNR additive noise. (D) Deconvolution result with corrupted MRF with 10 dB SNR of corruption. Note that noise limits the resolving capability and amplifies the noise floor of the deconvolution. (b) The equalizations of the two deconvolution examples shown in Figure 1a. (A) DFT of the Gaussian MRF function. (B) Equalization using a band-limited inverse Gaussian at 30 dB SNR. (C) Equalization after adding 10 dB SNR of corruption to the Gaussian MRF function. The ideal equalization is perfectly flat at magnitude 1, corresponding to an ideal band-limited reconstruction. However, the corrupted Gaussian equalization is distorted, corresponding to the amplification and attenuation of different content.

Regularization techniques can be used to avoid inverting small errors in the frequency domain, lessening the sensitivity of deconvolution to small errors. However, regularization inherently changes the assumed measurement response spectrum and distorts the reconstructed signal. To preserve the recoverable measurement response spectrum, the sampling matrix can instead be band-limited such that deconvolution only inverts a subset of the recoverable measurement response spectrum, minimizing the effects of small errors.

2.2.4. Band-limited Inverse

To mitigate the amplification of noise and reconstruction artifacts, band-limited filtering can be used to remove frequency content within the model before it is aliased in the calculation of the matrix inverse. One way to accomplish this is by band-limiting each row of \mathbf{M} (i.e., taking its DFT, applying a band-limiting filter by multiplication, then taking the inverse DFT) and then calculating the pseudo-inverse of \mathbf{M} . Assuming that the same band-limiting filter is applied to each row, we can express this process in linear algebra terms by recognizing that a band-limiting frequency filter \mathbf{h} can also be applied by

a band-limited matrix $\mathbf{D} = \mathbf{F}^* \text{diag}(\mathbf{h})\mathbf{F}$. We define the right-hand band-limited inverse of \mathbf{M} bandlimited by \mathbf{D} as

$$\mathbf{M}^\beta = (\mathbf{M}\mathbf{D})^\dagger, \quad (21)$$

where \mathbf{M}^β is the band-limited inverse, and \mathbf{D} is a band-limiting matrix that applies the frequency filter \mathbf{h} to each row of \mathbf{M} .

While there may not be a computational advantage to computing the band-limited inverse using Equation (21), since band-limiting with a matrix multiplication is $\mathcal{O}(n^3)$, and contrastingly, the DFT is $\mathcal{O}(n^2 \log n)$, Equation (21) can help us understand how a band-limited inverse is related to the ideal pseudo-inverse, for instance, $\mathbf{M}^\beta = \mathbf{M}^{-1}$ when \mathbf{M} is full rank and \mathbf{D} is an identity matrix (maximum discrete bandwidth). $\mathbf{M}^\beta = \mathbf{M}^\dagger$ when \mathbf{M} is not full rank and \mathbf{D} is a band-limited identity matrix (i.e., when \mathbf{D} is a Toeplitz Dirichlet kernel matrix) with a higher bandwidth than each row of \mathbf{M} . However, when \mathbf{D} is lower rank than \mathbf{M} , i.e., a lower bandwidth than the rows of \mathbf{M} , then $\mathbf{M}^\beta \neq \mathbf{M}^\dagger$.

Ideally, to preserve model content and remove noise, a band-limiting filter matching the recoverable model spectrum should be used. However, since the recoverable model spectrum depends on measurement sampling locations and MRF bandwidths, it can be hard to determine for variable aperture measurements a priori. One approach to teasing out the recoverable spectrum of variable aperture measurements is suggested in [3], which involves incrementally bandlimiting a sampling matrix and repetitively calculating its rank.

Instead of treating the bandlimiting and inverse steps sequentially, it becomes advantageous to compute a band-limited matrix and matrix inverse simultaneously. In this paper, we derive a method for calculating an iterative band-limited inverse, which we later show is equivalent to a partial QR decomposition using a frequency basis (see Appendix A).

2.3. Frequency QR

Here, we define a method for computing a band-limited inverse using single-rank updates. When computed in full, this method is equivalent to performing QR decomposition on a frequency basis (FQR). When computed sequentially, it can progressively construct a band-limited inverse from low to high frequencies, stopping when the magnitude of each frequency component falls below the assumed SNR noise floor in the frequency domain. This stops the reconstruction at the point where measurement noise theoretically overtakes the magnitude of reconstructable content. In the noise-free case, these criteria stop the reconstruction at the inherent sampling band limit without going over. An additional example comparing the result of FQR to that of a QR decomposition is provided in the following section.

2.3.1. Partitioned Band-Limited Matrix

Let \mathbf{f}_n represent an orthonormal basis vector of the frequency domain such that $\mathbf{f}_i^* \mathbf{f}_j = 0, \forall i \neq j$.

$$\mathbf{F} = [\mathbf{f}_0, \mathbf{f}_1, \dots, \mathbf{f}_{N-1}], \quad (22)$$

$$\mathbf{F}^* \mathbf{F} = \mathbf{I},$$

where \mathbf{F} is a matrix of the basis vectors \mathbf{f}_n and \mathbf{I} is the identity matrix. Typically, each \mathbf{f}_n corresponds to a DFT kernel within the calculation of the DFT. However, other frequency bases may be used when advantageous. More examples of frequency bases and their advantages are discussed in a later section.

A band-limiting frequency filter can be applied in the frequency domain by applying the filter weights h_n to each corresponding basis vector \mathbf{f}_n . Taking advantage of the orthogonal basis, \mathbf{MD} can be partitioned into a set of singular matrices:

$$\begin{aligned}\mathbf{MD} &= \mathbf{MF}\text{diag}(\mathbf{h})\mathbf{F}^*, \\ \mathbf{MD} &= (\mathbf{B}_0 + \mathbf{B}_1 + \cdots + \mathbf{B}_{N-1}),\end{aligned}\quad (23)$$

where $\text{diag}(\mathbf{x})$ denotes the formation of a diagonal matrix with \mathbf{x} along its diagonal, \mathbf{h} is a frequency filter, and each \mathbf{B}_n is a singular matrix defined by the outer product:

$$\begin{aligned}\mathbf{B}_n &= \mathbf{M}\mathbf{f}_n h_n \mathbf{f}_n^*, \\ \mathbf{B}_n &= \mathbf{b}_n \mathbf{f}_n^*,\end{aligned}\quad (24)$$

where h_n are the values of \mathbf{h} corresponding to each basis vector \mathbf{f}_n , and \mathbf{b}_n is a column vector of \mathbf{MF} weighted by h_n :

$$\mathbf{b}_n = \mathbf{M}\mathbf{f}_n h_n. \quad (25)$$

Note that the matrix product $\mathbf{B}_i \mathbf{B}_j^* = 0, \forall i \neq j$, because the row space of each \mathbf{B}_n is orthogonal. However, the column spaces of \mathbf{B}_n are not orthogonal, meaning that $\mathbf{B}_i^* \mathbf{B}_j \neq 0, \forall i \neq j$.

Using the partitioned representation of \mathbf{MD} , we can express the progression of band-limited sampling matrices $(\mathbf{MD})_n$ as

$$(\mathbf{MD})_n = \sum_{k=0}^n \mathbf{b}_k \mathbf{f}_k^*, \quad (26)$$

where each $(\mathbf{MD})_n$ progressively increases in bandwidth as n increases. The direct inversion of each $(\mathbf{MD})_n$ leads us to the existence of a progressive band-limited inverse:

$$\begin{aligned}\mathbf{M}_n^\beta &= (\mathbf{MD})_n^\dagger, \\ \mathbf{M}_n^\beta &= (\mathbf{B}_0 + \mathbf{B}_1 + \cdots + \mathbf{B}_n)^\dagger.\end{aligned}\quad (27)$$

Regrettably, while each \mathbf{B}_n may have orthogonal row spaces, the progressive band-limited inverse cannot be formed by the sum of \mathbf{B}_n^\dagger , i.e.,

$$\mathbf{M}_n^\beta \neq \sum_{k=0}^n \mathbf{B}_k^\dagger, \quad (28)$$

since each \mathbf{B}_n does not have orthogonal column spaces. However, given that sequential $(\mathbf{MD})_n$ only differ by singular matrix updates, it stands to reason that \mathbf{M}_n^β can also be partitioned into singular band-limited matrices.

2.3.2. Derivation

Our derivation of the progressive band-limited inverse \mathbf{M}_n^β begins with the following theorem, which relates the generalized inverse of the sum of two matrices to the generalized inverse of each individual matrix [12]:

$$\begin{aligned}\text{if } \text{rank}(\mathbf{X} + \mathbf{Y}) &= \text{rank}(\mathbf{X}) + \text{rank}(\mathbf{Y}), \\ \text{then } (\mathbf{X} + \mathbf{Y})^\dagger &= (\mathbf{I} - \mathbf{S})\mathbf{X}^\dagger(\mathbf{I} - \mathbf{T}) + \mathbf{S}\mathbf{Y}^\dagger\mathbf{T}, \\ \mathbf{S} &= (\mathbf{P}_Y^*[\mathbf{I} - \mathbf{P}_X^*])^\dagger, \\ \mathbf{T} &= ([\mathbf{I} - \mathbf{P}_X]\mathbf{P}_Y)^\dagger,\end{aligned}\quad (29)$$

where \mathbf{X} and \mathbf{Y} are two arbitrary matrices, \mathbf{P}_X denotes the projection onto the column space of \mathbf{X} , \mathbf{S} represents the disjoint portion of the row space of \mathbf{Y} from the row space of \mathbf{X} , and \mathbf{T} represents the disjoint portion of the column space of \mathbf{Y} from the column space of \mathbf{X} .

Recursively applying the theorem in Equation (29), we can find each progressive band-limited inverse. Applying the theorem to \mathbf{B}_n and the zero matrix, we trivially find the initial bandlimited inverse:

$$\mathbf{M}_0^\beta = \mathbf{B}_0^\dagger = \frac{\mathbf{f}_0 \mathbf{b}_0^*}{\mathbf{b}_0^* \mathbf{b}_0}. \quad (30)$$

The next bandlimited inverse can be found by performing a rank update to \mathbf{B}_0^\dagger using \mathbf{B}_1 :

$$\begin{aligned} \text{given } \text{rank}(\mathbf{B}_0 + \mathbf{B}_1) &= \text{rank}(\mathbf{B}_0) + \text{rank}(\mathbf{B}_1), \\ \text{then } (\mathbf{B}_0 + \mathbf{B}_1)^\dagger &= (\mathbf{I} - \mathbf{S}_1) \mathbf{B}_0^\dagger (\mathbf{I} - \mathbf{T}_1) + \mathbf{S}_1 \mathbf{B}_1^\dagger \mathbf{T}_1, \\ \mathbf{S}_1 &= (\mathbf{f}_1 \mathbf{f}_1^* (\mathbf{I} - \mathbf{f}_0 \mathbf{f}_0^*))^\dagger = \mathbf{f}_1 \mathbf{f}_1^*, \\ \mathbf{T}_1 &= \left(\begin{bmatrix} \mathbf{I} - \frac{\mathbf{b}_0 \mathbf{b}_0^*}{\mathbf{b}_0^* \mathbf{b}_0} & \mathbf{b}_1 \mathbf{b}_1^* \\ \mathbf{b}_1^* \mathbf{b}_0 & \mathbf{b}_1^* \mathbf{b}_1 \end{bmatrix} \right)^\dagger, \end{aligned} \quad (31)$$

where \mathbf{T}_1 is the disjoint portion of the column space of \mathbf{B}_1 from \mathbf{B}_0 , and \mathbf{S}_1 is the row space projection of \mathbf{B}_1 . As the row space of each \mathbf{B}_1 is already disjoint $\mathbf{S}_1 \mathbf{B}_1^\dagger = \mathbf{B}_1^\dagger$ and $(\mathbf{I} - \mathbf{S}_1) \mathbf{B}_0^\dagger = \mathbf{B}_0^\dagger$, upon simplifying using this result, we find the following:

$$\mathbf{M}_1^\beta = (\mathbf{B}_0 + \mathbf{B}_1)^\dagger = \mathbf{M}_0^\beta (\mathbf{I} - \mathbf{T}_1) + \mathbf{B}_1^\dagger \mathbf{T}_1, \quad (32)$$

where \mathbf{M}_1^β can be found using \mathbf{M}_0^β , \mathbf{T}_1 , and \mathbf{B}_1^\dagger . Applying the theorem to the n th rank update, we can form a recursive relationship, which iteratively constructs a band-limited inverse:

$$\begin{aligned} \mathbf{M}_n^\beta &= \mathbf{M}_{n-1}^\beta (\mathbf{I} - \mathbf{T}_n) + \mathbf{B}_n^\dagger \mathbf{T}_n, \\ \mathbf{T}_n &= \left(\begin{bmatrix} \mathbf{I} - \mathbf{P}_{(\mathbf{MD})_{n-1}} & \mathbf{b}_n \mathbf{b}_n^* \\ \mathbf{b}_n^* \mathbf{b}_n & \mathbf{b}_n^* \mathbf{b}_n \end{bmatrix} \right)^\dagger, \end{aligned} \quad (33)$$

where \mathbf{M}_n^β can be found by applying a rank update to \mathbf{M}_{n-1}^β with the disjoint column space projection of each \mathbf{B}_n^\dagger .

To find \mathbf{T}_n recursively, we use the fact that each \mathbf{T}_n 's column space is disjoint and $\{\mathbf{T}_0, \mathbf{T}_1, \dots, \mathbf{T}_n\}$ spans the column subspace of $(\mathbf{MD})_n$ to re-express $(\mathbf{I} - \mathbf{P}_{(\mathbf{MD})_{n-1}})$ with the sum of the disjoint column spaces of \mathbf{T}_n :

$$\begin{aligned} (\mathbf{I} - \mathbf{P}_{(\mathbf{MD})_{n-1}}) &= (\mathbf{I} - \mathbf{P}_{\mathbf{T}_{n-1}})(\mathbf{I} - \mathbf{P}_{\mathbf{T}_{n-2}}) \dots (\mathbf{I} - \mathbf{P}_{\mathbf{T}_0}), \\ (\mathbf{I} - \mathbf{P}_{(\mathbf{MD})_{n-1}}) &= \left(\mathbf{I} - \sum_{k=0}^{n-1} \mathbf{P}_{\mathbf{T}_k} \right). \end{aligned} \quad (34)$$

By substitution, we find a recursive relationship to calculate each singular matrix \mathbf{T}_n :

$$\begin{aligned} \mathbf{T}_n &= \left(\begin{bmatrix} \mathbf{I} - \sum_{k=0}^{n-1} \mathbf{P}_{\mathbf{T}_k} & \mathbf{b}_n \mathbf{b}_n^* \\ \mathbf{b}_n^* \mathbf{b}_n & \mathbf{b}_n^* \mathbf{b}_n \end{bmatrix} \right)^\dagger, \\ \mathbf{T}_n &= \mathbf{b}_n \left(\mathbf{b}_n - \sum_{k=0}^{n-1} \mathbf{P}_{\mathbf{T}_k} \mathbf{b}_n \right)^\dagger, \\ \mathbf{T}_n &= \mathbf{b}_n \mathbf{t}_n^*, \end{aligned} \quad (35)$$

where \mathbf{t}_n^* is the vector inverse of \mathbf{q}_n

$$\mathbf{q}_n = \mathbf{b}_n - \sum_{k=0}^{n-1} \frac{\mathbf{t}_k \mathbf{t}_k^*}{\mathbf{t}_k^* \mathbf{t}_k} \mathbf{b}_n, \quad (36)$$

and \mathbf{q}_n is a column space basis vector of \mathbf{M}_n^β . Note that this recursive relationship is equivalent to the Gram–Schmidt process used in QR factorization [13–15].

Substituting $\mathbf{T}_n = \mathbf{b}_n \mathbf{t}_n^*$ into the recursive relationship for \mathbf{M}_n^β in Equation (33), we have the following:

$$\mathbf{M}_n^\beta = \mathbf{M}_{n-1}^\beta (\mathbf{I} - \mathbf{b}_n \mathbf{t}_n^*) + \mathbf{B}_n^\dagger \mathbf{b}_n \mathbf{t}_n^*, \quad (37)$$

which reveals that the row space of \mathbf{M}_n^β is spanned by orthogonal \mathbf{t}_n . As each band-limited update is singular, we can solve for the column space vector \mathbf{m}_n^\dagger , which updates each \mathbf{M}_n^β . By multiplying both sides of Equation (37) by \mathbf{q}_n and noting that $\mathbf{B}_n^\dagger \mathbf{b}_n = \mathbf{f}_n$, we find an expression for \mathbf{m}_n^\dagger :

$$\begin{aligned} \mathbf{m}_n^\dagger &= \mathbf{M}_n^\beta \mathbf{q}_n, \\ \mathbf{m}_n^\dagger &= \mathbf{f}_n + \mathbf{M}_n^\beta [\mathbf{q}_n - \mathbf{b}_n], \\ \mathbf{m}_n^\dagger &= \mathbf{f}_n - \mathbf{M}_n^\beta \left[\sum_{k=0}^{n-1} \mathbf{q}_k \mathbf{t}_k^* \mathbf{b}_n \right], \\ \mathbf{m}_n^\dagger &= \mathbf{f}_n - \sum_{k=0}^{n-1} \mathbf{m}_k^\dagger \mathbf{t}_k^* \mathbf{b}_n. \end{aligned} \quad (38)$$

Note that Equation (38) is equivalent to the process of forward substitution used to solve a matrix inverse using a QR decomposition. Each $\mathbf{m}_n^\dagger \mathbf{t}_n^*$ forms a band-limited singular matrix partition of \mathbf{M}_n^β , which—when summed together—form progressive band-limited inverses:

$$\mathbf{M}_n^\beta = \sum_{k=0}^n \mathbf{m}_k^\dagger \mathbf{t}_k^*. \quad (39)$$

It is important to note that although the row space of each $\mathbf{m}_n^\dagger \mathbf{t}_n^*$ is orthogonal, and the column space (spanned by \mathbf{m}_n^\dagger) is not orthogonal, which is similar to the inequality in Equation (28). Specifically, each \mathbf{m}_n^\dagger may contain content from previous \mathbf{m}_k^\dagger , which are removed through forward substitution each band-limited update. This means that each band-limited inverse \mathbf{M}_n^β is not necessary equivalent to the pseudo-inverse of \mathbf{M} band-limited by each previous \mathbf{B}_n ,

$$\mathbf{M}_n^\beta \neq \left[\sum_{k=0}^n \mathbf{B}_k \right] \mathbf{M}^\dagger, \quad (40)$$

since \mathbf{M}^\dagger may contain content associated with higher-frequency features than contained in $[\sum_{k=0}^n \mathbf{B}_k]$. Therefore, $\mathbf{M}_n^\beta = [\sum_{k=0}^n \mathbf{B}_k] \mathbf{M}^\dagger$ only when \mathbf{B}_n has a higher bandwidth than the maximum bandlimit of \mathbf{M} .

As pointed out in the derivation, the vectors \mathbf{t}_n and \mathbf{m}_n^\dagger can be solved for using Gram–Schmidt orthogonalization and forward substitution methods, respectively. These are the basic steps for QR decomposition and inverse solutions. Appendix A shows how FQR is equivalent to applying a frequency rotation to a matrix and then solving for its partial inverse using QR iterations. A summary of the algorithm steps matching the notation given

in this paper is shown in Algorithm 1. Note that many algorithmic improvements can be made to make the orthogonalization and forward substitution methods more numerically stable and efficient [16,17]

Algorithm 1 Frequency-constrained QR.

$$\begin{aligned} \mathbf{f}_i^* \mathbf{f}_j &= 0, \quad \mathbf{f}_i^* \mathbf{f}_i = 1, \quad \forall i \neq j \\ \mathbf{b}_n &\leftarrow \mathbf{M} \mathbf{f}_n h_n \\ \mathbf{q}_n &\leftarrow \mathbf{b}_n - \sum_{k=0}^{n-1} \mathbf{q}_k \mathbf{t}_k^* \mathbf{b}_n \\ \mathbf{t}_n &\leftarrow \mathbf{q}_n / (\mathbf{q}_n^* \mathbf{q}_n) \\ \mathbf{m}_n^\dagger &\leftarrow \mathbf{f}_n - \sum_{k=0}^{n-1} \mathbf{m}_k^\dagger \mathbf{t}_k^* \mathbf{b}_n \\ \mathbf{M}_n^\beta &\leftarrow \mathbf{M}_{n-1}^\beta + \mathbf{m}_n^\dagger \mathbf{t}_n^* \end{aligned}$$

2.3.3. Observations

When computed over all N nonzero singular updates, this method is identical to taking a full inverse. In theory, this progressive band-limited inverse can be used as an inversion method for any arbitrary matrix \mathbf{X} . However, the computational complexity of this method is no better than using QR decomposition and forward substitution methods using $\mathcal{O}(n^3)$.

The main difference between this inversion method and other inversion methods is the order in which mutually independent information is inverted. For comparison, we describe how the SVD and QR decompositions break up a matrix inverse. The computation of the inverse using SVD decomposition breaks a matrix down starting with its largest eigenvectors, cutting off the inversion when eigenvalues become too small. The computation of the inverse using QR decomposition breaks a matrix down starting with its first column, cutting off the inversion when new column's contributions become too small. The computation of the inverse using FQR breaks a matrix down starting with its lowest frequency projection, cutting off the inversion when higher frequency's contributions become too small. All three methods (SVD, QR, and FQR) are essentially rotations of each other. In the noise-free case, all three of these methods are equivalent when run to completion, but in the presence of computation noise or model errors, the methods possess different sensitivities.

As a sampling matrix is inherently band-limited, the FQR inverse is advantageous to limit the inversion of noise content outside the sampling bandlimit.

2.3.4. SNR Filtering

When the inherent bandlimit of \mathbf{M} is unknown, but the expected noise floor or signal-to-noise ratio (SNR) is known, FQR can be modified to stop inversion when the magnitude of new \mathbf{t}_n drops below the expected noise level. This is very similar to the process of Wiener filtering [18], except that no regularization is added to the inversion process. As the expected SNR is decreased, the recoverable bandwidth shrinks and FQR iterations reach the raised expected noise level earlier. Further SNR filtering can be implemented by adding a weighting or taper on the magnitude of each \mathbf{t}_n . Such a taper can be used to mitigate the Gibbs phenomenon present at the edge of band-limited signals.

3. Results

To demonstrate the difference between QR and FQR decompositions, we iteratively performed reconstruction on two example 1D truth signals. To contrast both decompositions, two simple synthetic signals of 101 samples were generated and used as truth for each simulation. The first signal consisted of three low-bandwidth peaks, and the second

signal resembles a rect function with higher frequency trends. The average magnitude of the second signal was set to $100\times$ that of the first signal.

In order to form a variable aperture sample matrix \mathbf{M} , 101 MRFs were synthetically generated selecting by shifting three MRFs apertures centered at random indices. The three MRF apertures were one Gaussian centered with a standard deviation of 25 indices, one left skewed Gaussian approximately 13 indices wide, and one skewed right that mirrored the previous MRF. Each MRF was normalized to sum to 1. The same sampling matrix, a 101×101 matrix, was used on each input truth signal to create two corresponding sets of noise free measurements. To both of these measurement sets, additive Gaussian noise was added to achieve a 30 dB SNR. Example MRFs are shown in Figure 2 to visualize the centered and skewed apertures.

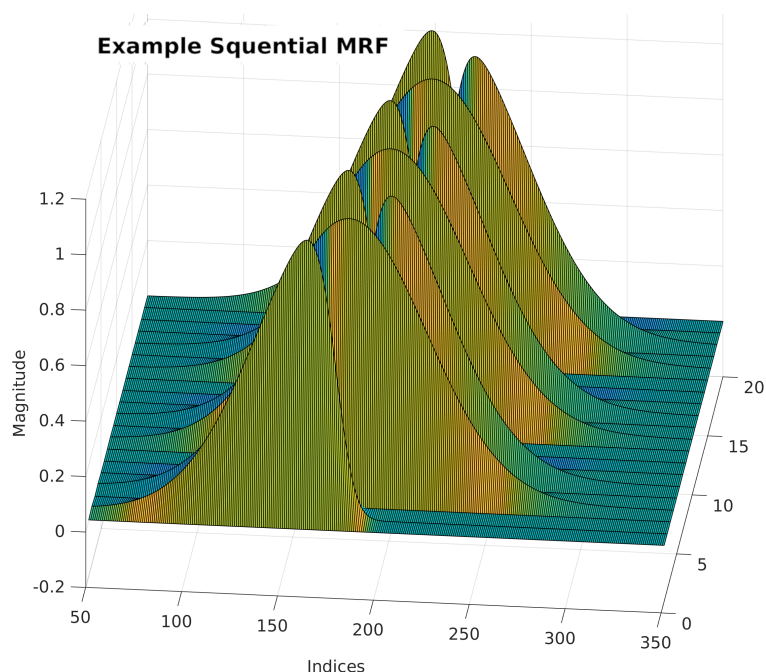


Figure 2. Selected portion of the sampling matrix using 3 randomly shifted MRFs. The MRFs were $4\times$ upsampled and renormalized to make them easier to visualize in this illustration. Each MRF was normalized to sum to 1.

In general the main difference between the reconstruction using QR vs. FQR is that QR reconstructs onto a basis of orthogonal MRFs, while FQR reconstructs onto an orthogonal frequency basis. To illustrate this, three selected MRFs are shown in Figure 3 along with the low-order truth signal for reference. Next to this plot in Figure 3, an example QR reconstruction using the same three MRFs is shown. Additionally, a FQR reconstruction using the first three discrete cosine transform (DCT) bins of all MRFs is shown. Not surprisingly, the FQR took fewer iterations than the QR to represent a low-bandwidth signal, since FQR uses all MRFs to reconstruct each frequency basis, while QR only uses one MRF for each iteration.

The full reconstruction for the low bandwidth input signal was performed using all measurements and frequency iterations. A few sample QR reconstructions using only the first 5, 11, and 21 MRFs are shown in Figure 4a. Similarly, sample FQR reconstructions using only the first 5, 11, and 21 DCT bins are also shown in Figure 4a. The same process was repeated for the higher bandwidth input signal in Figure 4b. Note that in the QR decomposition, a single MRF could greatly bias the reconstruction due to the mismatch between MRF shape and truth signal features. Theoretically, if all iteration steps are taken, the bias caused by a single measurement is removed in future iteration steps. Likewise,

in FQR, a single frequency basis can bias the reconstruction due to sampling density sensitivities.




Figure 3. (A) Example MRFs overlaid with the truth signal. (B) The reconstruction from the first 3 columns of a QR decomposition (first 3 MRFs). (C) The reconstruction from the first 3 FQR DCT bins.

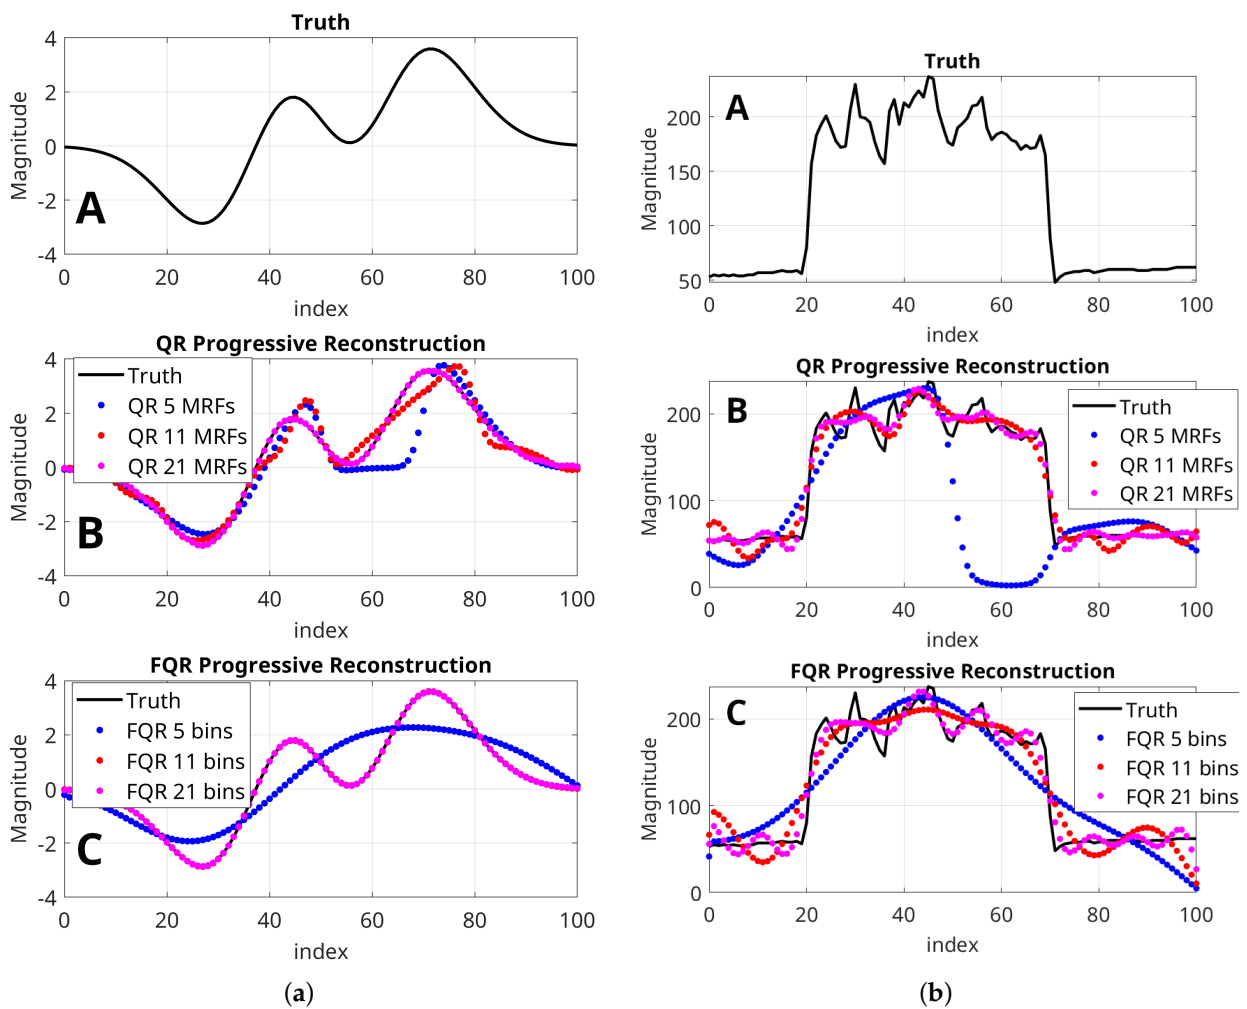


Figure 4. Partial reconstructions using QR and FQR decomposition. (a) (A) Example low bandwidth truth signal. (B) The reconstruction from the first 5, 11, and 21 columns of a QR decomposition. (C) The reconstruction from the first 5, 11, and 21 DCT bins of a FQR decomposition. Note that the band-limited QR reconstruction required less iterations to achieve similar results to QR. (b) (A) Example truth signal. (B) The reconstruction from the first 5, 11, and 21 columns of a QR decomposition. (C) The reconstruction from the first 5, 11, and 21 DCT bins of a FQR decomposition. Note that the FQR reconstruction uses all measurements' frequency response for each iteration, while conventional QR reconstructs one measurement at a time.

The total error with respect to each truth signal was calculated for each iteration for both the QR and FQR reconstruction methods. The total error per pixel for these four simulation are all shown on the same scale in Figure 5. Note that signal 2 is higher than signal 1 due to the intentional $100\times$ magnitude to plot both simulations on the same total error curves. In both cases, the FQR method was able to resolve more low-bandwidth content before the amplification of noise overtook the reconstruction process.

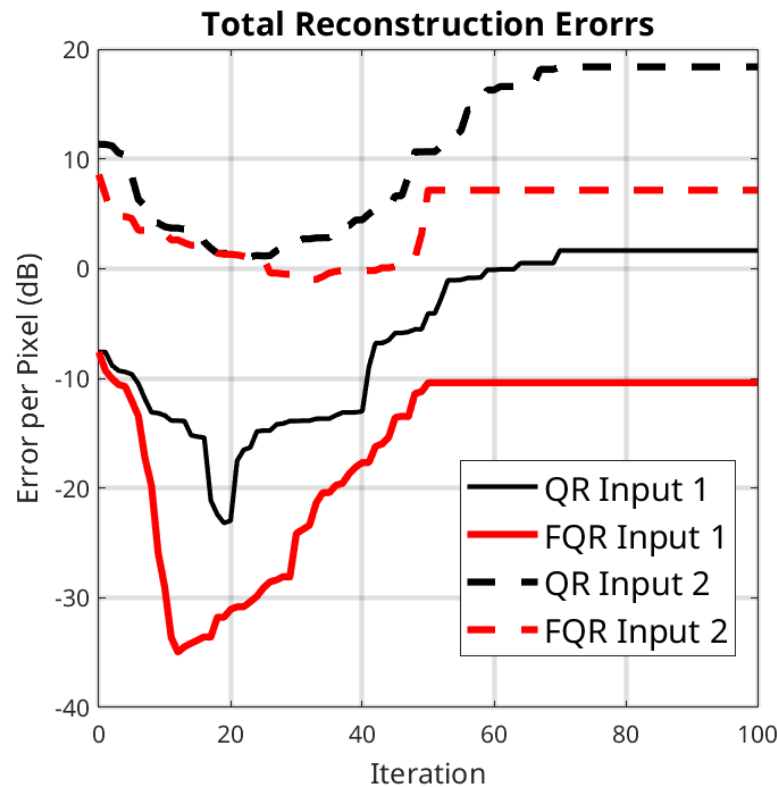


Figure 5. Total reconstruction error per pixel for each of the 4 simulations. Each simulation contained noise at a 30 dB SNR level. Input 1 is the low-bandwidth signal. Input 2 is the higher frequency signal, whose average magnitude was intentionally $100\times$ larger than Input 1 to plot these errors on the same scale. Note that in each error trend's error, there appears a minimum between reconstruction and noise amplification.

4. Discussion

4.1. Equivalence

Given that FQR is equivalent to performing QR decomposition after a frequency-basis rotation (see Appendix A), it stands to reason that other matrix inversion techniques can also be used to compute a band-limited inverse. For example, an iterative SVD decomposition could be used to compute band-limited inverses (see Appendix B). If inverting a noise-free and well-conditioned matrix, the frequency-basis rotation becomes redundant, as an identity matrix could be theoretically used as a trivial rotation. However, in the presence of noise or poor conditioning, the basis rotation can reorder and prioritize reconstructing signal content to minimize amplifying noise.

The prioritizing of reconstructed content in a QR inverse computation has been a wide field of study. Many rigorous and specialized algorithms have been developed to solve very large or sparse matrices [19,20]. To extend the discussion of how the choice of frequency basis can prioritize signal content differently, we review several frequency basis choices and their benefits in 1D and 2D sampling.

4.2. Choice of Frequency Basis

Theoretically, any basis that is chosen results in the same inverse if the sampling matrix is full rank and all basis vector iterations are taken. However, not all basis vectors are needed to compute a band-limited inverse. Thus, the choice of frequency basis can affect when and how much noise amplification occurs during the inversion process. Depending on the assumed structure of each inverse, a recombination or reordering of frequency components can be useful in computing progressively band-limited inverses.

4.2.1. DFT vs. DCT

The DFT is typically preferred in 1D sampling, as it works for both real and complex signals. To progressively reconstruct lowpass signals with an increasing bandwidth, reordering the DFT frequencies from low to high is useful.

If we define the DFT basis vector as

$$f_{nk} = \sqrt{\frac{1}{N}} \exp\left(\frac{-2\pi i k n}{N}\right), n, k \in [0, N - 1] \quad (41)$$

we find that the first half $n \in [1, \text{floor}(N/2)]$ increases from low- to high- frequency components (often called the positive frequencies), while the second half $n \in [\text{ceil}(N/2), N - 1]$ decreases from high- to low- frequencies (often called negative frequencies). To rearrange these basis vectors in order of increasing bandwidth, we need to alternate positive and negative frequencies. The ordering is $[0, 1, N - 1, 2, N - 2, \dots]$. For real signals, the positive and negative frequencies of the DFT are conjugates of each other, making one redundant in computing the inverse. Thus, for real signals, only the first half of the complex DFT basis vectors is needed to compute each band-limited inverse $[0, 1, 2, \text{floor}(N/2)]$. An equivalent real-only basis can be formed by using the discrete cosine transform (DCT), which takes advantage of conjugate symmetric pairs. The most common forms of the DCT are already ordered low to high frequencies (DCT-I, DCT-II, DCT-III, DCT-IV) by taking half-frequency steps [21].

4.2.2. 2D Frequency Basis

For the case of 2D sampling, similar bases can be used such as the 2D DFT for complex signals and 2D DCT for real signals. However, the extra dimensionality of the 2D basis introduces directional frequencies. While directional frequencies may not be a concern when inverting a uniform sampling scheme, some irregular 2D sampling schemes may be sensitive to directional sampling errors.

One method to mitigate these directional sensitivities is to combine the symmetry of the four quadrants of the 2D DFT. Just as the DCT in one dimension takes advantage of the positive and negative symmetry of DFT bins offset by $\pi/2$, it can also be used to pair symmetric quadrants in 2D. To take advantage of this symmetry for complex signals, the discrete sine transform (DST) can be used to fill in the other $\pi/4$ shifted quadrants, i.e., using the DCT and DST to respectively compute the real and imaginary portions of a symmetric DFT basis. Note that basis vectors still need to be ordered by their radial distance to perform progressively bandlimited reconstructions.

It is possible to further combine radially symmetric DFT bins by reducing the frequency step size and applying the appropriate 2D rotations to preserve the orthogonality of each unique radial DFT bin. Example 2D basis vectors for the DFT, DCT, and a shifted even DCT basis are shown in Figure 6.




Figure 6. Example 2D basis vectors using DFT, DCT-IV, and a shifted even DCT basis. The DFT basis vectors have directional components, the DCT has fractional frequency content and the shifted even DCT Basis has up to 4 radially symmetric orthogonal vectors.

4.3. Advantages

While there may be many practical uses for the incremental computation of a bandlimited inverse using FQR, we highlight three main advantages for its use in discrete reconstruction problems. Firstly, since it can be hard to determine the maximum reconstruction bandwidth a priori, FQR can be used to simultaneously reconstruct and find the maximum reconstruction bandwidth. Secondly, since FQR uses a frequency basis, modern filtering techniques can be leveraged to further constrain, condition, and regularize the FQR band-limited inverse. Thirdly, since each singular update to the FQR inverse uses all measurements, noisy measurement errors are mitigated when terminating the inversion process early.

4.4. Applications

Using a 2D frequency basis, the FQR or similar frequency-constrained reconstruction methods can be applied to enhance the spatial resolution of remote sensing systems, without the need to determine the maximum reconstruction bandwidth a priori. This is particularly useful for reconstructing irregular sampling schemes and measurements with variable apertures, often found in spaceborne radiometers and scatterometers. However, the requirement of FQR to reconstruct each individual frequency basis can be as computationally taxing as other pseudo-inverse methods. Iterative methods could be used to reduce the computational complexity by approximating a band-limited inverse, but may amplify more noise than an individual frequency-basis inverse. Furthermore, since a frequency basis implicitly assumes a spatial periodicity or band limit, additional spatial tapering or windowing may be required to preserve the temporal or spatial relevance between measurements.

5. Conclusions

A frequency-constrained QR decomposition improves reconstruction performance for bandlimited signals, particularly for measurement with variable apertures, SNR limitations, or 2D sampling considerations. Additionally, a frequency-constrained QR reconstruction can simultaneously reconstruct a signal and determine its effective bandwidth. While this paper focused on applying frequency constraints to QR reconstruction, in order to highlight the theoretical difference between standard matrix inverse techniques and a bandlimited inverse, similar frequency constraints can be applied to more practical and efficient inverse

methods. Further work is needed to demonstrate the benefits of a bandlimited inverse, i.e., applying frequency constraints to existing inversion methods and remote sensing datasets.

Author Contributions: Conceptualization, H.G. and D.G.L.; investigation, H.G.; writing—original draft preparation, H.G.; writing—review and editing, H.G. and D.G.L.; supervision, D.G.L.; funding acquisition, D.G.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded through NASA grant number 80NSSC19K0057.

Data Availability Statement: The original contributions presented in the study are included in the article, further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

DFT	Discrete Fourier transform
DCT	Discrete cosine transform
DST	Discrete sine transform
FQR	Frequency-constrained QR reconstruction
MRF	Measurement response function
SNR	Signal to noise ratio
SVD	Singular value decomposition
QR	QR matrix decomposition

Appendix A. Relation to QR

As noted previously, \mathbf{t}_n can be found using the Gram–Schmidt orthogonalization procedure. By comparing the computation of \mathbf{t}_n and \mathbf{m}_n^\dagger to the QR decomposition method, we can see that the derivation of FQR is a particular case of QR decomposition using a frequency basis.

By examining Equations (23)–(25), we can express (MDF) as a matrix with rows of \mathbf{b}_n :

$$(\mathbf{MDF}) = \begin{bmatrix} \mathbf{b}_0 & \mathbf{b}_1 & \dots & \mathbf{b}_{N-1} \end{bmatrix}. \quad (\text{A1})$$

An orthogonal basis matrix (\mathbf{T}) can be formed by grouping together the column basis vectors \mathbf{t}_n :

$$\mathbf{T} = \begin{bmatrix} \mathbf{t}_0 & \mathbf{t}_1 & \dots & \mathbf{t}_{N-1} \end{bmatrix}. \quad (\text{A2})$$

Based on the orthogonalization process, applying \mathbf{T}^* to (MDF) results in an upper-triangular matrix \mathbf{U} :

$$\begin{aligned} \mathbf{U} &= \mathbf{T}^*(\mathbf{MDF}), \\ &= \begin{bmatrix} 1 & \mathbf{t}_0^*\mathbf{b}_1 & \mathbf{t}_0^*\mathbf{b}_2 & \dots & \mathbf{t}_0^*\mathbf{b}_{N-1} \\ 0 & 1 & \mathbf{t}_1^*\mathbf{b}_2 & \dots & \mathbf{t}_1^*\mathbf{b}_{N-1} \\ \vdots & \vdots & \ddots & & \vdots \\ 0 & 0 & 0 & \dots & 1 \end{bmatrix}. \end{aligned} \quad (\text{A3})$$

Rearranging the definition of \mathbf{U} we find a factorization of (MDF):

$$(\mathbf{MDF}) = (\mathbf{T}^*)^\dagger \mathbf{U}. \quad (\text{A4})$$

Taking advantage of the orthogonality of \mathbf{T} , we know that $\mathbf{T}^*\mathbf{T} = \mathbf{L}$ is a diagonal matrix, which can be inverted and used to find $(\mathbf{T}^*)^\dagger$:

$$(\mathbf{T}^*)^\dagger = \mathbf{T}\mathbf{L}^\dagger, \quad (\text{A5})$$

where \mathbf{L}^\dagger is a diagonal matrix depicted as

$$\mathbf{L}^\dagger = \begin{bmatrix} (\mathbf{t}_0^*\mathbf{t}_0)^{-1} & 0 & 0 & \dots & 0 \\ 0 & (\mathbf{t}_1^*\mathbf{t}_1)^{-1} & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & & \vdots \\ 0 & 0 & 0 & \dots & (\mathbf{t}_{N-1}^*\mathbf{t}_{N-1})^{-1} \end{bmatrix}. \quad (\text{A6})$$

Using the substitution in Equation (A5), we can find the QR factorization of (\mathbf{MDF}) :

$$\begin{aligned} (\mathbf{MDF}) &= \mathbf{Q}\mathbf{R}, \\ \mathbf{Q} &= \mathbf{T}\mathbf{L}^{\dagger\frac{1}{2}}, \\ \mathbf{R} &= \mathbf{L}^{\dagger\frac{1}{2}}\mathbf{U}. \end{aligned} \quad (\text{A7})$$

where \mathbf{Q} and \mathbf{R} are the QR decomposition of (\mathbf{MDF}) , \mathbf{T} is a matrix of \mathbf{t}_n , \mathbf{U} is an upper-triangular matrix formed by applying \mathbf{T}^* to (\mathbf{MDF}) , and \mathbf{D}^\dagger is a diagonal matrix of the inverse squared norm of each \mathbf{t}_n . By closer inspection of the recursive definition of \mathbf{t}_n , it is clear that the FQR process is an un-normalized QR decomposition.

Further analysis reveals that \mathbf{m}_n^\dagger is also related to the QR decomposition of (\mathbf{MDF}) . By examining Equation (39), we can form a matrix of \mathbf{m}_n^\dagger by using Equation (A5):

$$(\mathbf{T}^*\mathbf{MD})^\dagger = \begin{bmatrix} \mathbf{m}_0^\dagger & \mathbf{m}_1^\dagger & \dots & \mathbf{m}_{N-1}^\dagger \end{bmatrix} \quad (\text{A8})$$

Taking advantage of the orthonormal frequency basis $\mathbf{F}\mathbf{F}^* = \mathbf{I}$ and substituting in the factorization from Equation (A4), we can reduce the matrix of \mathbf{m}_n^\dagger further:

$$\begin{aligned} \begin{bmatrix} \mathbf{m}_0^\dagger & \mathbf{m}_1^\dagger & \dots & \mathbf{m}_{N-1}^\dagger \end{bmatrix} &= \mathbf{F}\mathbf{F}^*(\mathbf{T}^*\mathbf{MD})^\dagger, \\ &= \mathbf{F}(\mathbf{T}^*\mathbf{MDF})^\dagger, \\ &= \mathbf{F}\mathbf{U}^\dagger, \end{aligned} \quad (\text{A9})$$

where $\mathbf{F}\mathbf{U}^\dagger$ is the normalized QR decomposition of the matrix of \mathbf{m}_n^\dagger .

Appendix B. Relation to SVD

Single-rank SVD decomposition updates can be used to incrementally find a band-limited inverse. However, the reorderings of new singular values of each update makes the SVD method more tedious. To relate the SVD decomposition of a band-limited inverse to its corresponding FQR decomposition, we show how the SVD of a band-limited inverse is a rotation of the FQR decomposition.

Define the band-limited SVD as

$$(\mathbf{MD}) = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*, \quad (\text{A10})$$

where \mathbf{U} and \mathbf{V} are both unitary matrices (i.e., $\mathbf{U}^*\mathbf{U} = \mathbf{I}$ and $\mathbf{V}^*\mathbf{V} = \mathbf{I}$), and $\mathbf{\Sigma}$ is a diagonal matrix of the singular values of (\mathbf{MD}) . With the FQR decomposition as $\mathbf{M}\mathbf{F}\text{diag}(\mathbf{h}) = \mathbf{Q}\mathbf{R}$ (see Appendix A), then

$$(\mathbf{MD}) = \mathbf{Q}\mathbf{R}\mathbf{F}^*, \quad (\text{A11})$$

where \mathbf{Q} is a normalized \mathbf{T} , and \mathbf{R} is a weighted upper-triangular matrix from the FQR decomposition.

By rotating the SVD of (\mathbf{MD}) we find the SVD of \mathbf{R} ,

$$\mathbf{R} = \mathbf{Q}^*(\mathbf{MD})\mathbf{F} = \mathbf{U}_r\boldsymbol{\Sigma}\mathbf{V}_r^*, \quad (\text{A12})$$

which, if we substitute back in to Equation (A11), we find

$$(\mathbf{MD}) = \mathbf{Q}\mathbf{U}_r\boldsymbol{\Sigma}\mathbf{V}_r^*\mathbf{F}^*, \quad (\text{A13})$$

where we compare it with Equation (A10) $\mathbf{U} = \mathbf{Q}\mathbf{U}_r$ and $\mathbf{V} = \mathbf{F}\mathbf{V}_r$. We compute the bandlimited inverse as

$$(\mathbf{MD})^\dagger = \mathbf{F}\mathbf{V}_r\boldsymbol{\Sigma}^{-1}\mathbf{U}_r^*\mathbf{Q}^*. \quad (\text{A14})$$

References

- Unser, M. Sampling-50 years after Shannon. *Proc. IEEE* **2000**, *88*, 569–587. [\[CrossRef\]](#)
- Gröchenig, K. Irregular sampling, Toeplitz matrices, and the approximation of entire functions of exponential type. *Math. Comput.* **1999**, *68*, 749–765. [\[CrossRef\]](#)
- Long, D.G.; Franz, R.O. Band-limited Signal Reconstruction from Irregular Samples with Variable Apertures. *IEEE Trans. Geosci. Remote Sens.* **2015**, *54*, 2424–2436. [\[CrossRef\]](#)
- Long, D.G. Discrete Band-Limited Signal Reconstruction from Irregular Samples. *IEEE Trans. Geosci. Remote Sens.* **2021**, *59*, 4033–4043. [\[CrossRef\]](#)
- Long, D.G.; Hardin, P.J.; Whiting, P.T. Resolution enhancement of spaceborne scatterometer data. *IEEE Trans. Geosci. Remote Sens.* **1993**, *31*, 700–715. [\[CrossRef\]](#)
- Migliaccio, M.; Gambardella, A. Microwave radiometer spatial resolution enhancement. *IEEE Trans. Geosci. Remote Sens.* **2005**, *43*, 1159–1169. [\[CrossRef\]](#)
- Nunziata, F.; Alparone, M.; Camps, A.; Park, H.; Zurita, A.M.; Estatico, C.; Migliaccio, M. An enhanced resolution brightness temperature product for future conical scanning microwave radiometers. *IEEE Trans. Geosci. Remote Sens.* **2021**, *60*, 5301812. [\[CrossRef\]](#)
- Zeiger, P.; Picard, G.; Richaume, P.; Mialon, A.; Rodriguez-Fernandez, N. Resolution enhancement of SMOS brightness temperatures: Application to melt detection on the Antarctic and Greenland ice sheets. *Remote Sens. Environ.* **2024**, *315*, 114469. [\[CrossRef\]](#)
- Gröchenig, K. A discrete theory of irregular sampling. *Linear Algebra Its Appl.* **1993**, *193*, 129–150. [\[CrossRef\]](#)
- Strohmer, T. Numerical analysis of the non-uniform sampling problem. *J. Comput. Appl. Math.* **2000**, *122*, 297–316. [\[CrossRef\]](#)
- Davis, P.J. *Circulant Matrices*; Wiley New York: New York, NY, USA, 1979; Volume 2.
- Fill, J.A.; Fishkind, D.E. The Moore–Penrose Generalized Inverse for Sums of Matrices. *SIAM J. Matrix Anal. Appl.* **2000**, *21*, 629–635. [\[CrossRef\]](#)
- Givens, W. Computation of plain unitary rotations transforming a general matrix to triangular form. *J. Soc. Ind. Appl. Math.* **1958**, *6*, 26–50. [\[CrossRef\]](#)
- Steinhardt, A.O. Householder transforms in signal processing. *IEEE ASSP Mag.* **1988**, *5*, 4–12. [\[CrossRef\]](#) [\[PubMed\]](#)
- Kelley, C.T. *Iterative Methods for Linear and Nonlinear Equations*; SIAM: Philadelphia, PA, USA 1995.
- Gander, W. Algorithms for the QR decomposition. *Res. Rep.* **1980**, *80*, 1251–1268.
- Trefethen, L.N.; Bau, D. *Numerical Linear Algebra*; SIAM: Philadelphia, PA, USA 2022.
- Wiener, N. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series, with Engineering Applications*; MIT Press: Cambridge, MA, USA 1964.
- Bai, Z.; Fahey, G.; Golub, G. Some large-scale matrix computation problems. *J. Comput. Appl. Math.* **1996**, *74*, 71–89. [\[CrossRef\]](#)
- Zhou, K.X.; Roumeliotis, S.I. A sparsity-aware QR decomposition algorithm for efficient cooperative localization. In Proceedings of the 2012 IEEE International Conference on Robotics and Automation, Saint Paul, MN, USA, 14–18 May 2012; pp. 799–806.
- Ahmed, N.; Natarajan, T.; Rao, K.R. Discrete cosine transform. *IEEE Trans. Comput.* **1974**, *100*, 90–93. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.