

Gridding

Nicholas Dwork

©2018

Abstract

This document attempts to explain the Gridding and Inverse Gridding algorithms. This document is heavily based on notes for Stanford's EE 369C class by Dr. John Pauly [1] and [2] by Beatty et al.

1 Background

For this document, we will use the following definitions for the Fourier Transform and inverse Fourier Transform:

$$F(k) = \mathcal{F}\{f\}(k) = \int_{-\infty}^{\infty} f(x) \exp(-i 2\pi k x) dx$$
$$f(x) = \mathcal{F}^{-1}\{F\}(x) = \int_{-\infty}^{\infty} F(k) \exp(i 2\pi k x) dk$$

where $i = \sqrt{-1}$.

The Discrete Fourier Transform (DFT) and the inverse Discrete Fourier Transform (DFT^{-1}) are linear operators; they can be derived as specific Riemann sums that approximate the Fourier Transforms where each element of the partition is equal in size, and the element in the summation (x_n or k_m) is the midpoint of the partition element. They are denoted as follows:

$$V[m] = \text{DFT}(v)[m] = \sum_{n=0}^{N-1} v[n] \exp\left(-i 2\pi \frac{m n}{N}\right)$$
$$v[n] = \text{DFT}^{-1}(V)[n] = \frac{1}{N} \sum_{m=0}^{N-1} V[m] \exp\left(i 2\pi \frac{m n}{N}\right).$$

The DFT and DFT^{-1} can be calculated very quickly using the Fast Fourier Transform (FFT) algorithm.

2 Introduction

This document is concerned with the following question: given a set of coordinates and corresponding values in the source domain how can we approximate the values at a set of coordinates in the destination domain? If the coordinates in both domains lie on the vertices of a cartesian grid then the answer is the DFT. But what do we do when the coordinates of one domain (or both) don't lie on a uniform grid?

The general problem, where there are no restrictions on the coordinates of either domain, is called the Non-uniform DFT. One might think that this problem is easily solved by interpolating in the source domain, performing a transform, and then interpolating in the destination domain. Unfortunately, this does not work out well.

We will look at two special cases of the Non-uniform DFT problem: "Gridding" and "Inverse Gridding".

3 Gridding

Gridding is an algorithm that approximates the Inverse Fourier Transform when the source coordinates (k_1, \dots, k_M) do not lie on a uniform grid but the destination coordinates (x_1, \dots, x_N) do. The question becomes, given the coordinates (k_1, \dots, k_M) and values $(F(k_1), \dots, F(k_M))$, how can we approximate the destination domain values $(f(x_1), \dots, f(x_N))$? Importantly, to answer this question we want to take advantage of the computational efficiency of the Inverse Discrete Fourier Transform; that is, Gridding should call the DFT^{-1} .

Define S as the “sampling function”, $S(k) = \sum_j \delta(k - k_j)$. It is a distribution that specifies which frequencies k have known Fourier values. It is assumed that the number of samples is finite, and therefore S has compact support. The “sampled data” is the distribution FS .

The trick of gridding will be to *push not pull*. That is, rather than interpolating from nearby values to determine the values of frequencies that lie on a Cartesian grid, we are going to *push* information from all known frequencies onto the nearby grid points. This is done by performing a convolution in the Fourier domain with a kernel C . Using this idea, the approximations to the Fourier values in the source domain are

$$\tilde{F} = (FS) * C, \quad (1)$$

where $*$ denotes convolution. Note that convolution is an integration, so (by the sifting property) convolving with C converts the distribution FS to finite (possibly complex) values at all points in the source domain.

Now that we have an approximation to the function F , we can approximate the values in the destination domain by using the Inverse Fourier Transform

$$\mathcal{F}^{-1}\{F\} \approx \mathcal{F}^{-1}\{FS * C\}. \quad (2)$$

This operation can't be performed on a computer. So we sample the source domain data on a uniform grid, modeled as multiplication with a comb function as follows¹:

$$[FS * C] \Delta k \text{III}_{\Delta k}, \quad (3)$$

where III is the comb (or Shah) function defined as $\text{III}_{\Delta k}(k) = \sum_{m=-\infty}^{\infty} \delta(k - m \Delta k)$.

Define \mathcal{T} as a mapping such that $\mathcal{T}([FS * C] \Delta k \text{III}_{\Delta k})[m]$ equals the weight of the m^{th} nonzero delta. Then \mathcal{T} is an isomorphism that maps the set of functions of the form in (3) (together with pointwise addition and scalar multiplication) onto \mathbb{C}^M .

We can approximate the values of (2) by performing an DFT^{-1} on this vector of sampled data; this yields N equally spaced approximations in the destination domain. For any $n \in \{0, 1, \dots, N - 1\}$,

$$\hat{f}(n) = \text{DFT}^{-1}\{\mathcal{T}([FS * C] \Delta k \text{III}_{\Delta k})\}[n] \quad (4)$$

$$\approx ((f * s) \cdot c) * \text{III}_{\text{FOV}}(n), \quad (5)$$

where s is the inverse Fourier Transform of S , and $\text{FOV} = 1/\Delta k$ where FOV stands for “Field of View”. Note that $((f * s) \cdot c) * \text{III}_{\text{FOV}}$ is periodic with period FOV ; we will estimate samples from one period, hence the name “Field of View.” The function $c = \mathcal{F}^{-1}\{C\}$ is called the “apodization function.” Note that (5) is only approximately equal to (4) since pointwise multiplication with the apodization function c in (5) equals circular convolution in (4).

Circular convolution with C is a linear transformation from $\mathbb{C}^M \rightarrow \mathbb{C}^N$, thus it can be represented by a matrix $\mathcal{C}[\mathbf{k}, \mathbf{k}'] \in \mathbb{R}^{N \times M}$, where \mathbf{k} is a vector of the sampled source domain coordinates and \mathbf{k}' is a vector of the source domain coordinates on a Cartesian grid. If the support of C is small then this matrix is sparse. The value $\mathcal{C}[\mathbf{k}, \mathbf{k}']_{nm}$ is the value of C evaluated at $\|k'_n - k_m\|_2$.

$$\underbrace{\begin{bmatrix} \mathcal{C}_{11} & \mathcal{C}_{12} & \cdots & \mathcal{C}_{1M} \\ \mathcal{C}_{21} & \mathcal{C}_{22} & \cdots & \mathcal{C}_{2M} \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{C}_{N1} & \mathcal{C}_{N2} & \cdots & \mathcal{C}_{NM} \end{bmatrix}}_{\mathcal{C}[\mathbf{k}, \mathbf{k}']} \begin{bmatrix} \hat{F}_1 \\ \hat{F}_2 \\ \vdots \\ \hat{F}_M \end{bmatrix} = \begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \vdots \\ \gamma_N \end{bmatrix}.$$

¹One might think that we also need to multiply by a rect function in (3). But, since S has compact support, $FS * C$ has compact support. So multiplication with a rect would be superfluous.

One can form the matrix $C[\mathbf{k}, \mathbf{k}']$. However, for images or volumes that matrix can consume a great deal of memory (even when it's stored as a sparse matrix). It may benefit us, then, to identify an algorithm to implement left multiplication by $C[\mathbf{k}, \mathbf{k}']$. A natural implementation of left multiplication by $C[\mathbf{k}, \mathbf{k}']$ is one that iterates over the Fourier estimates \hat{F} , multiplies each column of $C[\mathbf{k}, \mathbf{k}']$ by the corresponding estimate of \hat{F} , and sums all the results together. Note that this algorithm is embarrassingly parallelizable.

It is important to choose a good kernel function C . One might think to convolve with a rect function, which would be the simplest thing to do (it equates to nearest neighbor interpolation). This leads to multiplication by a sinc function in the destination domain. Since a sinc has significant values over a large segment of the domain (i.e. it dies out slowly), this leads to large aliasing artifacts. Additionally, the sinc function may be 0 in the Field of View which would make deapodization impossible. The theoretically ideal choice would be to convolve with a sinc in the source domain, which leads to multiplication by a rect in the destination domain. But this is extremely computationally intensive, and can't be accomplished theoretically since we won't retain an infinite number of values. A great choice of convolution kernel is the Kaiser Bessel function [2, 3]. A multidimensional kernel can be constructed as a separable product of 1D kernels.

Equation (4) represents the simplest form of Gridding: sample in the Fourier domain, perform a convolution with C , sample on a uniform grid, and perform an DFT^{-1} . Notice, though, that in equation (5) the reconstruction \hat{f} does not equal f ; the differences are all corruptions (visualized in Figure 1). These corruptions include convolution by s , apodization (or multiplication) by c , and aliasing due to convolution with III_{FOV} ; we'll now alter the gridding algorithm to address each of these corruptions.

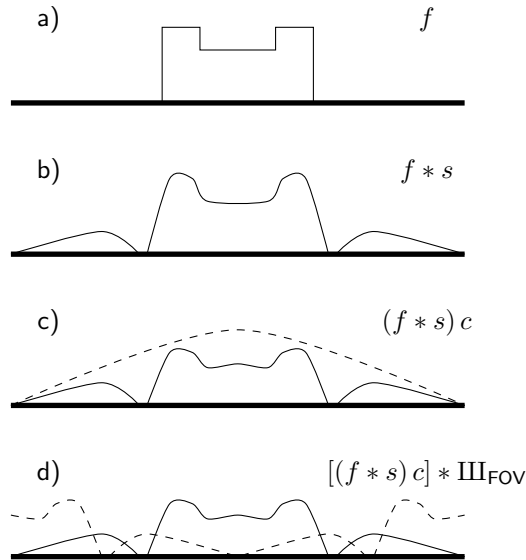


Figure 1: (a) The true function f . (b) f after convolution with s . The sharp edges become rounded and sidelobe artifacts appear. (c) Multiplication by the apodization function c . (d) Aliasing due to convolution with III_{FOV} . The final function would be the sum of the dashed and solid lines of subfigure (d).

3.1 Reducing Aliasing Artifacts

In (4), we chose to sample on a grid with spacing Δk , where this spacing leads to the Field of View in the destination domain that we desired. But this was not required. We could choose to reconstruct on a grid with smaller spacing, leading to a larger than desired Field of View, and then crop out the center region corresponding to our desired Field of View. This is called “oversampling”.

Suppose that, with oversampling, the number of samples is $\tilde{N} > N$; then the oversampling factor is $\alpha = \tilde{N}/N$. The gridding algorithm becomes

$$\begin{aligned} \hat{f}(n) &= \text{DFT}^{-1} \left\{ \mathcal{T} \left[(F S * C) (\Delta k / \alpha) \text{III}_{\Delta k / \alpha} \right] \right\} [n] \\ &\approx ((f * s) \cdot c) * \text{III}_{\alpha \text{FOV}}(n). \end{aligned}$$

The benefit can be seen in figure 2. The aliased sidelobes are pushed farther out and so they have less impact in the central region. After cropping, most of the aliasing artifacts are eliminated. Note that with

oversampling, the circular convolution operator has been changed; it now maps $\mathbb{C}^M \rightarrow \mathbb{C}^{\alpha N}$ and can be denoted $\mathcal{C}[\mathbf{k}, \mathbf{k}'']$, where \mathbf{k}'' is a vector of the k-space coordinates on the oversampled Cartesian grid (which is a larger vector than \mathbf{k}').




Figure 2: (a) Aliasing artifacts shown with oversampling factor $\alpha = 1$. (b) Aliasing artifacts moved farther away by increasing α .

3.2 Deapodization

We will now address the corruption of apodization by c . After reviewing (5), one would hope that simply dividing $\hat{f}(n)$ by $c(n)$ for all n would resolve the issue. This is a deconvolution in the source domain², and it's close to what we want, but not quite. Let $\mathbf{c} = (c(0), c(1), \dots, c(\alpha N - 1))$; then

$$\frac{\alpha N}{c(n)} ((f * s) \cdot c) * \text{III}_{\alpha \text{FOV}}(n) = \text{DFT}^{-1} \left\{ \text{DFT} \left\{ \frac{1}{\mathbf{c}} \right\} \otimes \mathcal{T} [(F S * C)(\Delta k / \alpha) \text{III}_{\Delta k / \alpha}] \right\} [n],$$

where $1/\mathbf{c}$ is a vector of the values of \mathbf{c} inverted. The problem is that in the source domain we are convolving by C . But, since we are using the DFT^{-1} to approximate the inverse Fourier Transform, dividing \hat{f} by c results in a circular convolution with $1/\mathbf{c}$. Thus, the deconvolution operation is not the inverse of the convolution operation. This works ok, but can be significantly improved. Instead of convolving by C , we alter the gridding algorithm to perform a circular convolution by C . That is, whenever convolution by C extends beyond the bounds $[-0.5, 0.5]$, we *wrap around*. Thus, the gridding algorithm becomes

$$\begin{aligned} \hat{f}(n) &= \frac{\alpha N}{c(n)} \text{DFT}^{-1} \left\{ \mathcal{T} [(F S \otimes C)(\Delta k / \alpha) \text{III}_{\Delta k / \alpha}] \right\} [n] \\ &\approx ([f * s] * \text{III}_{\alpha \text{FOV}})(n). \end{aligned} \quad (6)$$

Dividing by c is called “deapodization”; it boosts the intensities that were apodized by the convolution with C . The division by c and the multiplication by c approximately cancel each other out. However, any aliasing artifacts that remain and noise that is present will also get boosted up; hopefully these are small.

Since $\Delta k = 1/N$, we can simplify equation (6) by cancelling out the two αN as follows:

$$\hat{f}(n) = \frac{1}{c(n)} \text{DFT}^{-1} \left\{ \mathcal{T} [(F S \otimes C) \text{III}_{\Delta k / \alpha}] \right\} [n].$$

3.3 Density Compensation

We will now address the corruption due to the sampling function. The problem arises when the density of sample points in the source domain is not uniform; thus, the name of the algorithm to address this corruption is “Density compensation.” The key idea is to alter the sampling function; instead of making the sampling function a sum of impulses, we make the sampling function a sum of weighted impulses³: $S_w(k) = \sum_i w_i \delta(k - k_j)$. The question then becomes, how do we determine $w \in \mathbb{R}^M$?

Here, we will present the method of Pipe et al. [4]. The trick is to realize that if $f = \delta$ then $F = 1$ (i.e. F is the constant function 1). In this case, after convolution with C , we would like a value of 1 everywhere in

²Recall the convolution theorems: $\text{DFT}(a \otimes b) = \text{diag}(\text{DFT}(a)) \text{DFT}(b)$ and $\text{DFT}(\text{diag}(a)b) = (1/N)\text{DFT}(a) \otimes \text{DFT}(b)$, where $a, b \in \mathbb{C}^N$.

³This is called pre-density compensation. There also exists a post-density compensation which won't be discussed in this document. Refer to [1] for more details.

the source domain. Since we don't have samples everywhere, we desire, instead, that $S_w * C \approx 1$. With this goal in mind, Pipe et al. designed the following iteration:

$$S_w^{(i+1)} = \frac{S_w^{(i)}}{S_w^{(i)} * C} \quad (7)$$

where $S_w^{(i+1)} = 0$ whenever the denominator equals 0. Let $\mathcal{T}_S [S_w^{(i+1)}] = w^{(i+1)}$; i.e. it is a vector of the weights at the non-zero locations of the distribution. Note that \mathcal{T}_S is an isomorphism from the set of all possible S_w to \mathbb{R}^M (where we are assuming pointwise addition and scalar multiplication). If $S_w^{(i)} * C \approx 1$ at all nonzero values of $S_w^{(i)}$ then $S_w^{(i+1)} \approx S_w^{(i)}$; that is, the iteration has converged. Pipe et al. specify that this iteration should be run 10 times. Then, the final weight vector is $w = \mathcal{T}_S [S_w^{(15)}]$.

Hopefully, after density compensation, $s_w = \mathcal{F}^{-1}\{S_w\}$ approximates the Dirac delta δ .

3.4 Final Gridding Algorithm

The final gridding algorithm is shown in equation (8) and summarized in Algorithm 1.

$$\hat{f}(n) = \frac{1}{c(n)} \text{DFT}^{-1} \{ \mathcal{T} [(F S_w \otimes C) \text{III}_{\Delta k/\alpha}] \} [n]. \quad (8)$$

Algorithm 1: Gridding

Inputs: (k_1, \dots, k_M) , $(F(k_1), \dots, F(k_M))$, $\alpha > 1$

Step 1: Determine the weighting vector w using iteration (7).

Step 2: Determine the result of convolving the set $F S_w$ with the convolution kernel C and sampling on the grid points. Note that the grid points are the vertices of the oversampled grid.

Step 3: Perform an Inverse Discrete Fourier Transform.

Step 4: After the DFT^{-1} , the number of elements in the destination domain is larger than desired by a factor of α . Crop out the center region, which corresponds to the desired Field of View.

Step 5: Deapodize by c .

Gridding is a linear transformation \mathcal{G} , which can be expressed as

$$\mathcal{G} = \mathcal{K} \text{diag}(c)^{-1} \text{DFT}^{-1} \mathcal{C}[\mathbf{k}, \mathbf{k}'] \text{diag}(w),$$

where $\mathcal{K} : \mathbb{C}^{\alpha N} \rightarrow \mathbb{C}^N$ is the cropping transformation. A slight improvement in efficiency could be had by switching the order of cropping and deapodization; one would deapodize only those points that remain after cropping.

4 Inverse Gridding

Inverse Gridding is an algorithm that approximates the Fourier Transform when the source coordinates (x_1, \dots, x_N) lie on a uniform grid but the destination coordinates (k_1, \dots, k_M) do not⁴. The question becomes, given the coordinates (x_1, \dots, x_N) and values $(f(x_1), \dots, f(x_N))$, how can we approximate the destination domain values $(F(k_1), \dots, F(k_M))$?

The sampled data is $f \text{III} \Pi_D$ where the samples are contained within $(-D/2, D/2)$, where Π_D is a rect function of width D . Define t as the mapping such that $t(f \text{III} \Pi_D)[n]$ is the weight of the n^{th} delta. Then after sampling, the Discrete Fourier Transform yields

$$\mathcal{T}^{-1} \text{DFT} \{ t(f \text{III} \Pi_D) \} = (F * \text{III} * b) \text{III}_{\Delta k} \Pi$$

where $b(k) = D \text{sinc}(Dk)$ and $\Delta k = 1/N$. The convolution of F with III is aliasing and the convolution with the sinc is a blur.

⁴It's important to note that Inverse Gridding is not the inverse of Gridding. Instead, Inverse Gridding is an algorithm that approximates the inverse of Gridding.

We assume that f has compact support and that the support is a subset of $(-D/2, D/2)$. With this assumption, $F * b = F$. Thus, we can eliminate b from the above equation:

$$\mathcal{T}^{-1} \text{DFT}\{t(f \text{III} \Pi_D)\} = (F * \text{III}) \text{III}_{\Delta k} \Pi.$$

At this point we have uniformly spaced samples in the destination domain, but what we want is samples at a specific set of destination domain coordinates $\{k_1, \dots, k_M\} \subset [-1/2, 1/2)$. Again, we use the *push not pull* idea and circularly convolve with the kernel C :

$$\mathcal{T}^{-1} \text{DFT}\{t(f \text{III} \Pi_D)\} \otimes C = (F * \text{III}) \text{III}_{\Delta k} \Pi \otimes C.$$

Finally, we sample in the destination domain at the coordinates of interest:

$$[\mathcal{T}^{-1} \text{DFT}\{t(f \text{III} \Pi_D)\} \otimes C] S = [(F * \text{III}) \text{III}_{\Delta k} \Pi \otimes C] S.$$

Note that for some function h , hS is a distribution with M nonzero values at locations in the set $\{k_j : j = 1, \dots, M\}$.

We have arrived at the simplest form of inverse gridding:

$$\hat{F}(k_m) = \mathcal{T}_S \left([\mathcal{T}^{-1} \text{DFT}\{t(f \text{III} \Pi_D)\} \otimes C] S \right) [m].$$

The procedure is to compute the DFT of the data, circularly convolve with the kernel C , and sample on the desired coordinates.

Again, the estimates \hat{F} are not equal to F , and the differences are corruptions: aliasing due to the convolution with III , and blurring due to the circular convolution with C . Unfortunately, there's nothing we can do about the aliasing. In the next subsection, we will discuss how to address the corruption of the circular convolution with C .

Note that density compensation is not needed with inverse gridding; intuitively this makes sense since the samples that are convolved with C are evenly spread out in the destination domain.

4.1 Pre-emphasis

Like deapodization in Gridding, we wish to perform a deconvolution by C in the destination domain. This can approximately be done by performing a pointwise multiplication by $1/(c)$ in the source domain as follows:

$$\hat{F}(k_m) = \mathcal{T}_S \left(\left[\mathcal{T}^{-1} \text{DFT} \left\{ \text{diag}(N \mathbf{c})^{-1} t(f \text{III} \Pi_D) \right\} \otimes C \right] S \right) [m], \quad (9)$$

This algorithm is called pre-emphasis since the effect is to change the emphases (relative magnitudes) of different samples of f .

4.2 Final Inverse Gridding Algorithm

Putting it all together, we arrive at the final Inverse Gridding algorithm, summarized in Algorithm 2:

$$\hat{F}(k_m) = \mathcal{T}_S \left(\left[\mathcal{T}^{-1} \text{DFT} \left\{ \text{diag}(N \mathbf{c})^{-1} t(f \text{III} \Pi_D) \right\} \otimes C \right] S \right) [m].$$

Algorithm 2: Inverse Gridding

Inputs: $(x_1, \dots, x_N), (f(x_1), \dots, f(x_N))$

Step 1: Pre-emphasize by $1/(N \mathbf{c})$.

Step 2: Compute the DFT.

Step 3: Perform a circular convolution with C and sample on the points of interest.

Let the input be represented by $v = t(f \text{III} \Pi_D) \in \mathbb{C}^N$. Then the Inverse Gridding transformation can be written as

$$\mathcal{T}\hat{F} = \mathcal{C}[\mathbf{k}', \mathbf{k}] \text{DFT} \text{diag}(N \mathbf{c})^{-1} v,$$

where $\mathcal{C}[\mathbf{k}', \mathbf{k}] \in \mathbb{R}^{M \times N}$ is the circular convolution with the kernel C . The value of $\mathcal{C}[\mathbf{k}', \mathbf{k}]_{mn}$ is the kernel C evaluated at $\|k_m - k'_n\|_2$. Inverse Gridding is a linear transformation \mathcal{G}_I , which can be written as

$$\mathcal{G}_I = \mathcal{C}[\mathbf{k}', \mathbf{k}] \text{DFT} \text{diag}(N \mathbf{c})^{-1}. \quad (10)$$

4.3 Adjoint of Inverse Gridding

Note that $\mathcal{C}[\mathbf{k}', \mathbf{k}] = \mathcal{C}[\mathbf{k}, \mathbf{k}']^T$. From equation (10), it immediately follows that the adjoint of Inverse Gridding is⁵

$$\mathcal{G}_I^* = \text{diag}(\mathbf{c})^{-1} \text{DFT}^{-1} \mathcal{C}[\mathbf{k}, \mathbf{k}'].$$

Here we see a beautiful relationship; the Adjoint of Inverse Gridding is a substep of Gridding:

$$\mathcal{G} = \mathcal{K} \text{diag}(\mathbf{c})^{-1} \text{DFT}^{-1} \mathcal{C}[\mathbf{k}, \mathbf{k}'] \text{diag}(\mathbf{w}) = \mathcal{K} \mathcal{G}_I^* \text{diag}(\mathbf{w}).$$

Gridding is density compensation weighting, followed by the adjoint of Inverse Gridding, followed by a cropping. The algorithm for the adjoint of Inverse Gridding is shown in Algorithm 3.

Algorithm 3: Adjoint of Inverse Gridding

Inputs: (y_1, \dots, y_M)

Step 1: Compute the circular convolution with the kernel C and sample on grid points.

Step 2: Compute the DFT^{-1} .

Step 3: Deapodize by $1/\mathbf{c}$.

Consider left multiplication by $\mathcal{C}[\mathbf{k}', \mathbf{k}]^T$ as shown in equation (11). A natural implementation of this is to iterate over the Fourier coefficients, and compute the inner product between the rows of $\mathcal{C}[\mathbf{k}', \mathbf{k}]^T$ (or the columns of $\mathcal{C}[\mathbf{k}', \mathbf{k}]$) and the Fourier coefficients γ . Note that this algorithm is embarrassingly parallelizable.

$$\begin{bmatrix} \hat{F}_1 \\ \hat{F}_2 \\ \vdots \\ \hat{F}_M \end{bmatrix} = \underbrace{\begin{bmatrix} \mathcal{C}_{11} & \mathcal{C}_{12} & \cdots & \mathcal{C}_{1N} \\ \mathcal{C}_{21} & \mathcal{C}_{22} & \cdots & \mathcal{C}_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{C}_{M1} & \mathcal{C}_{M2} & \cdots & \mathcal{C}_{MN} \end{bmatrix}}_{\mathcal{C}[\mathbf{k}', \mathbf{k}]} \begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \vdots \\ \gamma_N \end{bmatrix}. \quad (11)$$

References

- [1] John M. Pauly. Medical image reconstruction. <http://web.stanford.edu/class/ee369c/>. Accessed: 2016-02-15.
- [2] Philip J Beatty, Dwight G Nishimura, and John M Pauly. Rapid gridding reconstruction with a minimal oversampling ratio. *Medical Imaging, IEEE Transactions on*, 24(6):799–808, 2005.
- [3] John I Jackson, Craig H Meyer, Dwight G Nishimura, and Albert Macovski. Selection of a convolution function for fourier inversion using gridding [computerised tomography application]. *Medical Imaging, IEEE Transactions on*, 10(3):473–478, 1991.
- [4] James G Pipe and Padmanabhan Menon. Sampling density compensation in mri: rationale and an iterative numerical solution. *Magnetic Resonance in Medicine*, 41(1):179–186, 1999.

⁵Recall that $\text{DFT}^* = N \text{DFT}^{-1}$ where the DFT^{-1} maps \mathbb{C}^N to \mathbb{C}^N .