Implementations of Shannon's sampling theorem, a time-frequency approach

Thomas Strohmer and Jared Tanner*

Abstract

Shannon's sampling theorem quantifies the Fourier domain periodization introduced by the equidistant sampling of a bandlimited signal when the sampling rate is at least as fast as the Nyquist rate dictated by the signal's bandwidth. If sampled faster than the Nyquist rate, i.e., oversampling, a reconstruction composed of highly localized atoms is possible, allowing for practical applications where only a truncated set of samples is available. More specifically, it is known that root-exponential accuracy can be achieved by constructing atoms whose Fourier transform (filter) is infinitely differentiable and compactly supported in the appropriate bandwidth. Unfortunately, there is no known compactly supported infinitely smooth filter whose corresponding atom has a known explicit representation. By considering filters with Gevrey regularity, we obtain root-exponential localization for the atom, and an efficient truncated Gabor approximation of the filter and atom. Furthermore, we present an alternative error decomposition that allows for the complete rigorous analysis of the error in truncating the signal, and of the error introduced in approximating the filter and atom. By scaling the approximation order appropriately, the root-exponential convergence rate is not adversely effected by the filter's approximation.

Keywords: Shannon's sampling theorem, oversampling, Gabor frames, Gevrey regularity.

1 Introduction

The classical Shannon sampling theorem plays a crucial role in signal processing and communications, indicating how to transfer between analog signals and discrete se-

*Department of Mathematics, University of California, Davis, CA 95616-8633, USA; The authors acknowledge partial support by NSF DMS grant 0208568. J.T. was supported by NSF DMS VIGRE grant 0135345. Correspondence: strohmer,jtanner@math.ucdavis.edu.

quences [29]. Shannon's sampling theorem states that if a function[1] belongs to the space of bandlimited functions B_σ, i.e.,

$$f(t) := \frac{1}{\sqrt{2\pi}} \int_{-\sigma}^{\sigma} e^{2\pi i w t} F(w) dw, \qquad F(w) \in L_0^2[-\sigma, \sigma], \tag{1}$$

then it can be recovered exactly from its equidistant samples

$$f(t) \equiv \sum_{k=-\infty}^{\infty} f\left(\frac{k}{2\sigma}\right) \frac{\sin(2\pi\sigma t - \pi k)}{2\pi\sigma t - \pi k} := \sum_{k=-\infty}^{\infty} f\left(\frac{k}{2\sigma}\right) \operatorname{sinc}(2\pi\sigma t - \pi k). \tag{2}$$

Sampling a function in the time domain introduces a periodization in the associated Fourier dual space, where sampling at the rate $1/2\sigma$ corresponds to a 2σ periodization. In equation (2) the reproducing atom (time domain), $\operatorname{sinc}(2\pi\sigma t - \pi k)$, removes the periodization introduced by sampling, through the action of its associated filter (Fourier dual space), $\chi_{[-\sigma,\sigma]}$, see Figure 1. For this critical, Nyquist sampling rate, $\operatorname{sinc}(\cdot)$ is the unique atom that can be used to remove the periodization.

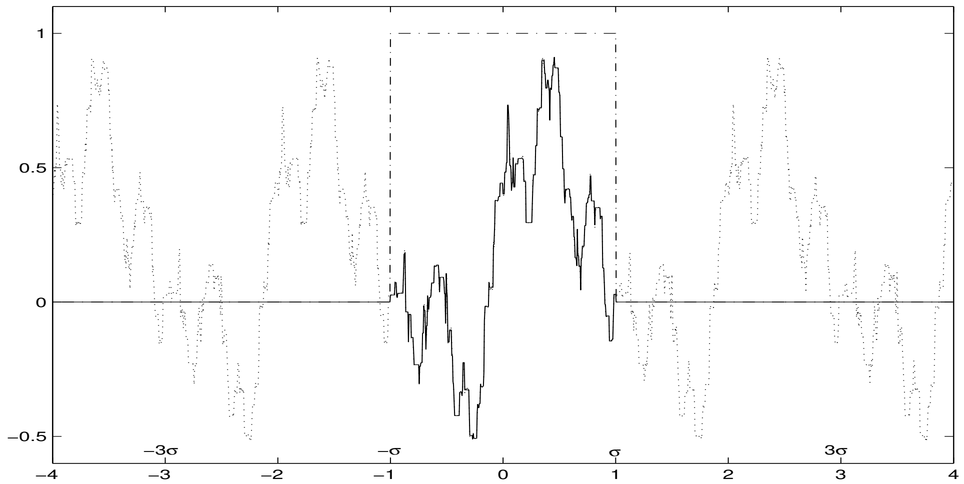

Figure 1: The dual space representation of a representative signal bandlimited to $[-\sigma, \sigma]$, solid. The periodization of the signal in the dual space, dotted, due to sampling at the Nyquist rate $1/2\sigma$, and the sinc's dual space representation, $\chi_{[-\sigma,\sigma]}$, dot-dash.

Although a critical observation, the formulation in equation (2) is unsuitable for practical applications, where only a finite number of samples is available, $\{f(k/2\sigma)\}_{|k|\leq L}$. For

[1] Throughout this paper we will use lower case letters to designate functions in the time domain, and upper case for the Fourier transform of the same function.

truncated samples the error, classically referred to as the truncation error, is controlled by the atom's localization

$$\epsilon(t,\sigma,L) := \left| f(t) - \frac{\sqrt{2\pi}}{2\sigma} \sum_{|k|\le L} f\left(\frac{k}{2\sigma}\right) \psi\left(t - \frac{k}{2\sigma}\right) \right| \le \frac{\sqrt{2\pi}}{2\sigma} \cdot \|f\|_{L_\infty} \sum_{|k|>L} \left| \psi\left(t - \frac{k}{2\sigma}\right) \right|. \tag{3}$$

In the case of the classical Shannon sampling theorem, the atom, $\psi(\tau) := \mathrm{sinc}(\tau)$, suffers from an unacceptably slow decay, $lim_{\tau\to\infty}\psi(\tau) \sim 1/\tau$, resulting in a first order convergence rate while moving from the sample boundaries, $\pm L/2\sigma$,to the interior. Moreover if the samples $f(k/(2\sigma))$ are replaced by noisy samples $f(k/(2\sigma)) + \epsilon_k$, then the corresponding approximation via the cardinal series in (2) may differ significantly from $f(t)$, cf. [4].

To remedy these problems in applications one usually introduces oversampling: If the bandlimited signal is sampled at a faster rate, $T := r/2\sigma$ where $r < 1$, then the dual space periodization is increased to σ/r, allowing a large family of reproducing filters. Specifically, any function satisfying[2]

$$\Psi(w) = \begin{cases} 1 & |w| \le \sigma \\ 0 & |w| > \sigma(2-r)/r =: \Omega \\ \text{anything} & \text{else,} \end{cases} \qquad \Rightarrow \qquad \psi \in B_\Omega \tag{4}$$

gives rise to a Shannon-type series expansion

$$f(t) \equiv \sqrt{2\pi}T \sum_{k=-\infty}^{\infty} f(kT)\, \psi(t - kT)\,. \tag{5}$$

For $r = 1$ the above filter reduces to $\chi_{[-\sigma,\sigma]}$, and the classical Shannon's sampling theorem. However, for $r < 1$ a gap is introduced between σ and $\sigma(2-r)/r$, allowing for a host of other reproducing filters, including those with a high degree of regularity, see Figure 2.

Asymptotically the atom's localization is reflected in the filter's smoothness; consequently, the filter's regularity controls the convergence rate of the truncation error, (3).

[2]It has been noted in [16] that the reproducing property is somewhat less strict than as stated in equation (4), in that the filter need not be zero for all $|w| \ge (2-r)/r$. Rather, the reproducing property is satisfied if the filter is one for $|w| \le \sigma$, and zero at the points where the periodic extension of the signal's dual representation is non-zero. However, this added flexibility can not increase the regularity of the filter, or decrease its regularity constants; and as such, can not improve the asymptotic convergence rate. Although this added flexibility can be used to increase the atom's immediate localization about the origin, it introduces substantial peaks away from the origin, [16], decreasing the overall convergence rate.

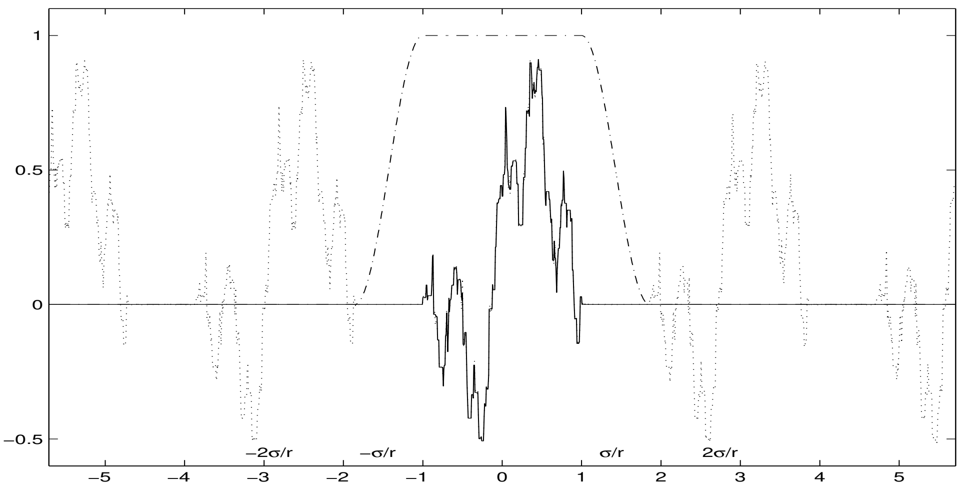

Figure 2: The signal's dual space representation, solid, and the signal's periodization due to sampling at the rate, $r/2\sigma$, for $r < 1$, dotted. With the gap between periodization, a smooth filter, dot-dash, can be used to remove the periodization introduced by sampling.

By constructing infinitely differentiable filters with precise regularity estimates, we obtain root-exponential accuracy for the approximation of a bandlimited signal, as the point to be approximated moves from the sampling boundary, $\pm LT$, to the interior, Section 2. Unfortunately, unlike classical finite regularity filters, such as the raised cosine, which have a closed form expression for their corresponding atoms; to the authors' knowledge, there is no known infinitely differentiable filter who's atom allows an explicit closed form expression. As such, we introduce an efficient exponentially accurate Gabor approximation to infinitely differentiable filters, which allows for explicit representations in both the time and Fourier dual space. Furthermore, we present an alternative error decomposition that allows for the complete rigorous analysis of the error in truncating the signal, and of the error introduced in approximating the filter and atom, Section 3. A direct numerical comparison of a classical finite regular filter, and our infinitely smooth filter, illustrates superior convergence in the interior of the samples, and a similar error near the sampling boundaries, Section 4. By adaptively scaling the filter's Gabor approximation order with the distance from the point to be recovered and the sampling boundaries, $LT - |t|$, the filter's approximation does not adversely effect the overall convergence rate. However, before proceeding, we briefly describe one of a variety of applications where high order accuracy is of importance.

1.1 Oversampling and channel estimation in communications

Channel estimation and equalization in wireless communications is an example where the problem of reconstructing an oversampled bandlimited signal from its samples plays an important role. Data transmission over wireless channels is impaired by multipath propagation [24], where the signal sent from the transmit antenna arrives at the receive antenna via multiple propagation paths, due to reflections from the ground and surrounding structures.

If the transmitter and receiver are not moving the channel is essentially a time-invariant system and can be modeled as convolution operator. More precisely, let $f \in B_\sigma$ be the transmitted signal in a single-carrier communication system[3], i.e., f is of the form

$$f(t) = \sum_{k \in \mathbb{Z}} c_k \psi(t - kR), \tag{6}$$

where $\{c_k\}_{k\in\mathbb{Z}}$ are the data to be transmitted, $\psi \in B_\sigma$ is called the pulse shaping function, and $R \leq \sigma$ is referred to as transmission rate. The received signal g can be written as

$$g(t) = \int_{-\infty}^{+\infty} f(s)h(t-s)\,ds + \eta(t),$$

where h is the impulse response of the channel and η represents additive Gaussian white noise [24].

The process of removing the interference caused by multipath propagation is referred to as equalization. We do not go into detail about the various kinds of equalizers, the interested reader may consult [24] and the references therein. In many equalization methods one solves a system of equations $Ax = y$, where the entries of the matrix A are samples of h and x and y are vectors containing samples of f and g. The impulse response h can be estimated at the receiver by sending pilot signals (ideally a delta-distribution, practically a so-called pseudo-noise signal of the same bandwidth as the transmitted signal) before the actual data transmission. The receiver first performs an analog anti-alias filtering of the pilot signal (i.e., a projection onto B_σ), and then samples the filtered signal at Nyquist rate or higher, to form an estimate of h. The error due to additive white noise can be greatly reduced by oversampling as well as by averaging over several such estimates.

Ideally the receiver would sample h at times $t = \{kR\}_{k\in\mathbb{Z}}$ and the samples $h(kR)$ would be sufficient for a successful equalization. However in practice the transmission rate is often independent from the sampling clock at the receiver [10]. Thus we may

[3]The restriction to single-carrier systems is just for convenience, the described problem is similar in multi-carrier systems

sample h at times $\{kR_s\}_{k\in\mathbb{Z}}$ instead, where the ratio R/R_s is assumed to be irrational. Hence we have to interpolate h at nodes $\{kR\}_{k\in\mathbb{Z}}$ from its oversampled values taken at times $\{kR_s\}_{k\in\mathbb{Z}}$. In theory such an interpolation can be carried out exactly by means of Shannon's sampling theorem; however, in practice one has to work with a finite number of samples, which means we can only compute an approximate interpolation of h from its samples $\{kR\}_{|k|\leq L}$.

Due to the truncation error introduced during the interpolation of h the entries of A are not known exactly. Therefore we are dealing with a perturbed matrix $\tilde{A} := A + E$, where E represents the perturbation caused by the truncation error. The role of this perturbation error becomes significant if the channel experiences so-called deep fades, which in turn means that $\|A^{-1}\|$ will be large (if A is invertible at all). Standard error estimates for perturbed linear systems tell us that the amplification of the error in the solution depends on the quantity $\|A^{-1}\| \cdot \|E\|$, cf. [15]. Therefore it is important to make $\|E\|$ and thus the truncation error as small as possible. This is where the highly localized reconstruction methods proposed in this paper come into play for this application.

2 Oversampling and Localization

The time domain localization of an atom, $\psi(\cdot)$, is reflected in the regularity[4] of its corresponding filter, $\Psi(\cdot)$,

$$|\psi(t)| \leq (2\pi t)^{-s}\|\Psi\|_{C^s} \cdot \frac{2\Omega}{\sqrt{2\pi}} \quad \forall s, \qquad \psi \in B_\Omega \tag{7}$$

where $\|\Psi\|_{C^s} := \|\Psi^{(s)}\|_{L^\infty}$.

Consequently, convergence is gained at the polynomial rate[5]

$$\epsilon(t, L, T) \leq Const \cdot (LT - |t|)^{1-s}(2\pi)^{-s}\|\Psi\|_{C^s} \cdot \left(\frac{\Omega}{\sigma}\right), \quad s \geq 2 \tag{8}$$

as t passes from the boundary, $\pm LT$, to the interior, where $T := r/2\sigma$.

Rather than improving the atom's localization by increasing its corresponding filter's regularity, attempts have been made to construct highly localized atoms by maximizing the atom's local amplitude, $\int_{-R}^{R}\psi^2(t)dt/\int_{-\infty}^{\infty}\psi^2(t)dt$. However, such approaches have resulted in discontinuous filters, [22], and atoms which do not decay globally, [16]. A

[4] This is achieved by s integration by parts, where the derivatives are transfered onto the filter.

[5] If the underlying filter possesses $\|\Psi\|_{C^{s+1}} < \infty$, then the below bound can be tightened by one order of $(LT - |t|)$ to the rate $(LT - |t|)^{-s}$, [12].

much more successful approach for polynomial order filters is to minimize the filter's regularity constant, $||\Psi||_{C^s}$. The classical raised cosine is such a filter, [24],

$$\Psi_{rc}(w) = \begin{cases} 1 & |w| \le \sigma \\ 0 & |w| > \sigma(2-r)/r \\ \frac{1}{2}(1+\cos\left(\frac{\pi}{2}\left(\frac{r}{1-r}\right)\left(\frac{w}{\sigma}-1\right)\right) & \sigma < w < \sigma\frac{2-r}{r} \\ \frac{1}{2}(1+\cos\left(\frac{\pi}{2}\left(\frac{r}{1-r}\right)\left(\frac{w}{\sigma}+1\right)\right) & -\sigma > w > -\sigma\frac{2-r}{r} \end{cases}, \tag{9}$$

where the bounded regularity constants are given explicitly by $||\Psi||_{C^1} = \frac{1}{2}\left(\frac{T}{2(1-r)}\right)$ and $||\Psi||_{C^2} = \frac{1}{2}\left(\frac{T}{2(1-r)}\right)^2$.

When a comparatively small number of sampling points is taken, low regularity constant polynomial order methods give extremely good reconstructions. However, when a larger number of samples is available, atoms with significantly improved asymptotic localization can be achieved by constructing infinitely regular filters, $\Psi \in C_0^\infty$. It should be noted that for infinitely differentiable functions, the optimal bound in equation (7) is not necessarily obtained for large s, as the regularity constant $||\Psi||_{C^s}$ grows rapidly in s. Rather, for functions where precise regularity estimates are known, the optimal s can be determined, resulting in an exponential decay without necessarily large constants. These claims will be realized in the numerical experiments presented in Section 4, contrary to assertion in [3] where it is claimed that the increased regularity does not improve numerical convergence. In a direct numerical comparison with the raised cosine filter, our infinitely differentiable filter achieves dramatically superior convergence in the interior of the samples, and quantitatively similar errors near the sampling boundaries.

2.1 Localization and Gevrey regularity

To achieve exponential accuracy, and satisfy the reproducing condition, (4), requires a filter which is infinitely differentiable and compactly supported. The natural space for infinitely differentiable compactly supported functions is the Gevrey class[6] which consists of functions satisfying the smoothness bound

$$||\rho||_{C^s} := ||\rho^{(s)}||_{L^\infty} \le Const \cdot \frac{(s!)^\alpha}{\eta^s} \Longleftrightarrow \rho \in G^\alpha \tag{10}$$

[6]Compact support is inconsistent with analyticity, G^1, so reproducing atoms can at most be in the space G^α, for $\alpha > 1$. The Gevrey class of functions is essentially similar to ultra-differentiable functions, [23].

where η is a constant independent of s. Incorporating the regularity information in the localization bound, (7), and minimizing over all admissible s, we conclude that Gevrey class filters satisfy a root-exponential localization decay[7],

$$|\psi(t)| \leq Const \cdot \sqrt{|t|} \exp(-\alpha(2\pi\eta|t|)^{1/\alpha}), \quad \Psi \in G^{\alpha}, \tag{11}$$

and root-exponential truncation error

$$\epsilon(t, L, T) \leq Const_{\alpha,\eta} \exp(-(2\pi\eta(LT - |t|))^{1/\alpha}) \tag{12}$$

where $Const_{\alpha,\eta} \sim \eta^{-2} \sum_{l=0}^{q} q!\eta^{l/2}/(q-l)!$, with q the smallest integer greater than or equal to $(3\alpha - 2)/2$.

A similarly localized atom was constructed in [11, 21] by multiplying the sinc function with the inverse Fourier transform of an appropriately dilated G^2 function. Alternatively, such G^{α} filters can be expressed explicitly in the dual space, such as

$$\Psi_{G2}(w) = \begin{cases} 1 & |w| \leq \sigma \\ 0 & |w| > \sigma(2-r)/r \\ \exp[\beta \cdot (w - \sigma(2-r)/r)^{-1} \, e^{-(w-\sigma)^{-1}}] & \sigma < w < \sigma\frac{2-r}{r} \\ \exp[\beta \cdot (-w - \sigma(2-r)/r)^{-1} \, e^{(w+\sigma)^{-1}}] & -\sigma > w > -\sigma\frac{2-r}{r} \end{cases} \tag{13}$$

which is in G^2, [20].

Although the filter $\Psi_{G2}(\cdot)$ and the one in [21] result in rapid convergence while approaching the interior, $|t| \leq LT$, their associated atoms lack an explicit construction. As a result, to reconstruct the bandlimited signal at an arbitrary point requires the costly implementation of a quadrature evaluation, or a global approximation of the atom, such as the Pade approximation proposed in [21]. In the following section we construct an alternative Gabor approximation well suited to the characteristic shape of filter, and which allows for both an efficient approximation, and a rigorous analysis of the overall error introduced in the approximation.

3 An implementation with exponential accuracy

Due to the structure of the desired filters, and the need to compute its inverse Fourier transform, we propose a particular Gabor approximation, discussed in section 3.2, as a highly efficient method to construct atoms based upon general G^{α} regular filters. Furthermore, unlike the Pade approximation, the Gabor decomposition allows a detailed analysis of the error introduced in the filter's approximation. Before preceding in constructing these atoms, we review those properties of Gabor frames which are relevant for our application.

[7] True exponential decay, i.e., $\alpha = 1$, can not be realized, as was shown in the classical paper [2].

3.1 Gabor frames and their time-frequency localization

We briefly review some results about Gabor frames and refer to [13] for details and references. A Gabor system consists of functions of the form

$$g_{na,mb}(t) = e^{2\pi imbt} g(t - na), \quad n, m \in \mathbb{Z}, \tag{14}$$

for some window $g \in \boldsymbol{L}^2(\mathbb{R})$ and $a, b \in \mathbb{R}$. The parameters a and b are the time- and frequency translation parameters respectively.
We say that the triple (g, a, b) generates a *Gabor frame* for $\boldsymbol{L}^2(\mathbb{R})$, for given a and b, if there exist constants (*frame bounds*) $A, B > 0$ such that

$$A\|f\|^2 \leq \sum_{n,m\in\mathbb{Z}} |\langle f, g_{na,mb}\rangle|^2 \leq B\|f\|^2,$$

for any $f \in \boldsymbol{L}^2(\mathbb{R})$, where throughout this section we use the usual $L^2(\mathbb{R})$ inner product. The Gabor frame operator S is given by

$$Sf = \sum_{n,m\in\mathbb{Z}} \langle f, g_{na,mb}\rangle g_{na,mb}, \qquad f \in \boldsymbol{L}^2(\mathbb{R}),$$

and satisfies $AI \leq S \leq BI$, where I is the identity operator of $\boldsymbol{L}^2(\mathbb{R})$. If $\{g_{na,mb}\}_{n,m\in\mathbb{Z}}$ is a Gabor frame for $\boldsymbol{L}^2(\mathbb{R})$, then any f in $\boldsymbol{L}^2(\mathbb{R})$ can be written in the form

$$f = \sum_{n,m\in\mathbb{Z}} \langle f, \gamma_{na,mb}\rangle g_{na,mb} = \sum_{n,m\in\mathbb{Z}} \langle f, g_{na,mb}\rangle \gamma_{na,mb}, \tag{15}$$

with the *canonical dual frame* $(\gamma_{na,mb})_{n,m\in\mathbb{Z}}$ given by

$$\gamma_{na,mb} = e^{2\pi imbt} \gamma(t - na), \quad n, m \in \mathbb{Z}, \tag{16}$$

where $\gamma = S^{-1}g$ is referred to as the *canonical dual window.* Gabor frames for $\boldsymbol{L}^2(\mathbb{R})$ can only exist if $ab \leq 1$ while Gabor Riesz bases can only exist if $ab = 1$. Thus the ratio $\frac{1}{ab}$ measures the redundancy of Gabor system.
Since one key motivation for considering Gabor frames is to obtain a joint time-frequency localized representation of functions and operators one usually attempts to choose the window g to be well localized in time and frequency. While the Balian–Low theorem precludes Gabor Riesz bases with good time-frequency localization, it is not difficult to design Gabor frames with excellent localization properties. For instance, if g is a Gaussian or a hyperbolic secant, we obtain a Gabor frame whenever $ab < 1$. It is clear that for a truly local time-frequency representation, not only g, but also its dual γ has to be well concentrated in the time-frequency plane. The problem under which

conditions γ will inherit the time-frequency decay properties of g have been studied in detail in, e.g., [8, 14, 17, 27].
For the efficient approximation of G^α filters, we select the Gaussian which is invariant under the Fourier transform defined in equation (1). This particular window and its dual, were investigated extensively by Janssen in [18].

Proposition 3.1 *Let $g(t) = 2^{\frac{1}{4}} \exp(-\pi t^2)$ and $(ab)^{-1} = 2q, q \in \mathbb{N}$. Then the canonical dual window γ is given by*

$$\gamma(t) = \frac{ab}{\vartheta_3(\pi t/a)} \sum_{k=-\infty}^{\infty} c_k g(t - k/b), \tag{17}$$

where $c(y) := \exp(-\pi/2y^2)$,

$$c_k = \frac{\sum_{m=0}^{\infty}(-1)^{m+k} c(b)^{(m+\frac{1}{2})(2|k|+m+\frac{1}{2})}}{\sum_{n=-\infty}^{\infty}(-1)^n(n+\frac{1}{2})c(b)^{(n+\frac{1}{2})^2}},$$

and

$$\frac{1}{\vartheta_3(z)} := \frac{2\sum_{k=-\infty}^{\infty}(-1)^k e^{2ikz}\sum_{m=0}^{\infty}(-1)^m c(a)^{(m+1/2)(2|k|+m+1/2)}}{\sum_{n=-\infty}^{\infty}(-1)^n(2n+1)c(b)^{(m+1/2)(2|k|+m+1/2)}},$$

with $\vartheta_3(z) := \vartheta_3(z, c(a))$ denoting the third theta function in the notation of [28]. Furthermore, γ is analytic in the strip $|\mathrm{Im}(t)| < \frac{\pi}{4a}$, and decays at the rate,

$$|\gamma(t)| \leq \exp(-\pi|t|/2b). \tag{18}$$

Equation (17) can be found in Section 6 of [18]. As noted by Janssen, [19], the analyticity follows from the fact that the $1/\vartheta_3(z)$ has its poles exactly at $(m+1/2)a + (n+1/2)i/2a$, see [28], Subsec. 21.12.

3.2 An efficient approximation and complete error analysis

Rather than constructing an approximation of the atom directly, as in [21], we construct an approximation to the filter, which allows for an explicit construction of the atom, and permits a complete error analysis. From this approach we can view the approximation error in terms of sacrificing the reproducing property in the dual space, (4), and adapt the approximation to carefully balance the truncation and reproducing errors.
In practice, only a finite number of sampling points are available, and as a result, the exact reproducing property of Shannon's sampling theorem is lost. The introduced error is quantified in the truncation error, (3), which we have shown to be root-exponentially small, (11), for the Gevrey-class filters, $\Psi \in G^\alpha$. By relaxing the reproducing property,

we construct a filter that allows for explicit representations of both the filter and its atom. Towards this end we introduce a new error decomposition for the recovery of a bandlimited signal with sampling rate, $T := r/2\sigma$, using an arbitrary atom $\psi(\cdot)$ and its associated filter $\Psi(\cdot)$, not necessarily satisfying the reproducing property, (4).

$$\begin{aligned} \mathrm{Error}(L,t,r) &:= f(t) - \sum_{|k|\le L} f(kT)\psi(t-kT) \\ &= \left(f(t) - \sum_{k=-\infty}^{\infty} f(kT)\psi(t-kT) \right) + \sum_{|k|>L} f(kT)\psi(t-kT) \end{aligned} \tag{19}$$

The first error component, the reproducing error, exactly vanishes for filters satisfying condition (4), and the second component is the traditional truncation error. With the goal of achieving minimal error in the approximation of the bandlimited signal, there is no reason to enforce the reproducing property, (4). Rather, the filter should be selected to minimize the overall error given in equation (19). Below we illustrate how a truncated Gabor decomposition allows for the explicit construction of both the filter and its associated atom, where the error introduced by truncating the Gabor expansion, is expressed in the reproducing error. We now turn our attention to the construction of such a non-reproducing filter.

To allow for the explicit representation of the filter and its associated atom we compute the Gabor expansion of an arbitrary filter where the synthesis window has a known inverse Fourier transform. Practical implementation will require that only a finite number of Gabor coefficients are used, making this an efficient approach if the analysis window is analytic, allowing the Gabor coefficient decay rate and the filters localization to be similar, as they are both determined by the filters regularity. In particular, root-exponential decay for filters, $\Psi \in G^{\alpha}$. To represent the filter and atom as finite combinations of Gaussians, we compute the Gabor coefficients from the Gaussian's canonical dual, equation (17). The resulting filter and atom are then given explicitly by

$$\Psi_{M,N}(w) := \sum_{|m|\le M} \sum_{|n|\le N} \langle \Psi, \gamma_{na,mb} \rangle \exp(-2\pi i m b w) g(w-na) \tag{20}$$

$$\psi_{M,N}(t) := \sum_{|m|\le M} \sum_{|n|\le N} \overline{\langle \Psi, \gamma_{na,mb} \rangle} e^{2\pi i m n a b} \exp(-2\pi i n a t) g(t-mb) \tag{21}$$

where $g(t) := 2^{1/4}\exp(-\pi t^2)$ is the Gaussian invariant under the Fourier transform defined in equation (1).

As the filter is in G^{α}, and the canonical dual is analytic, [19], the Gabor coefficients in equations (20,21) decay at the root-exponential rate in the modulation term m. Additionally, as the filter is compactly supported in $[-\sigma, \sigma]$, and the Gaussian's dual decays exponentially, equation (18), the coefficients decay exponentially in the translation term, n, once away from the filter, i.e., $|n| > \sigma/a$. More specifically,

$$\begin{aligned}|\langle\Psi, \gamma_{na,mb}\rangle| &\le (2\pi b|m|)^{-s}\|\gamma(t-na)\Psi(t)\|_{C^s}T^{-1} \\ &\le Const_{\Omega,b}|m|^{3/2}\cdot e^{-c(|n|+|m|^{1/\alpha})}, \qquad \Psi \in G^{\alpha}[-\Omega,\Omega] \qquad (22)\end{aligned}$$

where $Const_{\Omega,b} := \Omega\eta^{3/2}b^{3/2}\exp(\pi\Omega/2b)\|\Psi\|_{C^0}$ and $c := \min(\pi a/2b, \alpha(2\pi b\eta_{min})^{1/\alpha})$ with $\eta_{min} := \min(\eta_f, \eta_\Psi)$.
The reproducing error is then given by the truncation of the Gabor expansion,

$$\begin{aligned}\left|f(t) - \sum_{k=-\infty}^{\infty} f(kT)\psi_{M,N}(t-kT)\right| &\equiv \left|\int_{-\infty}^{\infty} F(w)[\Psi(w) - \Psi_{M,N}(w)]e^{iwt}dw\right| \\ &\le Const\cdot\|F\|_{L^\infty}\sum_{|m|>M}\sum_{|n|>N}|\langle\Psi,\gamma_{na,mb}\rangle| \qquad (23) \\ &\le C_\alpha\cdot\|F\|_{L^\infty}e^{-c(N+M^{1/\alpha})} \qquad (24)\end{aligned}$$

with $C_\alpha := Const_{\Omega,b}\cdot c^{-1}\sum_{l=0}^{q} q!c^{-l}/(q-l)!$, where q is the smallest integer greater than or equal to $\frac{5}{2}\alpha - 1$.
As the atom is composed of translated and dilated Gaussians, who's amplitude decay root-exponentially in the translation term, and as the Gaussians are only translated to $\pm 2\pi Mb$, the atom's localization is given by

$$|\psi(t)| \le Const_{\Omega,b}N\begin{cases}\exp(-c|t|^{1/\alpha}) & |t| \le Mb \\ e^{-cMb}\exp(-\pi(|t| - 2\pi Mb)^2) & |t| > Mb.\end{cases} \qquad (25)$$

when $\Psi \in G^\alpha$. Due to the superexponential decay for $|t| > Mb$, the composite error is controlled by the reproducing error for points away from the boundary, $LT - |t| > Mb$. The above results are summarized in the following theorem.

Theorem 3.2 *Let $f(t)$ be a signal bandlimited to $[-\sigma,\sigma]$, and $\Psi(\cdot)$ a filter in G^α satisfying equation (4). The atom $\psi_{M,N}(\cdot)$, given in (21) can be used to recover the signal, $f(t)$, from its sampled values, $\{f(kT)\}_{|k|\le L}$, where $T := r/2\sigma$ and $r < 1$, with the error bounded by*

$$\begin{aligned}|Error(L,t,r)| &:= \left|f(t) - \sum_{|k|\le L} f(kT)\psi_P(t-kT)\right| \qquad (26) \\ &\le C_\alpha N\|F\|_{L^\infty}\left(e^{-c(N+M^{1/\alpha})} + e^{-c(LT-|t|)^{1/\alpha}}\right) \qquad (27)\end{aligned}$$

where c and C_α are defined as before. Moreover, the signal can be recovered at the fractional power exponential rate, $\exp(LT - |t|)^{1/\alpha}$*, if the number of Gabor coefficients are selected at, at least the rates* $M \sim LT - |t|$ *and* $N \sim M^{1/\alpha}$*.*

4 Numerical Experiments

Using highly regular filters, $\Psi \in G^\alpha$, allow for atoms, $\psi(\cdot)$, with exponential localization as seen in equation (11). Consequently, a bandlimited signal, $f \in B_\sigma$, can be recovered with the exponential accuracy indicated in Theorem 3.2. We now verify this theorem, and contrast the exponential accuracy to the more common polynomial order accuracy obtained by filters possessing only a finite order regularity. To compare representative filters with finite and infinite regularity, we use the canonical raised cosine filter, (9), and the Gevrey order two filter given in equation (13), respectively.
Much of the success of the raised cosine filter is due to the optimally small first two regularity constants, $\|\Psi\|_{C^s}$ for $s = 1, 2$, which result in rapid initial localization. Infinitely regular filters posses bounded regularity constants for all s, but at the cost of necessarily larger regularity constants for small s. However, a great deal of freedom exists in the selection of G^α regular filters, for example the constant β used in the filter of equation (13). Although there is no closed form solution for the β which minimizes the first regularity constant, a good approximation can be obtained for a large range of oversampling by selecting β such that the filter's points of inflection are at the middle of the region connecting zero and one, i.e. $\Psi^{(2)}(\pm\sigma/r) = 0$. As such, for the numerical experiments involving the filter Ψ_{G2}, we use the β which satisfies this inflection condition

$$\beta := \frac{(1+2z^2)z}{(1+z)^2} e^{1/z} \quad \text{where} \ \ z := \sigma(r^{-1} - 1).$$

For the approximation of an arbitrary bandlimited signal, we form a test signal whose Fourier transform is composed of one hundred characteristic functions with random complex valued amplitudes normalized to unit l_2 norm; and with random widths and centers, normalized so that the largest magnitude bandwidth is σ. The resulting numerics shown, are characteristic of arbitrary complex valued bandlimited signals. The dual space representation, real portion, of such a function can be seen in Figure 2, which is the signal whose approximation error is shown in Figures 3 and 4. We now illustrate the convergence properties for the implementation described in Theorem 3.2. For a given signal, and oversampling rate, Theorem 3.2 indicates that exponential accuracy is gained as the point to be recovered moves from the sampling boundary to the interior, and as the number of terms in the truncated Gabor decomposition,

N and M, are increased. Ideally we would like to maximize the lattice dependent value of c given in equation (22), however, this requires knowledge of the filter and bandlimited signal's regularity coefficients, $\eta_{\Psi,f}$, which are traditionally not known. Alternatively, we use a tiling of the time-frequency plane most appropriate for our Gabor synthesis window. For our synthesis window, we use the Gaussian which is invariant under the Fourier transform, and which tiles the time-frequency plane with elements centered at the lattice points, and possessing circular decay in the plane. This suggests we select equal translation and modulation parameters, and for computational robustness a redundancy of two, $a = b := 2^{-1/2}$. It should also be noted that the Gaussian's dual, γ, can easily be computed with exponential accuracy as the terms involved in the infinite series decay exponentially, see equation (17), cf. [27] for a proof and [25] for details about a fast numerical method to compute γ. Furthermore, the dual need only be computed once, and its stored values used to compute the Gabor coefficients, $\langle \Psi, \gamma_{na,mb} \rangle$. For simplified notation we express the filter, $\Psi_{N,M}$, as simple Ψ_N, where M is given by $M := \text{round}(2^{-1/2}N^2)$ to balance the decay rates in the Gabor decomposition. Figure 3 illustrates the behavior of the error as a function

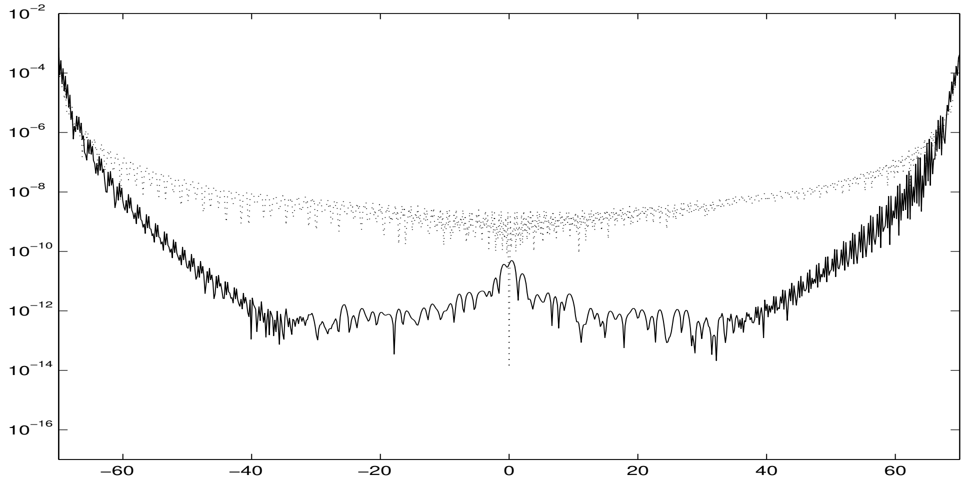

Figure 3: The $\log_{10}$ of the error in recovering a signal in B_1, with oversample rate $r = 0.7$. Note the log convergence exhibited by the reconstruction using the raised cosine(dotted), as compared to the root convergence obtained by the truncated Gabor expansion Ψ_{15} of Ψ_{G2}, (solid). In the interior of the sample, $|t| \lesssim 30$, the error is dominated by the reproducing error, $\|\Psi_{15} - \Psi_{G2}\|_{L^\infty} \leq 2.1^{-10}$, whereas near the boundaries, the error is dominated by the root-exponential localization error, see Theorem 3.2.

of the distance to the sampling boundaries, and the number of terms in the Gabor

decomposition, see Theorem 3.2. Away from the boundaries, $|t| \lesssim 30$, the error is dominated by the reproducing errors introduced in the truncated Gabor decomposition; whereas near the boundaries the error is controlled by the truncation errors.
By balancing the reproducing and truncation errors, as noted in Theorem 3.2, root-exponential accuracy can be achieved while moving away from the sampling boundaries, by selecting the number of terms in the Gabor decomposition as a function of the distance to the sampling boundaries. Figure 4 illustrates the error for recovering the same bandlimited signal used in Figure 3, but with a spatially adaptive number of Gabor coefficients, $N := 5 + (LT - |t|)/4$, and $M = 2^{-1/2}N^2$, where 5 is the minimum number of Gabor coefficients needed for a qualitatively accurate approximation of the filter. Figures 3 and 4 illustrate the superior accuracy obtained by using highly regular filters, as indicated in Theorem 3.2.

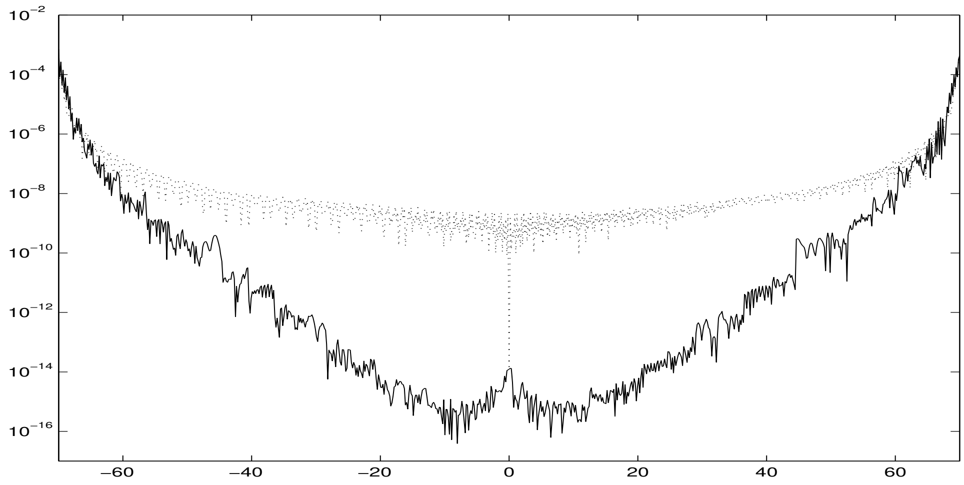

Figure 4: The $\log_{10}$ of the error in recovering a signal in B_1, with oversample rate $r = 0.7$. Note the log convergence exhibited by the reconstruction using the raised cosine(dotted), as compared to the root convergence obtained by the truncated Gabor expansion Ψ_N of Ψ_{G2}, (solid), where $N := 5 + (LT - |t|)/4$, and $M = 2^{-1/2}N^2$. Unlike the error in Figure 3, the root-exponential accuracy is not halted near the interior of the domain, and the similar accuracy near the sampling boundaries is achieved with significantly fewer Gabor coefficients.

5 Final remarks

It has been well understood that the oversampling of bandlimited signals allows for the reconstruction with smooth filters, and correspondingly well localized atoms. More precisely that root-exponential accuracy can be achieved with infinitely differentiable filters, [11, 21]. However, such filters have resisted explicit construction of their atoms, and as such practical implementation has been limited. Here we illustrated an efficient Gabor approximation of a large class of smooth filter, and a new error formulation that allows a rigorous analysis of the overall error introduced in the approximation of the atom, as well as the truncation of the samples. Although this implementation does result in root-exponential convergence, it remains computationally intensive to obtain the full reconstruction accuracy, scaling like $(LT-|t|)^3$. Even so, as discussed in section 1.1, there are applications which require such high resolution approximations.
With the goal of developing highly localized fast algorighms, future work should consider implementations where the periodization induced by sampling is removed through the direct action of the filter in the Fourier dual space, and hence avoiding the construction of an atom. Moreover, we have only considered uniform sampling, since if the sampling geometry can be chosen in advance, uniform sampling is optimal in many applications. However, in several applications one is confronted with nonuniformly spaced samples, see e.g. [1, 9] for theory and algorithms for the reconstruction of bandlimited functions from nonuniform sampling. The methods developed in [5, 7] are based on the convolution identity $f = f * \psi$ for $f \in B_\sigma$, where $\Psi(\omega) = 1$ for $\omega \in [-\sigma, \sigma]$ and $\Psi(\omega) = 0$ for $|\omega| > \sigma_0$ for properly chosen $\sigma_0 > \sigma$. While these methods are highly localized the truncation error analysis given in [6, 7] is based on the unrealistic assumption that in certain steps of the reconstruction all (infinitely many!) samples are available. Unlike for uniform sampling the truncation to finitely many samples can lead to severe instabilities of certain reconstruction methods, as pointed out in [26]. While approximation algorithms based on trigonometric polynomials are stable and fast [26], they do not achieve high convergence rate with respect to truncation error. A natural extension of the work presented in this paper would therefore be to derive a complete truncation error analysis for a highly localized reconstruction method for the case of nonuniform sampling.

References

[1] John J. Benedetto and Paulo J. S. G. Ferreira, editors. *Modern sampling theory*. Applied and Numerical Harmonic Analysis. Birkhäuser Boston Inc., Boston, MA, 2001. Mathematics and applications.

[2] A. Beurling and P. Malliavin. On the closure of characters and the zeros of entire functions. *Acta Math.*, 118:79–93, 1967.

[3] I. Daubechies. *Ten Lectures on Wavelets*, volume 61 of *Regional conference series in applied mathematics*. SIAM, 1992.

[4] I. Daubechies and R. DeVore. Approximating a bandlimited function using very coarsely quantized data: A family of stable sigma-delta modulators of arbitrary order. *Annals of Mathematics*, 158:679–710, 2003.

[5] H. G. Feichtinger and K. Gröchenig. Iterative reconstruction of multivariate bandlimited functions from irregular sampling values. *SIAM J. Math. Anal.*, 231:244–261, 1992.

[6] H. G. Feichtinger and K. Gröchenig. Error analysis in regular and irregular sampling theory. *Applicable Analysis*, 50:167–189, 1993.

[7] H. G. Feichtinger and S.S. Pandey. Error estimates for irregular sampling of bandlimited functions on a locally compact abelian group. *J. Math. Anal. Appl.*, 279(2):380–397, 2003.

[8] H.G. Feichtinger and K. Gröchenig. Gabor frames and time-frequency analysis of distributions. *J. Funct. Anal.*, 146(2):464–495, 1996.

[9] H.G. Feichtinger and K.H. Gröchenig. Theory and practice of irregular sampling. In J. Benedetto and M. Frazier, editors, *Wavelets: Mathematics and Applications*, pages 305–363. CRC Press, 1994.

[10] F.M. Gardner. Interpolation in digital modems—part I: Fundamentals. *IEEE Trans. Comm.*, 41(3):501–507, 1993.

[11] R. Gervais, Q.I. Rahman, and G. Schmeisser. A bandlimited function simulating a duration-limited one. In *Anniversary volume on approximation theory and functional analysis (Oberwolfach, 1983)*, volume 65 of *Internat. Schriftenreihe Numer. Math.*, pages 355–362. Birkhäuser, Basel, 1984.

[12] D. Gottlieb and S.A. Orszag. *Numerical Analysis of Spectral Methods: Theory and Applications*, volume 26 of *Regional conference series in applied mathematics*. SIAM, 1977.

[13] K. Gröchenig. *Foundations of Time-Frequency Analysis.* Birkhäuser, Boston, 2001.

[14] K. Gröchenig and M. Leinert. Wiener's lemma for twisted convolution and Gabor frames. *J. Amer. Math. Soc.*, to appear.

[15] R.A. Horn and C.R. Johnson. *Matrix analysis.* Cambridge University Press, Cambridge, 1990. Corrected reprint of the 1985 original.

[16] S. Huestis. Optimum kernels for oversampled signals. *J. Acoust. Soc. Amer.*, 92:1172–1173, 1992.

[17] A.J.E.M. Janssen. On rationally oversampled Weyl-Heisenberg frames. *Signal Process.*, 47:239–245, 1995.

[18] A.J.E.M. Janssen. Some Weyl-Heisenberg frame bound calculations. *Indag. Mathem.*, 7(2):165–182, 1996.

[19] A.J.E.M. Janssen. Personal communication, 2003.

[20] Fritz John. *Partial differential equations*, volume 1 of *Applied Mathematical Sciences.* Springer-Verlag, New York, fourth edition, 1991.

[21] Frank Natterer. Efficient evaluation of oversampled functions. *J. Comput. Appl. Math.*, 14(3):303–309, 1986.

[22] R. A. Niland. Optimum oversampling. *J. Acoust. Soc. Amer.*, 86(5):1805–1812, 1989.

[23] Stevan Pilipović and Nenad Teofanov. Wilson bases and ultramodulation spaces. *Math. Nachr.*, 242:179–196, 2002.

[24] T.S. Rappaport. *Wireless communications: Principles & Practice.* Prentice Hall, New Jersey, 1996.

[25] T. Strohmer. Numerical algorithms for discrete Gabor expansions. In H.G. Feichtinger and T. Strohmer, editors, *Gabor Analysis and Algorithms: Theory and Applications*, chapter 8, pages 267–294. Birkhäuser, Boston, 1998.

[26] T. Strohmer. Numerical analysis of the non-uniform sampling problem. *J. Comp. Appl. Math.*, 122(1–2):297–316, 2000.

[27] T. Strohmer. Approximation of dual Gabor frames, window decay, and wireless communications. *Appl. Comp. Harm. Anal.*, 11(2):243–262, 2001.

[28] E. T. Whittaker and G. N. Watson. *A course of modern analysis.* Cambridge University Press, Cambridge, 1996. Reprint of the fourth (1927) edition.

[29] A.I. Zayed. *Advances in Shannon's Sampling Theory.* CRC Press, Boca Raton, 1993.