Automated Filter Selection for Suppression of Gibbs Ringing Artefacts in MRI

Yue Wang[a], John J. Healy[a,*]

[a]*School of Electrical and Electronic Engineering, University College Dublin, Belfield, Dublin, Ireland*

Abstract

Gibbs ringing creates artefacts in magnetic resonance images that can mislead clinicians. Reconstruction algorithms attempt to suppress Gibbs ringing, or an additional ringing suppression algorithm may be applied post reconstruction. Novel reconstruction algorithms are often compared with filtered Fourier reconstruction, but the choices of filters and filter parameters can be arbitrary and sub-optimal. Evaluation of different reconstruction and post-processing algorithms is difficult to automate or subjective: many metrics have been used in the literature. In this paper, we evaluate twelve of those metrics and demonstrate that none of them are fit for purpose. We propose a novel metric and demonstrate its efficacy in 1D and 2D simulations. We use our new metric to optimise and compare 17 smoothing filters for suppression of Gibbs artefacts. We examine the transfer functions of the optimised filters, with counter-intuitive results regarding the highest-performing filters. Our results will simplify and improve the comparison of novel MRI reconstruction and post-processing algorithms, and lead to the automation of ringing suppression in MRI. They also apply more generally to other applications in which data is captured in the Fourier domain.

Keywords: image quality metrics, MR image reconstruction, dark rim artefact, Gibbs ringing artefact, magnetic resonance imaging, filtered Fourier reconstruction, filter comparison

1. Introduction

Gibbs ringing is an artifact that arises when a discontinuous signal is truncated in the Fourier domain. The reconstruction does not converge correctly near the discontinuity, but instead oscillates with peak error near the discontinuity of approximately 9% of the jump. Additional samples reduce the oscillation, but the peak error is inevitable from this kind of reconstruction. One situation in which Gibbs ringing arises is magnetic resonance imaging (MRI). An MR image is reconstructed from samples measured one at a time in the Fourier domain. Gibbs ringing is inevitable in MRIs, where many tissue boundaries are discontinuous, and a finite number of measurements must be taken to limit motion blur, maximise throughput of an expensive hospital resource, and minimise patient discomfort. Modern compressive sensing techniques result in the same number of samples in less time and don't impact this discussion. Gibbs oscillation creates artefacts at tissue boundaries that are known as dark rim artefacts (DRAs) and can lead to misdiagnosis. The traditional method for suppressing Gibbs artefacts is to apply a linear low-pass (i.e. smoothing) filter to the Fourier reconstruction. It is usually implemented by multiplying the data in the Fourier domain by the filter's transfer function prior to Fourier reconstruction. We refer to this as filtered Fourier reconstruction.

A number of modern techniques have been proposed for the suppression of Gibbs ringing, including reprojection methods [1, 2, 3, 4], sub-voxel shifts [5, 6], and machine learning methods [7, 8, 9, 10]. Most authors evaluate new methods by comparison with filtered Fourier reconstruction, but authors are inconsistent in terms of the specific filter or filters used in the comparison and in terms of the nature of the comparison. Some of the comparisons are partly or wholly qualitative, especially in the medical literature.

In this paper, we examine some of the most commonly used metrics for quantitative comparison. We reject twelve metrics in all as unfit for purpose, and present an argument for the use of the ℓ^0_ϵ sparsity metric. We also make a comparison of 17 commonly used filters, including those used in previous bench-marking exercises, and recommend

*john.healy@ucd.ie

six filters as the best performing filters. We anticipate that the adoption of the metric we recommend will make it significantly easier to benchmark reconstruction algorithms and post-processing algorithms consistently. Our work is also a major step towards the automation of the process of Gibbs suppression, currently performed by manual parameter search in a post-processing step by a radiologist.

The remainder of this paper is structured as follows. In Section 2, we 0investigate candidate metrics by evaluating their ability to answer the following question: which is the best Gaussian filter for a given test signal. In σ, we then have a one-parameter search space. We identify a leading candidate metric, and further examine the reliability of its recommendation. In Section 3, we compare 17 filters using the metric identified in Section 2. The filters are applied in the filtered Fourier reconstruction of different Fourier domain truncations of our test signal. For each truncation, the best parameters are determined for each filter and the score is recorded. We can then clearly identify the best performing filters. Finally, we present our conclusions.

2. A Metric to Compare Gibbs Ringing Suppression Techniques

In this section, we compare 13 metrics for quantifying the performance of Gibbs ringing suppression techniques. First, we introduce and define the metrics. Then, we introduce our test signal. Finally, we make a comparison of the metrics to identify which of them can discriminate between algorithms.

2.1. Candidate Metrics

The metrics that we are going to discuss and compare are given in Tab. 2.1. Also given in the table are citations for papers on Gibbs ringing that make use of those metrics. Not all of the papers mentioned in the Introduction appear on this table: the exceptions perform more qualitative comparisons. Some plot the logarithmic point-wise error, and evaluate the results by eye [1, 2, 3, 4]. The frequent resort to human evaluation of ringing suppression is indicative of the difficulty of the problem addressed in this Section. This has been noted elsewhere, e.g. "The frustrating disagreement between numerical metrics and image quality perception by human observers is not easy to resolve." [11]. Some computer vision researchers have even used Amazon's Mechanical Turk platform for such evaluations because of the inadequacy of common metrics [12, 13]. We also note that some authors have made use of metrics in their work which are non-trivial to cast in an 1D form, e.g. the Structural Similarity Index Measure (SSIM) [7, 9, 10]. We have neglected such metrics in this discussion.

Table 1: Table of metrics compared in Section 2, and references that use them for quantifying the performance of Gibbs ringing suppression in MRI reconstruction.

#	Name	References
1	Peak signal-to-noise ratio (PSNR)	[7, 9, 10, 14, 15]
2	Mean squared error (MSE)	[9, 14]
3	Entropy	[14, 15]
4	Variance	[15]
5	Variance of error	[14]
6	Maximum error	[16]
7	Signal-to-Noise Ratio (SNR)	[17]
8	Energy	[15]
9	Correlation	[6, 15]
10	High frequency error norm (HFEN)	[7]
11	Power spectral ratio	[8]
12	Edge Preservation Index (EPI)	
13	ℓ^0_ϵ	

The definitions of the metrics are given in the following equations. In these definitions, $g(t)$ is the signal to be reconstructed. It has Fourier transform $G(f)$. The samples of $g(t)$ are given by $g[n] = g(nT)$, where T is the sampling period, and the samples of $G(f)$ are similarly given by $G[m] = G(m/NT)$. The reconstruction is denoted by r, and

has discrete Fourier transform R. We specifically denote the Fourier (i.e. unfiltered) reconstruction r_F. $e = g - r$ is the reconstruction error. The mean of a function is depicted by an overline, e.g. $\overline{e}$ represents the mean of the reconstruction error. N refers to the number of samples.
Peak Signal-to-Noise Ratio (PSNR):

$$10 \cdot \log_{10}\left(\frac{(\max(|r|)^2}{\frac{1}{N}\sum_{n=1}^{N}|e[n]|^2}\right) \tag{1}$$

Mean Squared Error (MSE):

$$\frac{1}{N}\sum_{n=1}^{N}|e[n]|^2 \tag{2}$$

Entropy:

$$-\sum_{n=1}^{N} r[n]\log_2 r[n] \tag{3}$$

Variance:

$$\frac{1}{N-1}\sum_{n=1}^{N}|r[n]-\overline{r}|^2 \tag{4}$$

Variance of error:

$$\frac{1}{N-1}\sum_{n=1}^{N}|e[n]-\overline{e}|^2 \tag{5}$$

Maximum absolute error is simply given by $\max(|e|)$.
Signal-to-Noise Ratio (SNR):

$$10 \cdot \log_{10}\frac{\sum_{n=1}^{N} g[n]^2}{\sum_{n=1}^{N}|e[n]|^2} \tag{6}$$

Energy is given by the ℓ^2 norm of the reconstruction.

$$\sum_{n=1}^{N}|r[n]|^2 \tag{7}$$

Correlation:

$$\frac{\sum_{n=1}^{N}(g[n]-\overline{g})\cdot(r[n]-\overline{r})}{\sqrt{\sum_{n=1}^{N}(g[n]-\overline{g})^2\cdot\sum_{n=1}^{N}(r[n]-\overline{r})^2}} \tag{8}$$

High Frequency Error Norm (HFEN) is given by the ℓ^2 norm of the Laplacian-of-Gaussian of the error.

$$\sum_{n=1}^{N}|\mathrm{LoG}(e[n])|^2 \tag{9}$$

Power Spectral Ratio (PSR):

$$\frac{1}{N}\sum_{m=1}^{N}\left(\sqrt{\frac{R^2}{G^2}}\right) \tag{10}$$

The Edge Preservation Index (EPI), β, also known as the Edge Keeping Index has not been previously evaluated as a quantification of Gibbs ringing. It was first proposed by Sattar *et al.* as a measure of edge preservation in noise suppression algorithms for images [18]. Edges are discontinuities, and hence we considered EPI as a plausible candidate for the quantification of Gibbs ringing suppression also. In the following definition, Δ represents the Laplacian operation, and Γ is the gamma function.

$$\beta = \frac{\Gamma(\Delta g - \overline{\Delta g}, \Delta r - \overline{\Delta r})}{\sqrt{\Gamma(\Delta g - \overline{\Delta g}, \Delta g - \overline{\Delta g})\cdot\Gamma(\Delta r - \overline{\Delta r}, \Delta r - \overline{\Delta r})}} \tag{11}$$

The ℓ^0_ϵ metric was proposed in sparse signal processing [19]. It is defined as follows.

$$\|z\|_0 : |e[n]| < \epsilon \tag{12}$$

where $\|z\|_0 = \#\{z \neq 0, j = 1, ..., N\}$ is the ℓ^0 norm. We chose $\epsilon = \frac{\tilde{e_f}}{10}$, where $e_f = \left|g - r_f\right|$, the absolute error of the Fourier reconstruction for the least truncated case we considered, and $\tilde{e_f}$ is the median of e_f. Hence, $\ell^0_{\tilde{e_f}}$ counts the number of sample points with error smaller than the median of Fourier reconstruction's absolute error. This provides a signal-dependent threshold, so comparison of ℓ^0_ϵ for any ϵ has meaning only if the same test signal is used. The threshold could be modified or further refined in the future work, but our choice has proven adequate for our investigations in this paper.

2.2. Test Signal

We evaluate the metrics using a rectangular test function, $g(t) = \text{rect}(\frac{t}{4}) - \frac{1}{2}$. The signal is evaluated at 400 sampling points at a sampling frequency of 20Hz. The function has Fourier transform given by $G(f) = 4 * \text{sinc}(4\pi f)$ with a DC component of -120. We prefer to use an analytic test function over an MRI example because MRI examples have already been sampled and truncated, and so exhibit aliasing and Gibbs ringing, meaning that the ground truth is obscured. The function is plotted in Fig. 1. On the same figure, we plot the signal obtained by sampling $G(f)$, truncating it to remove all frequency components above 7Hz, and performing an inverse discrete Fourier transform.

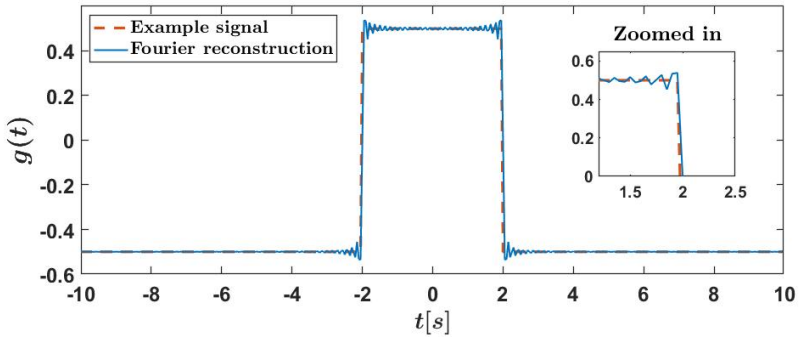

Figure 1: The test signal (red, dashed) is plotted as a function of time (equivalently, a spatial domain in an MRI). The reconstruction (blue) of the signal from truncated Fourier data exhibits Gibbs ringing. An insert zooms in on the ringing near the discontinuity.

In this Section, we limit our comparisons to the filtered Fourier reconstruction using a Gaussian filter, as the filter's standard deviation, σ, provides a 1D parameter space. Its transfer function is defined as $\exp(-f^2/2\sigma^2)$. Large values of σ represent gentle filtering, while small values represent more aggressive filtering. Gentle filtering will do less to reduce ringing, but also blur the signal less. Aggressive filtering will better suppress the ringing at the cost of more blurring. We aim to be able to identify a suitable trade-off between the two. We note that non-linear algorithms such as sub-voxel shifts and machine learning are not constrained by this trade-off in the same way.

2.3. Comparison of Metrics

We propose that when evaluating algorithms that tried to suppress Gibbs ringing, there are three key properties that an ideal metric should have.

- **Punishes deviation**. Large error (e.g. blurring) should result in a low score.
- **Punishes ringing**. Large ringing should result in a low score.
- **Consistency**. Given a continuous parameter, e.g. the cutoff frequency of the truncation, or the σ of the Gaussian filter, the metric should not change erratically with a small change in the parameter.

We evaluate the candidate metrics for filtered reconstructions of the test function (truncated to 7 Hz) for $0.05 < \sigma < 7.5$ Hz. A good metric that exhibits properties 1 and 2 should identify a best σ, which will be highly dependent on the test signal and will also depend on the truncation. This permits a researcher to optimise the filter type: knowing a suitable filter may not be enough, as parameter choice can be difficult and subjective. In fact, currently, clinical expertise is used to judge this by eye, "In our hospital, the technologists prepare and process the images for the physicians to read. The technologists use a slide-bar to adjust the filtering parameters until the images on the screen appear satisfactory" [20].

Most of the metrics exhibit a monotonic behaviour. Those metrics are plotted in Fig. 2. Where necessary, the metrics in this plot are normalized to allow easy comparison using two extremes: when the σ parameter is 10^{-5} (effectively the inverse Fourier transform of the D.C. term only) and the Fourier reconstruction (effectively $\sigma = \infty$), so that they all range from 0 to 1. Correlation, energy, and PSNR all exhibit a monotonic increase with increasing σ, while MSE, Variance, variance of error, max error, HFEN, and SNR all decrease with increasing σ. We have investigated products and linear combinations of these metrics, but none of these yielded a combined metric with a consistent peak. Also plotted is PSR, which shows a similar upward trend to, e.g., PSNR; it is not monotonic, but does not exhibit a useful global or local peak. Consequently, these metrics and their combinations do not have predictive power for the best choice of σ, and so cannot evaluate a Gibbs suppressing algorithm in general.

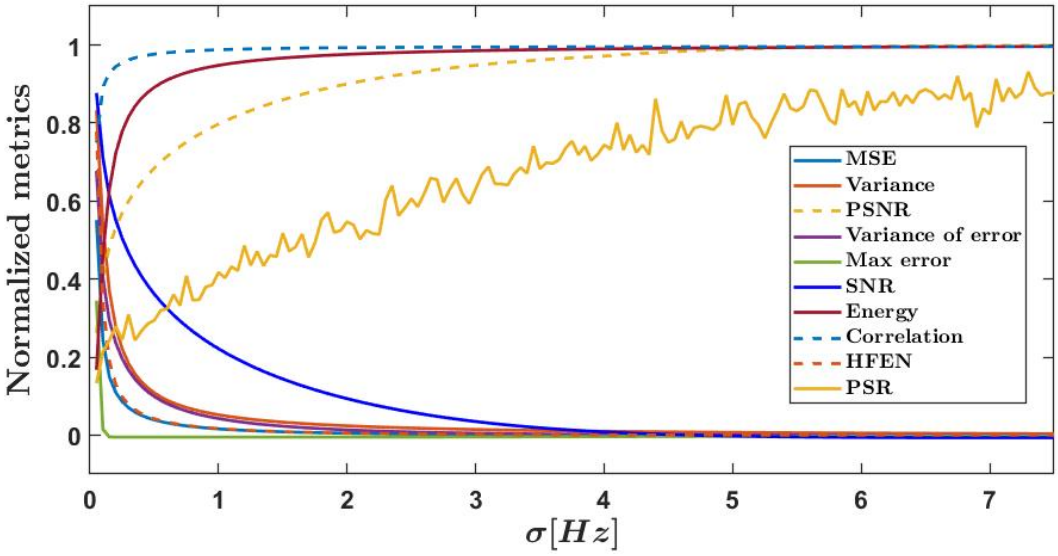

Figure 2: Evaluation of different Gaussian filters using a variety of metrics. We reject all of these metrics as they do not exhibit a peak/nadir.

In Fig. 3, we plot β and entropy as a function of σ, similarly to the metrics in Fig. 2. The behaviour is erratic for aggressively filtered signals; this is not ideal, but perhaps is manageable. Normalisation is performed a little differently comparing with the previous metrics because of this behaviour in small σ. The signal was normalised only by dividing the signal by the value obtained for the Fourier reconstruction. We observe that both metrics present shallow local maximum at plausible best values of σ for the test signal in Fig. 1 (plausible is necessarily subjective; we evaluated this manually); however, this does not appear consistently for other test signals, and hence these two metrics must also be rejected.

Finally, we turn to ℓ^0_ϵ. In Fig. 4, we plot ℓ^0_ϵ as a function of σ, similarly to the metrics in Fig. 2, and Fig. 3. We chose $\epsilon = \tilde{e_f}$. No normalisation is applied. The metric presents a global maximum at a plausible best value of σ. The maxima appear consistently across many test signals we have applied it to. We argue therefore that ℓ^0_ϵ meets the first two properties we set out at the beginning of Section 2.

In order to evaluate ℓ^0_ϵ under the third desirable property - consistency - we create a continuum of test signals by varying the truncation of our rectangular test signal in the Fourier domain. In Fig. 5, we plot the best value of

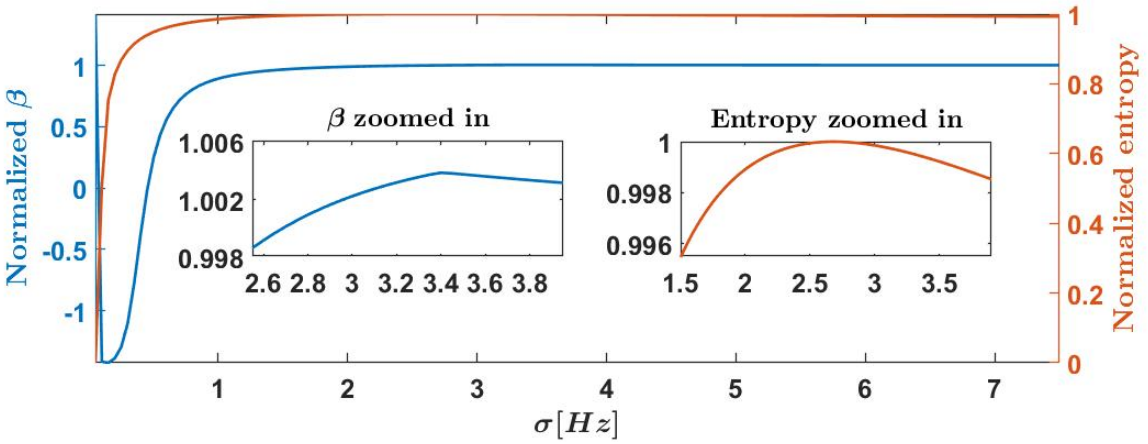

Figure 3: Evaluation of different Gaussian filters using β (EPI) and entropy. The results of EPI are shown in blue while the results of entropy are shown in red. Both metrics exhibit shallow peaks at plausible values of σ, but this does not consistently appear for different test signals, and so we reject these two metrics.

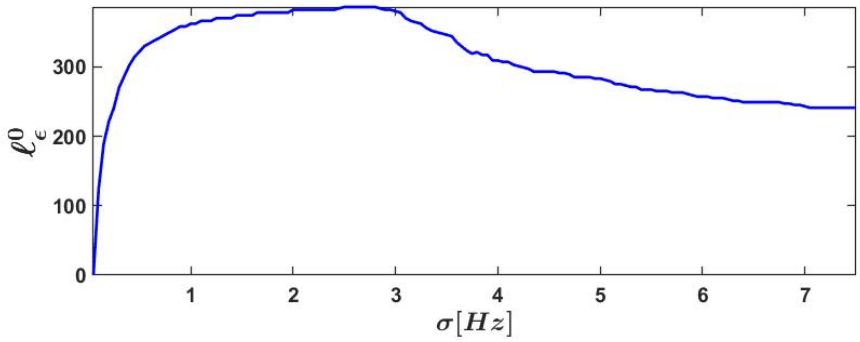

Figure 4: The ℓ^0_ϵ consistently exhibit peaks at plausible values of σ. Consequently, it can be used to rank different algorithms in terms of Gibbs suppression and fidelity to the original signal.

σ identified by ℓ^0_ϵ as a function of the cut-off frequency used in this truncation. We chose $\epsilon = \frac{\tilde{e_f}}{10}$. We argue that a consistent metric should exhibit a smooth monotonic behaviour in this test. Excluding extreme truncation (< 3 samples remaining in the Fourier domain), we observe approximately the desired behaviour in ℓ^0_ϵ. We conclude that ℓ^0_ϵ is highly suited to evaluating Gibbs suppression in a quantitative fashion so we will focus on it for the filters comparison in Section 3.

3. A comparison of filters for filtered Fourier reconstruction

There is very little consistency of choice of filter used in filtered Fourier reconstruction in the literature. Filters used in the literature include:

- Sharpened raised cosine [2]
- Triangle filter [2]
- Exponential filter [2, 4, 14, 16]
- Median filter [6, 15]

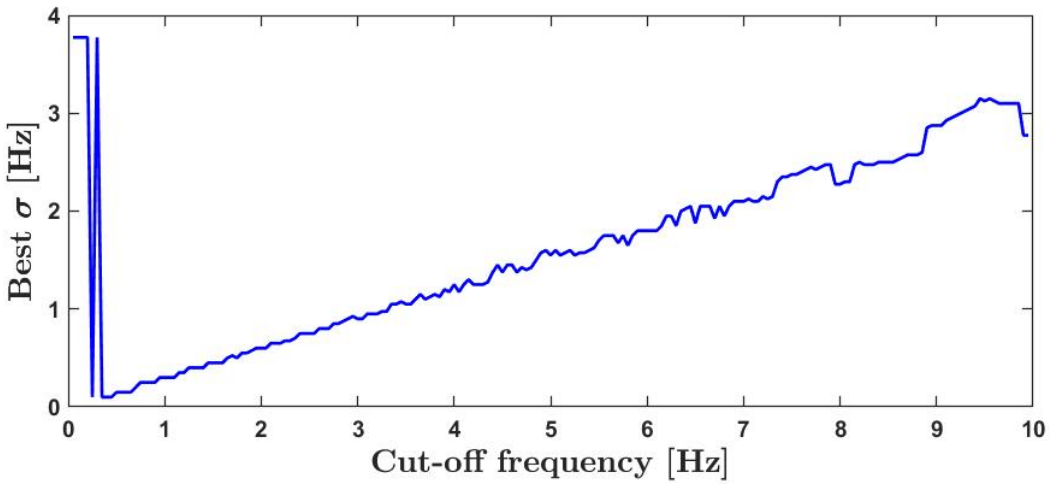

Figure 5: Best choice of σ for a Gaussian filter to use in the filtered Fourier reconstruction of our test signal truncated using different cut-off frequencies.

- Tukey window [7], including special case Hann/raised cosine window [2, 6, 17].
- Hamming window [9]
- Gaussian filter [15, 21]
- Butterworth filter [20]

Since filtered Fourier is the benchmark being used routinely in the literature against which to compare novel algorithms, this variation is a source of confusion for anyone attempting to compare algorithms reported in different papers, unless the filters have roughly comparable performance. In this Section, we apply the ℓ^0_ϵ metric to quantify the performance of 17 filters for Gibbs suppression, including the eight listed above, and nine others: Parzen window, Blackman window, Bohman window, Type-I and Type-II Chebyshev filters, Dolph-Chebyshev window, Flat top weighted window, Kaiser window, and Savitzky Golay filter.

The results of our evaluation are shown in Fig. 6. We generate a collection of versions of our test signal with different cut-off frequencies as in Fig. 5. Most of the filters have one or more parameters, e.g. cut-off frequency, or filter order. For each cut-off frequency, we identify the best performing parameters for each filter type by grid search, and plot its ℓ^0_ϵ score. We chose $\epsilon = \frac{\hat{e}_f}{10}$. We group the 17 filters into three plots for readability. The top figure shows the six best performing filters (identified by median ℓ^0_ϵ score). The middle figure shows the next five best filters, and the remaining six filters are collected in the low-performing plot.

It can be seen from Fig. 6 that the filtered Fourier reconstructions can consistently outperform Fourier reconstruction except for the triangular filter. When the truncation is relatively small (on the right hand side of the figure), most of the filters have similar results. This is a function of the choice of ϵ; we can decrease epsilon to obtain a more sensitive comparison, though we must be careful: epsilon can not be reduced to sensitivity. ϵ provides a noise floor below which all errors are accepted; reducing it demands more accuracy of our reconstruction, rejecting both blurring and ringing. The drawback is that all errors above that floor are treated as equivalent, which is not strictly true perceptually: an artefact will disturb a person far more than a constant offset.

Some of the filters, especially those which perform less well, and the Fourier reconstruction exhibit a periodic peak in Fig. 6. This is related to the sub-voxel shifting approaches that have been proposed for Gibbs suppression [5, 6]. Kellner *et al.* [6] is interested in still images; they use the sub-pixel shifts to obtain a low resolution reconstruction and then up-samples using linear interpolation, noting that other interpolation algorithms regenerate the ringing. Kellner's work is very interesting, and there are subtleties there that merit further investigation. While linear interpolation is easy to use, it is not accurate. The sub-pixel shift is different for each boundary, meaning their calculated image is sampled

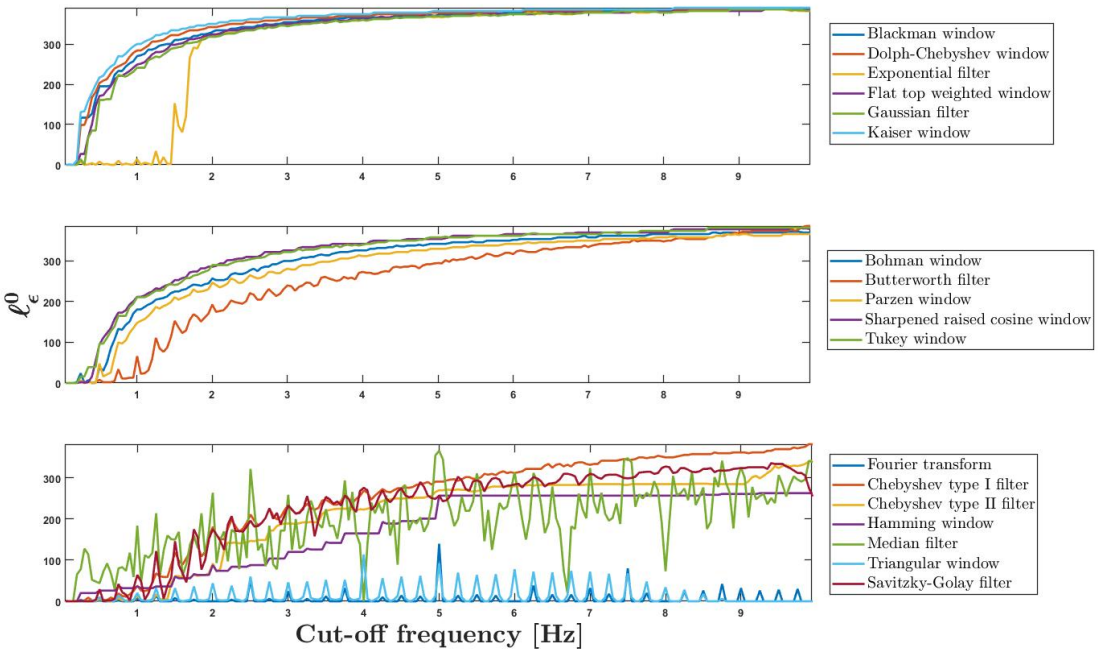

Figure 6: Quantitative evaluation of low-pass filters in filtered Fourier reconstruction for Gibbs ringing suppression for various truncations of our test signal. (top) **The Good**: Blackman window(median score 372), Dolph Chebyshev window(378), exponential filter (368), flat top weighted window(370), Gaussian filter (368) and Kaiser window (382). (middle) **The Bad**: Medium-performing filters, with median scores in the range 280-360. (bottom) **The Ugly**: Low-performing filters, with median scores in the range 3-280. The (unfiltered) Fourier reconstruction (median score 0) is included for reference.

on a non-cartesian grid and requires gridding before it can be displayed. This also is a source of error. However, it is likely an excellent solution within certain constraints. Our work is not in conflict: Kellner uses a measurement of high frequency power, not dissimilar to the HFEN we examine in this manuscript. Our metric could take the place of that in their algorithm. Kellner also evaluates the results entirely qualitatively. Our approach will allow a quantitative comparison of the sub-voxel shift approach with machine learning and Gegenbauer approaches. It is also increasingly relevant as automatic interpretation of MRIs matures. Ferreira *et al.* [5] examined the effect of different interpolation algorithms on dark rim artefacts in myocardial perfusion studies, concluding that sinc interpolation is best in that situation because sub-voxel shifts caused by motion between frames results in a modulation of the ringing that has the appearance of a dynamic artefact. Our work also uses sinc interpolation, which is the most consistent within the Fourier transform framework, though we implement it as zero-padding in the Fourier domain. This interpolation means that the image is reconstructed at far too many pixel locations to align with the zeros of the ringing, and hence we cannot zero the ringing using sub-voxel shifting. We can, however, modulate it, leading to the periodic peaks observed in Fig. 6.

The median filter has an erratic curve. We investigated this further. The median filter can be quite subjectively effective where the region between boundaries is relatively flat. For our test signal, we observe that it tracks the flat regions quite well, performing poorly only at the discontinuities. However, it scores poorly because the approximately constant values in the flat regions are incorrect, which is not true of other filters. This effect is sensitive to sub-sample shifts, explaining the erratic performance of the filter as measured by our metric; sub-sample shifts result in a modulation of the strength of ringing and so the median value is also varying. This constant offset matters little to human perception, but may be important in more quantitative fields. However, for even a slightly more complex test signal, the weaknesses of the median filter become apparent.

In Fig. 7, we plot a selection of the transfer functions of the filters analysed in this Section. We note that the most

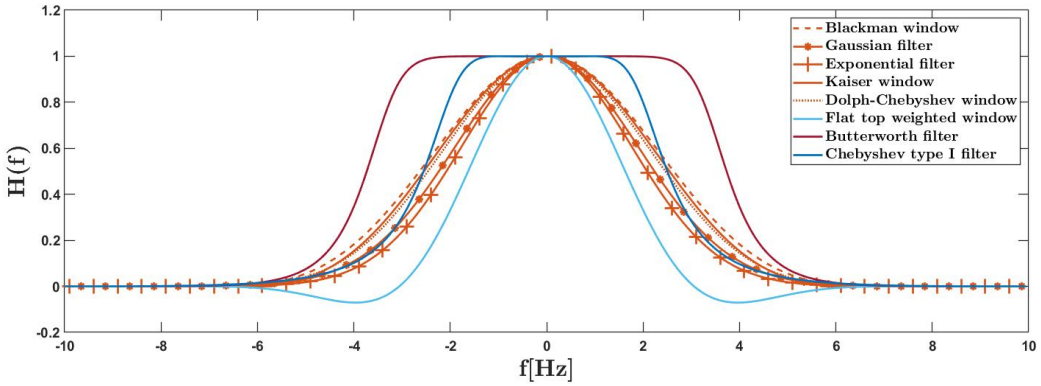

Figure 7: Transfer functions of a selection of the filters optimised for a test signal truncated at 4 Hz. In red, we have the four best performing filters. In other colors, we have three of the other filters including the Butterworth filter, flat top weighted window and Chebyshev type I filter.

successful filters (in red) have similar behaviours, varying smoothly from the peak at DC to their stop bands and with high attenuation near the cut-off frequency. Filters with flat pass-bands and those with more complex shapes are less successful. This is somewhat counter-intuitive, e.g. "To minimize unwanted smoothing, low-pass filters with a high cutoff-frequency seem preferable." [6]. We have observed that reducing the cut-off frequency of the Butterworth or Chebyshev type I filters, for example, noticeably deteriorated performance so it is also too simplistic to argue for more aggressive filtering.

In Fig. 8, we present the zoomed in 1D examples. The top figure shows the six best performing filters have suppressed Gibbs ringing quite well. The edges are preserved to an extent and the flat region have no visible ringing. The middle figure shows the results of six bad filters. The reconstructions present some artefacts around the edges after filtering and the ringing is still visible. The bottom figure shows the worst performing filters that couldn't suppress ringing well. We note that for this example, the Hamming window is grouped as a 'bad' performing one. This is because the cut-off frequency of the original signal is at 3.5 Hz where the Hamming window is not performing as well as overall.

In Fig. 9, we present the top left quarter of a 400×400 2D Shepp-Logan phantom example. On the top row, we show the Fourier reconstruction generated with 40×40 Fourier coefficients, causing visible ringing. This simulates a 10× zoom. We also show the outputs of three filters, one from each of the groups in Fig. 6, optimised as described in Section 3. On the bottom, the zoomed in slices demonstrate that the filters perform as expected.

4. Conclusion

The literature on novel MRI reconstruction and post-processing algorithms is plagued by qualitative evaluations and inconsistent comparisons. Novel algorithms continue to be developed in the literature. In particular, methods based on neural networks are rapidly becoming competitive with existing methods for speed and accuracy. Reconstruction and post-processing algorithms must suppress Gibbs ringing or risk introducing artefacts into MR images which could lead to misdiagnosis.

In this paper, we have demonstrated that many of the metrics previously used to evaluate the suppression of Gibbs ringing are unsatisfactory. We note that our simulations have been limited to linear algorithms. We cautiously suggest that non-linear algorithms, such as machine learning, may also benefit from the use of ℓ^0_ϵ. We have proposed a metric we believe to provide superior results, discriminating against reconstructions with excessive ringing or low fidelity. We used the metric to evaluate many potential filters for filtered Fourier reconstruction. We identified six high-performing filters. We note that some of the worst performing filters were amongst those used as benchmarks against which novel reconstruction algorithms have previously been judged in the literature.

We believe that the results presented in this paper provide a clear and consistent method of evaluating a novel algorithm: it should be bench-marked against one of the six filters we have identified, using ℓ^0_ϵ to select the parameters and to quantitatively compare the reconstructions.

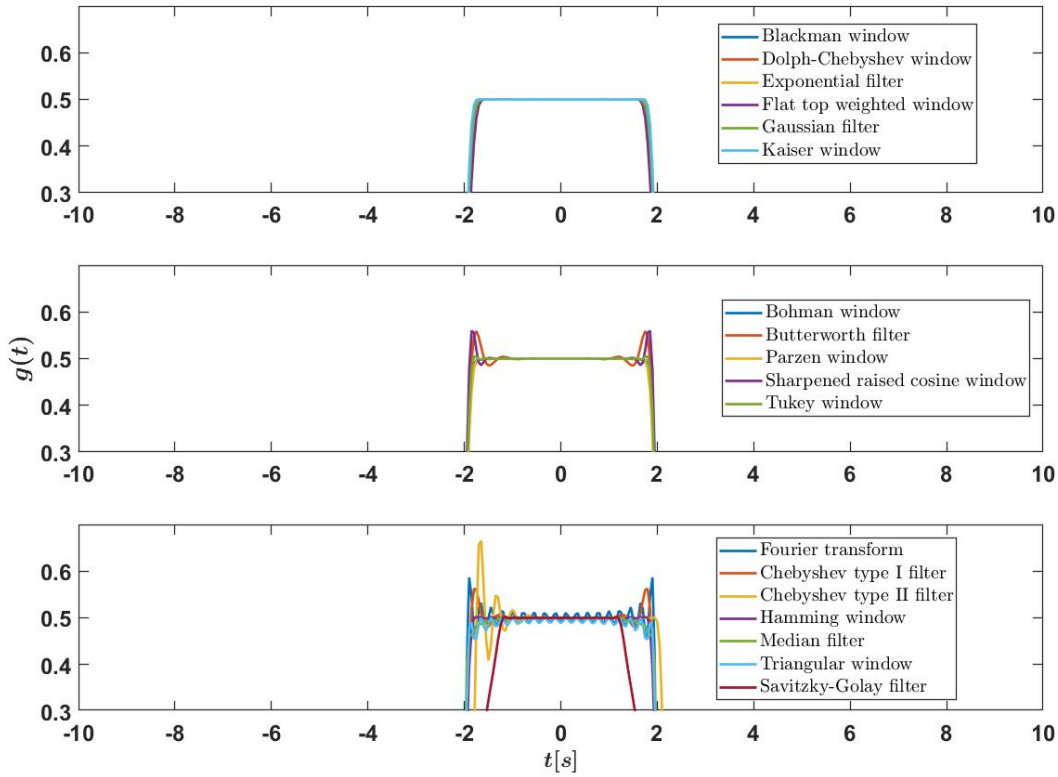

Figure 8: 1D example signal filtered reconstructions. On the top, six 'good' reconstructions are shown. In the middle, six 'bad' reconstructions. At the bottom, 6 'ugly' reconstructions including the Fourier reconstruction. We observe that the classification based on our novel metric is effective in terms of ranking the performance on suppression of ringing.

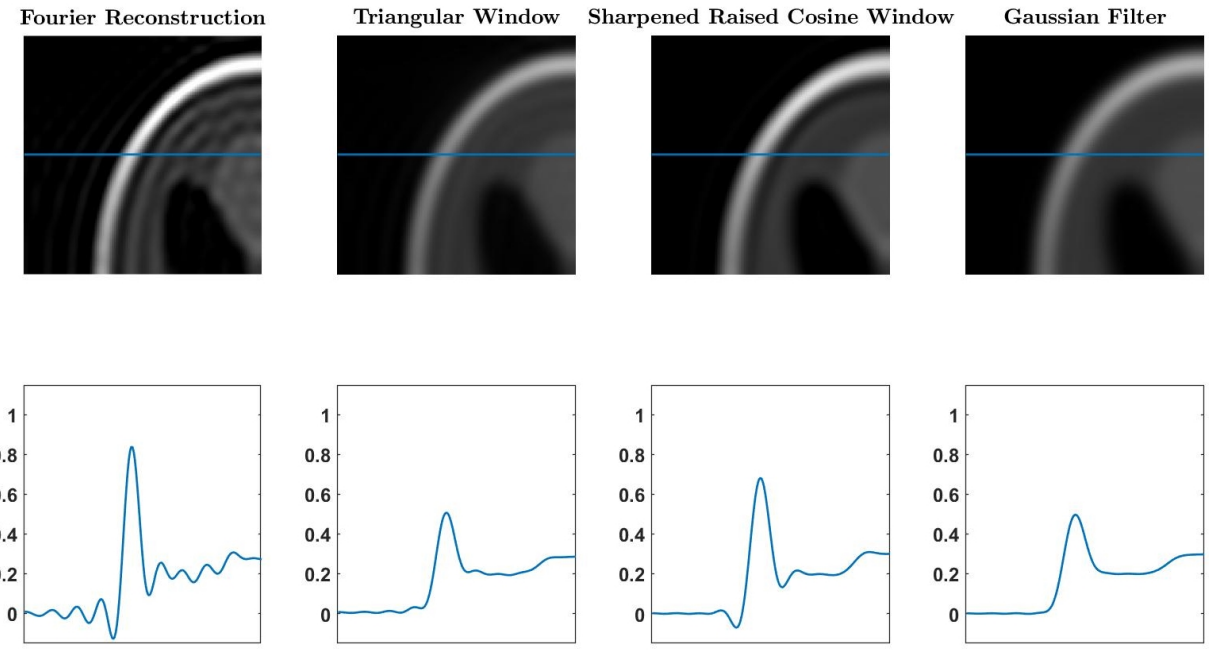

Figure 9: Results using a 2D Shepp-Logan phantom example. On the top row, we show the Fourier reconstruction of a 400×400 image generated with 40×40 Fourier coefficients, causing visible ringing. We also show the outputs of three filters, one from each of the groups in Fig. 6, optimised as described in Section 3. On the bottom, the zoomed in slices demonstrate that the filters perform as expected.

Acknowledgment

The authors would like to thank our colleague Dr Barry Cardiff for constructive conversations in relation to this work. The authors acknowledge the support of the GliMR EU COST Action CA18206. Yue Wang acknowledges the support of the Irish Research Council Government of Ireland Postgraduate Scholarship. The authors thank University College Dublin for their support.

References

[1] D. Gottlieb, C.-W. Shu, A. Solomonoff, H. Vandeven, On the Gibbs phenomenon i: recovering exponential accuracy from the Fourier partial sum of a nonperiodic analytic function, Journal of Computational and Applied Mathematics 43 (1-2) (1992) 81–98.
[2] D. Gottlieb, C.-W. Shu, On the Gibbs phenomenon and its resolution, SIAM review 39 (4) (1997) 644–668.
[3] A. Gelb, A hybrid approach to spectral reconstruction of piecewise smooth functions, Journal of Scientific Computing 15 (3) (2000) 293–322.
[4] R. Archibald, A. Gelb, A method to reduce the Gibbs ringing artifact in MRI scans while keeping tissue boundary integrity, IEEE Transactions on Medical Imaging 21 (4) (2002) 305–319.
[5] P. Ferreira, P. Gatehouse, P. Kellman, C. Bucciarelli-Ducci, D. Firmin, Variability of myocardial perfusion dark rim gibbs artifacts due to sub-pixel shifts, Journal of Cardiovascular Magnetic Resonance 11 (1) (2009) 1–10.
[6] E. Kellner, B. Dhital, V. G. Kiselev, M. Reisert, Gibbs-ringing artifact removal based on local subvoxel-shifts, Magnetic resonance in medicine 76 (5) (2016) 1574–1581.
[7] Y. Wang, Y. Song, H. Xie, W. Li, B. Hu, G. Yang, Reduction of Gibbs artifacts in magnetic resonance imaging based on convolutional neural network, in: 2017 10th international congress on image and signal processing, biomedical engineering and informatics (CISP-BMEI), IEEE, 2017, pp. 1–5.
[8] M. J. Muckley, B. Ades-Aron, A. Papaioannou, G. Lemberskiy, E. Solomon, Y. W. Lui, D. K. Sodickson, E. Fieremans, D. S. Novikov, F. Knoll, Training a neural network for gibbs and noise removal in diffusion mri, Magnetic resonance in medicine 85 (1) (2021) 413–428.
[9] Q. Zhang, G. Ruan, W. Yang, Y. Liu, K. Zhao, Q. Feng, W. Chen, E. X. Wu, Y. Feng, MRI Gibbs-ringing artifact reduction by means of machine learning using convolutional neural networks, Magnetic resonance in medicine 82 (6) (2019) 2133–2145.
[10] X. Zhao, H. Zhang, Y. Zhou, W. Bian, T. Zhang, X. Zou, Gibbs-ringing artifact suppression with knowledge transfer from natural images to mr images, Multimedia Tools and Applications (2019) 1–23.
[11] G. Barbastathis, A. Ozcan, G. Situ, On the use of deep learning for computational imaging, Optica 6 (8) (2019) 921–943.
[12] Q. Chen, V. Koltun, Photographic image synthesis with cascaded refinement networks, in: Proceedings of the IEEE international conference on computer vision, 2017, pp. 1511–1520.

[13] C. Chen, Q. Chen, J. Xu, V. Koltun, Learning to see in the dark, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 3291–3300.
[14] B. Roberts, M. Wan, S. P. Kelly, J. J. Healy, Quantitative comparison of gegenbauer, filtered Fourier, and Fourier reconstruction for MRI, in: Multimodal Biomedical Imaging XV, Vol. 11232, International Society for Optics and Photonics, 2020, p. 112320L.
[15] J. Seetha, S. S. Raja, Denoising of MRI images using filtering methods, in: 2016 International Conference on Wireless Communications, Signal Processing and Networking (WiSPNET), IEEE, 2016, pp. 765–769.
[16] E. Di Bella, D. Parker, A. Sinusas, On the dark rim artifact in dynamic contrast-enhanced MRI myocardial perfusion studies, Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine 54 (5) (2005) 1295–1299.
[17] D. L. Parker, G. T. Gullberg, P. R. Frederick, Gibbs artifact removal in magnetic resonance imaging, Medical physics 14 (4) (1987) 640–645.
[18] F. Sattar, L. Floreby, G. Salomonsson, B. Lovstrom, Image enhancement based on a nonlinear multiscale method, IEEE transactions on image processing 6 (6) (1997) 888–895.
[19] G. Rath, C. Guillemot, J.-J. Fuchs, Sparse approximations for joint source-channel coding, in: 2008 IEEE 10th Workshop on Multimedia Signal Processing, IEEE, 2008, pp. 481–485.
[20] G. L. Zeng, R. J. Allred, Partitioned image filtering for reduction of the Gibbs phenomenon, Journal of nuclear medicine technology 37 (2) (2009) 96–100.
[21] B. Sharif, R. Dharmakumar, T. LaBounty, R. Arsanjani, C. Shufelt, L. Thomson, C. N. Bairey Merz, D. S. Berman, D. Li, Towards elimination of the dark-rim artifact in first-pass myocardial perfusion MRI: removing Gibbs ringing effects using optimized radial imaging, Magnetic resonance in medicine 72 (1) (2014) 124–136.