High dynamic range fusion of magnetic resonance flow imaging data

Christopher M. Sandino
Stanford University
EE 367 Final Project Report
sandino@stanford.edu

Abstract

Time-resolved, volumetric phase-contrast magnetic resonance imaging (also known as 4D flow MRI) is a comprehensive, diagnostic tool used by radiologists to simultaneously assess cardiovascular anatomy, flow, and function. Its ability to resolve both slow and fast hemodynamics (dynamic range) is limited, however. To address this, post-processing techniques such as phase unwrapping have been used to extend the dynamic range, but are highly variable with respect to patient anatomy and physiology. Recently, a new technique called multi-venc 4D flow has gained traction due to its ability to simultaneously and robustly measure slow and fast flow. However, its use of the massive amount of data that it collects (10-20 GB) is inefficient. In this paper, I propose a high dynamic range (HDR) post-processing technique for reconstructing and compressing multi-venc data down into digestible images that allow for more precise measurement of typical hemodynamic parameters measured from this data such as peak flow. I applied this technique to both simulated 4D flow MRI data, and in-vivo data. This technique shows promise for allowing comprehensive abdominal and neurovascular imaging, in which there is a large dynamic range of blood velocity.

1. Introduction

Time-resolved, volumetric phase-contrast magnetic resonance imaging (also known as 4D flow MRI) is a powerful clinical tool used by radiologists to simultaneously assess cardiovascular anatomy, flow, and function [7]. With raw data acquired in one 8-10 minute scan, both anatomical and flow images can be reconstructed by encoding three-directional tissue velocity information into the phase of the raw MRI signal.

The velocimetric dynamic range of 4D flow MRI is inherently limited, however. Since signal is measured in quadrature (via real and imaginary channels), the principal value of the complex MRI signal can only take on values in the range $(-\pi, +\pi]$. More rigorously, the principal value of a complex number $z = a + ib$ for $a, b \in \mathrm{R}$ is defined as:

$$\mathrm{Arg}(z) = \begin{cases} \arctan\frac{b}{a} & a > 0 \\ \frac{\pi}{2} - \arctan\frac{b}{a} & b > 0 \\ -\frac{\pi}{2} - \arctan\frac{b}{a} & b < 0 \\ \arctan\frac{b}{a} \pm \pi & a < 0 \\ \text{undefined} & a = 0, b = 0 \end{cases}$$

Thus, there is a maximum encodable velocity also known as venc, which corresponds to the phase value $\pm\pi$. This is a prescribable parameter which determines the range of velocities $(-venc, +venc]$ which are directly measurable from a single scan. To avoid velocity aliasing (also known as phase wrapping), the venc is set higher than the highest velocity that the MRI technician expects to measure. However, scanning with higher venc's comes at the cost of velocity-to-noise ratio (VNR). This becomes detrimental in abdominal and neurovascular flow imaging due to the presence of a high dynamic range of flow velocities. The venc must be set high to capture fast flow dynamics, making VNR too low to resolve slow flow dynamics (See Figure 1).

In this work, I will propose a scheme where multiple 4D flow datasets are acquired with different venc's (multi-venc acquisition), and then fused together using a regularized high dynamic range (HDR) post-processing technique to precisely and accurately depict both fast and slow flow. This technique would not only streamline visualization of this high-dimensional data, but also increase precision and accuracy of biomarkers which are derived from flow data such as peak velocity, flow, wall shear stress, and turbulent kinetic energy.

2. Related Work

The velocimetric dynamic range of 4D flow MRI has previously been extended in three ways: (a) high VNR velocity encoding [4] (b) phase unwrapping techniques [6], and (c) multi-venc flow imaging [5, 10]. The performance of the first technique is highly dependent on spatial and

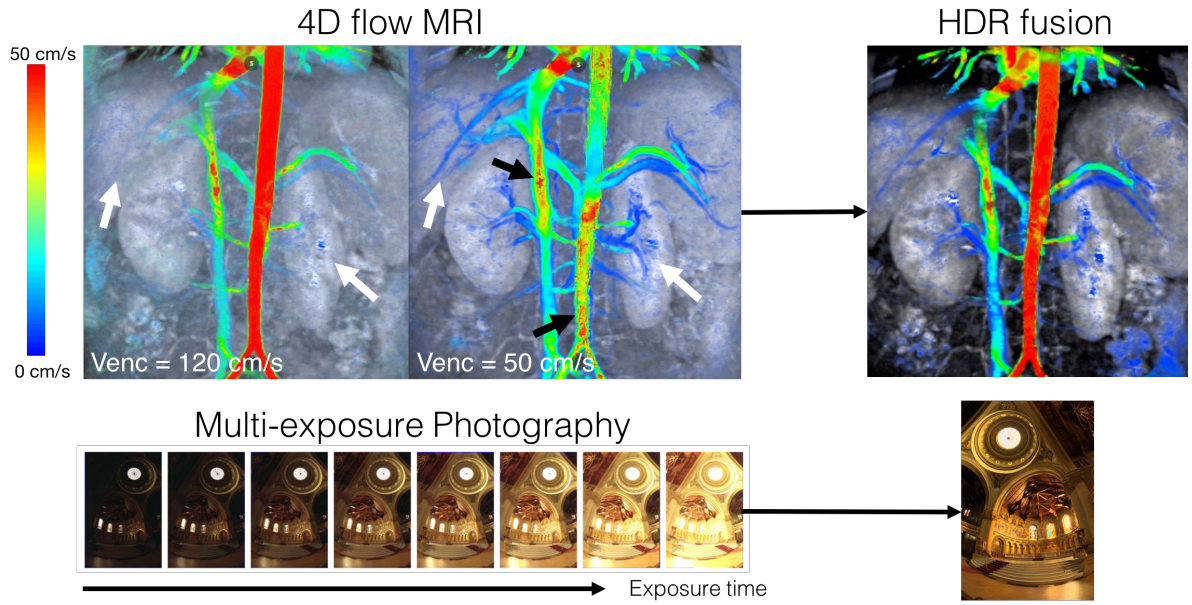

Figure 1. Two abdominal 4D flow datasets obtained with venc's of 120 cm/s and 50 cm/s. Color velocity data is scaled from 0 to 50 cm/s and overlaid on top of grayscale magnitude data. As shown here, low-venc data is prone to velocity aliasing in high velocity structures (black arrows), but offer much better low velocity resolution as seen by the renal and hepatic vasculature (white arrows). This problem is similar to multi-exposure photography, where short and long exposures only offer good resolution of bright and dark areas respectively. The idea behind this project is to use HDR techniques to combine two (or more) datasets like these into one composite dataset with good low velocity resolution and without velocity aliasing.

temporal resolution especially in fatty areas such as the abdomen. For this proposal, I will focus on the second and third techniques.

2.1. Phase unwrapping

Phase unwrapping is a classical image processing technique that has recently been applied to 4D flow in order to undo velocity aliasing in low venc acquisitions (See Figure 2). By estimating four-dimensional (space and time) gradients in the raw phase data, algorithms find and undo phase wraps by locating areas with local "jumps" of 2π. This can be done non-iteratively by analytically solving Poisson's equation [6], or iteratively by weighted least squares []. However, phase unwrapping in higher than one dimension has been shown to be intractable because the unwrapping result varies depending on the direction of the gradient-based heuristic. In general, phase unwrapping has been shown to be non-robust in 4D flow data, and its performance is highly dependent on patient anatomy and physiology.

2.2. Multi-venc techniques

Multi-venc 4D flow MRI is another technique that extends dynamic range by acquiring multiple datasets using different venc's and then using high venc measurements to

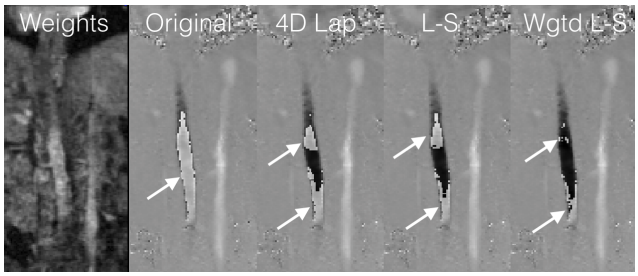

Figure 2. Performance of three phase unwrapping algorithms on low-venc abdominal data. 4D raw phase data is projected down onto a 2D coronal slice of the abdominal aorta. Phase wraps in each image are denoted by white arrows.

unwrap corresponding measurements in low venc datasets. Algorithms look at element-wise differences between low and high venc data, and locate phase wraps by thresholding voxels with differences above 2π [10]. However, acquiring multiple 4D flow datasets is not clinically feasible most of the time, so most techniques propose only acquiring two datasets with one high and one low venc [6]. In the end, most of the high venc data is thrown away since low venc data is replaced with high venc data only where wraps oc-

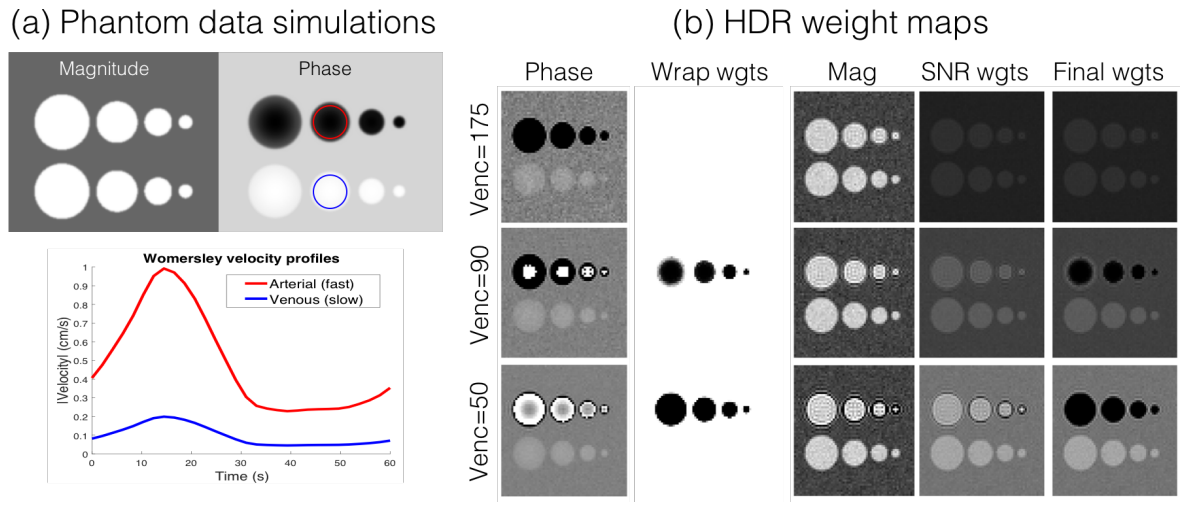

Figure 3. (Left) Both magnitude and phase data are simulated for phantom HDR experiments. Red and blue regions of interest are drawn around a simulated artery (red), and vein (blue), and then the velocity is plotted over time. (Right) Examples of weight maps for simulated phantom data with varying venc values are shown here. The wrapping weights are multiplied by the SNR weights to obtain the final weights shown on the far right column.

cur. While this technique has shown to be highly robust to velocity aliasing, its use of data is inefficient. Only high venc measurements corresponding to wrapped aliased voxels in the low venc data are used, while the rest is thrown away. I will propose a more efficient use of this data, that improves velocity precision.

3. Theory

The high dynamic range processing technique [2] is proposed here to reconstruct multi-venc 4D flow datasets using as much raw data as possible. Given magnitude images $M_i \in \mathrm{R}^m$ and velocity images $V_i \in \mathrm{R}^m$ acquired with venc_i for $i \in \{1, ..., n\}$, estimate composite HDR velocity image $\hat{V}$ by the following least absolute shrinkage and selection operator (LASSO) minimization:

$$\text{minimize} \sum_{i=1}^{n} ||W_i^{\intercal}(V_i - \hat{V})||_2^2 + \lambda ||\Psi(\hat{V})||_1$$

$$W_{ij} = w_{\text{VNR}}(M_{ij}) w_{\text{wrap}}(\text{venc}_i)$$

The data consistency term enforces least-squares agreement between the reconstructed images, and each of the venc acquisitions. Each measurement is weighed by an estimate of the local VNR, and local probability of phase wrapping. The regularization term is introduced as a denoising operation, but can also serve as a prior to estimate data that wasn't acquired (i.e. compressed sensing). This minimization problem is solved using a simple weighted average for $\lambda = 0$, or the alternating direction method of multipliers (ADMM) for $\lambda > 0$ [1]. The design of the cost function is decomposed into two problems: (a) design of the weights, and (b) choice of l_1-penalty.

3.1. Designing the weights

Assuming that the highest venc dataset has no velocity aliasing (as are most datasets acquired in the clinic), we can roughly estimate which voxels in the low venc datasets are wrapped and weigh against them. A soft-masking technique is implemented here using the sigmoid function:

$$w_{\text{wrap}}(v) = 1 - \text{sigm}(|v| - \text{venc}_i) = \frac{-e^{-a(|v| - \text{venc}_i)}}{1 - e^{-a(|v| - \text{venc}_i)}}$$

We can also weigh against measurements with low VNR, a parameter that can be estimated from corresponding magnitude data:

$$w_{\text{VNR}} = \frac{M_{ij}}{\text{venc}_i}$$

More precise data from low venc acquisitions is weighed more heavily than data from high venc acquisitions. This way of estimating VNR from magnitude data and the venc comes from statistical analysis of the complex MRI signal [7].

3.2. Incorporating a priori information

An l_1-regularization term is introduced to incorporate prior information into the reconstruction of composite HDR images. The total variation (TV) regularizer is used here in

both space and time:

$$\Psi_{\text{TV}}(\hat{V}) = \begin{pmatrix} \nabla_x \\ \nabla_y \\ \nabla_z \\ \nabla_t \end{pmatrix} \hat{V}$$

The TV penalty enforces continuity (small differences between neighboring pixels) in images, and may be a reasonable prior for 4D flow phase data. However, it may not be a good prior for cases where there are flow jets, such as vessel stenosis (narrowing), because it may bias peak velocities down to be consistent with surrounding static tissue.

The divergence penalty has previously been used to reconstruct and denoise 4D flow data [8, 9]:

$$\Psi_{\text{div}}(\hat{V}) = \begin{pmatrix} \nabla_x & \nabla_y & \nabla_z \end{pmatrix} \begin{pmatrix} \hat{V}_x \\ \hat{V}_y \\ \hat{V}_z \end{pmatrix}$$

It is a natural prior for flow fields due to a fluid dynamics result, which says that fluid flow is incompressible (divergence-free).

This idea derives from a classical fluid dynamics result that flow fields have zero divergence (incompressibility).

4. Analysis

4.1. Phantom experiments

Two-dimensional velocity profiles are simulated for cylindrical vessels of varying diameter (d=5,10,20,30 mm) using the Womersley method, a classic computational fluid dynamics model [11] (see Figure 3). These are used to synthesize 4D magnitude and phase data with varying venc and SNR by adding complex Gaussian white noise. Peak velocities were chosen to be 100 cm/s (arterial) and 5 cm/s (venous) to have a representation of the dynamic range that one would find in a typical abdomen.

The proposed HDR method was applied to phantoms with simulated vencs of 50, 75, and 150 cm/s. These are representative of the vencs that would be used to acquired slow and fast hemodynamic data. The data is merged using the weighted HDR method with $\lambda = 0$, and different values of λ to apply total variation regularization in space, time, and space/time. I did not apply the divergence penalty for the phantom data, because it only contains unidirectional flow, and thus the divergence would reduce down to a total variation penalty along the direction of flow.

Phantom HDR images fused with and without regularization are evaluated quantitatively with respect to PSNR, which was computed by comparing fused images with ground truth data. Images are also evaluated qualitatively by visual inspection of the data.

4.2. In-vivo experiments

In-vivo abdominal 4D flow MRI datasets were acquired on a 3T General Electric MR750 using a 32-channel cardiac coil. With IRB approval and informed consent, a pediatric patient was referred for a gadolinium-enhanced abdominal 4D flow MRI (Flip angle=15, TE/TR=2.49/7.84, Venc=120, 50 cm/s). A parallel imaging and compressed sensing based scheme, and soft-gated butterfly navigator motion compensation was used to reconstruct the data.

In-vivo HDR images fused with and without regularization are evaluated qualitatively by inspection of the data.

5. Results

5.1. Phantom experiments

See Figure 4 for resultant HDR fused images. Images that were fused with small λ generally achieved higher PSNR than for large values of λ. The best performance with respect to PSNR was achieved by HDR with TV regularization through space for $\lambda = 0.005$ with PSNR = 20.01 dB.

Qualitatitively, the best performance was achieved by HDR with TV regularization through time for $\lambda = 0.005$, which also achieved a PSNR of 19.77 dB. I chose this one because significant blurring is seen in the spatially-regularized images as the algorithm converges to a solution where the vessels seem to fuse into a singular blob. However, high amounts of temporal regularization causes an averaging effect on the velocity profiles through time.

5.2. In-vivo experiments

See Figure 5 for resultant HDR fused images. By visual inspection, the divergence-regularized HDR algorithm produces the most accurate and precise images. The high velocity values are close to what is measured in the high venc dataset (this is the clinical standard for abdominal imaging), while venous structures in the kidney and liver are well-resolved and visible. The TV regularized images again biased the absolute velocities down to both spatial and temporal averaging of the data.

6. Discussion

In this work, an HDR post-processing technique is proposed, implemented, and evaluated for fusing multi-"exposure" MRI data and allowing the visualization of slow and fast hemodynamics in abdominal 4D flow MRI scans. Based on the experiments in this paper, the technique seems to work robustly in the presence of noise (such as in the high venc data), and show reasonable resolution of low velocity structures. However, this technique has yet to be tested on other types of exams such as neurovascular data, where

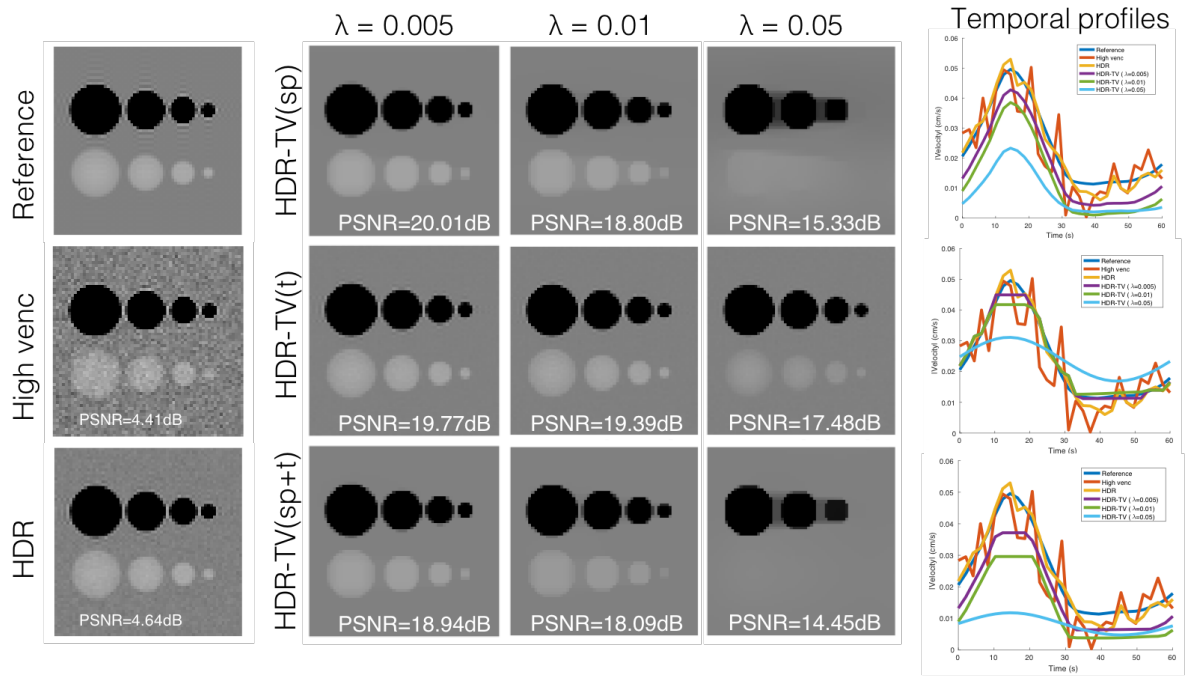

Figure 4. HDR fused images for $\lambda = 0$ (bottom left) and varying values of λ for TV regularization in space (sp), time (t), and space+time (sp+t). Temporal profiles of venous velocities are shown on the right for each case, which depict how each regularization can bias velocity measurements globally (TV in space), or locally (TV in time).

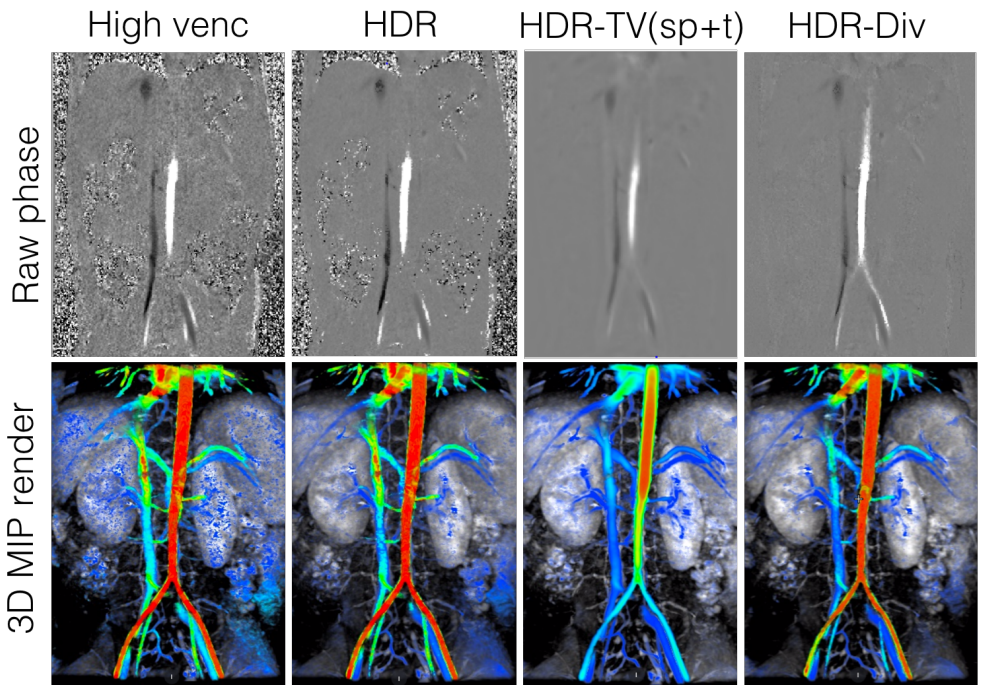

Figure 5. HDR fused in-vivo images for $\lambda = 0$ (second column) and $\lambda = 0.01$ TV regularization and divergence regularization. The top row is the raw phase data acquired on the scanner, the bottom row shows a rendered maximum intensity projection (MIP) of the 3D data at one time point. All velocity images are windowed to [-50, +50] cm/s.

there are three ranges of flow: cerebrospinal fluid, venous, and arterial.

One limitation of this study was that I did not compare the HDR fused images to any other images besides the high venc data. It would a fairer comparison between the HDR images and (1) denoised images and/or (2) images obtained through thresholding as is described in recent multi-venc literature [10].

Another limitation is the lack of ground truth data for which to evaluate HDR performance on in-vivo data. It is difficult to do a quantitative analysis of in-vivo data since a real ground truth dataset would take hours to acquire. Perhaps in the future, this method could be validated using computational fluid dynamic simulations performed on high-resolution anatomic data (computed tomography or MRI).

Future work will involve a more in depth study about how priors affect precision and accuracy in velocity quantitation. The weighted HDR algorithm without any regularization produced the most accurate velocity measurements, but still contained residual noise from high venc data. The phantom experiments showed that TV regularization improved PSNR, and visual quality by imposing a smoothness prior on the images. However, spatial regularization causes a blending of vessel structures and static tissue around it. This biases down the velocities on a global level, as the algorithm converges to images that try to make differences between vessel pixels and static tissue pixels small. Temporal regularization on the other hand causes averaging through time, which also biases velocities, but on a more local level. This may cause in accurate peak velocity values, which is an important diagnostic biomarker through which clinical decisions are made.

This work did not address one of the main problems with multi-venc 4D flow, and that is data acquisition time. Typically in the clinic, a cardiologist will order one 4D flow MRI scan to visualize either venous or arterial structures, because there may not be enough time to acquire both. I plan to use this HDR framework to reconstruct undersampled multi-venc data that is acquired in the same amount of time as a standard 4D flow dataset. One idea is to use superresolution techniques [3] to interpolate missing data from undersampled data. The cost function would instead be:

$$\text{minimize} \sum_{i=1}^{n} ||W_i^{\intercal}(AV_i - \hat{V})||_2^2 + \lambda||\Psi(\hat{V})||_1$$

where A is an interpolation matrix that upsamples each dataset to create a super-resolved dataset with good low and high velocity resolution.

7. Conclusion

A high dynamic range processing technique is presented for the reconstruction of high dynamic range 4D flow MRI data. This technique improves the comprehensiveness of 4D flow MRI as a clinical diagnostic exam. Information about slow and fast hemodynamics could potentially give clinicians more complete picture of patient pathologies, and a better understanding of the human cardiovascular system.

References

[1] S. Boyd, N. Parikh, E. Chu, and B. Peleato. Distributed optimization via alternating direction method of multipliers. *Foundations and Trends in* , 2010.

[2] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. *SIGGRAPH*, page 31, 1997.

[3] M. Elad and A. Feuer. Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images. *IEEE transactions on image processing*, 1997.

[4] K. M. Johnson and M. Markl. Improved SNR in phase contrast velocimetry with fivepoint balanced flow encoding. *Magn Reson Med*, 63:349–355, 2010.

[5] A. Lee, G. Pike, and N. Pelc. Three-Point Phase-Contrast velocity measurements with increased Velocity-to-Noise ratio. *Magn Reson Med*, 33:122–126, 1995.

[6] M. Loecher, E. Schrauben, K. Johnson, and O. Wieben. Phase unwrapping in 4D MR flow with a 4D singlestep laplacian algorithm. *J Magn Reson Imaging*, 43:833–842, 2016.

[7] M. Markl, F. Chan, M. Alley, K. Wedding, M. Draney, C. Elkins, D. Parker, R. Wicker, C. Taylor, R. Herfkens, and N. Pelc. Time-resolved three-dimensional phase-contrast MRI. *J Magn Reson Imaging*, 17:499–506, 2003.

[8] F. Ong, M. Uecker, U. Tariq, A. Hsiao, M. T. Alley, S. S. Vasanawala, and M. Lustig. Robust 4D flow denoising using divergencefree wavelet transform. *Magn Reson Med*, 73:828–842, 2015.

[9] C. Santelli, M. Loecher, J. Busch, O. Wieben, T. Schaeffter, and S. Kozerke. Accelerating 4D flow MRI by exploiting vector field divergence regularization. *Magn Reson Med*, 75:115–125, 2016.

[10] S. Schnell, S. Ansari, C. Wu, and J. Garcia. Accelerated dualvenc 4D flow MRI for neurovascular applications. *J Magn Reson Imaging*, 2017.

[11] J. R. Womersley. Method for the calculation of velocity, rate of flow and viscous drag in arteries when the pressure gradient is known. *J Physiol*, 127:553, 1955.