

Università degli studi di Napoli "Federico II"

Dipartimento di Scienze Biomediche Avanzate Direttore Prof. Claudio Buccelli

Dottorato di ricerca in Imaging Molecolare XXVII ciclo

Coordinatore Prof. Andrea Renda

Denoising in Magnetic Resonance Imaging: Theory, Algorithms and Applications.

Dottorando Dr. Pasquale Borrelli Relatore Dr. Giuseppe Palma

To my family

Abstract

In the context of medical image processing, denoising is widely considered as one of most fundamental postprocessing tasks. In this field, the non-local means (NLM) filter demonstrated to be a robust and performing approach respect to the previous state-ofart denoising methods. As the filtering strength must be tuned to obtain an optimized and customized restoring process, the estimation of image noise variance is an important issue. Althought in clinical practice noise estimation is performed on background (no signal area) of magnitude MR images, in case of parallel MR imaging (pMRI) techniques noise estimation from the image background produces biased results due to spatially varying noise distribution of the pMRI images. A novel NLM approach based on local noise estimation is introduced (hereafter indicated as SVN-NLM). Results show more accurate noise estimation in contrast to global noise calculation when spatially inhomogeneous noise was added. Denoising performances, measured by visual inspection and peak signal-to-noise ratio (PSNR), of the NLM are considerably improved using SVN-NLM in case of inhomogeneous noise. Furthermore, the SVN-NLM method produces similar results comparing with NLM when homogeneous noise was added, so as to make it an useful method for datasets with both spatially independent and dependent noise variance.

As second task, since the susceptibility-weighted imaging (SWI) suffers from reduced SNR due to the high resolution required to obtain a proper contrast generation, a novel pipeline (Multicomponent-Imaginary-Real-SWI, hereafter MIR-SWI) to obtain susceptibility-weighted images with higher SNR and improved conspicuity is proposed. In this context, the application of a denoising filter is non-trivial as the distributions of magnitude and phase noise may introduce biases during image restoration. Taking advantage of the potential multispectral nature of MR images, the multicomponent approach of the MIR-SWI approach performs better than a component-by-component image restoration method. Both qualitative and quantitative assessments showed that MIR-SWI fared consistently better than the other approaches. Noise removal with MIR-SWI also provided improvement in contrast-to-noise ratio (CNR) and vessel conspicuity at higher factors of phase mask multiplications than the one suggested in the literature for SWI vessel imaging.

Finally, a new strategy to address the computational demand of the NLM filter is investigated. Due to high computational complexity of the NLM denoising filter, in literature several 2D NLM implementations on Graphic Processor Unit (GPU) architectures were proposed. Here a fully 3D NLM implementation on a multi-GPU architecture is presented and its high scalability is suggested. Several configurations of thread block organization and data access are analyzed, thus identifying a set of optimal settings that guarantee high performance results for a wide spectrum of application scenarios. The reduction of running times shows that scalability is close to ideal one for most common dataset sizes, *e.g.* those typical of MRI clinical pratice.

Contents

Al	bstra	\mathbf{ct}	v
Co	onter	nts	vii
1	Intr	oduction	1
2	The	oretical background and basic methods	5
	2.1	Parallel MR Imaging	5
	2.2	Susceptibility-weighted Imaging	5
	2.3	Noise distribution in MRI	6
		2.3.1 Noise in pMRI	6
	2.4	Non-Local Means denoising filter	7
3	Mat	terials and Methods	11
	3.1	Noise estimation in MR images	11
		3.1.1 On the unbiased noise estimation in pMRI	11
	3.2	MNLM-based filter to denoise SWI images	13
		3.2.1 Acquisition protocol	13
		3.2.2 Preliminary investigation	13
		3.2.3 MIR-SWI method	14
	3.3	GPGPU implementation of the NLM filter	15
4	Res	ults	17
	4.1	SVN estimation	17
	4.2	Restoring SWI images	18
		4.2.1 Qualitative assessment (brain tissues)	19
		4.2.2 Quantitative evaluation (veins)	21
		4.2.3 Number of phase mask multiplications	23
	4.3	Multi-GPU implementation	23
		4.3.1 Consistency	23
		4.3.2 Performance	23
5	Disc	cussion and Conclusion	31
	5.1	Uniform and non-uniform unbiased noise estimation	31
	5.2	On the novel MIR-SWI approach	32
	5.3	On the Multi-GPU 3D NLM implementation	34

Contents	viii
Bibliography	35
Acknowledgements	39

Chapter 1

Introduction

In the context of medical image processing, denoising is widely considered as one of most fundamental post-processing tasks. The purpose of the denoising process is to estimate the true image by removing the noise and preserving edges and image structures at same time. In last decade, a wide variety of methods has been developed to address the problem of the noise removal. In this field, the non-local means (NLM) filter [1] demonstrated to be a robust and performing approach respect to the previous state-ofart denoising methods. The NLM filter is a non-linear neighborhood filter in which the voxel value to be restored is replaced by a weighted average of the voxel intensities in the entire noisy image. Unlike other neighborhood filters, the weights are determined by the neighborhood similarity on the basis of the intensities in patches surrounding each voxel of the image. In other words, the NLM algorithm can be seen as a patch-based mean filter in which the similarity between patches surrounding each voxel is used in the restoring process rather than the intensities of the voxel themselves. Furthermore, the region comparison is based on the radiometric proximity instead of geometrical distance between patches, therefore making it non-local. NLM scheme has been demonstrated both to preserve edges and fine structures from an excessive blurring and to correctly remove the noise thus avoiding the introduction of artifacts and spurious correlated signal [2-4]. In particular, has been proved that NLM filter is highly efficient on texture images as well as on natural images (e.g., on clinical datasets) by taking advantage of the redundancy between image patches. Moreover, in case of multispectral MR datasets where the contrasts between different acquisitions are strictly correlated on the base of the underlying anatomical structures, the redundant information between different components is used in the non-local metric as a way to better distinguish the noise from image features. Taking into account the multispectral nature of MR data, Manjon et al. [5] defined a multicomponent approach (hereafter indicated as MNLM) as a basis for the denoising process.

In the context of the denoising, as the filtering strength must be tuned in order to obtain an optimized and customized restoring process, the estimation of image noise variance is an important issue. In MR clinical practice, noise estimation is performed on Rayleigh-distributed background (no signal area) of magnitude images [6]. Althought noise variance in MR images is considered spatially independent, parallel MR imaging (pMRI) techniques as SENSE or GRAPPA generate spatially varying noise distribution [7]. In this scenario, noise estimation from background produces biased results [8]. To address these limitations a modified approach of the NLM filter based on a novel noise estimation (hereafter indicated as SVN-NLM) is introduced, based on a local statistic approach.

As other application field of the NLM denoising algorithm, a new strategy to restore susceptibility-weighted images was investigated. Susceptibility-weighted imaging (SWI) is an imaging technique based on the differences in tissue susceptibility to enhance the contrast in magnitude MR images [9]. In this technique, the local field inhomogeneities in the phase image are used as source of contrast in order to reveal important anatomical and physiological information about vessels and tissues [10-12]. Although used mainly in neuroimaging to study intracranial venous system (just to mention some of most clinical applications, see [13-17]), SWI images have recently been applyed in other parts of body as well [18-20]. However, the high resolution required to obtain sufficient phase information, which can be used for improved contrast, may lead to a reduced signal-to-noise ratio (SNR), thus compromising both postprocessing tasks and the overall visual inspection. In this context, the application of a denoising filter to produce images with higher SNR and still preserve small structures from an excessive blurring is extremely desirable. Taking into consideration the potential multispectral nature of complex MR images, the MNLM denoising filter may perform better than a component-by-component image restoration method. Moreover, as the distributions of magnitude and phase noise may introduce biases during image restoration, the application of a multispectral denoising filter is non-trivial. Here a new MNLM-based method (Multicomponent-Imaginary-Real-SWI, hereafter MIR-SWI) to produce SWI images with high SNR and improved conspicuity of elusive brain structures is presented, thus improving image quality for SWI data.

Finally, a new strategy to address the computational demand of the NLM filter was investigated. The use of accelerated hardware, in particular the Graphic Processor Units (GPUs), has proved reasonable running times of the NLM algorithm. In literature, several 2D NLM implementations on GPUs were proposed [21–23]. In case of 3D datasets, as in the context of MRI, the use of a fully 3D filter is more appropriate than a sliceby-slice filtering approach to exploit all the information contained in the image. Here a multi-GPU version of the 3D NLM filter based on Compute Unified Device Architecture (CUDA) [24] is presented.

The plan of the thesis is as follows. In §2 the theoretical background and basic methods are reviewed. In particular, the description of parallel imaging techniques and the susceptibility-weighted imaging are described in §2.1 and §2.2, respectively. Moreover, the noise distribution in MRI (§2.3) and the NLM denoising algorithm (§2.4) are presented. Then, in §3 the details of the proposed methods are described. Finally, in §4 and §5 the results are presented and discussed.

Chapter 2

Theoretical background and basic methods

2.1 Parallel MR Imaging

Parallel MRI (pMRI) is an emerged technique that increases the image acquisition rate by sampling a reduced amount of k-space data with an array of receiving coils [25, 26]. Generalized auto-calibrating partially parallel acquisition (GRAPPA) and sensitivityencoded (SENSE) MRI are most common image reconstruction schemes in pMRI. Both reconstruction algorithms share the incorporation of coil-sensitivity profiles into the image reconstruction process [27]. In GRAPPA algorithm missing k-space lines are computed before full-image is reconstructed for each receiver channel [28]. On the other hand, SENSE algorithm reconstructs complex image for each receiver channel and then final images are pixel-wise multiplied by appropriate coil sensitivity mask [29].

2.2 Susceptibility-weighted Imaging

In order to generate susceptibility-weighted magnitude images, the information from phase data is used to create a mask that enhances local changes in tissue magnetic susceptibility. As reported in [9], a high-pass filter (hp) is applied to phase image (ϕ) to obtain ϕ_{hp} and remove low spatial contributions from field variations due to external field inhomogeneities. Then, SWI image is computed as follow:

$$SWI(I_m, \phi_{hp}) = I_m \cdot \phi_{mask}^n(\phi_{hp}), \qquad (2.1)$$

where I_m is the magnitude image,

$$\phi_{mask}(\phi_{hp}) = \begin{cases} 1 - \frac{|\phi_{hp}|}{\pi}, & \text{if } -\pi < \phi_{hp} < 0, \\ 1, & \text{else.} \end{cases}$$
(2.2)

(for a right-handed system $\Delta \phi < 0$ corresponds to a paramagnetic behavior) and $n \sim 4$ is a parameter to tune in order to optimize the contrast-to-noise ratio (CNR) in the SWI image.

2.3 Noise distribution in MRI

The full k-space acquired in MRI is assumed to be corrupted with Gaussian white noise. After Fourier transform, real and imaginary images are still corrupted by uncorrelated Gaussian noise with same variance in both complex components [6]. The non-linear transformation leading to magnitude images modifies the noise distribution, which shows a Rician probability density function (PDF):

$$p_M(M|A,\sigma) = \frac{M}{\sigma^2} e^{-\frac{M^2 + A^2}{2\sigma^2}} I_o \frac{AM}{\sigma^2} \epsilon(M), \qquad (2.3)$$

where σ^2 is the noise variance in real and imaginary images; A is the noise-free image; M is the actual magnitude image; I_o is the 0th order modified Bessel function of first kind and ϵ is the Heaviside function. As shown in Fig. 2.1, in case of high SNR, the Rician distribution tends to the Gaussian one, while, in the opposite limit (A = 0), it becomes the Rayleigh distribution with PDF:

$$p_M(M|\sigma) = \frac{M}{\sigma^2} e^{-\frac{M^2}{\sigma^2}}.$$
(2.4)

2.3.1 Noise in pMRI

The application of multi-surface coil arrays and reconstruction filter can influence the statistical distribution of image noise [7]. In this scenario, variance of background regions will lead to innacurate estimations of the true local noise if a uniform Rayleigh distribution is erroneously assumed [8].

In pMRI raw data represent the subsampled k-spaces acquired from a multiple-coil system.

In case of GRAPPA technique, [30] shows that in a good approximation the magnitude image PDF is a non-central χ distribution, as if noise were distributed identically and



FIGURE 2.1: Rician PDF. Rician PDF plot at different A values with fixed $\sigma = 1$. In case of high SNR ($A \ge 3$) the Rician distribution starts to approximate to Gaussian one, whereas in image points where A = 0 (only noise is present) it corresponds to Rayleigh distribution.

independently in each coil:

$$p_{M_L}(M_L|A_L, \sigma_n, L) = \frac{A_L^{1-L}}{\sigma_n^2} M_L^L e^{-\frac{M_L^2 + A_L^2}{2\sigma_n^2}} I_{L-1} \frac{A_L M_L}{\sigma_n^2} \epsilon(M_L),$$
(2.5)

where L is the number of coils, M_L and A_L are respectively magnitude image and noisefree image reconstructed with sum-of-squares method and I_L is L^{th} order Bessel function of first kind.

On the other hand, as SENSE reconstructs sub-sampled acquisitions in the spatial domain, the noise distribution follows a non-stationary Rician PDF [31] whose variance is:

$$\sigma_s^2 = r\sigma^2 |W_i|^2, \tag{2.6}$$

where σ^2 is noise variance without subsampling, r is the reduction factor and W_i is the reconstruction matrix depending on sensitivity map for each coil.

2.4 Non-Local Means denoising filter

An *N*-D image X can be considered as a real function $X : \mathbb{R}^N \to \mathbb{R}$ with a bounded support $\Omega \subset \mathbb{R}^N$. The NLM filter [1] is a class of endomorphisms of the image space, identified by 2 parameters (a and h), that acts as follows:

$$[\operatorname{NLM}_{a,h}(X)](\vec{x}) = Y(\vec{x}) = \frac{\int_{\Omega} \exp\left[-\frac{d_a^2(\vec{x},\vec{y})}{h^2}\right] X(\vec{y}) d\vec{y}}{\int_{\Omega} \exp\left[-\frac{d_a^2(\vec{x},\vec{y})}{h^2}\right] d\vec{y}},$$
(2.7)

where

$$d_a^2(\vec{x}, \vec{y}) \equiv \int_{\mathbb{R}^N} \left| X(\vec{x} + \vec{t}) - X(\vec{y} + \vec{t}) \right|^2 \cdot \frac{\exp - \frac{\left\| \vec{t} \right\|^2}{2a^2}}{(2\pi)^{n/2} \cdot a} d\vec{t},$$
(2.8)

a represents the radius of the window centered on each point of the image and h rules the similarity measure in the window comparison.

Both computational issues and the convenience to introduce a geometric proximity criterion in addition to the pure radiometric distance measure led to a change in the original version of the NLM filter [3]. More specifically, given a search radius M, for each voxel i located at \vec{x}_i , a search box V_i is defined as

$$V_i \equiv \left\{ \vec{x}_j \in \Omega \left\| \|\vec{x}_j - \vec{x}_i\|_{\infty} < M \right\} .$$

$$(2.9)$$

The search box V_i defines the ensemble of voxels whose intensities will be available to restore the voxel located in \vec{x}_i , thus reducing the search freedom.

Likewise, given a similarity radius $d \sim a$, for each voxel \vec{x}_j within a given search box V_i , a similarity box $_jB_i$ is defined as

$${}_{j}B_{i} \equiv \left\{ \vec{x}_{k} \in \Omega \left\| \| \vec{x}_{k} - \vec{x}_{j} \|_{\infty} < d \right\} .$$
(2.10)

If the image is defined on a discrete grid, a suitable filter implementation is:

$$Y_{i} = \frac{\sum_{\vec{x}_{j} \in V_{i}} \exp\left[-\frac{\|jB_{i}-iB_{i}\|_{2}^{2}}{h^{2}}\right] X_{j}}{\sum_{\vec{x}_{j} \in V_{i}} \exp\left[-\frac{\|jB_{i}-iB_{i}\|_{2}^{2}}{h^{2}}\right]},$$
(2.11)

The filter strength, which is determined by h, can be automatically tuned to obtain an optimized denoising, independent of the search radius M and the standard deviation of noise σ :

$$h^2 = 2\beta\sigma^2 \left| V_i \right| \tag{2.12}$$

 $(\beta \sim 1 \text{ is an adimensional constant to be manually tuned}).$

For a multispectral framework approach, the filtering process can be improved by using intercomponent information to discriminate between noise and image features and reveal masked image details or discard false structures generated by noise. In this setting, the similarity measure in the NLM algorithm can be improved by combining not only the information of surrounding voxels within the image but also the information of different components in order to take advantage of the redundancy between MR series. According to [5], for each component c the filter implementation in Eq. (2.11) is extended to be used on a multispectral framework as follow:

$$Y_{i,c} = \frac{\sum_{\vec{x}_j \in V_i} \exp\left[-\frac{1}{C} \sum_{c=1}^{C} \frac{\|_j B_{i,c} - B_{i,c}\|_2^2}{2\beta \sigma_c^2 |V_{i,c}|}\right] X_{j,c}}{\sum_{\vec{x}_j \in V_i} \exp\left[-\frac{1}{C} \sum_{c=1}^{C} \frac{\|_j B_{i,c} - B_{i,c}\|_2^2}{2\beta \sigma_c^2 |V_{i,c}|}\right]},$$
(2.13)

where C represents the number of components and σ_c is the standard deviation of each component.

Chapter 3

Materials and Methods

3.1 Noise estimation in MR images

To avoid wrap-around artifacts, MR images are usually acquired with pretty large background area, so that the noise amplitude can be easily evaluated on no-signal image segment. Given the standard deviation $\hat{\sigma}$ of the image background, σ of Eqs. 2.3 and 2.4 is computed according to [6]:

$$\sigma^2 = (2 - \frac{\pi}{2})\hat{\sigma}^2.$$
(3.1)

3.1.1 On the unbiased noise estimation in pMRI

To handle potential inhomogeneities of noise power within the general context of MRI, a local approach was considered to estimate a mask of noise amplitude of input noisy image.

First, only-noise image (Fig. 3.1c) was extracted as the difference between input image (Fig. 3.1a) and the image denoised with a standard NLM algorithm (Fig. 3.1b): a pretty high filter streight ($\beta = 1.5$) was set in order to extract the noise as much as possible, still preserving the edges of the image structures from an excessive blurring. Then, a patch-based second order central moment of the only-noise image was calculated (Fig. 3.1d). To avoid patch-related effects and to wash out the spurious hyper-intensities (red arrows) in Fig. 3.1d around the image edges due to the unavoidable structure blurring introduced by NLM filter, a median filter (which is particularly effective in removing low-cardinality structures – pointed out by the red arrows in the figure) was applied to obtain the final mask of local variances (Fig. 3.1e).

(d)



FIGURE 3.1: Spatially varying noise estimation steps in the SVN-NLM method. Difference image (c) between noisy image (a) and NLM filtered image (b) was used to estimate local variances (d). Definitive noise mask (e) was obtained by applying a median filter to remove patch-related effects and residual signal structures (see the red arrows) from (d).

According to Eqs. 2.3–2.6, local noise mask would be underestimated by the local variance of magnitude image values in low SNR regions. To avoid biased results, fixed point formula proposed in [32] was used. Based on the first two moments of the Rician distribution, the correction factor ξ was evaluated as:

$$\xi(\theta_i) = 2 + \theta_i^2 - \frac{\pi}{8} \times \exp\left(\frac{\theta_i^2}{2}\right) \\ \times \left(\left(2 + \theta_i^2\right) I_0\left(\frac{\theta_i^2}{4}\right) + \theta_i^2 I_1\left(\frac{\theta_i^2}{4}\right)\right)^2, \tag{3.2}$$

(e)

where θ_i represents the local SNR and I_n is the n^{th} order Bessel function of first kind. Then, local noise variance σ_i^2 was computed as:

$$\sigma_i^2 = \frac{\hat{\sigma_i}^2}{\xi(\theta_i)},\tag{3.3}$$

where $\hat{\sigma}_i^2$ represents biased local variance.

3.2 MNLM-based filter to denoise SWI images

3.2.1 Acquisition protocol

Two axially-oriented fully flow-compensated spoiled gradient echo sequences were acquired on a Siemens Trio 3T scanner in 4 healthy volunteers using an 8-channel head receiver coil. Common acquisition parameters included a flip angle of 15 degrees, repetition time of 28 ms, echo time of 22.14 ms (in phase water and fat component), field of view of 230 x 194 x 166 mm³, acquisition matrix of 320 x 270 x 128, a GRAPPA factor of 2 and an acquisition time of 5 min and 8 s. The two acquisitions differed from each other in read-out bandwidth only, being 100 and 600 Hz/pixel, respectively. Unfiltered magnitude and phase reconstruction was enabled, thus obtaining a complex volume dataset for each acquisition.

Since in MRI the SNR is related to the square root of the bandwidth, 100 Hz/pixel bandwidth was selected for a high SNR acquisition (kind of "reference acquisition"). In fact, that value is indeed close to the lowest bandwidth limit compatible with clinical research protocols using SWI, as the echo time should not be increased beyond the values suggested in the literature [33] to keep the contrast unchanged.

On the other hand, using the higher bandwidth of 600 Hz/pixel ("noisy acquisition") yields a 2.5-fold decrease in SNR, but does not affect contrast in the image.

3.2.2 Preliminary investigation

Four different pipeline configurations were implemented over and above the proposed MIR-SWI approach.

For the first implementation, the standard NLM filter was applied downstream of the SWI image generation (hereafter as NLM-SWI).

Second, the complex dataset was processed according to IR-SWI denoising scheme described in [34]. For the next test, taking into account the multispectral nature of MR series, the MNLM algorithm was applied on magnitude and phase images before SWI generation (hereafter indicated as MNLM-SWI). In more detail, ϕ and I_m images were first restored with MNLM algorithm and subsequently the MNLM-SWI image was computed according to Eq. (2.2) and Eq. (2.1). Since the phase image is limited to the domain $[-\pi, +\pi)$, in the MNLM-SWI pipeline an unwrapping method must be used before the application of the MNLM denoising filter to recover the true phase on missing multiples of 2π . In that scheme, phase images were unwrapped according to [35].

As fourth test (referred to as MNLM-HP-SWI), the SWI pipeline was revised more deeply. In order to produce the restored SWI image, the MNLM filter was applied to magnitude and high-pass filtered phase images immediately before the phase mask generation, obtaining \hat{I}_m and $\hat{\phi}_{hp}$. The MNLM-HP-SWI image was then obtained applying Eq. (2.1) on \hat{I}_m and $\hat{\phi}_{hp}$.

3.2.3 MIR-SWI method

The proposed MIR-SWI scheme consists of a complex domain-based application of the MNLM filter. Since the full k-space acquired in MRI is assumed to be corrupted with Gaussian white noise, after Fourier transform real and imaginary images are still corrupted by uncorrelated Gaussian noise with the same variance in both complex components [6]. In this setting, unwanted low-frequency B_0 variation was removed by defining real and imaginary images as:

$$I_R = I_m \cdot \cos(\phi_{hp}) \tag{3.4}$$

$$I_I = I_m \cdot \sin(\phi_{hp}) \tag{3.5}$$

and applied MNLM algorithm to I_R and I_I .

After denoising \tilde{I}_R and \tilde{I}_I , the restored magnitude (\tilde{I}_m) and phase $(\tilde{\phi}_{hp})$ images are derived as

$$\tilde{I}_m = \sqrt{\tilde{I}_R^2 + \tilde{I}_I^2} , \qquad (3.6)$$

$$\tilde{\phi}_{hp} = \angle (\tilde{I}_R + i\tilde{I}_I) , \qquad (3.7)$$

and then processed according to Eq. (2.1).

3.3 GPGPU implementation of the NLM filter

The 3D NLM filter was implemented on the NVIDIA parallel computing architecture, which consists in a set of cores, or Scalar Processors (SPs), performing simple mathematical operations.

In the NVIDIA Fermi architecture, each SM has scheduler and dispatch units, execution units and a configurable memory of 64KB, which consists of a register file, an internal shared memory and an L1 cache. This memory is configurable in 16KB (or 48KB) for shared memory and 48KB (or 16KB) for L1 cache.

The Algorithm 1 is the pseudo-code of the NLM filter.

Algorithm 1 Pseudo-code of the NLM algorithm						
1: for each voxel (i_1, i_2, i_3) of the 3D image to be filtered do						
2: Initialize the cumulative sum of weights and the restored value to 0;						
3: for each voxel (j_1, j_2, j_3) of the search window $V_{(i_1, i_2, i_3)}$ do						
4: for each voxel (k_1, k_2, k_3) of the similarity window $(j_{1,j_2,j_3})B_{(i_1,i_2,i_3)}$ do						
5: Cumulate squared Euclidean distance;						
6: end for						
7: Calculate and cumulate the weight of the voxel in search window;						
8: Cumulate the restored value;						
9: end for						
10: Normalize restored value to the sum of the weights;						
11: end for						

In details, the statement at line 1 represents a nested iteration structure.

In the proposed GPU version, the loops on line 1 and line 3 in the Algorithm 1 are logically mapped onto the grid of thread blocks defined by means of the CUDA framework.

A first implementation in CUDA is presented in the Algorithm 2. Moreover, in order

 Algorithm 2 CUDA code of NLM algorithm

 1: int const i_1 = threadIdx.x + blockDim.x*blockIdx.x;

 2: int const i_2 = threadIdx.y + blockDim.y*blockIdx.y;

 3: /* local statements */

 4: if ((i_1 ≥ 0) && (i_1 < X_Dim) && (i_2 ≥ 0) && (i_2 < Y_Dim)) {</td>

 5: for(i_3=0; i_3 < Dim_Z; i_3++) {</td>

 6: /* do something on img[i_1 + i_2*X_Dim + i_3*X_Dim*Y_Dim] */}}

to make this algorithm compatible with multi-GPU architectures, some improvements are introduced. The number of GPU devices is returned by means of a CUDA library function and stored in the variable n_gpus. Then, the third dimension of the image is "splitted" between the available GPUs, setting the first (start_k) and the last (end_k) slices that each GPU has to manage. A sketch of the GPU implementation is reported in the Algorithm 3.

Algorithm 3 CUDA MULTI-GPU code of NLM algorithm

```
1: int const i_1 = threadIdx.x +
                                  blockDim.x*blockIdx.x;
2: int const i_2 = threadIdx.y +
                                  blockDim.y*blockIdx.y;
   /* split the image ''img'' between the ''n_gpus'' GPUs: each GPU works on the
3:
  section of the image ''my_img'' */
   /* local statements */
4:
5:
    for(i_3 = 0; i_3 < Z_Dim/n_gpus;
                                         i_3++) {
        /* do something on my_img[i_1 + i_2*X_Dim + i_3*X_Dim*Y_Dim/n_gpus] */
6:
7:
      } }
```

In order to explore different types of data access, several configurations are tested, both mono- and bi-dimensional, for the thread block size in which each slice is divided. Each thread processes sequentially the voxels along the third dimension. The workload is divided along the third dimension for multi-GPU configurations. Inside each GPU the workload is divided along the first and second dimensions, in strips (mono-dimensional configurations) and tiles (bi-dimensional configurations) of threads. Strip or tile is allowed to cover entirely or only partially the slice grid.

Also the impact of L1-cache on performance are tested, using the binary L1-prefer setting, which allows to choose between two possible configurations: 48KB of shared memory and 16KB of L1-cache (no L1-prefer), or 16KB of shared memory and 48KB of L1-cache (L1-prefer).

Chapter 4

Results

4.1 SVN estimation

In order to evaluate the performances of the NLM filter with actual SVN mask, noisefree MR images were corrupted with both Gaussian and Rician spatially varying noise and integrated the estimated noise mask in Eq. 2.12 to adapt voxel-by-voxel the filter strength as function of the local noise power. Due to high computational complexity of NLM algorithm, the multi-GPU implementation [36], as described in §3.3, in both pre-processing (standard NLM) and denoising (SVN-NLM) steps was used.

Visual inspection and residuals between noisy and denoised image have been used to rate the quality of denoising. As quality measure, Peak SNR (PSNR) was evaluated as:

$$PSNR(\hat{f}(x), f(x)) = 10 \log_{10} \frac{M^2}{\frac{1}{|\Omega|} \sum_{x \in \Omega} (f(x) - \hat{f}(x))^2},$$
(4.1)

where M is maximum value of noise-free image (f(x)) and $\hat{f}(x)$ is the denoised image.

SVN-NLM-denoised and residual images (Fig. 4.1(b)-(d)) show high performance denoising without sensible removal of image structures. Compared with standard NLM, the proposed SVN-NLM approach produces better results in terms of both visual inspection and PSNR, with a gain of ~ 1.3 dB (Fig. 4.1).

Moreover, to confirm the stability of the proposed SVN estimation scheme, the "ground truth" was corrupted with uniform noise. As shown in Fig. 4.2, SVN-NLM produces similar results in comparison with standard NLM.

Finally, the SVN-NLM method was compared with an implementation of the state-ofthe-art NLM algorithm for pMRI, proposed by [37]. The different strategies to estimate



FIGURE 4.1: Denoising results on corrupted image (average 10% non-uniform noise) with standard NLM (a) and NLM with SVN mask (b). PSNRs are 33.4162 and 34.7183 for NLM and SNV-NLM, respectively. Absolute values of residual images (enhanced by a factor 10) are shown in (c) and (d); the red ellipses highlight the image structures lost by standard NLM.

the only-noise image and to statistically handle the patch-related effects of the noise variance strongly influence denoising results in case of non-uniform noise distribution. In this context, the proposed method outperforms previous noise estimation and, accordingly, denoising in terms of both structure preservation and PSNR, with a gain of ~ 0.8 dB (see Fig. 4.3).

4.2 Restoring SWI images

In Fig. 4.4 were compared the denoising outcomes obtained by testing the different NLM pipelines described in $\S3.2.2$ in contrast to the proposed method (please refer to $\S3.2$)



FIGURE 4.2: Denoising results on corrupted image (10% uniform noise) with standard NLM (a) and NLM with SVN mask (b). 31.1033 and 31.1071 are PSNRs for NLM and SNV-NLM, respectively. Absolute values of residual images (enhanced by a factor 10) are shown in (c) and (d).

for the description of the different methods). In Fig. 4.5 the denoising results on highpass filtered phase images were show before the SWI image generation. The denoising performance of the proposed scheme was compared to the original SWI, to the NLM-SWI, to the MNLM-SWI, to MNLM-HP-SWI and to IR-SWI. To assess both the correct noise removal and the preservation of edges and tiny structures, we used a high-SNR SWI dataset acquired with a bandwidth of 100 Hz/pixel (labelled as SWI-100Hz) as the "reference SWI".

4.2.1 Qualitative assessment (brain tissues)

In order to evaluate the performance of MIR-SWI in comparison with original SWI, NLM-SWI, MNLM-SWI, MNLM-HP-SWI and the IR-SWI, the six sets of images from 4 healthy volunteers were randomly presented to two neuroradiologists with more than





FIGURE 4.3: Denoising results on "ground truth" (a) corrupted with an average 10% non-uniform noise. PSNRs of NLM proposed in [37] (b) and actual SVN-NLM (c) are 33.9606 and 34.7305, respectively. Red arrows highlight most relevant differences between images.

20-years experience in MR neuroimaging as 50 sextets of corresponding axial slices at different brain levels from the foramen magnum to the vertex. Semiquantitative assessment of the images was performed blindly and in consensus by rating the gray/white matter differentiation, the presence of artifacts at the brain/cerebrospinal fluid interface, or any other obvious artifacts, the confidence in detecting clinical relevant findings and the overall image quality, on a 0-5 scale, with 5 being the best representation of intracranial structures without relevant artifacts and 0 the worst and clinically inadequate display.

The MIR-SWI images scored 4 or 5 in 46 cases (92%), the SWI images scored 4 or 5 in 31 cases (62%), the IR-SWI images scored 4 or 5 in 41 cases (82%) while standard NLM-SWI scored 4 or 5 in only 16 cases (32%), and were never preferred over the corresponding SWI counterpart (Fig. 4.6).

MNLM-SWI and MNLM-HP-SWI schemes clearly produced images of systematically lower quality (e.g., sharpness, vein conspicuity, gray/white matter differentiation), as shown in Fig. 4.4 and in Fig. 4.5 and confirmed by the semi-quantitative analysis in



FIGURE 4.4: Results of different denoising pipelines on SWI image generation. Axial brain mIPs (corresponding to a volume of 20 mm) at the level of the lateral ventricles of SWI-100Hz (a), SWI (b), NLM-SWI (c), IR-SWI (d), MNLM-SWI (e), MNLM-HP-SWI (f), and MIR-SWI (g) images. The number of phase mask multiplications is set to 4. Enhanced visibility of venous structures without loss of tissue contrast is evident in (g) compared to (b-f).

Fig. 4.6. For this reason, the quality metrics computed in the subsequent analysis for these two schemes are not shown.

4.2.2 Quantitative evaluation (veins)

The denoising performance was quantitatively evaluated both by a vein-based contrastto-noise ratio (VB-CNR) comparison and by a vessel-profile analysis.

As proposed by Jang et al. [38], VB-CNR was defined as:

$$VB-CNR_{ab} = \frac{|S_a - S_b|}{\sigma_b}.$$
(4.2)



FIGURE 4.5: Influence of different denoising pipelines on high-pass filtered phase images. Argument of the phase mask function (somehow equivalent to highpass filtered phase) in the following pipelines: SWI-100Hz (a), SWI (b), IR-SWI (c), MNLM-SWI (d), MNLM-HP-SWI (e) and MIR-SWI (f). The tissues outside the brain were masked in order to focus on the denoising results. The image obtained with MIR-SWI scheme shows good noise suppression while preserving brain structures compared to both MNLM-SWI and MNLM-HP-SWI images.



FIGURE 4.6: Semiquantitative visual assessment. Frequency histogram of the semiquantitative scores for the display of the brain structures of the MNLM-SWI (gray), MNLM-HP-SWI (orange), SWI (green), NLM-SWI (yellow), IR-SWI (cyan) and MIR-SWI (red) images. Score values from 1 to 5 indicate increasing overall image quality (see text).

where S_a and S_b are the mean signal intensities in region a and b and σ_b is standard deviation of region b. Region a is defined as the set of pixels on a line passing through the center of the vein whereas region b is the set of pixels on a line surrounding the vein [38]. Using the same approach, a single observer manually segmented five veins: anterior septal vein (AS), thalamostriate vein (TS), internal cerebral vein (IC), lateral atrial vein (LA) and a silvian cortical vein (SC), and measured their VB-CNRs (Fig. 4.7a). In Fig. 4.7b the VB-CNR values for each SWI scheme (SWI, NLM-SWI, IR-SWI and MIR-SWI) and for each vein are compared in a healthy subject.

Another quantitative description of the denoising effects on the contrast changes between tissues and an estimate of possible edge blurring is given by vessel-profile analysis. As an example, in Fig. 4.8 the plotting of in-plane profiles of SWI, NLM-SWI, IR-SWI and MIR-SWI voxel intensities perpendicular to a cortical small vein was reported.

4.2.3 Number of phase mask multiplications

From Eq. (2.1) one sees that the number of phase mask multiplications may be optimized in order to enhance the contrast between veins or gray matter versus the surrounding tissues while keeping the noise level within reasonable limits. In this setting, different values of phase mask multiplications were tested both in terms of visual inspection (Fig. 4.9) and of VB-CNR (Fig. 4.10).

4.3 Multi-GPU implementation

4.3.1 Consistency

Since the aim is to produce a strictly equivalent GPU implementation of the sequential NLM algorithm, the implementation consistency was checked by comparing voxel-byvoxel the images obtained by one-core-CPU and GPU denoising. In Fig. 4.11 the 3D NLM filtering result was show on a real 3D knee MRI dataset. The difference between the GPU and CPU restored images falls within machine precision order of magnitude which are likely to be due to the arithmetic logic unit precision.

4.3.2 Performance

In order to investigate cache size impact on the execution time, several test runs were performed varying L1-prefer switch. Results are shown in Table 4.1. L1-prefer choice gives a benefit on lager dataset, with a performance improvement ranging from fraction



FIGURE 4.7: **VB-CNR analysis.** SWI-100HZ axial brain mIP (a) corresponding to a volume of 20 mm shows the five venous ROIs used for the quantitative evaluation of the MIR-SWI denoising scheme. Green lines represent the veins used for VB-CNR analysis while cyan lines are the background counterparts positioned on neighbooring tissues (anterior septal vein, AS; thalamostriate vein, TS; internal celebral vein, IC; lateral atrial vein, LA; silvian cortical vein, SC). The VB-CNR bar graph of each vein (b) shows an overall higher contrast between veins and background of the MIR-SWI (red bars) compared to the other schemes.

of percent in the smallest dataset to some 5% in the largest ones. These results suggest that the L1 miss rate is low enough to have high performance processing even with old generation cards having small amount of cache.

The strip or tile thread division influenced the performance of the filter in terms of computing time due to the different type of data access. Experimental results proved that the optimal configuration was given by the strip subdivision. In Table 4.1 running times of (128,1,1) configuration on 2-GPU were reported.



FIGURE 4.8: Vessel-profile comparison. Top: SWI, NLM-SWI, IR-SWI and MIR-SWI axial brain slices (from left to right respectively) in a healthy volunteer. The red lines represent the domain used to plot the in-plane profiles of the voxel intensities perpendicular to a small right frontal vein. Bottom: the comparison of the corresponding in-plane profiles of the SWI (green line), NLM-SWI (yellow line), IR-SWI (cyan line) and MIR-SWI (dotted red line) voxel intensities shows that MIR-SWI, IR-SWI and NLM-SWI schemes enhance the SNR of the parenchyma (depicted by the line plateau) compared to the SWI vessel profile, but only the MIR-SWI does not introduce a detrimental blurring between the vessel and surrounding tissues.



FIGURE 4.9: The effect of the n values on SWI images. mIPs of the same targeted volume of 20 mm at varying n values. In reference to the SWI-100Hz image, MIR-SWI shows both satisfactory noise removal and better vessel enhancement at increasing n values compared to the other SWI schemes.



FIGURE 4.10: Bar graph of VB-CNRs at different n values. VB-CNR analysis performed on the same veins of Fig. 4.7a (AS: anterior septal vein, TS: thalamostriate vein, IC: internal cerebral vein, LA: lateral atrial vein, SC: sylvian cortical vein) as they appeared in the three rightmost columns of Fig. 4.9. From each vein, the triplets of bars with the same color correspond to the images with n value of 6, 7 and 8, from left to right, respectively. Among the denoising schemes under evaluation, only MIR-SWI (red bars) consistently showed increased VB-CNR in all selected veins.

	Cache configuration			
	L1-prefer	no L1-prefer		
64x64x64	8.04	8.55		
128x128x64	9.43	9.42		
128x128x128	11.2	11.2		
256x256x128	19.9	20.2		
256x256x256	28.3	29.0		
512x512x128	39.5	40.4		
512x512x256	71.7	75.0		
512x512x512	138	148		

TABLE 4.1: Performance of L1-prefer switch. L1-prefer switch influence on execution times for (128,1,1) block size configuration and 3D random datasets.

In Table 4.2 a comparison between running times of CPU, single GPU and multi-GPU implementation of 3D NLM filter with various thread block size was shown. Reported running times included the overall data transfer between CPU and GPU and viceversa, which even for the biggest datasets appears to be negligible. Speed-up values suggested that the bigger the dataset to be filtered, the better the scalability of the implementation. Moreover, the optimal thread size seems to be strips of thread between 128 and 256 elements. This result was consistent with NVIDIA guidelines [24]. Finally, on large



FIGURE 4.11: **Implementation consistency.** From left to right and from top to bottom, the frames show a central slice of the original dataset, the GPU restored image, the CPU restored image and the difference between CPU and GPU filtered images (enhanced by a scaling factor of 10^6), respectively.

datasets strip configuration should be prefered to tile configuration of the same size because more sequential memory access of the former.

In Table 4.3, the behavior of running times against the search $(|V_i|)$ and similarity $(|_j B_i|)$ window cardinalities was investigated. An high and almost constant speed-up among the various experiments was noted, which made feasible large window filter testing in a reasonable time.

Finally, in Figure 4.12 the CPU, single GPU and multi-GPU GFlops for variable dataset sizes were outlined.

	Execution time/ Speed-up								
	Single GPU				Multi-GPU				CPU
Dataset size	(16, 16, 1)	(128,1,1)	(256,1,1)	(512,1,1)	(16, 16, 1)	(128,1,1)	(256,1,1)	(512,1,1)	
64^3	5.08/4.47	5.73/ 3.96	5.63/ 4.03	6.94/ 3.27	11.7/1.94	8.04/2.82	8.53/2.66	12.6/ 1.80	22.7
$128^2 \text{x}64$	6.48/ 13.6	7.30/ 12.1	7.08/ 12.5	10.2/8.61	8.97/ 9.83	9.43/ 9.35	9.31/9.47	10.8/8.13	88.2
128^{3}	9.06/ 19.3	10.8/16.2	10.4/16.9	16.7/ 10.5	10.3/17.0	11.2/15.7	10.9/16.0	14.0/ 12.5	175
$256^2 \text{x} 128$	22.1/ 31.8	24.3/28.9	25.0/ 28.0	31.0/22.6	16.9/ 41.6	19.9/35.2	18.6/37.8	21.1/33.2	702
256^{3}	40.5/ 34.6	44.7/ 31.3	47.2/ 29.6	59.4/ 23.6	26.0/53.8	28.3/49.4	29.5/47.5	35.1/ 39.9	1400
$512^2 x 128$	68.8/ 41.0	67.0/42.1	67.1/ 42.0	72.7/38.8	40.5/69.6	39.5/71.4	40.0/70.5	42.4/66.5	2820
$512^2 x 256$	136/41.2	132/42.6	131/42.8	142/39.5	73.8/ 76.3	71.7/78.5	72.0/78.2	77.5/ 72.6	5630
512^3	277/40.7	268/42.2	264/42.7	285/ 39.6	142/79.3	138/ 82.0	137/82.4	148/76.2	11300

TABLE 4.2: Execution times and speed-up values for several block size configurations and 3D random datasets. Search and similarity windows have been set according to $|V_i| = 11^3$ and $|_j B_i| = 3^3$.

	Execution time / Speed-up								
	Single GPU				Multi-GPU				CPU
$(V_i , _j B_i)$	(16, 16, 1)	(128,1,1)	(256,1,1)	(512,1,1)	(16, 16, 1)	(128,1,1)	(256, 1, 1)	(512,1,1)	
$(11^3, 3^3)$	68.8/ 41.0	67.0/42.1	67.1/ 42.0	72.7/ 38.8	40.5/ 69.6	39.5/71.4	40.0/ 70.5	42.4/66.5	2820
$(21^3, 3^3)$	447/ 44.4	434/45.7	434/45.7	467/42.4	228/ 87.0	222/89.4	221/ 89.6	239/82.8	19800
$(11^3, 5^3)$	235/ 34.6	221/ 36.7	223/ 36.6	255/31.9	123/ 61.1	116/69.9	117/ 69.3	130/ 62.4	8140
$(21^3, 5^3)$	1650/ 35.5	1510/38.8	1520/ 38.7	1820/ 32.2	817/ 72.0	757/ 77.6	764/ 77.0	906/ 64.8	58800

TABLE 4.3: Execution times and speed-up values for a 3D random dataset (size = $512 \times 512 \times 128$) for several window configurations.



FIGURE 4.12: Outline of the CPU, single GPU and multi GPU GFlops values for different dataset sizes. Please note the logarithmic scale of the axes.

Chapter 5

Discussion and Conclusion

5.1 Uniform and non-uniform unbiased noise estimation

A new method for MRI noise estimation has been presented to address the problem of biased estimation in case of spatially dependent noise distribution when backgroundbased variance extraction is performed. Visual inspection clearly proves better results compared to standard NLM filter where anatomical structures are visible in image residuals. Moreover, both quality measure and PSNR demonstrate the ability of SVN-NLM to remove noise adaptively according to local noise variance. As the actual SVN estimation does not need an *a priori* knowledge of sensitivity maps and subsampling factor, the derived noise mask is applicable on parallel MR images reconstructed with both SENSE and GRAPPA techniques. Moreover, the SVN estimation outperforms previous noise calculation in case of non-uniform noise distribution [37, 39, 40]. In particular, compared to denoising method proposed in [37], both better PSNR results and gain in contrast-to-noise ratio (CNR) are obtained.

On the other hand, as the proposed approach has been demonstrated to produce similar results as the standard NLM in case of uniform noise distribution, SVN-NLM is robust enough to be applicable also on general MRI datasets (both standard MRI and pMRI).

In conclusion, the application of SVN method in post-processing tasks of both standard and parallel MRI can clearly benefit not only visual inspection, but also quantitative techniques that rely on good quality of the data.

5.2 On the novel MIR-SWI approach

Noise removal algorithms in a SWI pipeline have received little attention to data, perhaps even the application of a robust denoising filter as the NLM algorithm produces poor results on SWI images. The excessive blurring and loss of anatomical structures shown in Fig. 4.4c is likely to be caused by an incomplete exploitation of the information contained within the two complex channels. In this context, the generation of a robust method to both enhance the vessels and increase the CNR of SWI images is a non-trivial operation.

A new strategy to restore SWI images based on a dedicated NLM scheme has been recently developed [34]. In that approach, the restoring process of SWI images was solved by applying the NLM filter on the complex image domain between the high-pass filtering step and the phase mask generation, thus taking advantage of the uncorrelated Gaussian distribution with the same noise variance in the real and imaginary parts of the complex image. In that study, while effective noise removal led to an evident noise reduction and a SNR increase, the restoring pipeline did not take into account the multispectral nature of the SWI scheme, thus producing blurred images with reduced differentiation between the brain regions and loss of tiny structures such as small vessels (Fig. 4.4d).

In the proposed approach, a new denoising pipeline based on a NLM algorithm to restore SWI images is presented. Several pipeline configurations of the multispectral version of the NLM denoising filter [5] were tested. As outlined in Fig. 4.5d, due to incorrect noise removal in phase image, lacking phase contribution is displayed in SWI image when MNLM-SWI method was applied, resulting in an extensive smoothing and loss of phase contrast (Fig. 4.4e). Even the MNLM-HP-SWI scheme similarly produces unsatisfactory outcomes: in fact although the phase information in the MNLM-HP-SWI scheme (Fig. 4.5e) is higher compared to the phase image processed with the MNLM-SWI method, poor anatomical information was obtained due to the different noise distribution between magnitude and high-pass filtered phase images. As outlined in Fig. 4.4f, the contrast in SWI images is generated by an additional phase contribution, but the contrast from some small vessels is missing.

The key in the MIR-SWI approach is to revisit the SWI pipeline by applying a MNLM filter on real and imaginary components of computed complex data between the phase high-pass filtering stage and the phase mask evaluation. The main benefit of the proposed method is to reduce the noise propagation in the non-linear SWI pipeline, thus avoiding the introduction of the biases due to a non-null first moment of the transformed zero-mean white Gaussian noise in complex images. Moreover, the noise removal in

ering also the magnitude information into

the phase image is addressed by considering also the magnitude information into the weighted averaging process, thus taking advantage of the multispectral property of the SWI scheme.

The results demonstrate that the proposed method clearly improves the SNR and properly preserves the brain structures (Fig. 4.4g - 4.5f). In fact, the visual assessment by two neuroradiologists showed that MIR-SWI images consistently displayed better gray/white matter differentiation, fewer artifacts and improved image quality (Fig. 4.6). Moreover, the visibility of vessels is enhanced due to an increased CNR between the vessels and the surrounding tissues (Fig. 4.7b - 4.8). Unlike NLM-SWI, MIR-SWI images show a better preservation of faint structures revealed by the SWI phase mask, still guaranteeing a high SNR. Finally, compared to IR-SWI images, both better vein visibility and increased contrast between brain structures were obtained when MIR-SWI is applied, thus producing a benefit for quantitative techniques that rely on good quality of the data (such as the quantitation of brain iron content).

Taking into account the noise propagation during the SWI image generation, Haacke et al. [9] proposed to set the number of phase mask multiplications to 4, in order to obtain the best compromise between vessel enhancement and overall image quality. In this setting, several phase mask multiplication values were tested to achieve the most satisfactory phase contrast in SWI images. As pointed out in $\S4.2.3$, with the proposed method better phase contribution was obtained at higher n levels without a significant increase of noise, thus resulting in a clear improvement in both CNR and vessel visibility. This is particularly true when the minimum Intensity Projection (mIP) is performed (Fig. 4.9), a reformat technique in which noise propagation is more evident. Beside being preferable in the qualitative assessment of SWI images at different n values processed with the SWI schemes under evaluation (Fig. 4.9), MIR-SWI quantitatively showed increased VB-CNR in all selected veins (Fig. 4.10). Although the optimal choice of the nvalue was not the primary aim of this work, it was demonstrated that the proper handling of the noise in the MIR-SWI scheme may be used to increase the contrast with different *n* values. In fact, only MIR-SWI guarantees a positive trend in CNR in every vein within the considered range of n values. Moreover, MIR-SWI, compared to SWI-100Hz, proved capable of largely filling a 6-fold gap between the two acquisition bandwidths, as shown in Fig. 4.10 where red bars were never shorter than the corresponding blue bars.

VB-CNR values could not be reasonably compared across different subjects, due to marked individual differences in vessel features (e.g., vein caliber, ferromagnetic load, haemodynamic parameters, etc.), that prevent a proper reproducibility analysis. However, the robustness of MIR-SWI can be inferred by evaluating the VB-CNR values of the different veins within the same subject, due to known intra-subject vessel variability. As shown in Fig. 4.7 and 4.10, VB-CNR values of different veins can be considered representative of the performance variability of the different SWI schemes under a wide range of analyzed subjects.

MIR-SWI takes advantage of the multi-GPU implementation of a generic NLM-scheme (see [36]), and is computationally feasible within clinically acceptable times (~ 3 minutes for a 3D 320x270x128 complex dataset with a typical denoising parameter setup [d = 1, M = 5] and 2 NVIDIA GeForce GTX 690).

In summary, the above-mentioned findings may prompt the application of the proposed algorithm to SWI processing in larger patient databases to assess whether clinical usefulness may be improved. Moreover, it may be used to enhance the overall quality of finer phase-image processing, such as Susceptibility Weighted Image Mapping (SWIM) or Quantitative Susceptibility Mapping (QSM) [41].

5.3 On the Multi-GPU 3D NLM implementation

NLM filter is a state-of-the-art denoising algorithm. However, the huge amount of computational load prevents the large-scale diffusion of its most common implementations.

Here, the first multi-GPU implementation of a fully 3D NLM filter in presented. Several configurations of thread block organization and data access were analyzed, thus identifying a set of optimal settings that guarantee high performance results for a wide spectrum of application scenarios. The reduction of running times showed that scalability is close to ideal one for most common dataset sizes, specially those typical of MRI clinical pratice. Speed-up high values encouraged the exploration of more sophisticated algorithm variants, and reduced the gap between the previous execution times and acceptable performance for real-time scenarios.

Bibliography

- A. Buades, B. Coll, and J. Morel. Image denoising methods. a new nonlocal principle. Siam Review, 52(1):113–147, 2010.
- [2] José V Manjón, José Carbonell-Caballero, Juan J Lull, Gracián García-Martí, Luís Martí-Bonmatí, and Montserrat Robles. Mri denoising using non-local means. *Medical image analysis*, 12(4):514–523, 2008.
- [3] Pierrick Coupé, Pierre Yger, Sylvain Prima, Pierre Hellier, Charles Kervrann, and Christian Barillot. An optimized blockwise nonlocal means denoising filter for 3-d magnetic resonance images. *Medical Imaging, IEEE Transactions on*, 27(4):425– 441, 2008.
- [4] Nicolas Wiest-Daesslé, Sylvain Prima, Pierrick Coupé, Sean Patrick Morrissey, and Christian Barillot. Rician noise removal by non-local means filtering for low signal-to-noise ratio mri: applications to dt-mri. In *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2008*, pages 171–179. Springer, 2008.
- [5] José V Manjón, Neil A Thacker, Juan J Lull, Gracian Garcia-Martí, Luís Martí-Bonmatí, and Montserrat Robles. Multicomponent mr image denoising. *Journal of Biomedical Imaging*, 2009:18, 2009.
- [6] Hákon Gudbjartsson and Samuel Patz. The rician distribution of noisy mri data. Magnetic Resonance in Medicine, 34(6):910–914, 1995.
- [7] Santiago Aja-Fernández, Gonzalo Vegas-Sánchez-Ferrero, and Antonio Tristán-Vega. About the background distribution in mr data: a local variance study. *Magnetic Resonance Imaging*, 28(5):739–752, 2010.
- [8] Santiago Aja-Fernández, Antonio Tristán-Vega, and Carlos Alberola-López. Noise estimation in single-and multiple-coil magnetic resonance data based on statistical models. *Magnetic resonance imaging*, 27(10):1397–1409, 2009.
- [9] E Mark Haacke, Yingbiao Xu, Yu-Chung N Cheng, and Jürgen R Reichenbach. Susceptibility weighted imaging (swi). Magnetic Resonance in Medicine, 52(3): 612–618, 2004.

- [10] EM Haacke, J Tang, J Neelavalli, and YCN Cheng. Susceptibility mapping as a means to visualize veins and quantify oxygen saturation. *Journal of Magnetic Resonance Imaging*, 32(3):663–676, 2010.
- [11] E Mark Haacke, S Mittal, Z Wu, J Neelavalli, and Y-CN Cheng. Susceptibilityweighted imaging: technical aspects and clinical applications, part 1. American Journal of Neuroradiology, 30(1):19–30, 2009.
- [12] S Mittal, Z Wu, J Neelavalli, and EM Haacke. Susceptibility-weighted imaging: technical aspects and clinical applications, part 2. American Journal of Neuroradiology, 30(2):232–252, 2009.
- [13] Vivek Sehgal, Zachary Delproposto, E Mark Haacke, Karen A Tong, Nathaniel Wycliffe, Daniel K Kido, Yingbiao Xu, Jaladhar Neelavalli, Djamel Haddar, and Jürgen R Reichenbach. Clinical applications of neuroimaging with susceptibilityweighted imaging. *Journal of Magnetic Resonance Imaging*, 22(4):439–450, 2005.
- [14] Vivek Sehgal, Zachary Delproposto, Djamel Haddar, E Mark Haacke, Andrew E Sloan, Lucia J Zamorano, Geoffery Barger, Jiani Hu, Yingbiao Xu, Karthik Praveen Prabhakaran, et al. Susceptibility-weighted imaging to visualize blood products and improve tumor contrast in the study of brain masses. *Journal of magnetic resonance imaging*, 24(1):41–51, 2006.
- [15] E Mark Haacke, Malek Makki, Yulin Ge, Megha Maheshwari, Vivek Sehgal, Jiani Hu, Madeswaran Selvan, Zhen Wu, Zahid Latif, Yang Xuan, et al. Characterizing iron deposition in multiple sclerosis lesions using susceptibility weighted imaging. *Journal of Magnetic Resonance Imaging*, 29(3):537–544, 2009.
- [16] EM Haacke, ZS DelProposto, S Chaturvedi, V Sehgal, M Tenzer, J Neelavalli, and D Kido. Imaging cerebral amyloid angiopathy with susceptibility-weighted imaging. *American journal of neuroradiology*, 28(2):316–317, 2007.
- [17] Matthew Schrag, Grant McAuley, Justine Pomakian, Arshad Jiffry, Spencer Tung, Claudius Mueller, Harry V Vinters, E Mark Haacke, Barbara Holshouser, Daniel Kido, et al. Correlation of hypointensities in susceptibility-weighted images to tissue histology in dementia patients with cerebral amyloid angiopathy: a postmortem mri study. Acta neuropathologica, 119(3):291–302, 2010.
- [18] Zhen Wu, Sandeep Mittal, Karl Kish, Yingjian Yu, J Hu, and E Mark Haacke. Identification of calcification with mri using susceptibility-weighted imaging: a case study. *Journal of Magnetic Resonance Imaging*, 29(1):177–182, 2009.
- [19] Yongming Dai, Mengsu Zeng, Ruokun Li, Shengxiang Rao, Caizhong Chen, Zachary DelProposto, EM Haacke, Jiani Hu, and Jerecic Renate. Improving detection of

siderotic nodules in cirrhotic liver with a multi-breath-hold susceptibility-weighted imaging technique. *Journal of Magnetic Resonance Imaging*, 34(2):318–325, 2011.

- [20] Qi Yang, Jiangtao Liu, Kuncheng Li, Yongming Dai, Mark Haacke, and Renate Jerecic. Visualization of carotid plaque calcification-a novel approach using susceptibility weighted mr imaging. *Journal of Cardiovascular Magnetic Resonance*, 12: 1–2, 2010.
- [21] Fernanda Palhano Xavier De Fontes, Guillermo Andrade Barroso, Pierrick Coupé, and Pierre Hellier. Real time ultrasound image denoising. *Journal of real-time image processing*, 6(1):15–22, 2011.
- [22] Bart Goossens, Hiep Luong, Jan Aelterman, Aleksandra Pižurica, and Wilfried Philips. A gpu-accelerated real-time nlmeans algorithm for denoising color video sequences. Advanced concepts for intelligent vision systems, pages 46–57, 2010.
- [23] Kuidong Huang, Dinghua Zhang, and Kai Wang. Non-local means denoising algorithm accelerated by gpu. Sixth International Symposium on Multispectral Image Processing and Pattern Recognition, pages 749711–749711, 2009.
- [24] David Kirk et al. Nvidia cuda software and gpu parallel computing architecture. ISMM, 7:103–104, 2007.
- [25] Anagha Deshmane, Vikas Gulani, Mark A Griswold, and Nicole Seiberlich. Parallel mr imaging. *Journal of Magnetic Resonance Imaging*, 36(1):55–72, 2012.
- [26] James F Glockner, Houchun H Hu, David W Stanley, Lisa Angelos, and Kevin King. Parallel mr imaging: A user's guide1. *Radiographics*, 25(5):1279–1297, 2005.
- [27] Klaas P Pruessmann. Encoding and reconstruction in parallel mri. NMR in Biomedicine, 19(3):288–299, 2006.
- [28] Mark A Griswold, Peter M Jakob, Robin M Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine*, 47 (6):1202–1210, 2002.
- [29] Klaas P Pruessmann, Markus Weiger, Markus B Scheidegger, Peter Boesiger, et al. Sense: sensitivity encoding for fast mri. *Magnetic resonance in medicine*, 42(5): 952–962, 1999.
- [30] Santiago Aja-Fernández, Antonio Tristán-Vega, and W Scott Hoge. Statistical noise analysis in grappa using a parametrized noncentral chi approximation model. *Magnetic Resonance in Medicine*, 65(4):1195–1206, 2011.

- [31] Santiago Aja-Fernández, Gonzalo Vegas-Sánchez-Ferrero, and Antonio Trsitan-Vega. Statistical noise analysis in sense parallel mri. arXiv preprint arXiv:1402.4067, 2014.
- [32] Cheng Guan Koay and Peter Basser. Analytically exact correction correction scheme for signal extraction from noisy magnitude mr signals. *Journal of Magnetic Resonance*, 179:317–322, 2006.
- [33] E Mark Haacke and Jürgen R Reichenbach. Susceptibility weighted imaging in mri: basic concepts and clinical applications. John Wiley & Sons, 2014.
- [34] Pasquale Borrelli, Giuseppe Palma, Enrico Tedeschi, Sirio Cocozza, Carmela Russo, Marco Comerci, Bruno Alfano, Mark E Haacke, and Marco Salvatore. Improving snr in susceptibility weighted imaging by a nlm-based denoising scheme. *Imaging* Systems and Techniques (IST), 2014 IEEE International conference on., pages 346– 350, 2014.
- [35] R Cusack and N Papadakis. New robust 3-d phase unwrapping algorithms: application to magnetic field mapping and undistorting echoplanar images. *Neuroimage*, 16(3):754–764, 2002.
- [36] Giuseppe Palma, Francesco Piccialli, Marco Comerci, Pasquale De Michele, Pasquale Borrelli, Salvatore Cuomo, and Bruno Alfano. 3d non-local means denoising via multi-gpu. Proceedings of the 2013 Federated Conference on Computer Science and Information Systems, pages 495—498, 2013.
- [37] José V Manjón, Pierrick Coupé, Luis Martí-Bonmatí, D Louis Collins, and Montserrat Robles. Adaptive non-local means denoising of mr images with spatially varying noise levels. *Journal of Magnetic Resonance Imaging*, 31(1):192–203, 2010.
- [38] Ung Jang, Yoonho Nam, Dong-Hyun Kim, and Dosik Hwang. Improvement of the snr and resolution of susceptibility-weighted venography by model-based multi-echo denoising. *Neuroimage*, 70:308–316, 2013.
- [39] O Dietrich, J G Raya, M Ingrisch, M F Reiser, and O Schoenberg. Noise characteristics in mri with multi-channel arrays, parallel imaging, and reconstruction filters. in Proc. International Society of Magnetic Resonance in Medicine, 2007.
- [40] Jelle Veraart, Jeny Rajan, Ronald Peeters, Alexander Leemans, Stefan Sunaert, and Jan Sijbers. Estimation of spatially variable rician noise map in diffusion mri. in Proc. International Society of Magnetic Resonance in Medicine, 2013.
- [41] E Mark Haacke, Saifeng Liu, Sagar Buch, Weili Zheng, Dongmei Wu, and Yongquan Ye. Quantitative susceptibility mapping: current status and future directions. *Magnetic resonance imaging*, 33(1):1–25, 2015.

Acknowledgements

The research work on which this thesis is based has been performed during my stay at the Institute of Biostructures and Bioimaging - National Research Council of Italy, Naples, Italy. This work has been accomplished under the supervision of Giuseppe Palma, Ph.D., which I am grateful for his highly professional guidance and for being so patient with me.

I wish to express my gratitude to the Collegio dei Docenti del Dottorato di Ricerca in Imaging Molecolare, and to Prof. Marco Salvatore, for giving me the chance to pursue my research studies.

Finally, I would like to thank the institution of University of Naples "Federico II" that, during the three years of doctoral studies, has provided me with the necessary financial support.