NeuroImage 118 (2015) 584–597

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

An efficient total variation algorithm for super-resolution in fetal brain MRI with adaptive regularization

Sébastien Tourbier [a,b,*], Xavier Bresson [c], Patric Hagmann [b], Jean-Philippe Thiran [d,b], Reto Meuli [b], Meritxell Bach Cuadra [a,b,d]

[a] Centre d'Imagerie BioMédicale (CIBM), Switzerland
[b] Radiology department, Lausanne University Hospital Center (CHUV), University of Lausanne (UNIL), Switzerland
[c] Signal Processing Laboratory (LTS2), Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
[d] Signal Processing Laboratory (LTS5), Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland

ARTICLE INFO

Article history:
Received 17 December 2014
Accepted 4 June 2015
Available online 10 June 2015

Keywords:
Fetal brain MRI
Super-resolution
Total variation
Fast convex optimization

ABSTRACT

Although fetal anatomy can be adequately viewed in new multi-slice MR images, many critical limitations remain for quantitative data analysis. To this end, several research groups have recently developed advanced image processing methods, often denoted by super-resolution (SR) techniques, to reconstruct from a set of clinical low-resolution (LR) images, a high-resolution (HR) motion-free volume. It is usually modeled as an inverse problem where the regularization term plays a central role in the reconstruction quality. Literature has been quite attracted by Total Variation energies because of their ability in edge preserving but only standard explicit steepest gradient techniques have been applied for optimization. In a preliminary work, it has been shown that novel fast convex optimization techniques could be successfully applied to design an efficient Total Variation optimization algorithm for the super-resolution problem. In this work, two major contributions are presented. Firstly, we will briefly review the Bayesian and Variational dual formulations of current state-of-the-art methods dedicated to fetal MRI reconstruction. Secondly, we present an extensive quantitative evaluation of our SR algorithm previously introduced on both simulated fetal and real clinical data (with both normal and pathological subjects). Specifically, we study the robustness of regularization terms in front of residual registration errors and we also present a novel strategy for automatically select the weight of the regularization as regards the data fidelity term. Our results show that our TV implementation is highly robust in front of motion artifacts and that it offers the best trade-off between speed and accuracy for fetal MRI recovery as in comparison with state-of-the art methods.

Introduction

Recent advances in clinical magnetic resonance imaging (MRI) provide an unprecedented opportunity to study the human brain growth *in-utero*. Concretely, T2-weighted (T2w) MR images have proven to provide the best ability to delineate structures and layers in fetal brain MRI (Prayer, 2011; Coakley et al., 2004; PC et al., 2006; Garel, 2006). Clinical MRI examination involves the use of ultra-fast multi-slices MR sequences, such as Half Fourier Acquisition Single-Shot Turbo Spin Echo (HASTE) or Single-Shot Fast Spin Echo (ssFSE) (Levine et al., 1996) to avoid as much as possible motion. In practice, these sequences are acquired as several stacks of thick slices, often in different orthogonal views (Fig. 1 b) in order to provide the *in-vivo* 3D fetal anatomy in a short acquisition time. This procedure has allowed a more regular use of fetal MR imaging in clinics, with an excellent in-plane spatial resolution (around 1 mm in-plane resolution with a slice thickness between 3 and 5 mm), very good tissue contrast and anatomical detail within the images while reducing motion artifacts (Fig. 1 a).

Although fetal anatomy can be adequately viewed in new multi-slice MR images, many critical limitations remain for quantitative data analysis. Despite these fast acquisition techniques, motion remains and mostly generates inter-slice artifacts (see top row, coronal view, in Fig. 1 b) but also in-plane degradation. This limits the use of computer-assisted methods for large-scale studies. To this end, several research groups have recently developed advanced image processing methods, often denoted by super-resolution (SR) techniques, to reconstruct from several clinical stacks, low-resolution (LR) images, a high-resolution (HR) motion-free volume (Rousseau et al., 2006; Jiang et al., 2007; Gholipour and Warfield, 2009; Gholipour et al., 2010; Kim et al., 2010; Rousseau et al., 2010; Rousseau et al., 2013; Kuklisova-Murgasova et al., 2012; Fogtmann et al., 2012). Such HR volume facilitates early and precise diagnosis (see Fig. 1 c) and offers, to the neuroscientist, the possibility of an automated quantitative study of the

* Corresponding author at: Department of Radiology, CIBM BH 07-081, Rue du Bugnon 46, CH-1011 Lausanne, Switzerland. Tel.: + 41 21 314 74 76.
E-mail address: sebastien.tourbier@unil.ch (S. Tourbier).

http://dx.doi.org/10.1016/j.neuroimage.2015.06.018
1053-8119/

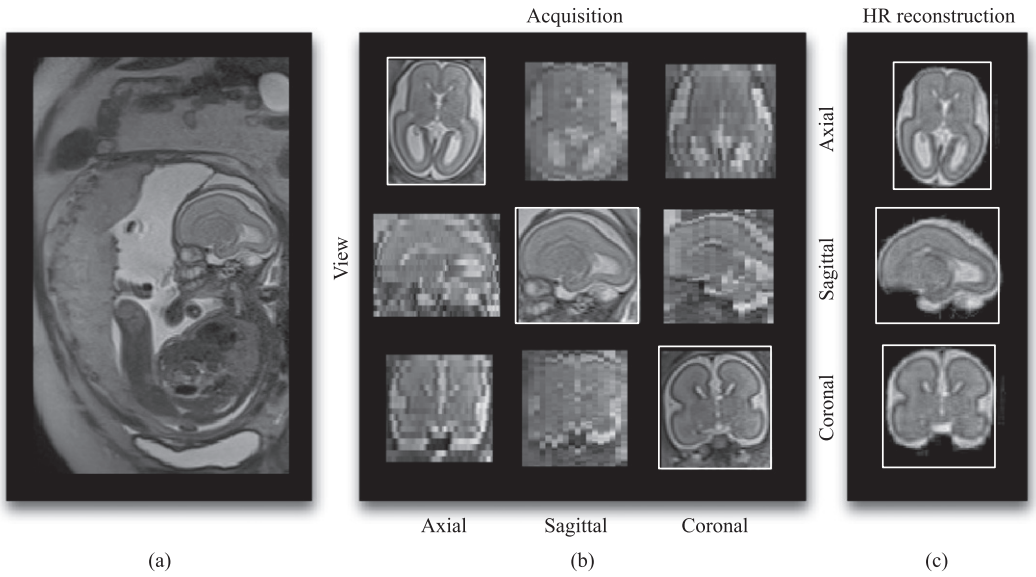

Fig. 1. Fetus of 25 weeks of gestational age: (a) Sagittal view of whole volume scan, (b) orthogonal stacks in brain region. Acquisitions were done with a Siemens Aera 1.5 Tesla, T2w HASTE sequence (TR 1200 ms, TE 90 ms, slice thickness 3.6 mm, in-plane resolution 1.13 mm), and (c) HR image of the fetal brain reconstructed by SR technique.

first stages of brain development (Gholipour et al., 2011; Corbett-Detig et al., 2011; Scott et al., 2011; P. A. Habas et al., 2012).

In the last years, SR algorithms have been successfully applied to MR imaging data (Van Reeth et al., 2012) and their advantage over single HR acquisition has been proven in terms of higher SNR for a given acquisition time (Plenge et al., 2012). SR techniques were initially dedicated to *static* subjects where scanning protocol can be adapted to add new information of the scene under controlled motion conditions. More recently, focus has been applied to SR of *moving* subjects, where motion is not controlled. In these cases, the need of advanced post-processing techniques to enhance spatial resolution is of great importance. Most of these works assumed a rigid motion and they are dedicated to the through-plane resolution improvement of fetal MRI images, providing us a fully isotropic 3D image only limited by the in-plane resolution of the acquisition (Rousseau et al., 2006; Jiang et al., 2007; Gholipour and Warfield, 2009; Gholipour et al., 2010; Kim et al., 2010; Rousseau et al., 2010; Rousseau et al., 2013; Kuklisova-Murgasova et al., 2012; Fogtmann et al., 2012). Thus, these methods consist of two common image processing steps: image registration (for motion compensation) and super-resolution[1] (for image recovery), that solves the so-called inverse problem (see Fig. 2 and Two-step reconstruction algorithm).

An objective of this work is to provide a short review of SR techniques in the case of fetal MRI. Along this line, our contribution differs from the recent review of Van Reeth et al. (2012) that aimed at discussing SR MRI techniques in various contexts. Here we focus on the literature of SR fetal MRI algorithms in order to emphasize the similarities and differences between the successful Bayesian and variational formulations for this specific problem. We built on this overview of SR techniques and recent advances in convex optimization to recall our optimization algorithm introduced in (Tourbier et al., 2014a) that offers simultaneously fast, accurate and robust solutions to the fetal image recovery problem.

Contributions of this work

- A short overview of existing SR reconstruction algorithms for fetal MRI within a unified mathematical formulation.
- An automatic process that estimates a near-visually optimal value for the regularization weight.
- An extensive validation of our SR algorithm introduced in Tourbier et al. (2014a) on simulated fetal data and on a set of ten clinical fetal datasets of normal and pathological brains under various acquisition conditions, including:
  - a developed comparison between our model, the baseline Tikhonov model and the open-source state-of-the-art method (Rousseau et al., 2013) and
  - a study of robustness of regularization terms w.r.t. motion error residuals.

The paper is organized as follows. The 2nd section provides A short review of fetal MRI reconstruction algorithms, starting with the description of the most generic observation model used in this context. The 3rd section Our Total Variation algorithm recalls our TV algorithm presented in (Tourbier et al., 2014a) and it describes the design and implementation of our reconstruction pipeline. The 4th section Results and

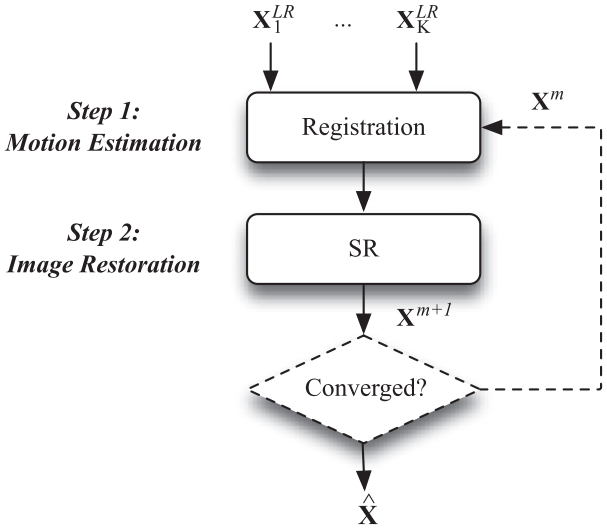

Fig. 2. General paradigm for direct and iterative fetal brain MR image reconstruction. Two iterative steps are considered until convergence: 1) a registration process for motion parameters estimation; 2) a restoration process for HR image recovery.

[1] All through the paper we will denote this step indistinctly by super-resolution, image recovery and restoration. Let us note that this step has to be differentiated from *reconstruction* that, in this work, refers to both motion estimation and image recovery steps.

validation develops an extensive validation of our algorithm, using simulated fetal data and clinically acquired fetal data. Finally, the Conclusion is presented in the 5th section.

A short overview of fetal MRI reconstruction algorithms

Notations and definitions

We consider the following discrete setting throughout the paper. Let $\mathbf{X}$ denotes the high-resolution (HR) image desired, $\mathbf{X}_{kl}^{LR}$ be the l-th observed slice of the k-th LR image, $\mathbf{n}_k$ is the observed noise in the k-th LR image. Let $S_{\mathbf{x}}$, $S_{\mathbf{y}}$ be two finite-dimensional real vector spaces with inner product $\langle \cdot, \cdot \rangle$, norm $\| \cdot \| := \langle \cdot, \cdot \rangle^{1/2}$, $dim S_{\mathbf{x}} = N$ where N is the total number of pixel voxels and $dim S_{\mathbf{y}} = m = N \cdot d$ where $d = \{2, 3\}$ for 2D 3D images. Let HR images $\mathbf{X} \in S_{\mathbf{x}}$ and $\mathbf{DX} \in S_{\mathbf{y}}$ their gradients. The gradient operator $\mathbf{D} : S_{\mathbf{x}} \to S_{\mathbf{y}}$ is a continuous linear operator with norm $\|\mathbf{D}\| := max\{\|\mathbf{DX}\| \mid \mathbf{X} \in S_{\mathbf{x}}, \ \|\mathbf{X}\| \leq 1\}$. Let $\frac{\partial}{\partial \mathbf{x}_i}$ denotes the derivative operator w.r.t. dimension $\mathbf{x}_i$.

A generic observation model

Given the literature, the most accepted image acquisition model in the context of fetal MRI is the linear model:

$$\mathbf{X}_{kl}^{LR} = \mathbf{H}_{kl}\mathbf{X} + \mathbf{n}_k, \tag{1}$$

where $\mathbf{H}_{kl}$ are linear operators that supposedly model the acquisition distortions such as noise, blurring, aliasing, intensity bias and motion (Fig. 3). In this process, the MR scanner acquires several LR stacks of 2D slices, referred as LR images, which are downsampled, degraded, and aliased from the HR original scene. A typical acquisition model accounting for motion, blurring and downsampling processes has been considered in (Rousseau et al. (2006), Jiang et al. (2007), Kim et al. (2008), Kim et al. (2010), Rousseau et al. (2010), Rousseau (2010) and Fogtmann et al. (2012)). In those formulations, the matrix $\mathbf{H}_{kl}$ is decomposed into three matrices accounting for motion $\mathbf{M}_{kl}$, blurring effects $\mathbf{B}_{kl}$, and a basic downsampling aliasing operator $\mathbf{A}_{kl}$, *i.e.*:

$$\mathbf{H}_{kl} = \mathbf{A}_{kl}\mathbf{B}_{kl}\mathbf{M}_{kl}, \tag{2}$$

where $\mathbf{B}_{kl}$ is the 3D Point-Spread-Function (PSF) of the system and $\mathbf{M}_{kl}$ is assumed to be 6-DOF 3D rigid motion. Typically, a good approximation of the PSF is a 3D Gaussian function with the full width at half maximum (FWHM) equal to the slice-thickness in the slice-select direction and 1.2× voxel size in-plane (Jiang et al., 2007; Kuklisova-Murgasova et al., 2012). In Gholipour and Warfield (2009) and Gholipour et al. (2010)), the slice acquisition process is fully modeled by decomposing $\mathbf{B}_{kl}$ into two matrices, that simulate (1) the in-plane 2D PSF ($\mathbf{P}_{kl}$) and (2) the slice profile ($\mathbf{S}_{kl}$):

$$\mathbf{H}_{kl} = \mathbf{A}_{kl}\mathbf{P}_{kl}\mathbf{S}_{kl}\mathbf{M}_{kl}, \tag{3}$$

where $\mathbf{P}_{kl}$ is a 2D Gaussian function with FWHM equal to 1.2× voxel size in-plane, $\mathbf{S}_{kl}$ corresponds to a 3D rotation operator that defines the orientation of the slice plane. In Kuklisova-Murgasova et al., 2012, the authors proposed to unify SR, regularization, robust statistics and intensity matching within a common expectation maximization framework. The model (1) is modified to consider intensity inhomogeneities (bias field and slice-dependent scaling factors) as follows:

$$\mathbf{X}_{kl}^{LR*} = \mathbf{H}_{kl}\mathbf{X} + \mathbf{n}_k, \ \mathbf{X}_{kli}^{LR*} = s_{kl} exp(-\mathbf{I}_{kli})\mathbf{X}_{kli}^{LR}, \tag{4}$$

where $\mathbf{X}_{kl}^{LR\,*}$ represents the l-th scaled, denoised and bias-corrected slice of the k-th LR image, $\mathbf{X}_{kli}^{LR\,*}$ is the intensity of the i-th voxel of $\mathbf{X}_{kl}^{LR\,*}$, $\mathbf{X}_{kli}^{LR}$ is the intensity of the i-th voxel of the l-th slice of the k-th original LR image, s_{kl} is the corresponding intensity scaling factor, and $\mathbf{I}_{kli}$ the corresponding bias field factor.

Finding the original HR image in Eqs. (1) and (4) that generated the LR images $\mathbf{X}^{LR}$ under the MRI acquisition is carried out with an inverse problem strategy, where Eqs. (1) and (4) are known to be the forward models. This means that the given measures $\mathbf{X}^{LR}$ are used to generate the original unknown HR image $\mathbf{X}$. The most natural way to combine the LR images is through a standard least-square problem:

$$\min_{\mathbf{X}} \sum_{kl} \left\| \mathbf{H}_{kl}\mathbf{X} - \mathbf{X}_{kl}^{LR} \right\|^2. \tag{5}$$

However, due to the presence of noise and insufficient number of acquired LR images, the above inverse problem is said to be ill-posed, meaning that it has generally no meaningful solutions. A natural solution of this issue is to add priors about $\mathbf{X}$, *i.e.* make use of known properties that $\mathbf{X}$ holds such as intensity smoothness.

Two-step reconstruction algorithm

Let us now introduce the most generic algorithm for fetal brain MRI reconstruction (Rousseau et al., 2006; Jiang et al., 2007; Gholipour and Warfield, 2009; Gholipour et al., 2010; Kim et al., 2010; Rousseau et al., 2010; Rousseau et al., 2013; Kuklisova-Murgasova et al., 2012; Fogtmann et al., 2012). The algorithm consists in iterating the following two steps until convergence:

Step 1: Motion estimation

$$\mathbf{M}_{kl}^{m+1} = arg \min_{\mathbf{M}_{kl}} \sum_{kl} \left\| \mathbf{A}_{kl}\mathbf{B}_{kl}\mathbf{M}_{kl}\mathbf{X}^{m} - \mathbf{X}_{kl}^{LR} \right\|_{M}, \tag{6}$$

where $\|.\|_M$ is a distance function that can be based on l_2-distance (Gholipour and Warfield, 2009; Gholipour et al., 2010; Kim et al., 2010), cross-correlation (Jiang et al., 2007), and NMI (Rousseau et al., 2006; Rousseau et al., 2010; Rousseau et al., 2013; Kuklisova-Murgasova et al., 2012).

Step 2: Image restoration

$$\mathbf{X}^{m+1} = arg \min_{\mathbf{X}} \ \frac{\lambda}{2} \sum_{kl} \left\| \mathbf{A}_{kl}\mathbf{B}_{kl}\mathbf{M}_{kl}^{m+1}\mathbf{X} - \mathbf{X}_{kl}^{LR} \right\|^2 + R(\mathbf{X}),$$

where $R(\mathbf{X})$ is a prior introduced to regularize the solution, the first term relates to data fidelity. The parameter λ balances the trade-off between $R(\mathbf{X})$ and data fidelity. The general reconstruction scheme is illustrated in Fig. 2. Our present work focuses on the image restoration step. It is beyond the scope of this paper to also tackle the motion compensation problem but we briefly summarize in the next section what the fetal MRI literature offered to solve this problem. We refer to (Malamateniou et al., 2013) for a detailed discussion about motion correction strategies in fetal applications.

Motion estimation algorithms

The motion estimation problem is addressed through an image registration task that aims at compensating mainly for motion occurring between slices of the LR images, typically generated by short and fast movements of the fetal head. All existing methodologies (Rousseau et al., 2006; Jiang et al., 2007; Gholipour and Warfield, 2009; Gholipour et al., 2010; Kim et al., 2010; Rousseau et al., 2010; Rousseau et al., 2013; Kuklisova-Murgasova et al., 2012; Fogtmann et al., 2012) have addressed the problem using voxel-based registration methods where fetal motion is modeled as a full 6 degree-of-freedom rigid 3D transformation (3 translations and 3 rotations). Typically,

they consist in (1) globally co-registering the LR images (volume-to-volume registration) and (2) hierarchically aligning every slice of the LR image to the reconstructed HR image (slice-to-volume registration), that is built using the current estimate of slice positions.

In practice, most registration methods differ only as regards their choice of similarity metric, the distance $\|.\|_M$ in Eq. (6), and the corresponding optimization algorithm used to estimate the spatial transformation. When the assumptions of similar levels of blur and identical contrasts between LR volumes are made, the mean square intensity difference (MSD) is adopted (Gholipour and Warfield, 2009; Gholipour et al., 2010). In order to force similar intensity values in the images being registered, a weighted MSD (Kim et al., 2010) and cross-correlation (CC) (Jiang et al., 2007) measures have also been proposed. To not be based on strong assumptions between the two images and to be less sensitive to changes in overlap, the Normalized Mutual Information (NMI) as cost function has also been used (Rousseau et al., 2006; Rousseau et al., 2010; Rousseau et al., 2013; Kuklisova-Murgasova et al., 2012). In all methods, a classical hierarchical strategy is applied in order to incorporate the interleaved aspect of acquisition and to reduce the risk to fall into local minima. Such hierarchical implementation involves decreasing step sizes of the optimization algorithm at the different stages of registration.

Rather than single slice-to-volume registration, the Slice Intersection Motion Correction (SIMC) (Kim et al., 2010) aims at solving a slice motion correction registration by seeking the collective alignment of all slices simultaneously, and considering the matching structure along all intersecting slice pairs.

While all approaches deal with the motion compensation as an independent step, a unified formulation for motion estimation and SR reconstruction was recently presented in Fogtmann et al. (2012). Their major contribution is a regularization term for the motion estimation that encodes the high correlation between temporal neighboring slices, increasing the robustness when large motion occurs.

Despite these motion compensation efforts, the HR image may still be corrupted by remaining registration errors and other types of motion artifacts not well-modeled by a 6 DOF transformation.[2] To treat these cases, outlier detection and rejection schemes are used, either implemented as a separate module preceding the SR (Rousseau et al., 2006) (Kim et al., 2010), or incorporated in the energy formulation of the restoration algorithm (Gholipour et al., 2010; Kuklisova-Murgasova et al., 2012).

Fetal MRI restoration algorithms

Original restoration algorithms were designed in the frequency domain, using the shifting and aliasing properties of the Fourier Transform in order to increase the image resolution. But these frequency approaches have shown to have an observation model limited to translation only. Nowadays, researchers in fetal MRI address this problem in the spatial domain, which allow models taking into rotation (Yang and Huang, 2011). Besides, desired edge-preserving properties are easier to deal with in the spatial domain.

Interpolation models

Pioneer works on fetal HR restoration are based on the interpolation SR model (Rousseau et al., 2006; Jiang et al., 2007; Kim et al., 2008; Kim et al., 2010). They consist in performing an iterative reconstruction procedure, interleaving rigid co-registration of the LR images and Scattered Data Interpolation (SDI) (interpolation onto an HR grid) steps.

[2] For instance, when fetus is displaced from the region being imaged and, due to the loss of signal, LR images have dark slices.

Based on the acquisition model Eq. (2), the development of such approaches have been motivated by assuming that $\mathbf{B}_{kl}$ is linearly spatial invariant and is the same for all slices and $\mathbf{M}_{kl}$ considers only rigid 6 DOF motions (translations + rotations). Thus, $\mathbf{B}_{kl}$ and $\mathbf{M}_{kl}$ can commute in Eq. (2), and after injection in Eq. (1), the acquisition model becomes:

$$\mathbf{X}_{kl}^{LR} = \mathbf{A}_{kl}\mathbf{M}_{kl}\underbrace{\mathbf{B}_{kl}\mathbf{X} + \mathbf{n}_k}_{\mathbf{Z}} = \mathbf{A}_{kl}\mathbf{M}_{kl}\mathbf{Z}. \tag{7}$$

A composite image on non-uniformly spaced sampling points is first obtained with the registration of LR images. Then, nonuniform interpolation is performed to get $\mathbf{Z}$, an image with uniformly spaced sampling points.

The first interpolation method that tackles the SR problem in fetal MRI was presented in (Rousseau et al., 2006). The method used an interpolation SR approach that comprises slice-to-volume registration interleaved with Scattered Data Interpolation (SDI). Once motion parameters are estimated, a bias field correction step is performed to correct the local relative intensity distortion between LR images. Finally, a computationally efficient local neighborhood Gaussian kernel SDI is employed for reconstruction but this reduces spatial frequency content and thus results in excessive blurring. In Jiang et al. (2007), SDI is improved with the use of cubic B-spline kernel SDI.

Introduction of spatial priors has been later adopted. They are reviewed in the next section.

Bayesian and Variational models

Two main categories of algorithms that introduce spatial regularizations in the restoration task can be distinguished: (1) Deterministic variational models and (2) stochastic Bayesian models. Typically, stochastic approaches rely on the formulation of the problem in a Bayesian framework such as *Maximum a Posteriori* (MAP), while deterministic approaches are based on the formulation of Constrained Least Squares (CLS) energy with functional space-based regularization. In both approaches, the HR image and motion among LR images can be considered as either stochastic or deterministic variables, relating the SR reconstruction steps stochastically or deterministically toward an optimal reconstruction.

Let us express the SR reconstruction problem into a full Bayesian formulation. Let $\mathbf{H}_{kl} = \mathbf{A}_{kl}\mathbf{B}_{kl}\mathbf{M}_{kl}$ be the matrix modeling the MRI acquisition. Let suppose $\mathbf{n}_k$ to be Gaussian. Then, SR reconstruction can be formulated as:

$$\hat{\mathbf{X}} = \arg\max_{\mathbf{X}} \prod_{kl} \Pr\left(\mathbf{X}|\mathbf{X}_{kl}^{LR}\right) \tag{8}$$

where $\Pr(\mathbf{X}|\mathbf{X}_{kl}^{LR})$ refers to the posterior probability. In fetal MRI, $\mathbf{X}$ and $\mathbf{H}_{kl}$ are assumed to be statistically independent and, if we suppose that $\mathbf{H}_{kl}$ is estimated beforehand, denoted as $\hat{\mathbf{H}}_{kl}$, Eq. (8) can be reformulated (using the Bayes rule) as:

$$\hat{\mathbf{X}} = \arg\max_{\mathbf{X}} \prod_{kl} \Pr\left(\mathbf{X}_{kl}^{LR}|\mathbf{X}, \hat{\mathbf{H}}_{kl}\right) \Pr(\mathbf{X}). \tag{9}$$

where probability $\Pr(\mathbf{X}_{kl}^{LR}|\mathbf{X}, \hat{\mathbf{H}}_{kl})$ corresponds to the data likelihood, $\Pr(\mathbf{X})$is the prior probability on the HR image desired. Eq. (9) describes the popular stochastic MAP formulation of SR. When no prior distribution over the image is incorporated, the *Maximum Likelihood* (ML) formulation could be obtained from Eq. (9). Such ML formulation has been adopted in the pioneer work presented in (Gholipour and Warfield, 2009; Gholipour et al., 2010), which aims at providing an optimum solution through maximizing the conditional probability density function $\Pr(\mathbf{X}_{kl}^{LR}|\mathbf{X})$ of the acquired slice $\mathbf{X}_{kl}^{LR}$ given the reconstructed volume $\mathbf{X}$. Standard stochastic optimization techniques, such as Monte Carlo (Metropolis et al., 1953; Hastings, 1970), Simulated annealing (Kirkpatrick et al., 1983), Iterated Conditional Modes (Besag, 1986) or

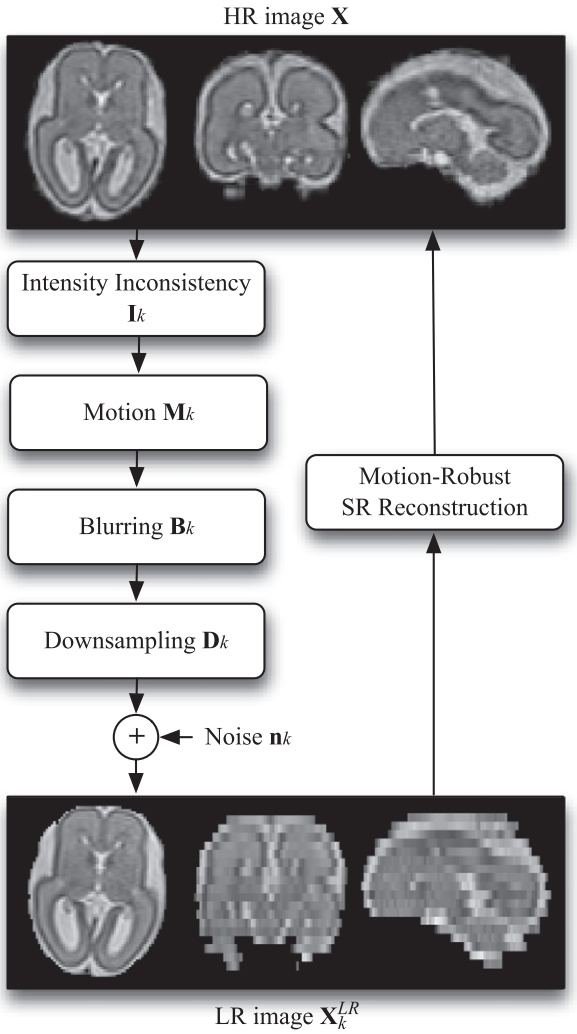

Fig. 3. The observation model of a real MR imaging system relating the low-resolution (LR) image $\mathbf{X}_k^{LR}$ (observed images) to an high-resolution (HR) image $\mathbf{X}$. The SR reconstruction step corresponds to the inverse problem, *i.e.*, finding $\mathbf{X}$ given the observations $\mathbf{X}_k^{LR}$.

stochastic Partial Differential Equations (Zabaras, 2010) can be used to find a solution to (9). However, using the assumption of independent slice acquisition, the stochastic formulation (9) can be cast into a deterministic optimization framework. If we suppose that $\Pr\left(\mathbf{X}_{kl}^{LR}|\mathbf{X},\hat{\mathbf{H}}_{kl}\right) = e^{-\lambda\left\|\hat{\mathbf{H}}_{kl}\mathbf{X}-\mathbf{X}_{kl}^{LR}\right\|^2}$, $\Pr(\mathbf{X}) = e^{-R(\mathbf{X})}$, when the noise residuals are presumed to be drawn from a Gaussian distribution, and using the logarithmic function, then (9) becomes

$$\begin{aligned}\hat{\mathbf{X}} &= -log\left[\arg\max_{\mathbf{X}}\prod_{kl}\Pr\left(\mathbf{X}_{kl}^{LR}|\mathbf{X},\hat{\mathbf{H}}_{kl}\right)\Pr(\mathbf{X})\right)\\ &= \arg\min_{\mathbf{X}}\sum_{kl}-log\left[\Pr\left(\mathbf{X}_{kl}^{LR}|\mathbf{X},\hat{\mathbf{H}}_{kl}\right)\right]-log[\Pr(\mathbf{X})]\\ &= \arg\min_{\mathbf{X}}\lambda\sum_{kl}\left\|\mathbf{H}_{kl}\mathbf{X}-\mathbf{X}_{kl}^{LR}\right\|^2+R(\mathbf{X}),\end{aligned} \tag{10}$$

where R($\boldsymbol{X}$) corresponds to the regularization term and λ controls the trade-off between regularization strength and data fidelity. As a result, it is equivalent to solve (9) and (10) under the slice independence hypothesis. Several proposed methods (Gholipour and Warfield, 2009; Gholipour et al., 2010; Rousseau et al., 2010; Rousseau et al., 2013; Kuklisova-Murgasova et al., 2012; Fogtmann et al., 2012) used this equivalence to formulate the problem in a Bayesian framework and compute a solution with variational optimization techniques. However, non-Gaussian noise and outliers might be present in the data due to possible intensity inhomogeneities and inaccurate slice motion estimation. In this case, the ℓ_2-norm is not robust and it has led to the development of modified versions of Eq. (10). They will be reviewed in the next section.

Variational terms

A more general formulation of Eq. (10) is:

$$\hat{\mathbf{X}} = \arg\min_{\mathbf{X}}\lambda F(\mathbf{X}) + R(\mathbf{X}), \tag{11}$$

where F($\boldsymbol{X}$) is a function that represents the data fidelity term. The choice of appropriate terms as well as an adequate optimization scheme is crucial for high reconstruction quality.

Data fidelity terms. In fetal MRI acquisition, motion estimation errors and intensity inhomogeneities usually occur. This is critical as these errors influence the quality of the restored image. Consequently, several error norms for the data fidelity term have been considered for dealing with outliers.

As noise might not be Gaussian, they have considered in (Gholipour and Warfield (2009) and Gholipour et al. (2010)) a modified error weight function $\mathbf{\Omega}^k$ (based on the Huber error function) to take into consideration more generic outliers such that:

$$F(\mathbf{X}) = \left\|\mathbf{\Omega}^k\left[\mathbf{H}_{kl}\mathbf{X}-\mathbf{X}_{kl}^{LR}\right]\right\|^2. \tag{12}$$

Such a formulation has shown to be a good balance between the most precise estimation in a Gaussian environment (ℓ_2-norm) and the most robust estimation in a non-Gaussian environment.

Alternatively in Kuklisova-Murgasova et al., 2012 the authors proposed to simultaneously maximize the fit between the estimated reconstruction, the estimated intensity inhomogeneities and the data acquisition s.t. $\hat{\mathbf{X}} = min_{\mathbf{X}}min_{\mathbf{I}_{kl}}F(\mathbf{X},\mathbf{I}_{kl})$, where I_{kl} is the underlying intensity inhomogeneity. They introduced probability density function in order to model outliers, where they consider the inlier class posterior probability $\mathbf{p}_{kl}^{slice}$ of the *l*-th slice of *k*-th LR image and the posterior probability $\mathbf{p}_{klj}^{voxel}$ of a voxel *j* being classified as an inlier. This corresponds to a hybrid approach, where probabilities are integrated in a deterministic energy minimization problem as weight functions:

$$F(\mathbf{X}) = \sum_{klj}\mathbf{p}_{kl}^{slice}\mathbf{p}_{klj}^{voxel}\mathbf{e}_{klj}^2, \tag{13}$$

where $\mathbf{e}_{klj} = \mathbf{H}_{klj}\hat{\mathbf{X}}_j - \mathbf{X}_{klj}^{LR*}$ is the estimation error of the *j*-th voxel of the *l*-th slice of the *k*-th intensity-corrected LR image. It has been demonstrated that this method performs better than the method with Huber statistics and exclusion of either intensity matching or robust statistics results in drop of performance compared to their full method.

Finally, the authors in Fogtmann et al. (2012) suggested reducing outliers by optimizing *F* with respect to both the underlying image $\mathbf{X}$ and the underlying slice motions M_{kl}. They also propose a unified formulation that simultaneously maximizes the fit between the estimated reconstruction, the estimated motion and the data acquisition data such that:

$$\hat{\mathbf{X}} = \arg\min_{\mathbf{X}}\min_{\mathbf{M}_{kl}}F(\mathbf{X},\mathbf{M}_{kl}). \tag{14}$$

This method has been compared against different formulations using different estimators and penalty terms for registration as well as different approaches such as the SIMC. On simulated data, the method designed with a ℓ_2 estimator and a Huber penalty has proved to provide the best overall performance.

Regularization terms. Due to the ill-posedness of the restoration problem it is essential to introduce regularization terms to constraint the solution to hold prior knowledge of the desired image. If the regularization term is not appropriate, SR could result in too blurry images or artifacts can be produced. Several local regularization terms have been considered in the literature s.a. Tikhonov regularization (Gholipour et al., 2010; Fogtmann et al., 2012), Total-Variation based regularization (Rousseau et al., 2010; Kuklisova-Murgasova et al., 2012), and non-local regularization terms such as non-local means (Rousseau et al., 2013). In general, TV energies can been seen as a measure of signal variability that penalizes only the total amount of gradient in the image, preserving edges during reconstruction, as opposed to Tikhonov which penalizes its distribution. For these reasons, they have been more widely adopted by the community.

Optimization via gradient descent

Let $F(\mathbf{X})$ be the simple l_2 error norm. The fetal brain reconstruction problem can be formulated as a first order convex optimization problem such that:

$$\min_{\mathbf{X}\in S_{\mathbf{X}}} \frac{\lambda}{2}\sum_{kl}\left\|\mathbf{H}_{kl}\mathbf{X}-\mathbf{X}_{kl}^{LR}\right\|^2 + \|\mathbf{X}\|_{TV} s.t. \mathbf{X}\geq 0 \tag{15}$$

where $\|\mathbf{X}\|_{TV} := \|\mathbf{DX}\|_1$ is the TV semi-norm (Rudin et al., 1992). Problem (15) is a convex but also non-smooth optimization problem. Non-smooth minimization has been challenging since the introduction of TV in imaging problems in the 90s. The most common methodology that optimizes TV consists in using a smooth ε-regularization of TV such that:

$$\|\mathbf{X}\|_{TV_\varepsilon} := \sum_{i,j}\sqrt{|\mathbf{DX}_{i,j}|^2+\varepsilon} \;\rightarrow\; \|\mathbf{X}\|_{TV} \;\; for\; \varepsilon\rightarrow 0. \tag{16}$$

The natural main motivation to introduce a quadratic approximation of TV is to apply smooth optimization tools like gradient descent or conjugate gradient schemes. In the context of fetal reconstruction, this approach was considered with $\varepsilon = 1$ in Kuklisova-Murgasova et al., 2012; Rousseau et al. (2010) based on the edge-preserving technique introduced in Charbonnier et al. (1997). In Gholipour et al. (2010), the authors considered large ε that links to the Dirichlet Laplacian energy, *i.e.* $\|\mathbf{X}\|_{TV_\varepsilon}\rightarrow\|\mathbf{X}\|_{Dir} = \|\mathbf{DX}\|^2$ for $\varepsilon\rightarrow\infty$. Standard optimization techniques for ε-TV have been carried out with the calculus of variation followed by a forward explicit steepest gradient descent scheme, such as in Kuklisova-Murgasova et al., 2012 and Rousseau et al. (2010), or followed by a quasi-Newton optimization scheme, such as in Fogtmann et al. (2012). The non-local reconstruction approach suggested in Rousseau et al. (2013), is solved by iteratively performing a basic image restoration algorithm, based on the Dirichlet energy, and a NLM denoising algorithm, as in Manjón et al. (2010). The energy is optimized using the Fletcher–Reeves form of conjugate gradient descent algorithm.

TV energies have been largely adopted in the literature because of their powerful ability in edge preservation. However, only standard steepest gradient techniques have been applied to optimize fetal-based TV energies. But, such standard PDE-based optimization schemes may be slow because they are restricted by the Courant–Friedrichs–Lewy's condition (Courant et al., 1967), that basically sets an upper bound on the time step of the iterative flow s.t. $\Delta t \leq \varepsilon\Delta x/max_{i,j}EL_{i,j}$, where *EL* corresponds to the Euler–Lagrange equations (speed of the flow) and Δx is the spatial step. It is known that the asymptotic rate of convergence of forward steepest descent techniques for smooth energy is $O(1/n)$ where n is the number of iteration steps and the iterative rate is $O(1/\varepsilon)$, see *e.g.* (Beck and Teboulle, 2009). Besides the speed, there are two other limitations with ε-TV: (1) the TV term is not exactly solved because we compute an ε-solution, which does not preserve as well image contrasts and small-scale structures as the exact TV, and (2) there is an extra parameter, ε, to select (unlike exact TV).

Our total variation algorithm

Algorithm description

In the last years, fast TV-based algorithms based on convex optimization theory have been developed to solve sparse reconstruction problems such as Compressed Sensing. Major classes of TV optimization methods are (1) Alternating Direction Method of Multipliers (ADMM) (Glowinski and Tallec, 1989; Goldstein and Osher, 2009), (2) - Forward–Backward algorithms (Nesterov, 2005; Combettes and Wajs, 2006; Combettes and Pesquet, 2011; Beck and Teboulle, 2009), and (3) Uzawa-based Primal-Dual methods (Arrow et al., 1958; Zhu and Chan 2008; Chambolle and Pock, 2011a). For the sake of clarity, we recall in this section our contribution (Tourbier et al., 2014a) where we reformulated Eq. (15) by using an accelerated primal-dual hybrid gradient (PDHG) method based on (Chambolle and Pock, 2011b) to design a fast, robust algorithm that offers accurate solutions and is guaranteed to converge to a global solution for Eq. (15) (*i.e.* solution is independent of the initialization). Specifically, PDHG consists in introducing the dual variable $\mathbf{P}$ in Eq. (15) that splits our initial complex problem on the primal variable $\mathbf{X}$ into two simpler problems easy to solve. It corresponds to rewrite Eq. (15) as the saddle point problem or equivalently as a primal-dual optimization problem:

$$\min_{\mathbf{X}\in S_X}\underbrace{\max_{\mathbf{P}\in S_Y}\langle\mathbf{DX},\mathbf{P}\rangle - F^*(\mathbf{P})}_{\|\mathbf{X}\|_{TV}} + G(\mathbf{X}) + \underbrace{\delta_C(\mathbf{X})}_{\mathbf{X}\geq 0} \tag{17}$$

where $\mathbf{X}$ corresponds to the original primal variable and $\mathbf{P}$ corresponds to the dual variable introduced, $G(\mathbf{X}) = \frac{\lambda}{2}\sum_{k=1}^{K}\left\|\mathbf{H}_k\mathbf{X}-\mathbf{X}_k^{LR}\right\|^2$, the convex function $F^\star$ denotes the barrier function of the ℓ_∞ unit ball, that is $F^\star(\mathbf{P}) = 0$ *if* $|\mathbf{P}_i|\leq 1$ *for* $1\leq i\leq n$, otherwise $F^\star(\mathbf{P}) = +\infty$ and $\delta_C(\mathbf{X})$ is a barrier function of the convex set $C := \{\mathbf{X}\geq 0\}$. As G is uniformly convex, we may therefore apply (Chambolle and Pock, 2011b) to solve Eq. (17). The proposed algorithm consists in iterating

$$\mathbf{P}^{n+1} = prox_{\sigma^n F^\star}\left(\mathbf{P}^n + \sigma^n\mathbf{D}\overline{\mathbf{X}}^n\right) \tag{18}$$

$$\mathbf{X}^{n+1} = prox_{\tau^n G+\delta_C}\left(\mathbf{X}^n - \tau^n\mathbf{D}^t\mathbf{P}^{n+1}\right) \tag{19}$$

$$\begin{aligned}\theta^{n+1} &= 1/\sqrt{1+2\rho\tau^n}, \quad \tau^{n+1} = \theta^{n+1}\tau^n,\\ \sigma^{n+1} &= \sigma^n/\theta^{n+1}\end{aligned} \tag{20}$$

$$\overline{\mathbf{X}}^{n+1} = \mathbf{X}^{n+1} + \theta^{n+1}\left(\mathbf{X}^{n+1}-\mathbf{X}^n\right) \tag{21}$$

where $prox_E$ the proximal operator of E defined as $prox_E(\mathbf{X}) := arg\,min_{\mathbf{Y}}\; E(\mathbf{Y}) + \frac{1}{2}\|\mathbf{Y}-\mathbf{X}\|^2$. Solution of the inner problem Eq. (18) is given by $\left(prox_{\sigma^n F^\star}(\mathbf{Z})\right)_i = \mathbf{Z}_i/max\{1,|\mathbf{Z}_i|\}$ where $\mathbf{Z} = \mathbf{P}^n + \sigma^n\mathbf{D}\overline{\mathbf{X}}^n$. The solution of the least-square problem Eq. (19), $min_{\mathbf{X}\geq 0}\;\frac{\lambda}{2}\sum_{k=1}^{K}\left\|\mathbf{H}_k\mathbf{X}-\mathbf{X}_k^{LR}\right\|^2 + \frac{1}{2\tau^n}\|\mathbf{X}-\mathbf{W}\|^2$ with $\mathbf{W} = \mathbf{X}^n - \tau^n\mathbf{D}^t\mathbf{P}^{n+1}$ can be computed with several approaches. We use a semi-implicit gradient descent scheme that provides fast good approximate minimizing solutions. More specifically, the Euler–Lagrange solution of Eq. (19) is $\lambda\tau^n(\mathbf{HX}-\mathbf{X}^{LR}) + \mathbf{X} - \mathbf{W} = 0$ where $\mathbf{H} := \sum_{k=1}^{K}\mathbf{H}_k^t\mathbf{H}_k$ and $\mathbf{X}^{LR} := \sum_{k=1}^{K}\mathbf{H}_k^t\mathbf{X}_k^{LR}$ (note that $\mathbf{H}$ and $\mathbf{X}^{LR}$ are computed only once) and the iterative semi-implicit scheme is defined as $\mathbf{X}^{l+1} = \mathcal{P}_C\left(\mathbf{X}^l - \Delta t\lambda\tau^n\left(\mathbf{HX}^l-\mathbf{X}^{LR}\right) + \Delta t\mathbf{W}\right)/(1+\Delta t)$ where $\mathcal{P}_C$ is the projection operator onto the set $\mathbf{X}\geq 0$ and $\Delta t = 0.1$ in all experiments.

Algorithm defined by Eqs. (18)–(21) is guaranteed to converge to a saddle point ($\mathbf{X}^\star$, $\mathbf{P}^\star$) to Eq. (17) (where $\mathbf{P}^\star$ is the maximizer of the dual problem and $\mathbf{X}^\star$ is the minimizer of the primal problem) and therefore a solution $\mathbf{X}^\star$ of the fetal reconstruction problem Eq. (15) as long as the initial time steps are chosen to be $\sigma^0\tau^0 \leq 1/\|\mathbf{D}\|^2$, see (Chambolle and Pock, 2011b) for more details.

We observe that TV-based models have been proposed in the literature for the fetal MRI reconstruction problems such as Charbonnier et al. (1997), Rousseau et al. (2010) and Kuklisova-Murgasova et al., 2012. However, these works have considered a smooth approximation of the TV semi-norm, that is Eq. (16). Although our technique and the ones in (Charbonnier et al., 1997; Rousseau et al., 2010; Kuklisova-Murgasova et al., 2012) target at optimizing the TV energy, there is a subtle but essential difference between our approach and the other published TV-based techniques. If one wants the exact TV for best image edge recovery in ε-TV approaches, then taking ε as small as possible is required, but the smaller ε, the slower the convergence to the steady state solution (and $\varepsilon = 0$ can never be chosen). More importantly, our approach solves the *exact* TV with a new optimization scheme that is optimal for this class of problems as proved by Nesterov in (Nesterov, 2005), as the asymptotic speed of convergence of the algorithm is quadratic, *i.e.* $O(\frac{1}{n^2})$, while optimization techniques based on ε-TV are restricted to $O(\frac{1}{n})$. We will carry out a numerical experiment in section Numerical comparison of optimization schemes: ε-TV v.s. exact TV to validate these asymptotic speeds of convergence.

Adaptive regularization parameter λ

Determination of an appropriate weight λ that controls the strength of the regularization terms is the key for successful regularized image reconstruction. Intuitively, the amount of regularization would depend on the image resolution and the number of LR images. In this sense, the more the number of available LR images is increased, the less ill-posed the problem will be, and thus, a lower level of regularization will be required. To our knowledge, existing SR techniques in fetal MRI set the amount of regularization *arbitrarily* based on visual perception (Gholipour and Warfield, 2009; Gholipour et al., 2010; Rousseau et al., 2010; Rousseau et al., 2013; Kuklisova-Murgasova et al., 2012; Fogtmann et al., 2012). However, many regularization parameter choice methods have been proposed in image restoration and reconstruction problems (Galatsanos and Katsaggelos, 1992; Karl, 2005). In general, they can be classified into two major categories. Methods in the first category s.a. those based the discrepancy principle (Galatsanos and Katsaggelos, 1992; Karl, 2005) seek to estimate the optimal value based on *a priori* knowledge of the image and/or statistics of the noise while methods in the second category s.a. those based on *L*-curve (Hansen and O'Leary, 1993), generalized cross-validation (GCV) (Mersereau and Reeves, 1990) and estimation of mean squared error (MSE) (Stein, 1981; Ramani et al., 2012) search to estimate the optimal value directly from the data available. We refer to (Galatsanos and Katsaggelos, 1992; Karl, 2005) for more details about regularization parameter choice. In this section, we propose a first attempt for adaptive regularization parameter setting in fetal MRI SR reconstruction. Two data-driven strategies are presented here to perform a fair comparison between different regularization terms carried out in section Results and validation.

In fetal MRI, quantitative evaluation of the quality of the reconstruction is challenging as *a priori* knowledge about the ground-truth HR image is not known. To overcome this limitation, previous works on fetal brain SR reconstruction adopted two different approaches. On one side, the authors in (Gholipour et al. (2010) and Rousseau et al. (2010) and Kuklisova-Murgasova et al., 2012) simulated fetal data from a known HR image to evaluate their algorithms under controlled conditions. On the other side, the authors in Kuklisova-Murgasova et al., 2012 suggested to evaluate their algorithms on clinical fetal data by performing a leave-one-out analysis. For all experiments, a standard evaluation of the quality of the reconstructed images in fetal MRI was to use either the normalized root mean squared error (*NRMSE*) (Gholipour and Warfield, 2009; Kuklisova-Murgasova et al., 2012; Tourbier et al., 2014a) or the peak signal-to-noise ratio (*PSNR*) (Rousseau et al., 2010; Gholipour et al., 2010; Kuklisova-Murgasova et al., 2012), related through $PSNR = 10 * log(1/NRMSE^2)$. Therefore, we have developed two strategies, as regards as the kind of experiments, where the PSNR is considered as the criterion for determining the optimal regularization weight, namely λ_{PSNR}, for comparison between different regularization terms.

For simulated experiments in section Quantitative analysis of simulated fetal images, we propose a strategy which automatically sets the optimal value of the regularization parameter using *the ground-truth HR image* $\mathbf{X}^{GT}$. Specifically, for each regularization term, we select λ that provides the best reconstruction quality in terms of the highest PSNR with respect to $\mathbf{X}^{GT}$:

$$\lambda_{PSNR} = arg\ \max_{\lambda} PSNR\left(\hat{\mathbf{X}}, \mathbf{X}^{GT}\right) \tag{22}$$

For *in-vivo* fetal experiments in section Quantitative analysis of real fetal images, as we do *not* know the ground truth HR image, we propose a strategy which automatically sets the optimal regularization weight using *the set of available ground-truth LR images*. Specifically, for each patient and for each regularization term, we will select λ that provides the best reconstruction quality in terms of the highest mean PSNR over the set of acquired LR images:

$$\lambda_{PSNR} = arg\ \max_{\lambda} \frac{1}{N}\sum_{k} PSNR\left(\hat{\mathbf{X}}_k^{LR}, \mathbf{X}_k^{LR}\right) \tag{23}$$

where $\hat{\mathbf{X}}_k^{LR} = \mathbf{H}\hat{\mathbf{X}}_k$ and $\hat{\mathbf{X}}_k$ corresponds to the HR image reconstructed in a leave-one-out fashion by excluding $\mathbf{X}_k^{LR}$ from the reconstruction. For both strategies, an exhaustive search is performed. To validate the quality of the estimate, we present in section Perceptual evaluation of the adaptive regularization a perceptual evaluation of our adaptive regularization parameter in clinical practice using a multiple-alternative forced-choice approach, as in where a radiologist expert was required to choose one image as his preference from a set a reconstructed images with different levels of regularization in a range around λ_{PSNR}.

Practical implementation

Data preprocessing

A first manual reorientation step is performed with Slicer (Pieper et al., 2004) (Step 1 in Fig. 4). Then, we perform semi-manual brain masking (aided by region growing segmentation algorithm) to ensure good results of the subsequent image processing steps, as maternal tissue surrounding the brain may changes and could consequently corrupt them (Step 2 in Fig. 4). Note that few works have addressed this problem automatically in the last past years, either through template-based segmentation (Anquez et al., 2009; Taleb et al., 2013; Tourbier et al., 2014b) or through machine learning (Ison et al., 2012; Keraudren et al., 2014). Finally, we standardize intensities of the LR stacks through *N*4 bias field correction, intensity scale standardization (Nyúl et al., 2000), and rescaling intensities into [0,255] (Step 3 in Fig. 4).

Motion estimation

The estimation of motion parameters is done with BTK. Firstly, global stack registration is used for initialization of the transforms. Then, rigid 6 *degrees-of-freedom* slice-to-volume registration is performed where normalized correlation is used as metric for optimization. Prior to motion estimation, LR images are filtered by NLM denoising. This reduces the chance of the registration process to fall into local minimas. Finally,

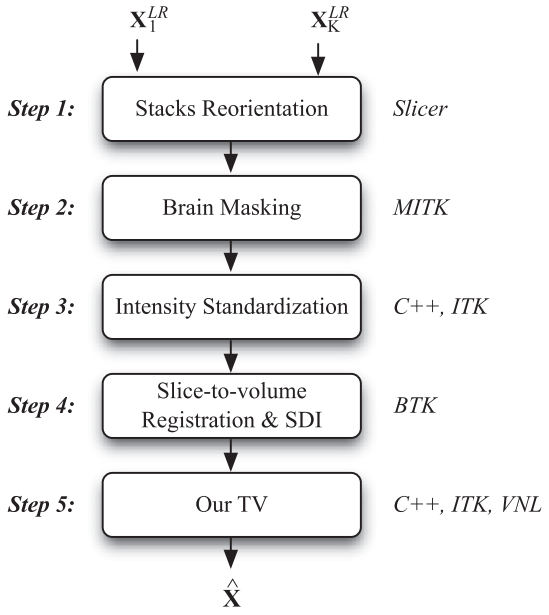

Fig. 4. Summary of the overall fetal brain HR reconstruction pipeline. Softwares, libraries and programming languages used are listed on the right-hand side in *italic*.

an initial HR image is estimated using SDI, as in Rousseau et al. (2006). It corresponds to Step 4 in Fig. 4.

Super-resolution

Our Total Variation algorithm is implemented in C++ with Insight Toolkit (Yoo et al., 2002). A first step is the computation of matrices H_k in Eq. (15). It consists of positioning and orienting Gaussian kernels according to the transformation between each slice and the HR image followed by sampling to the grid of the HR image. Calculations and operations inherited from our TV formulation are implemented using the numerical library VNL for algebra purposes. Note that input LR images to the SR algorithm are not filtered by NLM denoising. In addition, we prefer to not update motion estimation during super-resolution. We know that motion estimation updated during SR can further improve the quality of registration and consequently, the quality of image restoration. However, we noticed that the running time was dramatically increased. We decided to focus only on the optimization of the image restoration problem, which is the scope of this work. It corresponds to Step 5 in Fig. 4.

Results and validation

As described in section Variational terms, most published works in the literature of fetal brain MRI have focused and stayed to these three most important regularizers, *i.e.* Tikhonov, TV and NLM. In our knowledge, evaluations between different ε-TV formulations have been performed by the authors in (Rousseau et al., 2010) showing that they led to similar results in terms of PSNR. However, there is no existing comparison either with a classical Tikhonov as used in the pioneer work (Gholipour et al., 2010), which is very fast to converge, or with NLM as provided by the open-source toolbox BTK (Rousseau et al., 2013) (considered as one of the *state-of-the-art* method freely available).

In this section, we present an extensive validation of our algorithm. First, the convergence speed of our optimization scheme is validated as regards ε-TV optimization scheme on a 2D brain image of a mature newborn (section Numerical comparison of optimization schemes: ε-TV v.s. exact TV). Then, we carry out quantitative evaluations on 1) *simulated* fetal acquisitions from a newborn data and 2) normal and pathological fetal brain MRIs, where we compare our TV with Tikhonov and NLM. We study the robustness of the proposed TV reconstruction algorithm to different amounts of residual motion error (section Quantitative analysis of simulated fetal images). We also evaluate the ability of our algorithm to be conformed with the MRI stacks acquired (section Quantitative analysis of real fetal images). Eventually, we qualitatively evaluate the proposed algorithm on fetal MRIs (section Perceptual evaluation of the adaptive regularization).

As BTK, the Tikhonov method is implemented in C + + with Insight Toolkit (Yoo et al., 2002), where the Fletcher–Reeves version of the conjugate gradient is used for optimization. Projection onto the set $\{\mathbf{X} \geq 0\}$ is also applied to guarantee positive intensities. Global rigid stack registration is used for the newborn dataset (section Quantitative analysis of simulated fetal images). Global rigid stack registration followed by slice-to-volume registration is employed in the case of the clinical datasets (section Quantitative analysis of real fetal images). Tests are run on a 3.4 GHz Quad-core i7-3770 CPU. Reported running times correspond to the computational time of the SR problem optimization without motion estimation.

Material

Simulated fetal dataset

We use a T2-weighted TurboFLASH image of a mature newborn to simulate fetal brain MRIs, as in (Gholipour et al., 2010; Kuklisova-Murgasova et al., 2012). The image was acquired on a 3 T Siemens Tim Trio with $TR = 4000$ ms, $TE = 3.7$ ms, slice thickness of 1.2 mm and in-plane resolution of 0.78 mm. We consider the ground-truth image as the original image that was bias field corrected (Tustison and Gee, 2009) and isotropically interpolated to a resolution of 1 mm using B-splines. Six LR images, two per acquisition direction, with in-plane resolution of 1 mm, slice thickness of 3 mm, were simulated from the ground-truth image by applying downsampling and blurring operations. Shifts of 1 mm were also introduced for images with the same acquisition direction. Note that a number of six LR images were adopted in compliance with previous works Rousseau et al. (2010) and Kuklisova-Murgasova et al., 2012 showing that the

Table 1
Clinical datasets.

Patient	GA	Diagnostic	Brain Volume	Sequence	Stacks	Resolution	Mfactor
P1	22 weeks	Normal	112643 mm^3	HASTE	6	$1.13 \times 1.13 \times 3.6$	3.2
P5	22 weeks	Important malformation	118169 mm^3	HASTE	4	$1.13 \times 1.13 \times 3.6$	3.2
P2	25 weeks	Agnesis of the Corpus Colossum	204882 mm^3	HASTE	6	$1.13 \times 1.13 \times 3.6$	3.2
P4	26 weeks	Limited Gyration	290421 mm^3	HASTE	4	$1.13 \times 1.13 \times 3.6$	3.2
P10	30 weeks	Bilateral ventriculomegaly + kyste + hemorrhage	275914 mm^3	SSFSE	4	$1.09 \times 1.09 \times 5.5$	5.1
P3	32 weeks	Normal	470108 mm^3	HASTE	5	$1.13 \times 1.13 \times 4.8$	4.3
P7	32 weeks	Normal but small	324868 mm^3	SSFSE	8	$1.09 \times 1.09 \times 5.5$	5.1
P8	33 weeks	Asymmetrical ventriculomegaly	503682 mm^3	SSFSE	7	$1.09 \times 1.09 \times 5.5$	5.1
P6	34 weeks	Normal	344547 mm^3	SSFSE	5	$1.09 \times 1.09 \times 5.5$	5.1
P9	36 weeks	Bilateral ventriculomegaly + Cornes frontales	612998 mm^3	SSFSE	4	$1.09 \times 1.09 \times 5.5$	5.1
				Total	53		

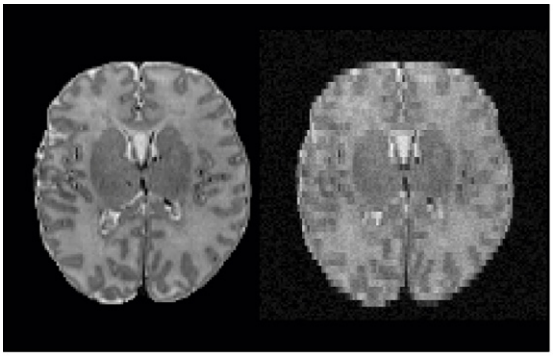

(a) 2D T2-weighted image of the mature newborn; Left: original HR image; Right: simulated LR image.

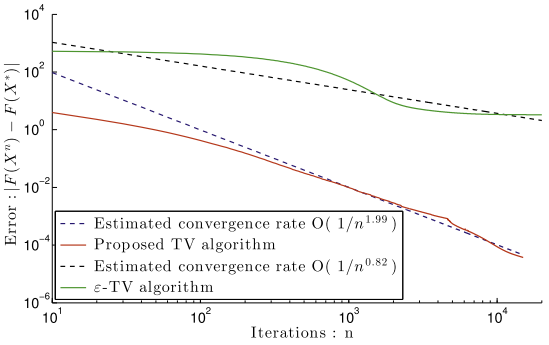

(b) Numerical asymptotic convergence rates.

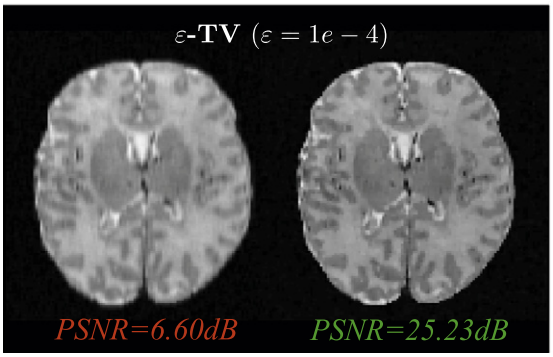

(c) Fetal MRI reconstruction with ε-TV with small ε; Left: solution after 20 iterations; Right: final steady state solution.

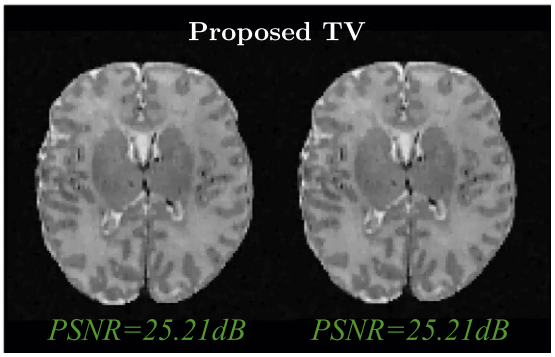

(d) Fetal MRI reconstruction with exact TV; Left: solution after 20 iterations; Right: final steady state solution.

Fig. 5. Numerical comparison of optimization schemes: ε-TV vs exact TV.

reconstruction quality marginally improves when more than 4–5 stacks LR images are used, in similar acquisition conditions.

Clinical fetal dataset

Our clinical fetal dataset is formed by 53 LR images coming from the acquisition of 10 fetus, aged between 22 and 36 weeks GA (see Table 1). The acquisition of each fetus consists of a set of 4 to 8 stacks, where at least one stack is available in each anatomical direction. Stacks were acquired using two different MRI scanners, including 1) a 1.5 T Siemens Aera using a T2-weighted HASTE sequence with a resolution of $1.13 \times 1.13 \times 3.6$ mm^3 (TE TR = 90/1200 ms), $1.13 \times 1.13 \times 4.8$ mm^3 (TE TR = 89/1000 ms) and 2) a 1.5 T Philips with a resolution of $1.09 \times 1.09 \times 5.5$ mm^3 (TE TR = 180/7000 ms). Our dataset is thus heterogeneous as acquisitions come from fetus with different age, performed on different MRI scanner, and with different slice-thickness-to-in-plane-resolution ratio, also known as magnification factor (Mfactor). It was shown by the authors in Lin and Shum 2004 on 2D images that when Mfactor is an integer, the sufficient number of LR images is Mfactor^2. A similar number of required LR images was found in Rousseau et al. (2010); Tourbier et al. (2014a) in the case of fetal 3D MRI.

This study has been approved by the Cantonal Research Ethics Committee of Vaud, Switzerland. The patient information from all data used in our study was anonymized and de-identified prior to our analysis.

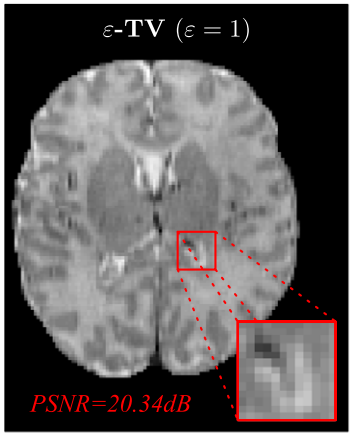

(a) Solution with ε-TV for *large* ε.

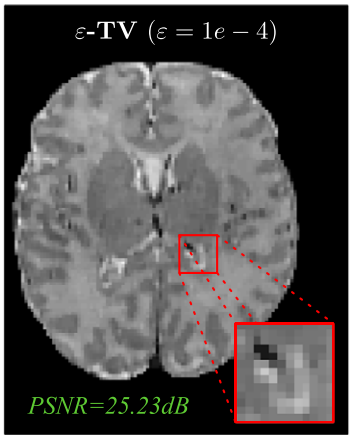

(b) Solution with ε-TV for *small* ε.

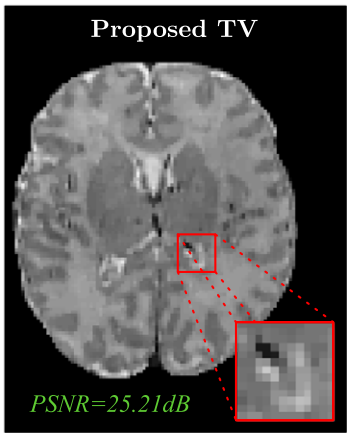

(c) Solution with exact TV.

Fig. 6. Impact of the choice of ε in ε-TV algorithms.

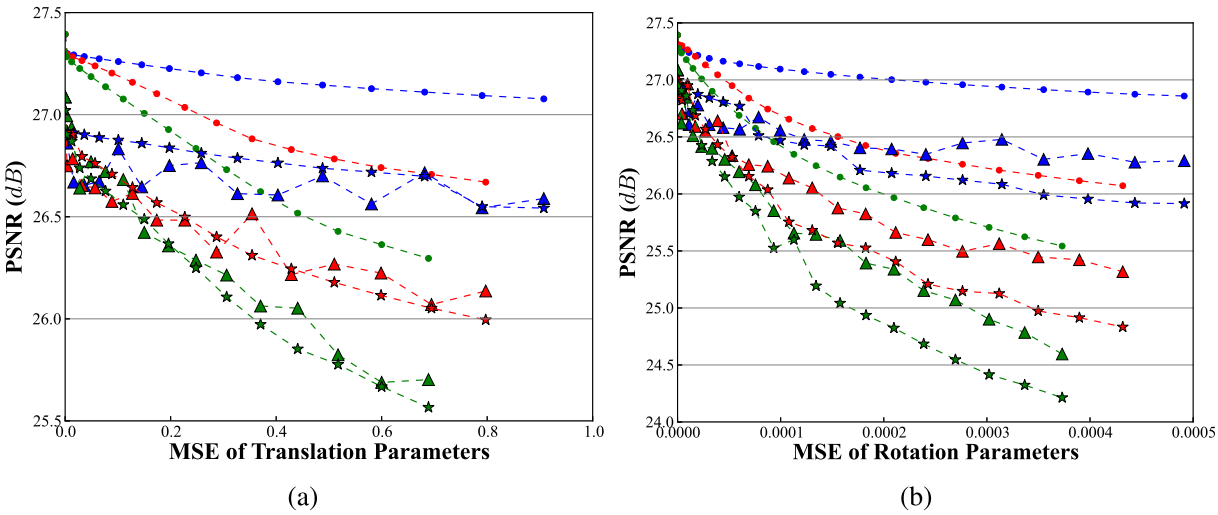

Fig. 7. Robustness of SR algorithms to motion error residuals when ★, ★, and ★ of the slices are affected. ★ denotes results obtained using the classical Tikhonov algorithm. ▲ denotes results obtained using the BTK algorithm based on NLM. ● denotes results obtained using our TV algorithm.

Numerical comparison of optimization schemes: ε-TV v.s. exact TV

The goal of this section is to show that (i) for *small* ε (here we take $\varepsilon = 1e-4$), ε-TV produces the same quality for reconstructed SR images than exact TV but it requires more time to converge, and (ii) for *large* ε (here we take $\varepsilon = 1$), ε-TV does not produce good image quality as the TV approximation is too smooth to recover sharp image discontinuities, *i.e.* edges.

For (i), we carry out a standard numerical experiment on Fig. 5 to validate the theoretical asymptotic rates of convergence for ε-TV with $\varepsilon = 1e-4$ (using an explicit forward gradient flow algorithm as in (Rousseau et al., 2010; Kuklisova-Murgasova et al., 2012)) and exact TV (using the proposed algorithm Eqs. (18)–(21) in section Our Total Variation algorithm). In order to speed up the computational time, we test both algorithms on a 2D HR image corresponding to an axial slice of the T2-weighted TurboFLASH image of the mature newborn (simulated fetal dataset). Four LR images were generated by applying subsampling operations and introducing noise but without adding any motion, see Fig. 4. The numerical asymptotic convergence rates are given in Fig. 4. The computed slope of our optimization algorithm, -1.99, is close to the theoretical one, *i.e.* -2, and the slope of ε-TV, -0.82, is also close to the theory, *i.e.* -1. To illustrate the impact of the difference of convergence rates, we present Figs. 4 and 4. At the 20th iteration, we can see that the solution given by our algorithm, Fig. 4, is very close to the final steady state solution with a PSNR value of 25.21 dB. However, the solution provided by the ε-TV algorithm, Fig. 4, is far away at the 20th iteration of its final solution as the PSNR value is 6.60 dB. It illustrates the fact that more iterations are required for the ε-TV algorithm to converge. Eventually, at convergence, both solutions have the same reconstruction quality, around 25.21 − 25.23 dB. This confirms that ε-TV for *small* ε offers same solution as exact TV, but at a higher computational cost.

For (ii), we present Fig. 6 that compares ε-TV for *large* ε (Fig. 5), ε-TV for *small* ε (Fig. 5), and exact TV (Fig. 5). The solution provided by small ε and exact TV have almost the same PSNR value of 25.23 dB, while the PSNR of large ε is lower at 20.34 dB. A lower PSNR for large ε is expected as a smoother version of TV is less able to recover image edges.

Quantitative analysis of simulated fetal images

We assess the performance of our algorithm w.r.t. the residual motion error by randomly affecting 1/5, 2/5, 3/5 of the slices in each LR images. Different amounts of residual motion errors are randomly added to one of the transform parameters (translation only and rotation only), following (Rousseau et al., 2006; Jiang et al., 2007; Gholipour et al., 2010). Translation errors are selected in $[-t, t]$ for $t = [0.1, 0.2, \ldots, 1.5$ mm$]$ and rotation errors are selected in $[-r, r]$ for $r = [0.1, 0.2, \ldots, 2.0^{\circ}]$. The amount of the overall residual motion error is measured by computing the mean square error of the transform parameters. The reconstruction quality is evaluated with the Peak-

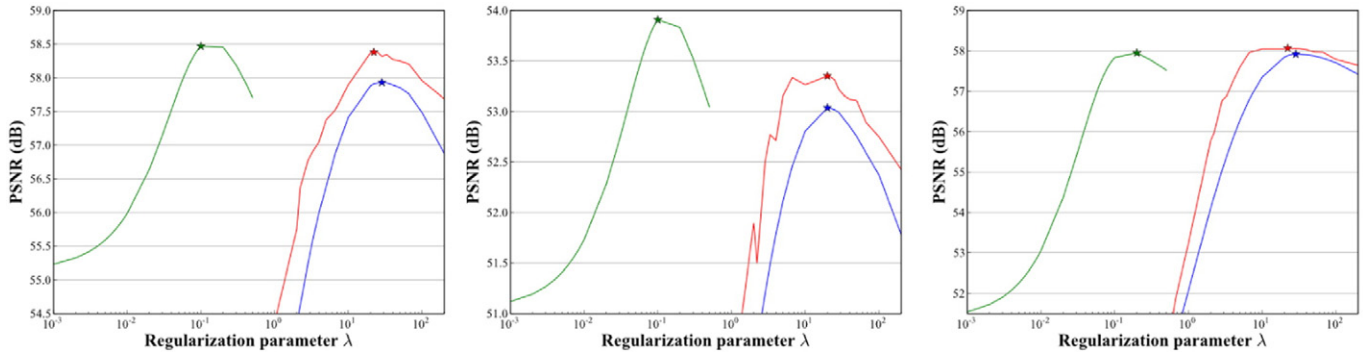

Fig. 8. Optimization of the regularization parameter λ for: a) patient P3 (5/18.5), b) patient P6 (5/26), and c) patient P7 (8/26). Numbers in parenthesis (X Y) correspond respectively to the number X of stacks of slices and to the slice-thickness-to-in-plane-resolution ratio Y = Mfactor^2 (defined in section Material). Both numbers indicate the "degree" of ill-posedness of the SR problem. ★ denotes the optimal regularization weight for Tikhonov. ★ denotes the optimal regularization weight for BTK. ★ denotes the optimal regularization weight for our TV. The estimated optimal values are further used for the leave-one-out analysis presented in Table 2.

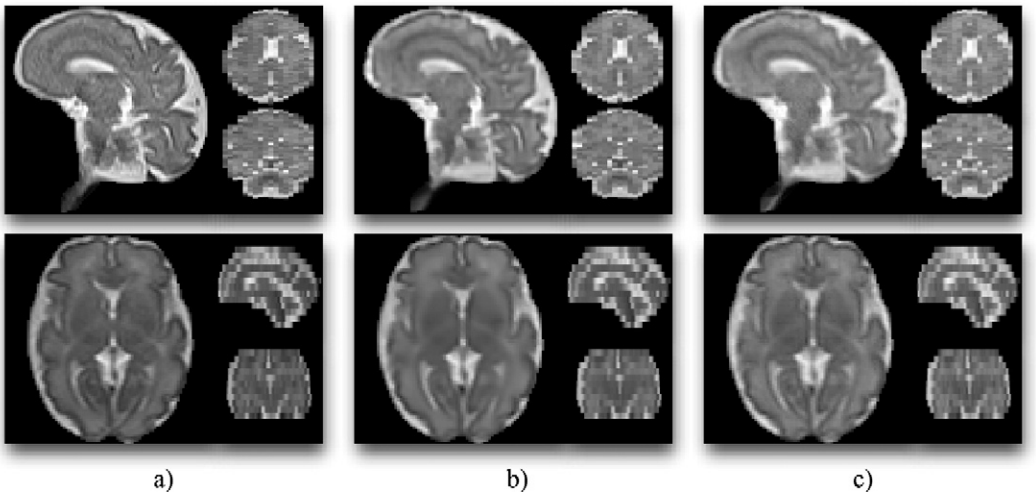

Fig. 9. Illustration of the leave-one-out analysis results for fetus P3 (row 1) and fetus P7 (row 2); a) Original LR Image $\mathbf{X}_k^{LR}$ left for evaluation; b) $\hat{\mathbf{X}}_k^{LR}$ using BTK; b) $\hat{\mathbf{X}}_k^{LR}$ using our TV.

Signal-To-Noise Ratio (PSNR). Parameters for Tikhonov, BTK and our algorithm are selected to have the best reconstructed HR image in the sense of the highest PSNR with respect to the original isotropic image of the newborn.

Results are presented in Fig. 7. We observe that the reconstruction quality always degrades independently of the regularization term as error motion residual is important. In addition, the TV regularization of our algorithm is less sensitive to residual errors of the motion estimation than the Tikhonov regularization and the NLM regularization employed in BTK. Finally, we observe that for similar MSE, the restoration quality of all methods decreases as the number of affected slice increases.

Quantitative analysis of real fetal images

We also assess the performance of our algorithm on clinical fetal brain MRI in terms of both fidelity with the original LR images and contrast in the reconstructed HR image. To do so, we perform a leave-one-out analysis, as in Kuklisova-Murgasova et al., 2012 that evaluates the ability of estimating the LR volume left out from the SR reconstruction based on all the other available LR volumes. We use the strategy described in section Adaptive regularization parameter λ to automatically set, for each patient, an optimal regularization weight in terms of PSNR. Fig. 8 illustrates the PSNR curves obtained for three different patients and for Tikhonov, BTK and TV. Obviously, λ is selected for each patient as the one providing the best reconstruction quality in terms of the highest mean PSNR (over the set of LR images). For adaptive regularization setting, exhaustive searches on the λ parameter are performed between the range of [0.001, 0.002, …, 0.01, 0.02, …, 0.1, 0.2, …, 0.5] for TV and between the range of [2/0.01, 2/0.02, …, 2/0.1, 2/0.2, …, 2/1.0, 2/2.0, …, 2/5.0] for BTK and Tikhonov.

Once λ_{PSNR} is set, for each patient and each regularization term, we assess both quality and speed of reconstruction with a leave-one-out analysis. Qualitative results on one left-out ground-truth LR volume and its estimation by BTK and TV regularization are shown in Fig. 9. Reconstruction results in terms of PSNR and running time are reported in Table 2. In terms of fidelity to the original LR stacks, both BTK and TV outperform Tikhonov and our TV algorithm performs better than BTK (0.34 dB in average for all subjects). In terms of speed, Tikhonov regularization is the fastest. TV regularization is about 4 times slower in average than Tikhonov (due to an extra inner problem optimisation), and TV is about 2–3 times faster than BTK. We also notice that BTK uses a parallel implementation of the non-local means algorithm, while the proposed TV algorithm has no such speed-up here. A parallel implementation of TV optimization may reduce the presented computational time.

It was observed that some reconstructed HR images suffer from reduced within-tissue contrast despite a high PSNR. Consequently, we decide to also measure the sharpness in the reconstructed HR image in complement to the PSNR. Sharpness validation was carried out using the energy of the gradient magnitude image (**M2**), as proposed in (Gholipour and Warfield, 2009). We computed **M2** by integrating the magnitude of the gradient of the HR reconstructed image at all voxels. The rational behind this study is that sharper structures would be observed if the motion-corrected images are more accurately fused in the restoration process. Sharpness results are summarized in Table 3. In most of all the patients, the HR images reconstructed by our TV-based restoration algorithm obtained higher **M2** values. Supported

Table 2
Leave-one-out analysis: similarity w.r.t. the original LR images.

	PSNR (dB)			Run Time (s)		
Patient	Tikhonov	BTK	Our TV	Tikhonov	BTK	Our TV
P1	53.64 ± 1.12	53.71 ± 1.29	**53.85 ± 1.19**	**3.21 ± 3.89**	10.76 ± 0.36	4.85 ± 0.06
P5	42.32 ± 2.45	43.44 ± 2.44	**44.64 ± 2.52**	**4.17 ± 0.60**	14.54 ± 0.60	4.66 ± 0.08
P2	48.31 ± 2.31	48.56 ± 2.44	**48.90 ± 2.37**	**6.66 ± 2.82**	21.52 ± 1.71	8.68 ± 0.11
P4	45.91 ± 3.21	46.38 ± 3.44	**47.14 ± 3.50**	**12.14 ± 5.45**	43.04 ± 5.04	15.08 ± 0.30
P10	53.35 ± 4.47	54.22 ± 4.46	**54.84 ± 4.40**	**2.79 ± 0.48**	36.53 ± 2.41	12.76 ± 0.16
P3	57.93 ± 1.11	58.39 ± 1.16	**58.47 ± 1.18**	**4.31 ± 0.67**	82.61 ± 6.80	30.21 ± 0.91
P7	57.91 ± 8.23	**58.07 ± 8.30**	57.94 ± 8.13	**4.23 ± 0.62**	52.56 ± 4.51	23.79 ± 0.39
P8	57.40 ± 3.06*s*	57.72 ± 3.19	**57.74 ± 3.06**	**4.92 ± 0.60**	104.74 ± 5.45	39.05 ± 0.46
P6	53.04 ± 4.24	53.39 ± 4.15	**53.91 ± 4.09**	**2.83 ± 0.40**	44.99 ± 1.91	15.88 ± 0.17
P9	52.90 ± 1.49	54.52 ± 1.11	**54.94 ± 1.00**	**4.09 ± 0.21**	76.33 ± 9.40	7.91 ± 0.18

Table 3
Leave-one-out analysis: sharpness of the reconstructed HR images. $\mathbf{M2}_{method}$ corresponds to the energy of the gradient magnitude image of the HR image reconstructed by the respective *method*.

Patient	$\mathbf{M2}_{Tikhonov}$	$\mathbf{M2}_{BTK}$	$\mathbf{M2}_{TV}$
P1	1437602	1428423	**1501358**
P5	959734	963221	**1327880**
P2	1869633	1696500	**2154680**
P4	2234225	2277655	**3040140**
P10	4272910	4291232	**4522360**
P3	**5918378**	5791398	5840018
P7	3988919	4053390	**4792452**
P8	5785973	5441901.43	**6407542**
P6	4732112	**5072980**	4757786
P9	4233893	3988113	**4247992**

by better PSNR, this suggests that our algorithm with TV regularization provides the best performance in terms of stronger contrasts between and within tissues.

Perceptual evaluation of the adaptive regularization

We show in this section the capability of our TV algorithm to restore HR images for diagnosis purposes in fetal MRI. In order to evaluate the optimality of the regularization weight λ_{PSNR} automatically set in terms of PSNR, see section Adaptive regularization parameter λ, we adopted a multiple-alternative forced-choice approach, inspired from (Reeves and Higdon, 1995). For each patient, a total of 6 HR images were reconstructed using a regularization weight of $\lambda_{PSNR}/8$, $\lambda_{PSNR}/4$, $\lambda_{PSNR}/2$, λ_{PSNR}, $2\lambda_{PSNR}$ and $4\lambda_{PSNR}$. Then, a radiologist expert proceeded to a visual evaluation where he indicated the best HR image in terms of image quality. Here, we consider only the reconstructed HR images of patients P1, P2, and P6 (diagnosed as normal brain), and the reconstructed HR images of patients P8 and P10 (diagnosed as abnormal brain) to minimize the number of images evaluated by the expert. Reconstruction results, running time for image restoration and the amount of regularization preferred by the expert are summarized in Fig. 10. Results for only 3 patients are visually shown due to space limitation. In general, HR reconstructed images with regularization weight of $\lambda_{PSNR}/4$ were selected for a majority of the patients (P1, P2, P8 and P10), except for patient P6, where an HR image with regularization of $\lambda_{PSNR}/2$ was chosen. This shows that the expert does prefer to visualize a little more regularized HR image than the one provided by the optimal PSNR (section Quantitative analysis of real fetal images). Let us notice that our PSNR image quality measure may be not fully correlated with the visual image quality. Especially, outliers in the LR images are still used in the reconstruction process which might emphasize the noise in the reconstructed HR image. Let us now recall that our PSNR-based parameter choice method works (on a leave-one-out fashion) on LR images simulated from the reconstructed HR image to compute the PSNR w.r.t. the excluded LR image. This simulation process reduces the noise in the LR image simulated as regards the real noise observed in the original HR reconstructed image. Therefore when we say "a little more regularized image" we here talk regularization w.r.t. PSNR and not w.r.t. visual quality measure. So even if it is more regularized w.r.t. PSNR, the image may not be regularized w.r.t. visual quality.

This suggests that our method, originally proposed to automatically set the regularization level for a fair comparison between our TV, Tikhonov and BTK, could indeed give an upper bound for regularization setting in clinical practice. This also suggests that the integration of a robust outlier removal scheme could give even a more optimal estimate of the regularization visually perceived by the radiologist expert.

Discussion

In this paper, we have revisited the existing SR techniques that address the fetal MRI brain reconstruction problem. Concretely, we have focused on the Bayesian and variational dual formulations with the goal of reviewing and unifying current state-of-the-art methods. Built on this formulation, we have recalled our TV-based optimization algorithm (Tourbier et al., 2014a) and we have presented the design and the implementation of our reconstruction pipeline. We have also numerically recovered the theoretical speed of convergence of the proposed algorithm. Precisely, we have shown that the proposed

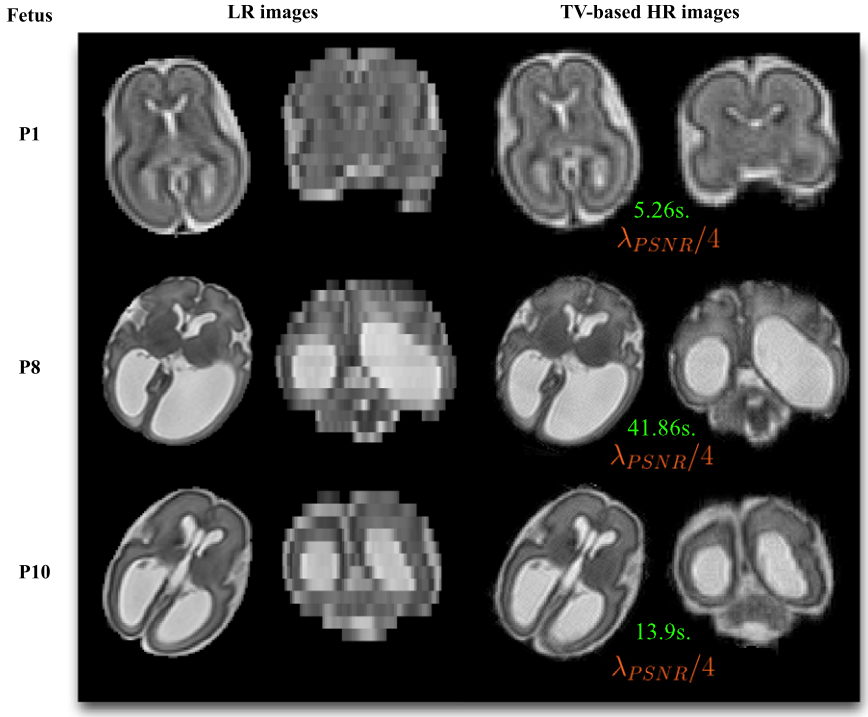

Fig. 10. Illustration of the qualitative analysis for fetus P1, P8 and P10.

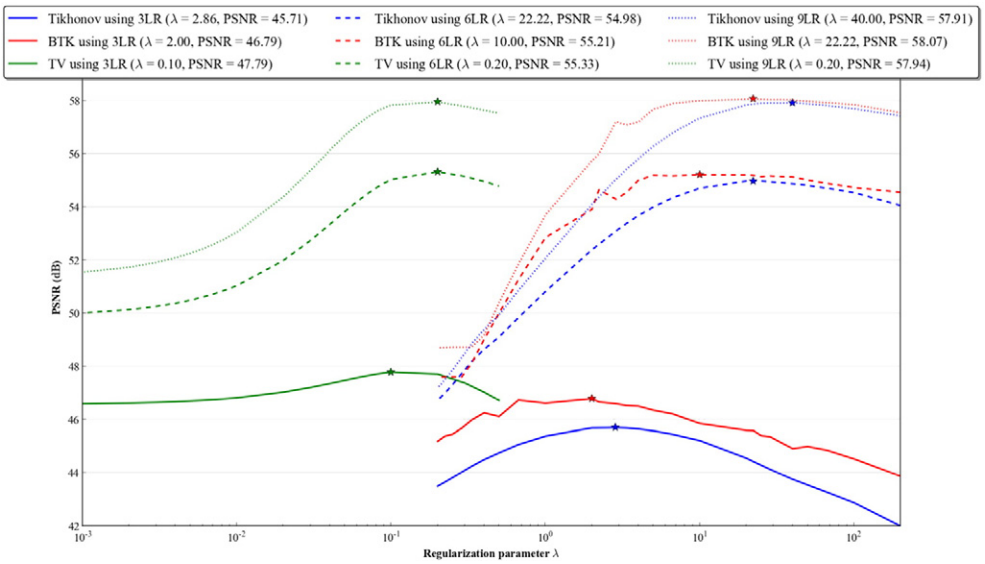

Fig. 11. Regularization *versus* number of input LR images.

algorithm solves the exact TV problem and outperforms the ε-TV optimization scheme in terms of convergence speed.

The second contribution is the extensive study of the convex TV regularization in comparison with classical Tikhonov and BTK. Concretely, we have presented for the first time a quantitative analysis of robustness of regularization w.r.t. residual registration errors (see section Quantitative analysis of simulated fetal images). Our experiments have clearly shown that TV is more robust to motion artifacts than Tikhonov and BTK. Moreover, this experiment confirms that registration accuracy is very crucial for the success of super-resolution algorithms. It also supports the development of outliers rejection schemes, as suggested in (Gholipour et al. (2010), Kim et al. (2010) and Kuklisova-Murgasova et al., 2012). Obviously, our approach would perform even better with the inclusion of these schemes. Using the automated setting of λ for each patient independently and inspired by (Kuklisova-Murgasova et al., 2012), we have performed a leave-one-out analysis to compare our TV regularization strategy to classical Tikhonov regularization and to NLM regularization on real fetal data (see section Quantitative analysis of real fetal images). We have shown that TV outperforms Tikhonov and BTK. We have observed that such improvement is higher in the cases where less LR volumes are available and where the Mfactor is higher *i.e.* in the cases where the problem is more ill-posed. To validate this hypothesis we have conducted an experiment that analyzes, for a given subject, the behavior of the regularization in function of the number of LR volumes (3, 6 and 9) used in the reconstruction (see Fig. 11). This confirms that TV is more robust when only few LR volumes are available. Such results go in favor of our TV, since in a clinical acquisition setting we will often have few LR available (between 3 and 6).

The third contribution is the adaptive setting of the amount of regularization w.r.t. each subject to be reconstructed and thus adapted to the ill-posedness nature of the reconstruction problem. We propose a first attempt to automatically select the optimal λ w.r.t. each subject and each algorithm in terms of PSNR (see section Quantitative analysis of real fetal images). Note that only an exhaustive search on λ was performed here. This could be improved by the formulation of the problem in terms of quadratic energy minimization. This will allow a faster convergence towards to optimal value of λ and it will guarantee that the solution is optimal and unique. We could also suggest integrating information about the sharpness of the reconstructed HR image into the energy to be minimized. This should even better estimate the appropriate regularization for diagnosis purposes. Moreover, we have adopted a simple heuristic approach as our goal was to fairly compare different regularization terms but the estimation of the optimal regularization weight could also be performed in a Bayesian framework, at the price of solving a non-convex optimization problem and computationally more expenses.

A last contribution is to show the capability of TV to reconstruct 3D volumes for diagnosis. We propose to proceed to a blind qualitative evaluation of the 3D reconstructed HR images conducted by an expert radiologist, including subjects diagnosed with both normal and abnormal brains (see section Perceptual evaluation of the adaptive regularization). It has been shown that in general, his choice was 4 times smaller than the estimated λ, optimal w.r.t. the PSNR. It demonstrates that the method, originally developed to determine the optimal regularization weight used by each algorithm with the leave-one-out analysis, is able to provide an upper bound to the regularization level in clinical application. In practice, we can imagine that this value could still be manually adjusted to fit end user expectations.

Conclusion

In summary, the proposed approach is a simple framework with an optimal and efficient TV optimization algorithm that has shown to be well adapted to solve the HR reconstruction problem in fetal MRI. It has been noticed that TV regularization generally outperforms Tikhonov and BTK in terms of reconstruction error w.r.t. the original LR images and in terms of sharpness of the reconstructed HR images. However, more sophisticated schemes could be integrated into the framework, such as complete outlier rejection scheme and bias field correction scheme in order to enhance its robustness to several kinds of clinical scenarios.

Acknowledgment

This work was supported by the Swiss National Science Foundation under Grant SNSF-141283 and by the CIBM of Geneva-Lausanne Universities and EPFL, as well as the Fondation Leenaards and Fondation Louis-Jeantet. The authors would like to thank Prof. Laurent Guibaud for providing access to the SSFSE data. We also would like to thank the reviewers for their time refereeing our manuscript and for providing us constructive comments.

References

Anquez, J., Angelini, E.D., Bloch, I., 2009. Automatic segmentation of head structures on fetal MRI. ISBI, IEEE, pp. 109–112.

Arrow, K., Hurwicz, L., Uzawa, H., 1958. Studies in Linear and Non-Linear Programming, Stanford University Press.

Beck, A., Teboulle, M., 2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM J. Imag. Sci. 2 (1), 183–202.

Besag, J., 1986. On the statistical analysis of dirty pictures. J. R. Stat. Soc. Ser. B 48 (3), 259–302.

Chambolle, A., Pock, T., 2011a. A first-order primal-dual algorithm for convex problems with applications to imaging. J. Math. Imaging Vision 40 (1), 120–145.

Chambolle, A., Pock, T., 2011b. A first-order primal-dual algorithm for convex problems with applications to imaging. JMIV 40 (1), 120–145.

Charbonnier, P., Blanc-féraud, L., Aubert, G., Barlaud, M., 1997. Deterministic edge-preserving regularization in computed imaging. IEEE Trans. Image Process. 6, 298–311.

Coakley, F.V., Glenn, O.A., Qayyum, A., Barkovich, A.J., Goldstein, R., Filly, R.A., 2004. Fetal MRI: a developing technique for the developing patient. Am. J. Roentgenol. 182 (1), 243–252.

Combettes, P., Pesquet, J., 2011. Proximal Splitting Methods in Signal Processing. Fixed-Point Algorithms for Inverse Problems in Science and Engineering. 185–212.

Combettes, P., Wajs, V., 2006. Signal recovery by proximal forward–backward splitting. Multiscale Model. Simul. 4 (4), 1168–1200.

Corbett-Detig, J., Habas, P., Scott, J., Kim, K., Rajagopalan, V., McQuillen, P., Barkovich, A., Glenn, O., Studholme, C., 2011. 3D global and regional patterns of human fetal subplate growth determined in utero. Brain Struct. Funct. 215, 255–263.

Courant, R., Friedrichs, K., Lewy, H., 1967. On the partial difference equations of mathematical physics. IBM J. Res. Dev. 11 (2), 215–234. http://dx.doi.org/10.1147/rd.112.0215.

Fogtmann, M., Seshamani, S., Kim, K., Chapman, T., Studholme, C., 2012. A unified approach for motion estimation and super resolution reconstruction from structural magnetic resonance imaging on moving objects. MICCAI workshop on Perinatal and Paediatric Imaging (PaPI 2012).

Galatsanos, N., Katsaggelos, A., 1992. Methods for choosing the regularization parameter and estimating the noise variance in image restoration and their relation. IEEE Trans. Image Process. 1 (3), 322–336. http://dx.doi.org/10.1109/83.148606.

Garel, C., 2006. New advances in fetal MR neuroimaging. Pediatr. Radiol. 36, 621–625.

Gholipour, A., Warfield, S.K., 2009. Super-resolution reconstruction of fetal brain MRI. MICCAI Workshop on Image Analysis for the Developing Brain (IADB'2009), pp. 45–52.

Gholipour, A., Estroff, J., Warfield, S., 2010. Robust super-resolution volume reconstruction from slice acquisitions: application to fetal brain MRI. IEEE Trans. Med. Imaging 29 (10), 1739–1758.

Gholipour, A., Estroff, J., Barnewolt, C., Connolly, S., Warfield, S., 2011. Fetal brain volumetry through mri volumetric reconstruction and segmentation. Int. J. Comput. Assist. Radiol. Surg. 6, 329–339.

Glowinski, R., Tallec, P.L., 1989. Augmented Lagrangian and Operator-Splitting Methods in Nonlinear Mechanics. SIAM.

Goldstein, T., Osher, S., 2009. The split Bregman method for L1-regularized problems. SIAM J. Imag. Sci. 2 (2), 323–343.

Habas, P.A., Scott, J.A., Roosta, A., Rajagopalan, V., Kim, K., Rousseau, F., Barkovich, A.J., Glenn, O.A., Studholme, C., 2012. Early folding patterns and asymmetries of the normal human brain detected from in utero mri. Cerebral Cortex 13–25.

Hansen, P., O'Leary, D., 1993. The use of the l-curve in the regularization of discrete ill-posed problems. SIAM J. Sci. Comput. 14 (6), 1487–1503. http://dx.doi.org/10.1137/0914086.

Hastings, W., 1970. Monte Carlo sampling methods using Markov chains and their applications. Biometrika 57, 97–109. http://dx.doi.org/10.1093/biomet/57.1.97.

Ison, M., Dittrich, E., Donner, R.e., Kasprian, G., Prayer, D., Langs, G., 2012. Fully Automated Brain Extraction and Orientation in Raw Fetal MRI. MICCAI Workshop on Perinatal and Paediatric Imaging: PaPI 2012pp. 17–24.

Jiang, S., Xue, H., Glover, A., Rutherford, M., Rueckert, D., Hajnal, J., 2007. MRI of moving subjects using multislice snapshot images with volume reconstruction (svr): application to fetal, neonatal, and adult brain studies. IEEE Trans. Med. Imaging 26 (7), 967–980.

Karl, W.C., 2005. Regularization in Image Restoration and Reconstruction, Handbook of Image and Video Processing. 2nd ed. pp. 183–202.

Keraudren, K., Kuklisova-Murgasova, M., Kyriakopoulou, V., Malamateniou, C., Rutherford, M., Kainz, B., Hajnal, J., Rueckert, D., 2014. Automated fetal brain segmentation from 2D {MRI} slices for motion correction. NeuroImage 101, 633–643.

Kim, K., Hansen, M., Habas, P., Rousseau, F., Glenn, O., Barkovich, A., Studholme, C., 2008. Intersection based registration of slice stacks to form 3D images of the human fetal brain. Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on, 2008, pp. 1167–1170.

Kim, K., Habas, P., Rousseau, F., Glenn, O., Barkovich, A., Studholme, C., 2010. Intersection based motion correction of multislice MRI for 3-D in utero fetal brain image formation. IEEE Trans. Med. Imaging 29 (1), 146–158.

Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P., 1983. Optimization by simulated annealing. Science 220, 671–680.

Kuklisova-Murgasova, M., Quaghebeur, G., Rutherford, M.A., Hajnal, J.V., Schnabel, J.A., 2012. Reconstruction of fetal brain mri with intensity matching and complete outlier removal. Medical Image Analysis 16 (8), 1550–1564. http://dx.doi.org/10.1016/j.media.2012.07.004.

Levine, D., Hatabu, H., Gaa, J., Atkinson, M., Edelman, R., 1996. Fetal anatomy revealed with fast mr sequences. AJ. Am. J. Roentgenol. 167, 905–908.

Lin, Z., Shum, H.Y., 2004. Fundamental limits of reconstruction-based superresolution algorithms under local translation. IEEE Transactions on Pattern Analysis and Machine Intelligence 26 (1), 83–97 (URL http://research.microsoft.com/apps/pubs/default.aspx?id=69073).

Malamateniou, C, Malik, S.J., Counsell, S.J., Allsop, J.M., McGuinness, A.K., Hayat, T., Broadhouse, K., Nunes, R.G., Ederies, A.M., Hajnal, J.V., Rutherford, M.A., 2013. Motion-Compensation Techniques in Neonatal and Fetal MR Imaging. AJNR American journal of neuroradiology 34 (6), 1124–1136. http://dx.doi.org/10.3174/ajnr.A3128.

Manjón, J., Coupé, P., Buades, A., Fonov, V., Collins, D.L., Robles, M., 2010. Non-local {MRI} upsampling. Med. Image Anal. 14 (6), 784–792. http://dx.doi.org/10.1016/j.media.2010.05.010 (URL http://www.sciencedirect.com/science/article/pii/S1361841510000630).

Mersereau, R.M., Reeves, S.J., 1990. Optimal estimation of the regularization parameter and stabilizing functional for regularized image restoration. Opt. Eng. 29 (5), 446–454. http://dx.doi.org/10.1117/12.55613.

Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E., 1953. Equation of state calculations by fast computing machines. J. Chem. Phys. 21, 1087–1092.

Nesterov, Y., 2005. Smooth minimization of Non-smooth functions. Math. Program. 103, 127–152.

Nyúl, L.G., Udupa, J.K., Zhang, X., 2000. New variants of a method of mri scale standardization. IEEE Trans. Med. Imaging 19 (2), 143–150.

PC, B., S. F., C, L., D, P., 2006. Eur. Radiol. 57, 172–181.

Pieper, S., Halle, M., Kikinis, R., 2004. 3D Slicer, ISBI. pp. 632–635.

Plenge, E., Poot, D.H.J., Bernsen, M., Kotek, G., Houston, G., Wielopolski, P., van der Weerd, L., Niessen, W.J., Meijering, E., 2012. Super-resolution methods in MRI: can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? Magn. Reson. Med. 68 (6), 1983–1993. http://dx.doi.org/10.1002/mrm.24187.

Prayer, D., 2011. Fetal MRI. Springer.

Ramani, S., Liu, Z., Rosen, J., Nielsen, J., Fessler, J., 2012. Regularization parameter selection for nonlinear iterative image restoration and mri reconstruction using gcv and sure-based methods. IEEE Trans. Image Process. 21 (8), 3659–3672. http://dx.doi.org/10.1109/TIP.2012.2195015.

Reeves, S., Higdon, A., 1995. Perceptual evaluation of the mean-square error choice of regularization parameter. IEEE Trans. Med. Imaging 4 (1), 107–110. http://dx.doi.org/10.1109/83.350808.

Rousseau, F., 2010. A non-local approach for image super-resolution using intermodality priors. Med. Image Anal. 14 (4), 594–605. http://dx.doi.org/10.1016/j.media.2010.04.005 (URL http://www.sciencedirect.com/science/article/pii/S1361841510000411).

Rousseau, F., Glenn, O.A., Iordanova, B., Rodriguez-Carranza, C., Vigneron, D.B., Barkovich, J.A., Studholme, C., 2006. Registration-based approach for reconstruction of high-resolution in utero fetal mr brain images. Acad. Radiol. 13 (9), 1072–1081.

Rousseau, F., Kim, K., Studholme, C., Koob, M., Dietemann, J., 2010. On super-resolution for fetal brain MRI. In: Jiang, T., Navab, N., Pluim, J., Viergever, M. (Eds.), Medical Image Computing and Computer-Assisted Intervention (MICCAI) 2010, Vol. 6362 of Lecture Notes in Computer Science. Springer Berlin, Heidelberg, pp. 355–362.

Rousseau, F., Oubel, E., Pontabry, J., Schweitzer, M., Studholme, C., Koob, M., Dietemann, J.-L., 2013. Btk: An open-source toolkit for fetal brain {MR} image processing, Computer Methods and Programs in Biomedicine. 109 (1) pp. 65–73. http://dx.doi.org/10.1016/j.cmpb.2012.08.007 URL http://www.sciencedirect.com/science/article/pii/S0169260712001897).

Rudin, L.I., Osher, S., Fatemi, E., 1992. Non-linear total variation based noise removal algorithms. Physica D 60 (14), 259–268. http://dx.doi.org/10.1016/0167-2789(92)90242-F (URL http://www.sciencedirect.com/science/article/pii/016727899290242F).

Scott, J.A., Habas, P.A., Kim, K., Rajagopalan, V., Hamzelou, K.S., Corbett-Detig, J.M., Barkovich, A.J., Glenn, O.A., Studholme, C., 2011. Growth trajectories of the human fetal brain tissues estimated from 3D reconstructed in utero MRI. Int. J. Dev. Neurosci. 29 (5), 529–536.

Stein, C.M., 1981. Estimation of the mean of a multivariate normal distribution. Ann. Stat. 9 (6), 1135–1151. http://dx.doi.org/10.1214/aos/1176345632.

Taleb, Y., Schweitzer, M., Studholme, C., Koob, M., Dietemann, J.-L., Rousseau, F., Jun. 2013. Automatic Template-based Brain Extraction in Fetal MR Images, organization for Human Brain Mapping (OHBM) conference. (URL http://hal.archives-ouvertes.fr/hal-00879089).

Tourbier, S., Bresson, X., Hagmann, P., Thiran, J.-P., Meuli, R., Bach Cuadra, M., 2014a. Efficient total variation algorithm for fetal brain MRI reconstruction. In: Golland, P. (Ed.), MICCAI Part II, no. 8674 in Lecture Notes in Computer Science. Springer, pp. 252–259.

Tourbier, S., Bresson, X., Hagmann, P., Cagneaux, M., Schaer, M., Guibaud, L., Thiran, J.-P., Meuli, R., Cuadra, M.B., 2014b. Automated brain extraction in fetal MRI by multi-atlas fusion strategy: study on healthy and pathological subjects. Joint Annual Meeting ISMRM-ESMRMB.

Tustison, N., Gee, J., 2009. N4ITK: ITK For MRI Bias Field Correction.

Van Reeth, E., Tham, I.W.K., Tan, C.H., Poh, C.L., 2012. Super-resolution in magnetic resonance imaging: a review. Concepts in Magnetic Resonance Part A 40A. 6, pp. 306–325. http://dx.doi.org/10.1002/cmr.a.21249.

Yang, J., Huang, T., 2011. Image super-resolution: historical overview and future challenges, from the book: super-resolution imaging. In: Milanfar, Peyman (Ed.), CRC Press (Taylor & and amp and Francis Group).

Yoo, T.S., Ackerman, M.J., Lorensen, W.E., Chalana, V., Aylward, S., Metaxas, D., Whitaker, R., 2002. ITK — the insight toolkit. Medicine Meets, Virtual Reality, pp. 586–592.

Zabaras, N., 2010. Solving Stochastic Inverse Problems: A Sparse Grid Collocation Approach. John Wiley & Sons, Ltd, pp. 291–319 http://dx.doi.org/10.1002/9780470685853.ch14.

Zhu, M., Chan, T., 2008. An Efficient Primal-Dual Hybrid Gradient Algorithm for Total Variation Image Restoration. UCLA CAM Report 08–34.