

Bayesian Parallel Imaging With Edge-Preserving Priors

Ashish Raj,^{1,2*} Gurmeet Singh,³ Ramin Zabih,^{4,5} Bryan Kressler,⁵ Yi Wang,⁵ Norbert Schuff,^{1,2} and Michael Weiner^{1,2}

Existing parallel MRI methods are limited by a fundamental trade-off in that suppressing noise introduces aliasing artifacts. Bayesian methods with an appropriately chosen image prior offer a promising alternative; however, previous methods with spatial priors assume that intensities vary smoothly over the entire image, resulting in blurred edges. Here we introduce an edge-preserving prior (EPP) that instead assumes that intensities are piecewise smooth, and propose a new approach to efficiently compute its Bayesian estimate. The estimation task is formulated as an optimization problem that requires a non-convex objective function to be minimized in a space with thousands of dimensions. As a result, traditional continuous minimization methods cannot be applied. This optimization task is closely related to some problems in the field of computer vision for which discrete optimization methods have been developed in the last few years. We adapt these algorithms, which are based on graph cuts, to address our optimization problem. The results of several parallel imaging experiments on brain and torso regions performed under challenging conditions with high acceleration factors are shown and compared with the results of conventional sensitivity encoding (SENSE) methods. An empirical analysis indicates that the proposed method visually improves overall quality compared to conventional methods. Magn Reson Med 57:8–21, 2007. © 2006 Wiley-Liss, Inc.

Key words: parallel imaging; edge-preserving priors; Bayesian reconstruction; SENSE; graph cuts; regularization

The use of multiple coils in MRI to reduce scan time (and thus motion artifacts) has become quite popular recently. Current parallel imaging techniques include, *inter alia*, sensitivity encoding (SENSE) (1–3), simultaneous acquisition of spatial harmonics (SMASH) (4,5), and generalized autocalibrating partially parallel acquisitions (GRAPPA) (6). A good comparative review was presented in Ref. 7. While all of these methods are mathematically similar, SENSE is the reconstruction method that performs exact matrix inversion (8) and is the focus of this work. These schemes use multiple coils to reconstruct (unfold) the unaliased image from undersampled data in SnFourier or

k-space. Successful unfolding relies on receiver diversity (i.e., each coil “sees” a slightly different image because each coil has a different spatial sensitivity profile).

Unfortunately, the conditioning of the encoding system becomes progressively worse with increasing acceleration factors. Therefore conventional parallel imaging methods, especially at accelerations above 3, suffer from a fundamental noise limitation in that unfolding is achieved at the cost of noise amplification. This effect depends on the coil geometry and acceleration factor, and is best captured in terms of the *g*-map, which is a spatial mapping of the noise amplification factor. Reconstructed data can be further degraded in practice by inconsistencies between encoding and decoding sensitivity due to physiological motion, misalignment of coils, and insufficient resolution of sensitivity calibration lines. In this paper we propose a novel Bayesian approach whereby an edge-preserving spatial prior is introduced to reduce noise and improve the unfolding performance of parallel imaging reconstruction. The resulting estimation task is formulated as an optimization problem whose solution is efficiently obtained by graph-based algorithms.

Methods to reduce noise and artifacts can be grouped into two classes: 1) those that handle sensitivity errors using a maximum-likelihood (9) or total least squares (10) approach, and 2) those that exploit some prior information about the imaging target via regularization. While regularization is generally effective for solving ill-posed problems, existing methods rarely exploit the spatial dependencies between pixels. Most techniques either impose minimum norm solutions (as in regularized SENSE (11–13)), or require a prior estimate of the target. Temporal priors for multiframe imaging have been reported (14–16), and a generalized series model was developed with the use of reduced-encoding imaging with generalized-series reconstruction (RIGR) (17,18). The limitations of these regularization techniques are clear: regularized SENSE makes unrealistic assumptions about the image norm, while methods that rely on a prior estimate of the imaging target (called the mean or reference image) (12,13,17,18) must be carefully registered to the target. In practice, the use of such strong reference priors is vulnerable to errors in their estimation, leading to reconstruction artifacts. Temporal priors are obviously restricted to dynamic imaging. Even though the minimum norm prior reduces noise, it is unsatisfactory for de-aliasing in parallel imaging because it can be shown by simple algebra that it favors a solution with equally strong aliases. For example, suppose we have an underdetermined aliasing system $y = x_1 + x_2$. Then the minimum norm solution is $x = x_2 = y/2$, which amounts to alias energy being equal to the desired signal

¹Department of Radiology, University of California–San Francisco, San Francisco, California, USA.

²Veterans Affairs Medical Center, San Francisco, California, USA.

³Department of Electrical Engineering, Cornell University, Ithaca, New York, USA.

⁴Department of Computer Science, Cornell University, Ithaca, New York, USA.

⁵Department of Radiology, Weill Medical College of Cornell University, New York, New York, USA.

*Correspondence to: Ashish Raj, Center for Imaging of Neurodegenerative Diseases (CIND), VA Medical Center (114M), 4150 Clement St., San Francisco, CA 94121. E-mail: ashish.raj@ucsf.edu

Received 4 November 2005; revised 1 June 2006; accepted 7 June 2006.

DOI 10.1002/mrm.21012

Published online in Wiley InterScience (www.interscience.wiley.com).

© 2006 Wiley-Liss, Inc.

energy. This is why conventional regularization techniques are good at reducing noise but countereffective for removing aliasing. The introduction of spatial priors is essential for the latter task. This is possible within the Tikhonov regularization framework (13), as long as it is assumed that intensities vary smoothly over the entire imaging target.

Our approach uses a spatial prior on the image that makes much more realistic assumptions regarding smoothness. Our prior model is quite general and has very few parameters; hence, little or no effort is required to find this prior, in contrast to image-based or temporal priors. The primary challenge in our formulation is a computational one: Unlike regularized SENSE, there is no closed-form solution, and we need to minimize a nonconvex objective function in a space with thousands of dimensions. However, we developed an efficient algorithm to solve this problem by relying on some powerful discrete optimization techniques that were recently developed in the computer-vision community (19,20,39).

We apply an edge-preserving prior (EPP) that assumes that voxel intensity varies slowly within regions but (in contrast to smoothness-enforcing Tikhonov regularization) can change discontinuously across object boundaries (21). Since our piecewise smooth model imposes relationships only between neighboring voxels, it can be used for a very wide range of images (for instance, it is applicable to any MR image regardless of contrast or modality). EPPs have been studied quite extensively in the fields of statistics (22), computer vision (21,23,24), and image processing (25), and are widely considered to be natural image models.

The computational challenge of EPPs is too difficult for conventional minimization algorithms, such as conjugate gradients or steepest descent. However, EPPs have become widely used in the field of computer vision in the last few years, due primarily to the development of powerful optimization techniques based on graph cuts (19,20). Graph cuts are discrete methods whereby the optimization task is reformulated as the problem of finding a minimum cut on an appropriately constructed graph¹. The minimum-cut problem, in turn, can be solved very efficiently by modern graph algorithms (26).

While graph-cut algorithms can only be applied to a restricted set of problems (19), they have proved to be extremely effective for applications such as stereo matching (20), where they form the basis for most of the top-performing algorithms (27). Although standard graph-cut algorithms (19,20) cannot be directly applied to minimize our objective function, we have developed a graph-cut reconstruction technique based on a subroutine from Hammer et al. (36) that is quite promising. Preliminary parallel imaging results indicate the potential and promise of this algorithm, which we call “edge-preserving parallel imaging with graph-cut minimization” (EPIGRAM).

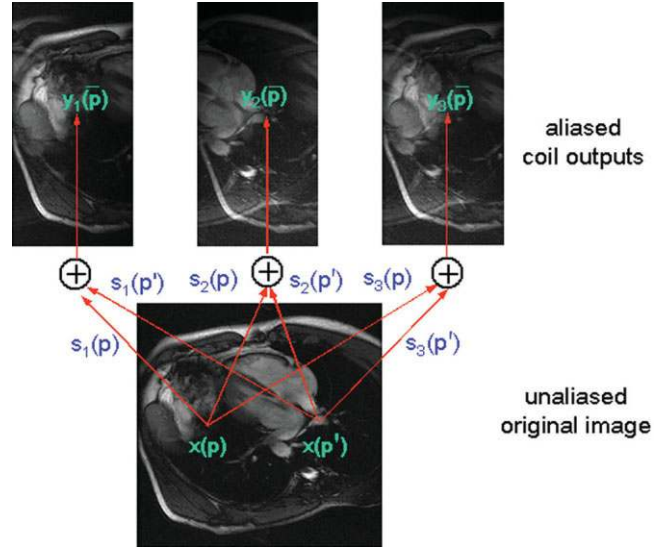


FIG. 1. Schematic of the pixelwise aliasing process for a single pair of aliasing pixels p and p' , for $2\times$ acceleration using three coils. The aliased observations Y_l are obtained by a weighted sum of the aliasing pixels, weighted by coil sensitivity values S_l . To simplify the figure, aliasing is shown in the horizontal direction. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

THEORY

Summary of Parallel Imaging

Suppose the target image is given by $X(\mathbf{r})$, where \mathbf{r} is the 2D spatial vector, and \mathbf{k} is a point in k -space. The imaging process for each coil l (from L coils) is given by

$$Y_l(\mathbf{k}) = \int d\mathbf{r} e^{-i2\pi\mathbf{r}\mathbf{k}} S_l(\mathbf{r}) X(\mathbf{r}) \quad [1]$$

where y_l is the (undersampled) k -space data seen by the l -th coil, and S_l is its sensitivity response. For Cartesian sampling it is well known that Eq. [1] reduces to folding in image space, such that each pixel p in Y_l , the Fourier transform of y_l , results from a weighted sum of aliasing pixels in X . If phase-encoding steps are reduced by R times, the $N \times M$ image will fold over into $N/R \times M$ aliased coil outputs. This can be easily verified by discretizing Eq. [1] and taking the Fourier transform. For $i=1, K, N; \bar{i}=1, K, N/R; j=1, K, M$; define $p=(i, j)$, $\bar{p}=(\bar{i}, j)$. We then have

$$Y_l(\bar{p}) = \sum_{[p]=\bar{p}} S_l(p) X(p) \quad [2]$$

where we denote $[p] = (\text{mod}(i, N/R), j)$. Note that as defined, \bar{p} spans the pixels of Y_l , and p spans the pixels of X . This process is depicted in Fig. 1. Over the entire image this has a linear form:

$$\mathbf{y} = \mathbf{E}\mathbf{x} \quad [3]$$

where vector $\mathbf{x} = \{x_p | p \in \mathbf{P}\}$ is a discrete representation of the intensities of the target image $X(\mathbf{r})$, p indexes the set \mathbf{P}

¹A graph is a set of objects called “nodes” joined by links called “edges.” The relationship between graph cuts and optimization problems is described in more detail at the end of this section, and in Refs. 19 and 20.

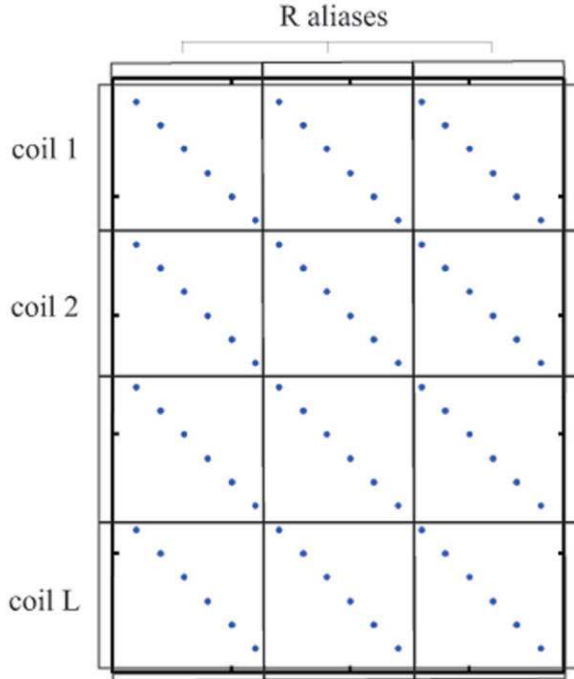


FIG. 2. The $L \times R$ block-diagonal structure of matrix \mathbf{E} . Each sub-block sensitivity-encodes one aliasing band in the image. Matrix \mathbf{E} is a concatenation over all coil outputs, as shown by the coil labels. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

of all pixels of the image, and vector \mathbf{y} contains the aliased images “seen” by the receiver coils. Matrix \mathbf{E} encodes coil sensitivity responses and is a $L \times R$ block-diagonal matrix of the form shown in Fig. 2.

SENSE takes a least-squares approach via the pseudo-inverse of \mathbf{E} :

$$\hat{\mathbf{x}}_{SENSE} = (\mathbf{E}^H \mathbf{E})^{-1} \mathbf{E}^H \mathbf{y} \quad [4]$$

This is the maximum-likelihood estimate under the assumption of additive white Gaussian noise (28). Unfortunately, inverse problems of this form become progressively ill-posed with increasing acceleration, leading to noise amplification and insufficient de-aliasing in many cases. To reduce these effects and stabilize the inversion in SENSE, a Tikhonov-type regularization is introduced (10,13):

$$\hat{\mathbf{x}}_{regSENSE} = \arg \min_{\mathbf{x}} \{ \|\mathbf{E}\mathbf{x} - \mathbf{y}\|^2 + \mu^2 \|\mathbf{A}(\mathbf{x} - \mathbf{x}_r)\|^2 \} \quad [5]$$

where the first term enforces agreement with observed data, and the second penalizes nonsmooth solutions through an appropriate matrix \mathbf{A} and some prior reference image \mathbf{x}_r . Its closed-form solution is

$$\hat{\mathbf{x}}_{regSENSE} = \mathbf{x}_r + (\mathbf{E}^H \mathbf{E} + \mu^2 \mathbf{A}^H \mathbf{A})^{-1} \mathbf{E}^H (\mathbf{y} - \mathbf{E}\mathbf{x}_r) \quad [6]$$

Equations [4] and [6] can both be computed very quickly for the Cartesian case because they readily break up into independent $L \times R$ sub-problems (for the standard choice

of $\mathbf{A} = \mathbf{I}$, the identity matrix). If there is no reference image (i.e., $\mathbf{x}_r = 0$), then with $\mathbf{A} = \mathbf{I}$ this computes the minimum norm solution (11), while more general Tikhonov forms of \mathbf{A} impose global smoothness.

Bayesian Reconstruction

Even with regularization there is a noise/unfolding limit. If μ is too small, there will be insufficient noise reduction. If μ is too high, noise will be removed but residual aliasing will occur. This fundamental aliasing/noise limit cannot be overcome unless more information about the data is exploited. This naturally suggests a Bayesian approach, which was the subject of a recent work (29). Given the imaging process (Eq. [3]), observation \mathbf{y} , and the prior probability distribution $\Pr(\mathbf{x})$ of the target image \mathbf{x} , Bayesian methods maximize the posterior probability:

$$\Pr(\mathbf{x}|\mathbf{y}) \propto \Pr(\mathbf{y}|\mathbf{x}) \cdot \Pr(\mathbf{x}) \quad [7]$$

The first right-hand term, called the likelihood function, comes from the imaging model (Eq. [3]), and the second term is the prior distribution. In the absence of a prior, this reduces to maximizing the likelihood, as performed by SENSE. Assuming that $\mathbf{n} = \mathbf{y} - \mathbf{E}\mathbf{x}$ is white Gaussian noise, Eq. [3] implies a simple Gaussian distribution for $\Pr(\mathbf{y}|\mathbf{x})$:

$$\Pr(\mathbf{y}|\mathbf{x}) \propto e^{-\|\mathbf{y} - \mathbf{E}\mathbf{x}\|^2} \quad [8]$$

Similarly, we can write the prior, without loss of generality, as

$$\Pr(\mathbf{x}) \propto e^{-G(\mathbf{x})} \quad [9]$$

Depending on the term $G(\mathbf{x})$, this form can succinctly express both the traditional Gaussianity and/or smoothness assumptions, as well as more complicated but powerful Gaussian or Gibbsian priors modeled by Markov random fields (MRFs) (23,24). The posterior is maximized by

$$\hat{\mathbf{x}}_{MAP} = \arg \min_{\mathbf{x}} (\|\mathbf{y} - \mathbf{E}\mathbf{x}\|^2 + G(\mathbf{x})) \quad [10]$$

which is the maximum a posteriori (MAP) estimate. Conventional image priors impose spatial smoothness; hence, this can be viewed as the sum of a data penalty and a smoothness penalty. The data penalty forces \mathbf{x} to be compatible with the observed data, and the smoothness penalty $G(\mathbf{x})$ penalizes solutions that lack smoothness. Traditionally, only smoothness penalties of the kind $G(\mathbf{x}) = \|\mathbf{A}\mathbf{x}\|^2$ have been used, where \mathbf{A} is a linear differential operator. This corresponds to Eq. [6] if the reference image $\mathbf{x}_r = 0$, and is commonly known as Tikhonov regularization (28).

However, this smoothness penalty assumes that intensities vary smoothly across the entire image. Such an assumption is inappropriate for most images because although most image data change smoothly, they have discontinuities at object boundaries. As a result, the Tikhonov smoothness penalty causes excessive edge blurring, while we seek an edge-preserving G . To illustrate

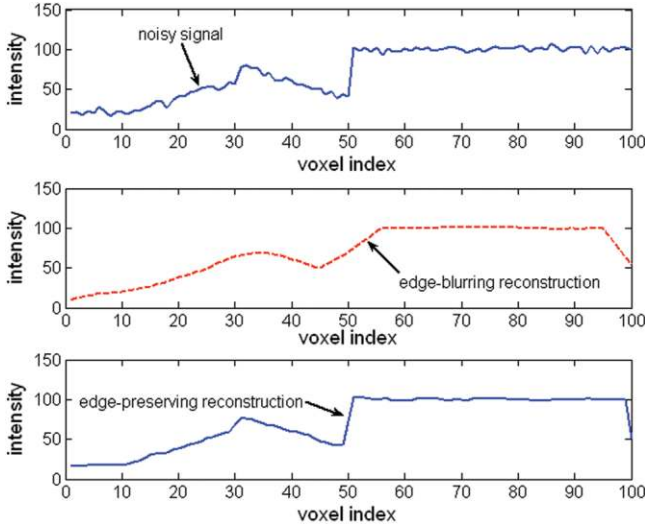


FIG. 3. Typical performance of edge-preserving and edge-blurring reconstruction on a noisy image row. Image **b** was obtained by globally smoothing via a Gaussian kernel, and image **c** was obtained with an edge-preserving median filter. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

this, we show in Fig. 3 a single noisy image row and two possible strategies to denoise it. The difference in performance between global smoothing (obtained by Gaussian blurring) and edge-preserving smoothing (via median filtering) is obvious: although both denoise the signal, one oversmooths sharp transitions while the other largely preserves them.

EPPs in MRI

A natural class of edge-preserving smoothness penalties is

$$G_{EP}(\mathbf{x}) = \sum_{(p,q) \in N_s} V(x_p, x_q) \quad [11]$$

The spatial neighborhood system N_s consists of pairs of adjacent pixels, usually the eight-connected neighbors. The separation cost $V(x_p, x_q)$ gives the cost to assign intensities x_p and x_q to neighboring pixels p and q , and the

form of this prior can be justified in terms of MRFs (23). Typically, V has a nonconvex form, such as $V(x_p, x_q) = \lambda \min(|x_p - x_q|, K)$, for some metric $|\cdot|$ and constants K, λ . Such functions effectively assume that the image is piecewise smooth rather than globally smooth. Figure 4 shows two possible choices of V : the right one preserves edges, and the left one does not. For MR data the truncated linear model appears to work best, and seems to present the best balance between noise suppression (due to the linear part) and edge-preservation (due to truncation of penalty function). Therefore, neighboring intensity differences within the threshold K will be treated as noise and penalized accordingly. However, larger differences will not be further penalized, since they occur, most likely, from the voxels being separated by an edge. Note that this is very different from using a traditional convex distance, such as the L_2 norm, which effectively forbids two adjacent pixels from having very different intensities, for the separation cost. Although the L_2 separation cost does not preserve edges, it is widely used because its convex nature vastly simplifies the optimization problem.

A possible problem with the truncated linear penalty is that it can lead to some loss of texture, since the Bayesian estimate will favor images with piecewise smooth areas over those with textured areas. In the Discussion we point out examples of this feature and suggest ways to mitigate it.

Parallel Imaging As Optimization

The computational problem we face is to efficiently minimize

$$\|\mathbf{y} - \mathbf{E}\mathbf{x}\|^2 + G_{EP}(\mathbf{x}) \quad [12]$$

From Ref. 4 we know that \mathbf{E} has a diagonal block structure and decomposes into separate interactions between R aliasing voxels, according to Eq. [2]. Let us first define for each pixel $\bar{p} = (\bar{i}, j)$ in Y_j the set of aliasing pixels in X that contribute to $Y_j(\bar{p})$, as follows: For image X of size $M \times N$ undergoing R -fold acceleration, aliasing occurs only in the phase-encode direction, between aliasing pixels. Then

$$\|\mathbf{y} - \mathbf{E}\mathbf{x}\|^2 = \sum_{\bar{p}} \sum_l \left[Y_l(\bar{p}) - \sum_{[p]=\bar{p}} S_l(p)x_p \right]^2 \quad [13]$$

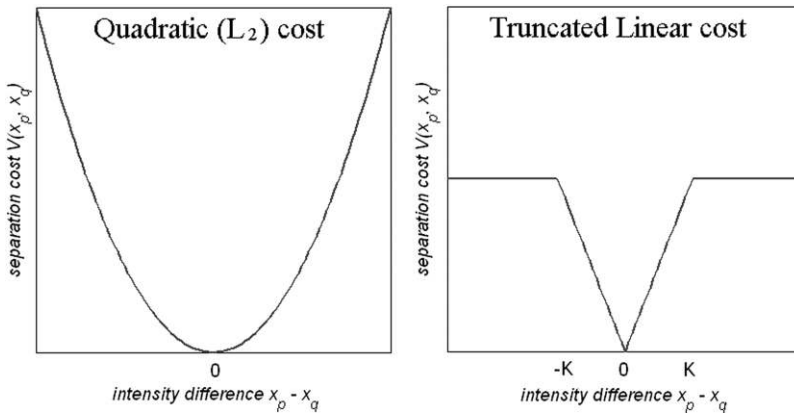


FIG. 4. Two natural separation cost functions for spatial priors. The L_2 cost on the left usually causes edge blurring due to an excessive penalty for high-intensity differences, whereas the truncated linear potential on the right is considered to be edge-preserving and robust. For MR data, the truncated linear model appears to work best. While the L_2 separation cost does not preserve edges, its convex nature vastly simplifies the optimization problem.

This can be intuitively understood by examining the aliasing process depicted in Fig. 1. After some rearrangement, this expands to

$$\begin{aligned} \|\mathbf{y} - \mathbf{E}\mathbf{x}\|^2 &= \sum_{\bar{p}} \sum_l Y_l^2(\bar{p}) + \sum_p \left(\sum_l S_l^2(p) \right) x_p^2 \\ &\quad - 2 \sum_p \left(\sum_l S_l(p) Y_l([p]) \right) x_p + 2 \sum_{(p,p') \in N_a} \left(\sum_l S_l(p) S_l(p') \right) x_p x_{p'} \end{aligned} \quad [14]$$

where we define the aliasing neighborhood set $N_a = \{(p, p'), [p] = [p'], p \neq p'\}$ over all aliasing pairs. Grouping terms under single pixel and pairwise interactions, we get

$$\begin{aligned} \|\mathbf{y} - \mathbf{E}\mathbf{x}\|^2 &= a^2 + \sum_p b(p)x_p^2 - 2 \sum_p c(p)x_p \\ &\quad + 2 \sum_{(p,p') \in N_a} d(p, p')x_p x_{p'} \end{aligned}$$

for appropriately chosen functions $b(p)$, $c(p)$, and $d(p, p')$. The first term is a constant and can be removed from the objective function, the next two terms depend only on a single pixel, and the last term depends on two pixels (both from the aliasing set) at once. This last term, which we will refer to as a ‘‘cross term,’’ arises due to the nondiagonal form of our system matrix \mathbf{E} .

To perform edge-preserving parallel imaging, we need to minimize our objective function:

$$\begin{aligned} \mathbf{E}\mathbf{x} &= a^2 + \sum_p b(p)x_p^2 - 2 \sum_p c(p)x_p + 2 \sum_{(p,p') \in N_a} d(p, p')x_p x_{p'} \\ &\quad + \sum_{(p,q) \in N_s} V(x_p, x_q) \end{aligned} \quad [15]$$

Let us first consider the simpler case of our objective function that would arise if \mathbf{E} were diagonal. In this case there would be no cross terms (i.e., $d(p, p') = 0$), which appears to simplify the problem considerably. Yet even this simplification results in a difficult optimization problem. There is no closed-form solution, the objective function is highly nonconvex, and the space over which we are minimizing has thousands of dimensions (one dimension per pixel). Worse still, minimizing such an objective function is almost certain to require an exponentially large number of steps².

If \mathbf{E} were diagonal, however, the objective function would be in a form that has been extensively studied in the computer-vision field (10,20,27), where significant recent progress has been made. Specifically, a number of powerful methods have been designed that employ a discrete optimization technique called ‘‘graph cuts’’ (20), which we briefly summarize in the next section. Graph cuts are a powerful means of minimizing $\mathbf{E}(\mathbf{x})$ in Eq. [15], and can be

easily applied as long as \mathbf{E} is diagonal (19). The presence of off-diagonal entries in \mathbf{E} gives rise to cross terms in our objective function, making traditional graph-cut algorithms (19,20) inapplicable, and requires an extension, as described in Materials and Methods.

Optimization With Graph Cuts

One can minimize objective functions similar to Eq. [15] by computing the minimum cut in an appropriately defined graph using the graph-cut technique. This technique was first used for images by Greig et al. (30), who used it to optimally denoise binary images. A recent series of papers (19,20) extended the method significantly, and it can now be used for problems such as stereo matching (20,31) and image/video synthesis (32) in computer vision, as well as medical image segmentation (33) and fMRI data analysis (34).

The basic idea is to first discretize the continuous pixel intensities x_p into a finite discrete set of labels $L = \{1, K, N_{\text{labels}}\}$. Since we focus on MR reconstruction problems, we will assume that the labels are always intensities and use the terms interchangeably; however, graph-cut algorithms are employed for a wide variety of problems in computer vision and graphics, and often use labels with a more complex meaning. Then instead of minimizing over continuous variables x_p , we minimize over individual labels $\alpha = L$, allowing any pixel in the image to take the label α . In practice the dynamic range of intensities may have to be reduced for computational purposes, although this is not a requirement of our technique. The most powerful graph-cut method is based on expansion moves. Given a labeling $\mathbf{x} = \{x_p | p \in P\}$ and a label α , an α -expansion $\chi = \{\chi_p | p \in P\}$ is a new labeling whereby χ_p is either x_p or α . Intuitively, one constructs χ from \mathbf{x} by giving some set of pixels the label α . The expansion-move algorithm picks a label α , finds the lowest cost χ , and moves there. This is pictorially depicted in Fig. 5.

The algorithm converges to a labeling where there is no α -expansion that reduces the value of the objective function \mathbf{E} for any α . The key subroutine in the expansion move algorithm is to compute the α -expansion χ that minimizes \mathbf{E} . This can be viewed as an optimization problem over binary variables, since during an α -expansion each pixel either keeps its old label or moves to the new label α . This is also shown in Fig. 5. An α -expansion χ is equivalent to a binary labeling

$$\chi_p = \begin{cases} x_p & \text{iff } b_p = 0 \\ \alpha & \text{iff } b_p = 1 \end{cases} \quad [16]$$

Just as for a labeling χ there is an objective function \mathbf{E} , for a binary labeling \mathbf{b} there is an objective function B . More precisely, assuming χ is equivalent to \mathbf{b} , we define B by

$$B(\mathbf{b}) = \mathbf{E}(\chi)$$

We have dropped the arguments \mathbf{x} , α for clarity, but the equivalence between the α -expansion χ and the binary labeling \mathbf{b} clearly depends on the initial labeling \mathbf{x} and on α .

²More precisely, it was shown in Ref. 20 to be NP-hard, which means that it is in a class of problems that are widely believed to require exponential time in the worst case.




FIG. 5. The expansion move algorithm. Start with the initial labeling of the image shown by different shades “light”, “medium” and “dark” in (a), assuming only 3 labels. Note that here labels correspond to intensities, although this need not be the case in general. First find the expansion move on the label “dark” that most decreases objective function \mathbf{E} , as shown in (b). Move there, then find the best “light” expansion move, etc. Done when no α -expansion move decreases the cost, for any label α . Corresponding to every expansion move (b) is a binary image (c), where each pixel is assigned 1 if its label changes, and 0 if it does not.

In summary, the problem of computing the α -expansion that minimizes \mathbf{E} is equivalent to finding the \mathbf{b} that minimizes the binary objective function B . The exact form of B will depend on \mathbf{E} . The minimization of \mathbf{E} proceeds via successive binary minimizations corresponding to expansion moves. The binary minimization subroutine is somewhat analogous to the role of line-searching in the conjugate-gradient algorithm, where a local minimum is repeatedly computed over different 1D search spaces. With graph cuts, however, the binary subroutine efficiently computes the global minimum over $2^{|P|}$ candidate solutions, where $|P|$ is the number of pixels in the image. Therefore, in contrast to traditional minimization algorithms, such as conjugate gradients, trust region, simulated annealing, etc. (28), graph cuts can efficiently optimize highly nonconvex objective functions that arise from edge-preserving penalties (19).

Consider a binary objective function of the form

$$B(\mathbf{b}) = \sum_p B_1(b_p) + \sum_{p,q} B_2(b_p, b_q) \quad [17]$$




FIG. 6. Graph construction to minimize a given binary objective function. The graph consists of nodes corresponding to image voxels, as well as two special “terminal” nodes \mathbf{S} and \mathbf{T} . Edges between nodes represent single-voxel (B_1) and pairwise (B_2) cost terms. A minimum cut on this graph solves the binary minimization problem. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Here B_1 and B_2 are functions of binary variables. The difference is that B_1 depends on a single pixel, while B_2 depends on pairs of pixels. Graph-cut methods minimize B by reducing the computation of a minimum cut on an appropriately constructed graph. The graph consists of nodes that are voxels of the image as well as two special terminal nodes, as shown in Fig. 6. The voxel nodes are labeled p, q, r , etc., and terminal nodes are indicated as S and T . All nodes are connected to both terminals via edges, each of which have weights obtained from the B_1 terms above. Nodes are also connected to each other via edges with weights obtained from the pairwise interaction term B_2 .

One can solve the binary optimization problem by finding the minimum cut on this graph (20). A cut is defined as a partition of the graph into two connected subgraphs, each of which contains one terminal. The minimum cut minimizes the sum of the weights of the edges between the subgraphs. Fast algorithms to find the minimum cut using max-flow methods (26) are available. It was shown in Ref. 19 that the class of B that can be minimized exactly by computing a minimum cut on such a graph satisfies the condition

$$B_2(0, 0) + B_2(1, 1) \leq B_2(1, 0) + B_2(0, 1) \quad [18]$$

If $B_2(x, y)$ satisfies Eq. [18], then it is said to be “submodular” with respect to x and y , and a function B is called submodular if it consists entirely of submodular terms³. Single-variable terms of the form of B_1 are always submodular. We will refer to the set of all pixel pairs for which $B_2(b_p, b_q)$ are submodular as the submodular set S .

Previous applications of graph cuts were designed for diagonal \mathbf{E} . This leads to no cross terms, and thus B_1 comes solely from the data penalty and B_2 comes only from the smoothness penalty. It was shown in Ref. 20 that if the separation cost is a metric, then B_2 satisfies Eq. [18]. Many edge-preserving separation costs are metrics, including the truncated cost V used here and shown in Fig. 2b (see Ref. 20 for details). However, the situation is more complicated in parallel MR reconstruction, where \mathbf{E} is nondiagonal. As

³Some authors (21) use the term “regular” instead of “submodular.”

a result, the data penalty has pairwise interactions due to the presence of the cross terms $d(p, p')$ in Eq. [15]. This also follows from Fig. 1, which shows how this data penalty arises from the joint effect of both aliasing voxels p and p' in the image.

It was previously shown (35) that the binary optimization problem arising from Eq. [15] is in general submodular only for a small subset of all cross terms. This necessitates the use of a subroutine (from Ref. 36) to accommodate cross terms arising in MR reconstruction, as described in the next section.

MATERIALS AND METHODS

New MR Reconstruction Algorithm Based on Graph Cuts

The subroutine we use to find a good expansion move is closely related to relaxation methods for solving integer programming problems (26). In these methods, if the linear programming solution obeys the integer constraints, it solves the original integer problem. We compute an expansion move by applying the algorithm of Hammer et al. (36), which was introduced in the field of computer vision in early 2005 by Kolmogorov and Rother (39). The theoretical analysis and error bounds presented in Ref. 40 help explain the strong performance of this construction for MR reconstruction.

For each pixel p , we have a binary variable b_p that is 1 if p acquired the new label α , and 0 otherwise. We introduce a new binary variable \tilde{b}_p , which has the opposite interpretation (i.e., it will be 0 if p acquires the new label α , and 1 otherwise). We call a pixel “consistent” if $\tilde{b}_p = 1 \mp b_p$. Instead of our original objective function $B(\mathbf{b})$, we minimize a new objective function $\tilde{B}(\mathbf{b}, \tilde{\mathbf{b}})$, where $\tilde{\mathbf{b}}$ is the set of new binary variables $\tilde{\mathbf{b}} = \{\tilde{b}_p | p \in P\}$. $\tilde{B}(\mathbf{b}, \tilde{\mathbf{b}})$ is constructed so that $\tilde{\mathbf{b}} = \mathbf{1} - \mathbf{b} \Rightarrow \tilde{B}(\mathbf{b}, \tilde{\mathbf{b}}) = B(\mathbf{b})$ (in other words, if every pixel is consistent, the new objective function will be the same as the old one). Specifically, we define our new objective function by

$$\begin{aligned} 2 \cdot \tilde{B}(\mathbf{b}, \tilde{\mathbf{b}}) = & \sum_p (B_1(b_p) + B_1(1 - \tilde{b}_p)) + \sum_{(p,q) \in \mathcal{G}} (B_2(b_p, b_q) \\ & + B_2(1 - \tilde{b}_p, 1 - \tilde{b}_q)) + \sum_{(p,q) \notin \mathcal{S}} (B_2(b_p, 1 - \tilde{b}_q) \\ & + B_2(1 - \tilde{b}_p, b_q)) \quad [19] \end{aligned}$$

Here the functions $B_1(\cdot)$ and $B_2(\cdot)$ come from our original objective function B in Eq. [17].

Importantly, our new objective function $\tilde{B}(\mathbf{b}, \tilde{\mathbf{b}})$ is submodular. The first summation only involves $B_1(\cdot)$, while for the remaining two terms simple algebra shows that

$$\begin{aligned} B_2(b, b') \text{ is submodular} & \Rightarrow B_2(1 - b, 1 - b') \text{ is submodular} \\ B_2(b, b') \text{ is non-submodular} & \Rightarrow \text{both } B_2(b, 1 - b') \\ & \text{and } B_2(1 - b, b') \text{ are submodular} \end{aligned}$$

As a result, the last two summations in Eq. [19] contain only submodular terms. Thus $\tilde{B}(\mathbf{b}, \tilde{\mathbf{b}})$ is submodular, and can be easily minimized using the binary graph-cut subroutine. In summary, minimizing $\tilde{B}(\mathbf{b}, \tilde{\mathbf{b}})$ is exactly equiv-

alent to minimizing our original objective function B , as long as we obtain a solution in which every pixel is consistent. We note that our technique is not specific to MR reconstruction, but can compute the MAP estimate of an arbitrary linear inverse system under an edge-preserving prior G_{EP} .

While we cannot guarantee that all pixels are consistent, in practice this is true for the vast majority of pixels (typically well over 95%). In our algorithm we simply allow pixels that are not consistent to keep their original labels rather than acquire the new label α . However, even if there are pixels that are not consistent, this subroutine has some interesting optimality properties. It is shown in Ref. 36 that any pixel that is consistent is assigned its optimal label. As a result, our algorithm finds the optimum expansion move for the vast majority of pixels.

Convergence Properties

We investigated the convergence properties of the proposed technique using simulated data from a Shepp-Logan phantom, with intensities quantized to integer values between 0 and 255. We computed the objective function (Eq. [15]) achieved after each iteration for $3 \times$ acceleration and eight coils.

In Vivo Experiments: Setup and Parameters

High-field-strength (4 Tesla) structural MRI brain data were obtained using a whole-body scanner (Bruker/Siemens Germany) equipped with a standard birdcage, eight-channel, phased-array, transmit/receive head coil localized cylindrically around the superior-inferior (S-I) axis. Volumetric T_1 -weighted images ($1 \times 1 \times 1 \text{ mm}^3$ resolution) were acquired using a magnetization-prepared rapid gradient-echo (MPRAGE) sequence with TI/TR = 950/2300 ms timing and a flip angle of 8° . The total acquisition time for an unaccelerated data set was about 8:00 min. In a separate study, images of the torso region were acquired using a gradient-echo sequence with a flip angle of 60° and TE/TR of 3.3/7.5 ms on a GE 1.5T Excite-11 system. Several axial and oblique slices of full-resolution data (256×256) were acquired with an eight-channel upper body coil arranged cylindrically around the torso.

To allow quantitative and qualitative performance evaluations, we acquired all data at full resolution and with no acceleration. The aliased images for acceleration factors of 3–5 were obtained by manually undersampling in k -space. In each case we also computed the full rooted sum of squares (RSOS) image after dividing the coil data by the relative sensitivity maps obtained from calibration lines. We used the self-calibrating strategy for sensitivity estimation, whereby the center of k -space is acquired at full density and used to estimate low-frequency (relative) sensitivity maps. We used the central 40 densely sampled calibration lines for this purpose. These lines were multiplied by an appropriate Kaiser-Bessel window to reduce ringing and noise, zero-padded to full resolution, and transformed to the image domain. We estimated the relative sensitivity by dividing these images by their RSOS. To avoid division by zero, we introduced a small threshold in the denominator, which amounted to 5% of the maximum

intensity. This also served effectively to make the sensitivity maps have zero signal in background regions. However, further attempts to segment background/foreground from these low-frequency data proved unreliable in some cases, and we did not implement background segmentation.

Algorithmic parameters were chosen empirically. It was sufficient to quantize intensity labels to $N_{labels} = 256$, since the resulting quantization error is much smaller than observed noise. Since the computational cost of EPIGRAM grows linearly with N_{labels} , fewer labels are preferable. Model parameters were varied (geometrically) over the range $K \in [N_{labels}/20, N_{labels}/2]$, $\lambda \in [0.01 \cdot \max(x), 1 \cdot \max(x)]$ to find the best values. However, we found that the performance was rather insensitive to these choices, and therefore used the same parameters for all cases shown in this paper. Graph-cut algorithms are typically insensitive to initialization issues, and we chose the zero image as an initial guess. All reconstructions were obtained after 20 iterations. Regularized SENSE reconstruction was for comparison with our method. We chose the regularization factor μ after visually evaluating image quality with a large range of values in the region $\mu \in [0.01, 0.6]$. The images that gave the best results are shown in the next section, along with those obtained with a higher regularization. We obtained the latter to observe the noise vs. aliasing performance of SENSE.

Quantitative Performance Evaluation

In addition to the visual evidence presented in the next section, we conducted a quantitative performance evaluation of the reconstructed in vivo data. For in vivo data the problem of ascertaining noise estimates or other performance measures is challenging due to the nonavailability of an accurate reference image. Unfortunately, none of the reconstruction methods we implemented (RSOS, regularized SENSE, and EPIGRAM) are unbiased estimators of the target. This makes it difficult to directly estimate noise performance, and the traditional root mean square error (RMSE) becomes inadequate. We follow instead a recent evaluation measure for parallel imaging methods proposed by Reeder et al. (37), which provides an unambiguous and fair comparison of SNR and geometry factor. Two separate scans of the same target with identical settings are acquired, and their sum and difference are obtained. The local signal level at each voxel is computed by averaging the sum image over a local window, and the noise level is obtained from the standard deviation (SD) of the difference image over the same window. Then

$$SNR = \frac{\text{mean}(\text{Sum image})}{\sqrt{2} \text{stdev}(\text{Diff image})}$$

Here the mean and SD are understood to be over a local window (in this case a 5×5 window around the voxel in question). This provides unbiased estimates that are directly comparable across different reconstruction methods. We perform a similar calculation, with a crucial difference: instead of acquiring two scans, we use a single scan but add random uncorrelated Gaussian noise in the

coil outputs to obtain two noisy data sets. This halves the acquisition effort without compromising estimate quality, since uncorrelated noise is essentially what the two-scan method also measures. Reeder et al. (37) explained that their method can be erroneous for in vivo data because motion and other physiological effects can seriously degrade noise estimates. While our modification achieves the same purpose, it does not suffer from this problem, and hence should be more appropriate for in vivo data. The SNR calculation also allows the geometry factor, or g -factor, maps for each reconstruction method to be obtained. For each voxel p and acceleration factor R :

$$g(p, R) = \frac{SNR_{RSOS}(p)}{SNR(p) \sqrt{R}}$$

where SNR_{RSOS} is the SNR of the RSOS reconstruction. The SNR and g -map images appeared quite noisy due to the division step and estimation errors, and we had to smooth them for better visualization.

Comparing EPIGRAM and Regularized SENSE

SNR and g -factor maps obtained by the method described by Reeder et al. (37) and above are well suited for unbiased reconstructions like unregularized SENSE because they provide comparable measures of noise amplification. Reeder et al. (37) demonstrated conclusively that the two-scan method gives comparable estimates to those obtained voxelwise by using hundreds of identical scans. Unfortunately, neither the SNR nor the g -factor can serve as a single measure of performance in the case of biased reconstructions like regularized SENSE and EPIGRAM, or indeed any Bayesian estimate. This is because there are now two quite different sources of degradation: noise amplification and aliasing. While the former can be measured accurately by SNR and g -maps using the above technique, the latter cannot. All regularization or Bayesian approaches come with external parameters, which can be thought of as “knobs” available to the user (μ for SENSE and λ for EPIGRAM). Depending on how much the user tweaks those knobs, it is possible to achieve any desired degree of noise performance. To illustrate this, we show in the Results section the effect of μ , the regularization parameter in SENSE. We demonstrate that any desired g -factor can be achieved by making μ large enough. Therefore, for a proper comparison of reconstruction methods, it is important to fix either aliasing quality or noise quality, and then compare the other measure. In this section we adopt the following approach: we turn the “knobs” until both EPIGRAM and regularized SENSE produce approximately similar g -maps and mean g -values. We then visually evaluate how each method performed in terms of aliasing and overall reconstruction quality. In each case we tabulate the mean SNR and g -values achieved by the given settings. The results are given below.

RESULTS

Convergence Result

The convergence behavior of the objective function (Eq. [15]) against the number of iterations for $R = 3$, $L = 8$ is




FIG. 7. Convergence behavior of the modified graph-cut algorithm (EPIGRAM) for acceleration factor $R = 3$, $L = 8$ coils, on GRE torso data. The vertical axis shows the value of objective function (Eq. [15]) achieved after each outer iteration, which represents a single cycle through $N_{labels} = 256$ expansion moves. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

shown in Fig. 7. As is typical of graph-cut algorithms (31), most of the improvement takes place in the first few (approximately five) iterations. In our initial implementation of EPIGRAM, which was written primarily in MATLAB, an iteration takes about 1 min on this data set. Since EPIGRAM is almost linear in the number of nodes, the

total running time approximately scales as N^2 , the image size.

In Vivo Results

The best parameter values were consistent across all in vivo data sets we tried: $K = N_{labels}/7$, $\lambda = 0.04 \cdot \max(x)$. The results of several brain imaging experiments with these parameter values are displayed in Figs. 8 and 9. Figure 8 shows the reconstruction of a MPRAGE scan of a central sagittal slice, with an undersampling factor $R = 4$ along the anterior–posterior (A-P) direction. The RSOS reference image is shown in Fig. 8a, regularized SENSE with (empirically obtained optimal) $\mu = 0.08$ is shown in b, regularized SENSE with $\mu = 0.16$ is shown in c, and EPIGRAM is shown in d. Reduced noise is visually noticeable in the EPIGRAM reconstruction compared to both SENSE reconstructions. Higher regularization in SENSE caused unacceptable aliasing, as observed in Fig. 8c. We note that the unregularized (i.e., standard) SENSE results were always worse than those of regularized SENSE, and consequently are not shown. Another sagittal scan result is shown in Fig. 9, this time from the left side of the patient. Image support is smaller, allowing $5\times$ acceleration. The optimally regularized SENSE output (b) with $\mu = 0.1$ is noisy at this level of acceleration, and $\mu = 0.2$ (c) introduced significant aliasing, especially along the central brain region. EPIGRAM (d) exhibits some loss of texture, but on the whole appears to outperform SENSE.

A set of torso images acquired on a GE 1.5T scanner using a GRE sequence and acceleration factor of 3 (along





FIG. 8. Brain A: In vivo brain result with $R = 4$, $L = 8$. Views were acquired vertically. **a**: Reference image. **b**: SENSE regularized with $\mu = 0.08$. **c**: SENSE regularized with $\mu = 0.16$. **d**: EPIGRAM reconstruction.

FIG. 9. Brain B: In vivo brain result with $R = 5, L = 8$. Views were acquired vertically. **a**: Reference image. **b**: SENSE regularized with $\mu = 0.1$. **c**: SENSE regularized with $\mu = 0.2$. **d**: EPIGRAM reconstruction.



the A-P direction) were resolved from 40 sensitivity calibration lines (Fig. 10). Various slice orientations (both axial and oblique) were used. These data also show the practical limitation of SENSE when an inadequate number of calibration lines are used for sensitivity estimation. The reconstruction quality of SENSE is poor as a result of the combination of ill-conditioning of the matrix inverse and calibration error. SENSE exhibits both high noise and residual aliasing. In fact, what appears at first sight to be uncorrelated noise is, upon finer visual inspection, found to arise from unresolved aliases, as the background in Fig. 10b clearly indicates. EPIGRAM was able to resolve aliasing correctly and suppress noise, without blurring sharp edges and texture boundaries. To demonstrate the performance of these methods more clearly, we show in Fig. 10d–f zoomed-in versions of the images in Fig. 10a–c.

Next we demonstrate the trade-off between noise and aliasing performance. Figure 11 shows another torso slice along with associated g -maps computed as specified in Materials and Methods. We investigate the effect of various regularizations of SENSE. The leftmost column (a and e) shows the SENSE result and its g -map for $\mu = 0.1$. Clearly, there is inadequate noise suppression (some regions of the image have a g -factor as high as 6). It is easy to reduce noise amplification by increasing regularization. In

the next column, results for $\mu = 0.3$ are shown. The g -map has become correspondingly flatter, but the reconstruction indicates that this was achieved at the cost of introducing some aliasing in the image. The rightmost column shows EPIGRAM results, which indicate significantly lower g -values and lower aliasing artifacts. In the third column we show that with an appropriate choice of μ it is possible to match the EPIGRAM g -map (c.f., Fig. 11g and h). However, with this choice of $\mu = 0.5$, SENSE yields unacceptable aliasing. Table 1 shows the mean SNR and g -factor values for EPIGRAM and various regularizations of SENSE.

We observe that the regularization needed in SENSE to match the EPIGRAM g -values (approximately $\mu = 0.5$ in almost all cases) yields unacceptable aliasing. Instead, a more modest μ must be chosen empirically. The standard RMSE criterion does not really help here. In Fig. 12 we plot the average over all our torso data of the RMSE for various regularization factors. For this data set, the optimum μ in terms of RMSE was found to be around 0.25, although in practice the best visual quality was observed below this value, at around 0.1–0.15. This suggests that the RMSE measure does not capture reconstruction quality very accurately, and in particular seems to underemphasize the effect of residual aliasing. In all experimental data shown in this section, the optimum regularization for




FIG. 10. Cardiac A: In vivo result with $R = 3$, $L = 8$. Views were acquired horizontally. **a**: Reference image. **b**: SENSE regularized with $\mu = 0.16$ (optimal). **c**: EPIGRAM reconstruction. Zoomed-in portions of a–c are shown in **d–f**.

SENSE was obtained empirically for each data set by visual inspection as much as possible. To give an idea of the visual quality corresponding to a certain μ value, Fig. 12 also shows a portion of the resulting SENSE reconstruction. The mean SNRs and g -factors for all reconstruction examples presented here are summarized in Table 2. Non-regularized SENSE data are not shown, because they were always worse than regularized SENSE data.

In all of these examples there are regions in the EPIGRAM data with some loss of texture, as well as regions of low signal that appear to have been “washed out” by a uniform value. Both effects result from the piecewise smooth assumption imposed by the prior. We discuss ways to mitigate this problem below, but note that it is typical of most applications in which MRF priors are applied.




FIG. 11. Cardiac B: In vivo result with $R = 3$, $L = 8$, showing the effect of regularization on g -factor maps. Views were acquired horizontally. **a** and **e**: SENSE reconstruction and its g -map for $\mu = 0.1$. There is inadequate noise suppression, with a g -factor as high as 6. Results for $\mu = 0.3$ are shown in **b** and **f**. The g -maps has become correspondingly flatter, but at the cost of aliasing. The rightmost column (**d** and **h**) shows EPIGRAM results, which indicate significantly lower g -values as well as lower aliasing artifacts. SENSE with $\mu = 0.5$ (**c** and **g**) matches the EPIGRAM g -map but yields unacceptable aliasing.

Table 1
Mean Reconstructed SNR and Mean g -Factor Values of Reconstructed Data for Various Regularization Factors, with $R = 3$ and $L = 8^*$

	Cardiac B		Brain A	
	Mean SNR	Mean g	Mean SNR	Mean g
RSOS	42	1.0	59	1.0
SENSE, $\mu = 0.1$	15	3.3	21	3.7
SENSE, $\mu = 0.3$	18	2.3	26	2.5
SENSE, $\mu = 0.5$	24	1.4	30	1.9
EPIGRAM	36	1.4	40	1.85

*Two datasets, Cardiac B and Brain A were used. The means are over object foreground. RSOS mean g value is 1 by definition.

DISCUSSION AND CONCLUSIONS

The experimental data shown above indicate that EPIGRAM shows promise for providing improved performance and visual quality. This is a direct result of our use of a powerful EPP model, and the efficient implementation of the associated Bayesian problem using graph-cut algorithms. An exhaustive clinical evaluation is needed to further characterize the method and identify specific clinical applications, and such work is ongoing. It is also desirable to develop appropriate quantitative measures for comparing the reconstruction performance of various methods. Conventional measures, such as SNR and g -factor maps, do not serve this purpose well. We have argued that looking solely at the reduction in noise as a measure of regularization performance is inadequate because it does not account for the introduction of concomitant aliasing artifacts (e.g., see Fig. 11). In fact, g -factor maps on their own might mislead one to conclude that large regularizations are always good because they produce small g -factors. Further work is therefore needed to devise more

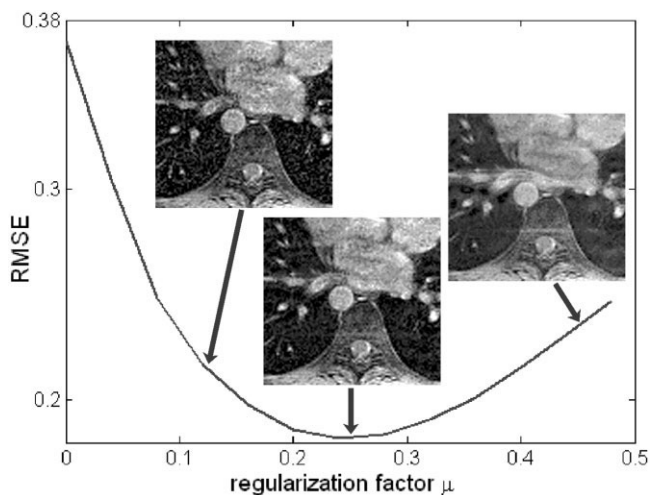


FIG. 12. Average RMSE of regularized SENSE reconstruction of torso data. The optimum was observed at around 0.25, although visual quality was found to be best for smaller values, in the region [0.1–0.2]. Three example torso reconstructions are also shown. Observe that as regularization increases, residual aliasing artifacts and blurring both become worse. Also note that nonregularized SENSE (i.e., $\mu = 0$) gives a substantially worse RMSE.




FIG. 13. Plot of the SNRs achieved by EPIGRAM and SENSE for various data sets. SNR values were taken from Table 2. Observe that large μ gives higher SNR, but, as demonstrated by the displayed images, this comes at the cost of unacceptable aliasing. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

suitable quantitative measures of performance, such as a statistically sound diagnostic score obtained blindly from a panel of experts.

We currently do not consider sensitivity errors or temporal information, but such information can be easily incorporated within our objective function. In particular, the use of graph-cut minimization for temporal or multidimensional data is very attractive. The graph-cut-based cost-minimization approach should be readily generalizable to other parallel imaging methods, such as GRAPPA. Since GRAPPA also performs essentially a least-squares fit via the optimization of some objective function, its implementation using graph-cut-based minimization appears possible. Currently the method is implemented only for Cartesian sampling. Extending EPIGRAM to non-Cartesian trajectories is an interesting but challenging task.

A shortcoming of our method is the possible loss of texture due to the piecewise smooth assumption. In this work the model parameters K, λ were determined heuristically and assumed to be constant over the entire image. An inappropriate choice of either of these parameters can adversely affect the reconstructed data. If λ is too large, it will force low signal areas to become uniformly zero, or remove informative texture from nonsmooth regions. Both of these effects can be observed in our example reconstructions. Decreasing λ can remove this to some extent, but the most effective approach may be an adaptive and locally varying choice of parameters. Although our experimentation produced the best results for the chosen truncated linear prior model, other models (e.g., truncated quadratic or piecewise linear) are possible. The power of the truncated linear model appears to be well substantiated by research in computer vision; however, future work in MR could conceivably yield better edge-preserving models that preserve textures and low signal regions. Further work on different prior models and their parameters will most likely improve the results shown here. However, it is important to note that the EPIGRAM algorithm can continue to be used in its current form because it is formulated for arbitrary objective functions of the kind shown in Eq. [15].

Table 2
Quantitative Performance Measures for Reconstruction Examples*

	R	SENSE, optimal μ		SENSE, large μ		EPIGRAM	
		Mean SNR	Mean g	Mean SNR	Mean g	Mean SNR	Mean g
Brain A	4	8	4.6	13	2.4	23	1.7
Brain B	5	10	3.5	13	2.4	17	2.2
Cardiac A	3	20	2.3	25	1.6	33	1.5
Cardiac B	3	15	3.3	18	2.3	36	1.4

*Both mean reconstructed SNR and mean g -factor values are shown. The means are over object foreground wherever the latter was available. Corresponding reconstructed images were presented in Figs. 8–12. Note that numbers for large μ shown below are closer to EPIGRAM numbers, but produce images with considerably more aliasing.

It is also noteworthy that Bayesian reconstruction can sometimes produce images that appear different from or even better than the best available reconstruction (in our case the RSOS). This is an interesting case of a Bayesian estimate fitting the prior better than the best RSOS estimate. This is a well known phenomenon in Bayesian estimation, and can be explained by the fact that our MRF prior models the target image—not the RSOS reconstruction. Another example of this can be seen in the g -factor maps of EPIGRAM, which show some regions with $g < 1$. This counterintuitive phenomenon results from the fact that the prior imposes smoothness to such an extent that the Bayesian estimate becomes less noisy than the unaccelerated image in some places. In the extreme case when $\mu, \lambda = \infty$, both SENSE and EPIGRAM will simply reproduce their respective priors and we will have $g = 0$ everywhere! Of course, such a result will be of no real use since it is independent of the observed data.

Comparison With Other Prior-Driven Approaches

It is instructive to compare our EPP model with other prior models. Some very interesting and impressive results from dynamic imaging models (e.g., RIGR) were presented in Refs. 17 and 18, which used a powerful generalized series to model k -space truncation, and dynamic and static parts of the MR data. Our Bayesian approach is an interesting counterpoint to that body of work. RIGR imposes a priori constraints in the form of a detailed generalized series model. This approach works best for dynamic imaging and requires a good, well-registered, high-resolution reference image. This technique certainly has significant value for many imaging situations, but it suffers from high sensitivity to modeling errors, patient motion, and the difficulty of obtaining correct model parameters from partial data. Our prior model is simpler and does not need a reference image. The advantage of our approach is that it is basically parameter-free, apart from the particular choice of separation cost. Since the EPP model relies solely on spatial coherence, it is effective under widely varying imaging conditions and should be robust against artifacts caused by geometric distortion, B_1 intensity variations, excessive noise, and even small amounts of motion. Furthermore, in theory, it appears to be possible to incorporate RIGR-type priors within the minimization framework proposed here. Another interesting approach that has the potential to improve conventional regularization schemes is the total variation (TV) method (38). Both EPIGRAM and TV are novel and basically unproven techniques for MR, and it

will be interesting to analytically and empirically compare the two.

In conclusion, we have argued that Bayesian estimation with edge-preserving image priors is an effective means of breaking out of the fundamental noise vs. aliasing trade-off in conventional parallel imaging without causing over-smoothing. We showed that this Bayesian estimate is a minimization problem that is closely related to problems that have been addressed in the computer-vision community. We demonstrated that standard graph-cut optimization methods (19,20) cannot be applied directly to our problem, and instead proposed a new MR reconstruction technique relying on a subroutine based on Ref. 36. The resulting MR reconstruction algorithm, which we call EPIGRAM, appears promising both visually and quantitatively. The resulting gains can potentially be exploited to further accelerate acquisition speed in parallel imaging, or to achieve higher reconstruction quality at the same speed.

ACKNOWLEDGMENTS

We thank Vladimir Kolmogorov for spending a considerable amount of time explaining the algorithm of Ref. 36; his help was invaluable in clarifying this important work. We also thank Lara Stables and Xiaoping Zhu for help with data acquisition.

REFERENCES

1. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med* 1999;42:952–962.
2. Pruessmann KP, Weiger M, Boernert P, Boesiger P. Advances in sensitivity encoding with arbitrary k -space trajectories. *Magn Reson Med* 2001;46:638–651.
3. Weiger M, Pruessmann KP, Boesiger P. 2D SENSE for faster 3D MRI. *Magn Reson Mater Biol Phys Med* 2002;14:10–19.
4. Sodickson DK, Manning WJ. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magn Reson Med* 1997;38:591–603.
5. Sodickson DK, McKenzie CA, Ohliger MA, Yeh EN, Price MD. Recent advances in image reconstruction, coil sensitivity calibration, and coil array design for SMASH and generalized parallel MRI. *Magn Reson Mater Biol Phys Med* 2002;13:158–163.
6. Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med* 2002;47:1202–1210.
7. Blaimer M, Breuer F, Mueller M, Heidemann RM, Griswold MA, Jakob PM. SMASH, SENSE, PILS, GRAPPA: how to choose the optimal method. *Top Magn Reson Imaging* 2004;15:223–236.
8. Wang Y. Description of parallel imaging in MRI using multiple coils. *Magn Reson Med* 2000;44:495–499.

9. Raj A, Zabih R. An optimal maximum likelihood algorithm for parallel imaging in presence of sensitivity noise. In: Proceedings of the 12th Annual Meeting of ISMRM, Tokyo, Japan, 2004 (Abstract 1771).
10. Liang ZP, Bammer R, Ji J, Pelc N, Glover G. Making better SENSE: wavelet denoising, Tikhonov regularization, and total least squares. In: Proceedings of the 10th Annual Meeting of ISMRM, Honolulu, HI, USA, 2002 (Abstract 2388).
11. Lin F, Kwang K, Belliveau J, Wald L. Parallel imaging reconstruction using automatic regularization. *Magn Reson Med* 2004;51:559–567.
12. Bammer R, Auer M, Keeling SL, Augustin M, Stables LA, Prokesch RW, Stollberger R, Moseley ME, Fazekas F. Diffusion tensor imaging using single-shot SENSE-EPI. *Magn Reson Med* 2002;48:128–136.
13. Ying L, Xu D, Liang ZP. On Tikhonov regularization for image reconstruction in parallel MRI. *Proc IEEE EMBS* 2004:1056–1059.
14. Tsao J, Boesiger P, Pruessmann KP. K-T BLAST and K-T SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations. *Magn Reson Med* 2003;50:1031–1042.
15. Hansen MS, Kozerke S, Pruessmann KP, Boesiger P, Pedersen EM, Tsao J. On the influence of training data quality in K-T BLAST reconstruction. *Magn Reson Med* 2004;52:1175–1183.
16. Tsao J, Kozerke S, Boesiger P, Pruessmann KP. Optimizing spatiotemporal sampling for K-T BLAST and K-T SENSE: application to high-resolution real-time cardiac steady-state free precession. *Magn Reson Med* 2003;53:1372–1382.
17. Chandra S, Liang ZP, Webb A, Lee H, Morris HD, Lauterbur PC. Application of reduced-encoding imaging with generalized-series reconstruction (RIGR) in dynamic MR imaging. *J Magn Reson Imaging* 1996; 6:783–797.
18. Hanson JM, Liang ZP, Magin RL, Duerk JL, Lauterbur PC. A comparison of RIGR and SVD dynamic imaging methods. *Magn Reson Med* 1997; 38:161–167.
19. Kolmogorov V, Zabih R. What energy functions can be minimized via graph cuts? *IEEE Trans Patt Anal Machine Intel* 2004;26:147–159.
20. Boykov Y, Veksler O, Zabih R. Fast approximate energy minimization via graph cuts. *IEEE Trans Patt Anal Machine Intel* 2001;23:1222–1239.
21. Poggio T, Torre V, Koch C. Computational vision and regularization theory. *Nature* 1985;317:314–319.
22. Besag J. On the statistical analysis of dirty pictures (with discussion). *J R Stat Soc Ser B* 1986;48:259–302.
23. Li S. Markov random field modeling in computer vision. Berlin: Springer-Verlag; 1995.
24. Geman S, Geman D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans Patt Anal Machine Intel* 1984;6:721–741.
25. Chellappa R, Jain A. Markov random fields: theory and applications. New York: Academic Press; 1993.
26. Cook WJ, Cunningham WH, Pulleyblank WR, Schrijver A. Combinatorial optimization. New York: Wiley; 1998.
27. Scharstein D, Szeliski R. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int J Comput Vis* 2002;47:7–42.
28. Press W, Teukolsky S, Vetterling W, Flannery B. Numerical recipes in C. 2nd ed. Cambridge: Cambridge University Press; 1992.
29. Raj A, Zabih R. MAP-SENSE: maximum a posteriori parallel imaging for time-resolved 2D MR angiography. In: Proceedings of the 13th Annual Meeting of ISMRM, Miami Beach, FL, USA, 2005 (Abstract 1901).
30. Greig D, Porteous B, Seheult A. Exact maximum a posteriori estimation for binary images. *J R Stat Soc Ser B* 1989;51:271–279.
31. Boykov Y, Kolmogorov V. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans PAMI* 2004;26:1124–1137.
32. Kwatra V, Schofield A, Essa I, Turk G, Bobick A. Graphcut textures: image and video synthesis using graph cuts. *ACM Trans Graph* 2003; 22:277–286.
33. Boykov Y, Jolly MP. Interactive organ segmentation using graph cuts. *Proc Med Image Comput Comput Assist Interv* 2000:276–286.
34. Kim J, Fisher J, Tsai A, Wible C, Willsky A, Wells W. Incorporating spatial priors into an information theoretic approach for fMRI data analysis. *Proc Med Image Comput Comput Assist Interv* 2000:62–71.
35. Raj A, Zabih R. A graph cut algorithm for generalized image deconvolution. *Proc Int Conf Comput Vis* 2005:1901.
36. Hammer P, Hansen P, Simeone B. Roof duality, complementation and persistency in quadratic 0-1 optimization. *Math Programming* 1984;28: 121–155.
37. Reeder SB, Wintersperger BJ, Dietrich O, Lanz T, Greiser A, Reiser MF, Glazer GM, Schoenberg SO. Practical approaches to the evaluation of signal-to-noise ratio performance with parallel imaging: application with cardiac imaging and a 32-channel cardiac coil. *Magn Reson Med* 2005;54:748–754.
38. Persson M, Bone D, Elmqvist H. Total variation norm for three-dimensional iterative reconstruction in limited view angle tomography. *Phys Med Biol* 2001;46:853–866.
39. Kolmogorov V, Rother C. Minimizing non-submodular functions with graph cuts: a review. Microsoft Research Technical Report MSR-TR-2006-100, July 2006. Redmond, WA: Microsoft; 2006.
40. Raj A, Singh G, Zabih R. MRIs for MRFs: Bayesian reconstruction of MR images via graph cuts. *IEEE Computer Vision and Pattern Recognition Conference*, June 17–22, 2006. IEEE Computer Society, p 1061–1069.