This paper, as originally published in CVPR 2006, contained a serious error. The authors inadvertently did not acknowledge Vladimir Kolmogorov for a significant contribution: in early 2005, he introduced the 1984 construction of Hammer *et al.* into computer vision, and explained it to us.

For a very clear description of the Hammer *et al.* construction, see:

> V. Kolmogorov and C. Rother, “Minimizing non-submodular functions with graph cuts - a review”. Microsoft Research Technical Report MSR-TR-2006-100, July 2006.

The enclosed revision of our paper has been updated to provide proper credit. We apologize for the oversight.

Ramin Zabih
July, 2006

MRF's for MRI's: Bayesian Reconstruction of MR Images via Graph Cuts

Ashish Raj
UC San Francisco
ashish.raj@ucsf.edu

Gurmeet Singh
Cornell University
gs244@cornell.edu

Ramin Zabih
Cornell University
rdz@cs.cornell.edu

Abstract

Markov Random Fields (MRF's) are an effective way to impose spatial smoothness in computer vision. We describe an application of MRF's to a non-traditional but important problem in medical imaging: the reconstruction of MR images from raw fourier data. This can be formulated as a linear inverse problem, where the goal is to find a spatially smooth solution while permitting discontinuities. Although it is easy to apply MRF's to the MR reconstruction problem, the resulting energy minimization problem poses some interesting challenges. It lies outside of the class of energy functions that can be straightforwardly minimized with graph cuts. We point out that a graph construction due to Hammer et al. *[3] allows graph cuts to be applied to this problem, and provide some theoretical analysis of the properties of our algorithm. Experimentally, our method gives very strong performance, with a substantial improvement in SNR when compared with widely-used methods for MR reconstruction.*

1. Introduction

Magnetic resonance (MR) imaging has great importance, both for clinical and for research applications, due to its safety and flexibility. The MR imaging process, however, imposes a fundamental tradeoff between image quality and scan time [7]. It is very important to reduce scan time, for a number of reasons. The primary issue is that MR is very sensitive to motion artifacts, so reducing scan times decreases the odds of obtaining a poor-quality image. Some parts of the body, notably the heart and lungs, undergo periodic motion; as a result, only sufficiently fast scans can be used to image these organs, or any other adjacent tissues. In addition, there are obvious issues of patient comfort as well as cost that arise from lengthy scans. Finally, if the basic image acquisition process were faster, even for a fixed scan duration a higher signal to noise ratio could be obtained [7]. This is important for many important modalities, such as fMRI, perfusion, diffusion or time-resolved angiography, where the physical process being imaged takes place over short time periods.

There are many opportunities for better algorithms to result in faster (and hence better) MR. The image formation process in MR is quite unlike a conventional camera, and this in turn leads to a number of interesting computational challanges. In MR, an image is acquired as its fourier transform (usually called "k-space"), and this acquisition takes place serially. Typically one row in k-space is acquired at a time, and each row takes tens of milliseconds to acquire (the details depend on the particular study being performed).

A particularly important technique for accelerating MR scans is called *parallel imaging*, which uses multiple receiver coils. According to a recent survey article "Parallel imaging is one of the most promising recent advances in MRI technology and has revolutionized MR imaging" [1]. While parallel imaging requires solving a linear inverse problem, existing methods [8, 9] either do not incorporate spatial priors, or assume that the image is globally smooth. In computer vision, of course, Markov Random Fields (MRF's) are commonly used to encode better priors.

In this paper we address the parallel imaging problem by applying discontinuity-preserving priors. We begin by giving a precise definition of the parallel imaging problem, and describing the regularization-based methods that have been used to solve it. In section 3 we adopt a Bayesian view of this problem and formulate a discontinuity-preserving prior as an MRF. The energy function for computing the MAP estimate has some interesting properties, but unfortunately it is known to lie outside the class that can easily be minimized via graph cuts [5, 11]. In section 5 we apply the graph cut construction of Hammer *et al.* [3] to effectively minimize the energy. While our approach is similar to that of [11], our algorithm has better theoretical properties, and gives significantly more impressive experimental results as demonstrated in section 6. More images, as well as some additional technical details, are available on the project web page at http://www.cs.cornell.edu/~rdz/SENSE.htm.

2. The parallel imaging problem

The parallel imaging problem is illustrated in figure 1. We can cut the scan time in half by dropping every other column in k-space. The resulting image, when reconstructed from a conven-

tional RF receiver coil, is an aliased image like those shown in the top row of figure 1. Such an image is half the width of the original (unobserved) image shown in the bottom row. It is formed by multiplying the original image by a slowly varying function $S(p)$ and then adding the left half of the resulting image to the right half. The multiplication by $S(p)$ comes from the spatial response of the coil, while the addition of the two image halves is a consequence of dropping alternate k-space columns. Thus, each pixel p in the aliased image is the weighted sum of two pixels (p, p') in the original image, where p' is in the same row but half the width of the image away. The weights come from the function S.

If we only had a single aliased image it would be impossible to reconstruct the original image. However, we can use multiple coils, each of which has a different S, without increasing the scan time. In our example, we have a simple linear system for each pair of pixels p, p':

$$\begin{bmatrix} Y_1(p) \\ Y_2(p) \\ Y_3(p) \end{bmatrix} = \begin{bmatrix} S_1(p) & S_1(p') \\ S_2(p) & S_2(p') \\ S_3(p) & S_3(p') \end{bmatrix} \begin{bmatrix} X(p) \\ X(p') \end{bmatrix}. \quad (1)$$

The functions S are called sensitivity maps, and can be computed in advance. In addition, the sum over the coils is usually normalized: $\sum_{l=1}^{L} S_l^2(p) \approx 1$.

Formally, assuming cartesian sampling in k-space we have an (unobserved) true image X, which we must reconstruct from the sensitivity maps $S_1, \ldots S_l, \ldots S_L$ and the coil outputs $Y_1, \ldots Y_l, \ldots Y_L$. While X and S_l are of size $N \times M$, Y_l is of size $\frac{N}{R} \times M$, where the acceleration factor is R. By stacking the receiver coil outputs into the column vector $\mathbf{y}$ and the image into $\mathbf{x}$, we can write the image formation process as

$$\mathbf{y} = H\mathbf{x}, \quad (2)$$

where the encoding matrix H contains the sensitivity maps.

Least squares can be used to compute the maximum likelihood estimate from equation (2)

$$\hat{\mathbf{x}}_{\text{ML}} = (H^t H)^{-1} H^t \mathbf{y}. \quad (3)$$

This is a famous MR reconstruction algorithm called SENSE [9], which has been widely applied in research and clinical systems. Tikhonov-style regularization methods are also commonly used to handle the ill-conditioning of H. The most general form of these [8] is

$$\hat{\mathbf{x}}_{\text{reg}} = \arg\min_{\mathbf{x}} \left\{ \|\mathbf{y} - H\mathbf{x}\| + \mu^2 \|A(\mathbf{x} - \mathbf{x}_r)\|^2 \right\}. \quad (4)$$

Here, the second term penalizes that are far from a prior reference image $\mathbf{x}_r$, and that are not spatially smooth according to the regularizer matrix A. The regularization parameter μ controls the tradeoff between aliasing and local noise in the output. Note that since A is a linear operator, global spatial smoothness is imposed. There is a simple closed-form solution for this equation,

$$\hat{\mathbf{x}}_{\text{reg}} = \mathbf{x}_r + (H^t H + \mu^2 A^t A)^{-1} H^t (\mathbf{y} - H\mathbf{x}_r). \quad (5)$$

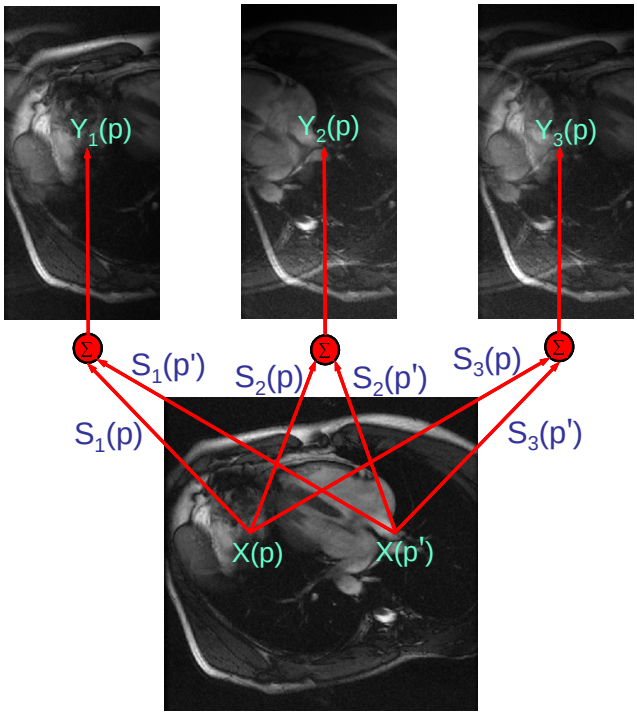

Figure 1. The parallel imaging problem, for 2× acceleration with 3 coils. The unobserved image is shown at bottom, and the aliased image observed from each coil is shown at top. Each pixel in an aliased image is the weighted sum of two pixels in the unobserved image. The weights are different for each pixel and each aliased image. The linear system that relates the unobserved image to the observations is given in equation (1)

2.1. Related work

Every clinical system, and nearly every research paper in MR, computes some variant of $\hat{\mathbf{x}}_{\text{reg}}$ [1]. The original paper on SENSE [9] did not perform regularization (i.e., $\mu = 0$). For numerical reasons, authors such as [8] proposed variants of equation (4), typically using the identity matrix for A. While this does not impose spatial smoothness, it has signficant computational advantages in terms of the structure of H. Even the most general variants of this method, however, can only impose spatial smoothness globally, and give poor performance for acceleration factors significantly greater than 2. This is well-known in the MR community [1], and is confirmed by our experiments in section 6 with SENSE and its variants.

As we will see in section 4, imposing a better model of spatial smoothness leads to an intractable energy minimization problem. Very little research has been done on this topic. The one exception is the work of [11], who address a generalized variant of our problem. After showing that it cannot be solved with previous graph cut energy minimization methods, they present a new graph cut algorithm to solve the problem. While we defer a more detailed discussion of their method until section 6, our method provides better theoretical guarantees and vastly better experimental results.

3. Parallel imaging with MRF priors

In a Bayesian view of the MR reconstruction problem, we are given an observation $\mathbf{y}$, a prior probability distribution over images $\Pr(\mathbf{x})$, and a likelihood $\Pr(\mathbf{y}|\mathbf{x})$ that models the image formation process (2). We seek the estimate that maximizes the posterior probability,

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg\max_{\mathbf{x}} \Pr(\mathbf{x}|\mathbf{y}) \propto \Pr(\mathbf{y}|\mathbf{x})\Pr(\mathbf{x}). \quad (6)$$

Assuming i.i.d. Gaussian noise in the imaging process, the likelihood can be written

$$\Pr(\mathbf{y}|\mathbf{x}) \propto \exp(-\|\mathbf{y} - H\mathbf{x}\|^2). \quad (7)$$

Assuming the prior can be written as $\Pr(\mathbf{x}) \propto \exp(-G(\mathbf{x}))$ for some function G, the estimate we seek minimizes the negative log of the posterior probability,

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg\min_{\mathbf{x}} \|\mathbf{y} - H\mathbf{x}\|^2 + G(\mathbf{x}). \quad (8)$$

In the Tikhonov-style regularization methods shown in equation (4), the second term can be interpreted as a particular choice of G.

3.1. Likelihood function

We can rewrite the first term in equation (8) as

$$\|\mathbf{y} - H\mathbf{x}\|^2 = \sum_p \sum_l \Big[Y_l(p) - \sum_{(p')\in\mathcal{N}_a(p)} S_l(p')X(p') \Big]^2. \quad (9)$$

Here, we have written the set of pixels in X that contribute to $Y_l(p)$ as $\mathcal{N}_a(p)$ (in the example shown in figure 1, $\mathcal{N}_a(p)$ contains p and p', where p' is horizontally separated from p by half the image width). With some algebra, this in turn can be rewritten as a constant plus

$$\sum_{p=1}^{n} f(p)X^2(p) + \sum_{p=1}^{n} g(p)X(p) + 2 \cdot \sum_{(p,p')\in\mathcal{N}_a} h(p,p')X(p)X(p').$$

We have written $\mathcal{N}_a = \{\, (p,p') \mid p' \in \mathcal{N}_a(p), p' \neq p \,\}$. We will refer to these last terms as *cross terms*, and they play an important role in this paper. We defined f, g, h by

$$f(p) = \sum_{l=1}^{L} |S_l(p)|^2, \quad g(p) = -2\sum_{l=1}^{L} |S_l(p)Y_l(p)|,$$
$$h(p,p') = \sum_{l=1}^{L} |S_l(p)S_l(p')|. \quad (10)$$

3.2. Choice of prior

The natural prior on images is a Markov Random Field [6], which flexibly enforces spatial coherence. The most widely used MRF prior in vision is of the form

$$G(\mathbf{x}) = \sum_{(p,q)\in\mathcal{N}_s} V(X(p) - X(q)). \quad (11)$$

Here, $\mathcal{N}_s$ is the neighborhood system over which spatial smoothness is imposed (typically 4- or 8-connected pixel pairs). The separation cost V provides a penalty for adjacent pixels to have different intensities. If V is convex, the prior imposes global smoothness and will blur edges, but $\hat{\mathbf{x}}_{\text{MAP}}$ is much easier to compute.

Our goal is to apply an MRF prior that imposes spatial smoothness while preserving edges. Such priors are usually called robust, or discontinuity-preserving, and rely on some non-convex separation cost. While there are many good choices for V, the primary challenge is to ensure that we can efficiently compute $\hat{\mathbf{x}}_{\text{MAP}}$.

4. Energy minimization via graph cuts

To find a discontinuity-preserving solution to the parallel imaging problem, we need to minimize the energy

$$E(\mathbf{x}) = \sum_{p=1}^{n} f(p)X^2(p) + \sum_{p=1}^{n} g(p)X(p) + 2 \cdot \sum_{(p,p')\in\mathcal{N}_a} h(p,p')X(p)X(p') + \sum_{(p,q)\in\mathcal{N}_s} V(X(p),X(q)). \quad (12)$$

The last summation comes from the prior (we have encorporated the standard regularization parameter into the function V), while the second-to-last summation contains the cross terms.

This is, on the surface, an extremely intractable energy minimization problem, since it involves minimizing a non-convex energy function in a high-dimensional space. The natural approach is to minimize the energy with graph cuts, which have proven very effective for similar problems with MRF's. For example, the majority of the top-ranked stereo algorithms on the Middlebury benchmarks use graph cuts for energy minimization [14].

The most powerful graph cut method is the expansion move algorithm of [2]. Given an input image $\mathbf{x}$ and a label (i.e. intensity) α, an α-expansion is a new image where each pixel has either its old label x_p or the new label α. The expansion move algorithm proceeds by repeatedly finding the lowest energy α-expansion. This in turn involves minimizing a restricted version of the original energy function, where each pixel makes a binary decision.

In vision applications, the original energy function generally involves terms with single pixels and with pairs of pixels. The expansion move algorithm can find the lowest energy α-expansion, as long as the terms involving pairs of pixels satisfy a certain constraint. Following the notation of [13], we can express a term involving a pair of pixels p, q as a function $\mathcal{E}_{p,q}$ of two binary variables. The binary variable b_p associated with the pixel p will be 0 if p keeps its old label x_p, and 1 if p gets the new label α. It is shown in [5] that graph cuts can be used to find the lowest energy expansion move as long as $\mathcal{E}_{p,q}$ obeys the constraint

$$\mathcal{E}_{p,q}(0,0) + \mathcal{E}_{p,q}(1,1) \leq \mathcal{E}_{p,q}(1,0) + \mathcal{E}_{p,q}(0,1). \quad (13)$$

If $\mathcal{E}_{p,q}$ obeys this constraint we will call it *submodular* ([5] refers to this as "regular").

The key question, then, is whether the terms $\mathcal{E}_{p,q}$ in the binary restriction of our energy function E from (12) are submodular. These terms arise from the last two summations in E. If V is a metric, [5] shows that the resulting $\mathcal{E}_{p,q}$ terms are submodular.

There are several metrics that preserve edges, such as the Potts model or the truncated L_1 distance, where graph cuts are known to give good results [2]. Unfortunately, the $\mathcal{E}_{p,q}$ terms that arise from cross terms are problematic: [11] shows that these terms are only submodular at pixels p where $(p,p') \in \mathcal{N}_a \implies X(p) \leq \alpha \leq X(p')$. On any realistic image almost never occurs.

5. Our energy minimization method

We will make use of a graph construction due to Hammer *et al.* [3], which was introduced into computer vision by Vladimir Kolmogorov in early 2005 (see [4] for a particularly clear review). [3] shows how to compute a good (but not necessarily optimal) solution to this class of binary energy functions, even though some of the terms are not submodular. This in turn will allow to compute a good expansion move. To explain the construction of [3], let us divide the $\mathcal{E}_{p,q}$ terms into the submodular ones R and the non-submodular ones S. Assuming that V is a metric, the non-submodular terms $\mathcal{E}_{p,q}$ will have $(p,q) \in \mathcal{N}_a$. We are minimizing an energy function $\mathcal{E}$ over n binary variables b_p, one per pixel, where

$$\mathcal{E}(b_1,\ldots,b_n) = \sum_{p=1}^{n} \mathcal{E}_p(b_p) + \sum_{(p,q)\in R} \mathcal{E}_{p,q}(b_p,b_q) + \sum_{(p,q)\in S} \mathcal{E}_{p,q}(b_p,b_q). \quad (14)$$

For every pixel p we introduce an additional binary variable $\bar{b}_p$. We will call p *inconsistent* if it is in the set $\mathcal{I} = \{\, p \mid \bar{b}_p \neq 1 - b_p \,\}$, and consistent if $p \notin \mathcal{I}$.

Instead of the original energy function $\mathcal{E}(b_1,\ldots,b_n)$, [3] minimizes the new energy function over both the new and old variables:

$$\tilde{\mathcal{E}}(b_1,\ldots,b_n,\bar{b}_1,\ldots,\bar{b}_n) = \frac{1}{2}\Big(\sum_{p=1}^{n} \mathcal{E}_p(b_p) + \sum_{p=1}^{n} \mathcal{E}_p(1-\bar{b}_p) + \sum_{(p,q)\in R} [\mathcal{E}_{p,q}(b_p,b_q) + \mathcal{E}_{p,q}(1-\bar{b}_p,1-\bar{b}_q)] + \sum_{(p,q)\in S} [\mathcal{E}_{p,q}(b_p,1-\bar{b}_q) + \mathcal{E}_{p,q}(1-\bar{b}_p,b_q)]\Big). \quad (15)$$

We are interested in $\tilde{\mathcal{E}}$ because of the following two lemmas:

Lemma 1 *If all pixels p are consistent, then $\tilde{\mathcal{E}}(b_1,\ldots,b_n,\bar{b}_1,\ldots,\bar{b}_n) = \mathcal{E}(b_1,\ldots,b_n)$*

Lemma 2 *Every term in the energy $\tilde{\mathcal{E}}$ is submodular.*

Both these lemmas can be easily checked.

Thus, instead of minimizing $\mathcal{E}$, [3] proposes minimizing $\tilde{\mathcal{E}}$. In addition [3] proves that the solution has an important property called *persistence*, which means that the pixels that are consistent have the optimal labels.[1]

Given a labeling of the old and new binary variables, we need to compute an expansion move. Consistent pixels get their labels according to the old binary variable b_p. If a pixel p is inconsistent, we can leave it with its old value x_p, or assign it some new value such as $\frac{x_p+\alpha}{2}$. In practice, the number of inconsistent pixels $|\mathcal{I}|$ is very small, and the experimental results show that our method works quite well. Due to the persistence property, this means that the optimal α-expansion is computed at nearly all the pixels. Often $|\mathcal{I}| = 0$, and we compute the optimal α-expansion.

[1] The proof of this fact is non-trivial.

5.1. Error analysis

While persistence is a powerful property, it says nothing about how many pixels are consistent, and as a result shows nothing about how close we are to computing the optimal expansion move. Since our algorithm is a variant of the expansion move algorithm, if we always compute the optimal expansion move we would obtain this algorithm's guarantees, which are described in [2].

Our error analysis, which is concerned with how close we get to computing the optimal expansion move, has two parts. First, we can bound how far we are from computing the optimal α expansion; the bound depends on various parameters, but an important one is the number of inconsistent pixels $|\mathcal{I}|$. Second, we can give some conditions on inconsistent pixels which suggest that $|\mathcal{I}|$ is small, and that an inconsistent pixel p has a label x_p that is close to α (and therefore would not change very much in an α-expansion move).

Our error analysis is also backed up by experimental evidence in section 5.1.3. In practice it appears that both the bound and $|\mathcal{I}|$ are small, which in turn helps explain our algorithm's strong performance.

5.1.1 Bounding the expansion move error

To analyze how far we are from computing the optimal expansion move, will need to introduce some notation. If $|\mathcal{I}| > 0$, let us define the consistent pixels as $\mathcal{C} = \mathcal{I}^c$, and the number of inconsistent neighbors as $|\mathcal{N}(I)| = |\{\,(p,q) \in \mathcal{N}_s \cup \mathcal{N}_a \mid p \in \mathcal{I}\,\}|$. An assignment of 0 or 1 to all the binary variables $\{b_p\}$ corresponds to a labeling of all pixels p, either with the old value x_p or with the new one α.

Similarly, an assignment to the larger set of binary variables $\{b_p\} \cup \{\bar{b}_p\}$ can be viewed as assigning p a pair of labels $(\chi_p, \bar{\chi}_p)$, defined as follows. $(\chi_p,\bar{\chi}_p) = (x_p,x_p)$ if $(b_p,\bar{b}_p) = (0,1)$, (α,α) if $(b_p,\bar{b}_p) = (1,0)$, (α,x_p) if $(b_p,\bar{b}_p) = (1,1)$ and (x_p,α) if $(b_p,\bar{b}_p) = (0,0)$. Under this view, p is inconsistent if the pair of labels disagree, $\chi_p \neq \bar{\chi}_p$.

When we minimize $\tilde{\mathcal{E}}$, some pixels will be consistent and some will be inconsistent. We will write the assignment of pairs of labels to pixels that minimizes $\tilde{\mathcal{E}}$ as $\chi_{\min}$, and can divide the pixels into the consistent and inconsistent set. Our algorithm computes $\chi_{\min}$ and then creates a consistent labeling by giving inconsistent pixels p the value $\frac{x_p+\alpha}{2}$. We will write this new labeling as $\hat{\chi}$.

Our goal is to bound the difference between the energy of $\hat{\chi}$ and the energy of the optimal α expansion, which we will write χ^*. To do this, we need to generalize E from taking an assignment of one label per pixel, in order to handle two labels per pixel. We use the natural generalization $\tilde{E}$, which has the same relationship to the binary energy function $\tilde{\mathcal{E}}$ that E has to $\mathcal{E}$. Thus, if χ is consistent then $E(\chi) = \tilde{E}(\chi)$ (using the obvious equivalence between a consistent labeling of each pixel, and a labeling that assigns each pixel a pair of identical labels).

Our error bound will use the constant C, defined as

$$-\sum_{p\in\mathcal{I}}(\frac{x_p-\alpha}{2})^2+\frac{1}{2}\sum_{p,q\in\mathcal{I}}|(x_p-\alpha)(x_q-\alpha)|+\lambda|\mathcal{N}(\mathcal{I})|.$$

Here, $\lambda=\max V(\cdot)$, which is essentially the standard regularization parameter.

Theorem 1 *$\tilde{E}(\hat{\chi})-\tilde{E}(\chi_{\min})$, is no larger than C.*

PROOF: After a great many terms cancel, we are left with the equality $\tilde{E}(\hat{\chi})-\tilde{E}(\chi_{\min})=$

$$\begin{aligned}
&-\sum_{p\in\mathcal{I}}f(p)(\frac{x_p-\alpha}{2})^2+-\sum_{p,q\in\mathcal{I}}\frac{h(p,q)}{2}(x_p-\alpha)(x_q-\alpha)\\
&+\sum_{p,q\in\mathcal{I}}\frac{h(p,q)}{2}(x_p-\alpha)(x_q-\alpha)\\
&+\frac{1}{2}\sum_{p,q\in\mathcal{I}}[2V(\frac{x_p-x_q}{2})-V(x_p-x_q)]\\
&+\frac{1}{2}\sum_{p,q\in\mathcal{I}}[2V(\frac{x_p-x_q}{2})-V(x_p-\alpha)-V(x_q-\alpha)]\\
&+\frac{1}{2}\sum_{p\in\mathcal{I},q\in\mathcal{C}}[2V(\frac{x_p+\alpha}{2}-\chi_q)-V(x_p-\chi_q)-V(\alpha-\chi_q)].
\end{aligned}$$

Since the sum over the coils is normalized, $f(p)=1$. In addition, $|h(p,q)|\leq 1$ (this follows from the definition of h). Now the second and third summations have opposite signs, depending on the particular arrangement of the pairwise labels. In the worst case, all terms end up having a positive sign, and in the bound we replace both by their absolute values (although note that most practical situations are considerably better than the worst case).

All summations involving separation costs $V(\cdot)$ can be further simplified using the triangle inequality since they are metrics. Thus the last three summations are no larger than

$$\frac{1}{2}\sum_{p,q\in\mathcal{I}}V(\frac{x_p-x_q}{2})+\sum_{p,q\in\mathcal{I}}V(\frac{x_p-x_q}{2}-\alpha)+\sum_{p\in\mathcal{I},q\in\mathcal{C}}V(\frac{x_p-\alpha}{2}). \quad (16)$$

Since the value of V is no larger than λ, we obtain the desired bound. ■

Note that $\tilde{E}(\chi_{\min})\leq\tilde{E}(\chi^*)$, which in turn is the energy of the optimal α-expansion $E(\chi^*)$. This is natural, since by allowing inconsistent pixels we can potentially achieve lower energy. This fact, together with the theorem, gives us our bound:

Corollary 1 $E(\hat{\chi})-E(\chi^*)\leq C$.

The quantity C, of course, depends on the image, the parameters of the energy function E, and the value of α. We will present evidence in section 5.1.3 that in practice its value appears to be quite small.

5.1.2 Limiting the number of inconsistent pixels

We can associate an undirected graph with the energy function $\tilde{\mathcal{E}}$ as follows. For every pixel, there will be two nodes, which we denote $p,\bar{p}$. We can represent the two-argument terms in the definition of $\tilde{\mathcal{E}}$ by an edge in this graph. Now consider the configuration $\chi_{\min}$ as above, which minimizes $\tilde{\mathcal{E}}$ but which suffers from inconsistent pixels $\mathcal{I}$. We define the induced graph $\mathcal{G}_{\mathcal{I}}(\chi_{\min})$ to contain the pair of nodes $p,\bar{p}$ for every pixel $p\in\mathcal{I}$, and to have an edge between two nodes when there is a term involving the associated pair of binary variables in $\tilde{\mathcal{E}}$. We will use the structure of $\mathcal{G}_{\mathcal{I}}$ to give some conditions that must be met by inconsistent pixels, and hence argue that such pixels are rare. (Recall that if there are no inconsistent pixels, our algorithm computes the optimal expansion move.)

Lemma 3 *We can efficiently construct a solution $\chi'_{\min}$ where $\tilde{\mathcal{E}}(\chi'_{\min})=\tilde{\mathcal{E}}(\chi_{\min})$ such that for every pixel p, there is a path in $\mathcal{G}_{\mathcal{I}}(\chi'_{\min})$ between p and $\bar{p}$.*

PROOF: We can find all connected components of $\mathcal{G}_{\mathcal{I}}(\chi_{\min})$ in linear time, and hence identify any pixels p such that there is no path between p and $\bar{p}$ in the graph $\mathcal{G}_{\mathcal{I}}(\chi_{\min})$. If $(p,q)\in\mathcal{N}_s\cup\mathcal{N}_a$, there is no path between between q and $\bar{q}$, since this would provide a path between p and $\bar{p}$. This in turn implies that $\mathcal{G}_{\mathcal{I}}(\chi_{\min})$ contains two disjoint sub-graphs $\mathcal{G}_A$ and $\mathcal{G}_{\bar{A}}$, such that $p\in\mathcal{G}_A$ and $\bar{p}\in\mathcal{G}_{\bar{A}}$ Due to the structure of $\tilde{\mathcal{E}}$, the energy contributed by $\mathcal{G}_A$ and $\mathcal{G}_{\bar{A}}$ is identical. As a result, we can make p consistent, by giving it either the label x_p or α, without changing the energy. ■

Theorem 2 *Suppose we have obtained a minimum solution $\chi'_{\min}$ satisfying lemma 3. Then*

1. *If all terms of $\mathcal{E}$ are submodular, then there are no inconsistent pixels in $\chi'_{\min}$.*
2. *Inconsistent pairs cannot occur in clusters of 1 or 2.*
3. *Inconsistent pairs $(p,\bar{p})$ cannot occur if they are members of acyclic sub-graphs.*

PROOF: These cases imply the existence of pairs $p,\bar{p}$ which do not have a path between them in $\mathcal{G}_{\mathcal{I}}(\chi'_{\min})$, which is impossible according to lemma 3. ■

This shows that inconsistent pixels can only occur under limited circumstances. However, for our specific energy function we can find further bounds on x_p when p is inconsistent. Let n_s be the number of neighbors that each pixel has in $\mathcal{N}_s$. Assume that $\epsilon=n_s\max(V)/x_{max}\ll 1$, i.e., the regularization factor is much smaller than the maximum image intensity. This is reasonable, since we usually have $\epsilon<0.1$. Next we define

$$\triangle_R=g(p)+\sum_{(p,p')\in\mathcal{N}_a\cap R}h(p,p')x_{p'},$$

$$\gamma_R=\min(\alpha,-\alpha-\triangle_R),\ \gamma'_R=\max(\alpha,-\alpha-\triangle_R).$$

Similarly define $\triangle_S$, γ_S and γ'_S by replacing R by S.

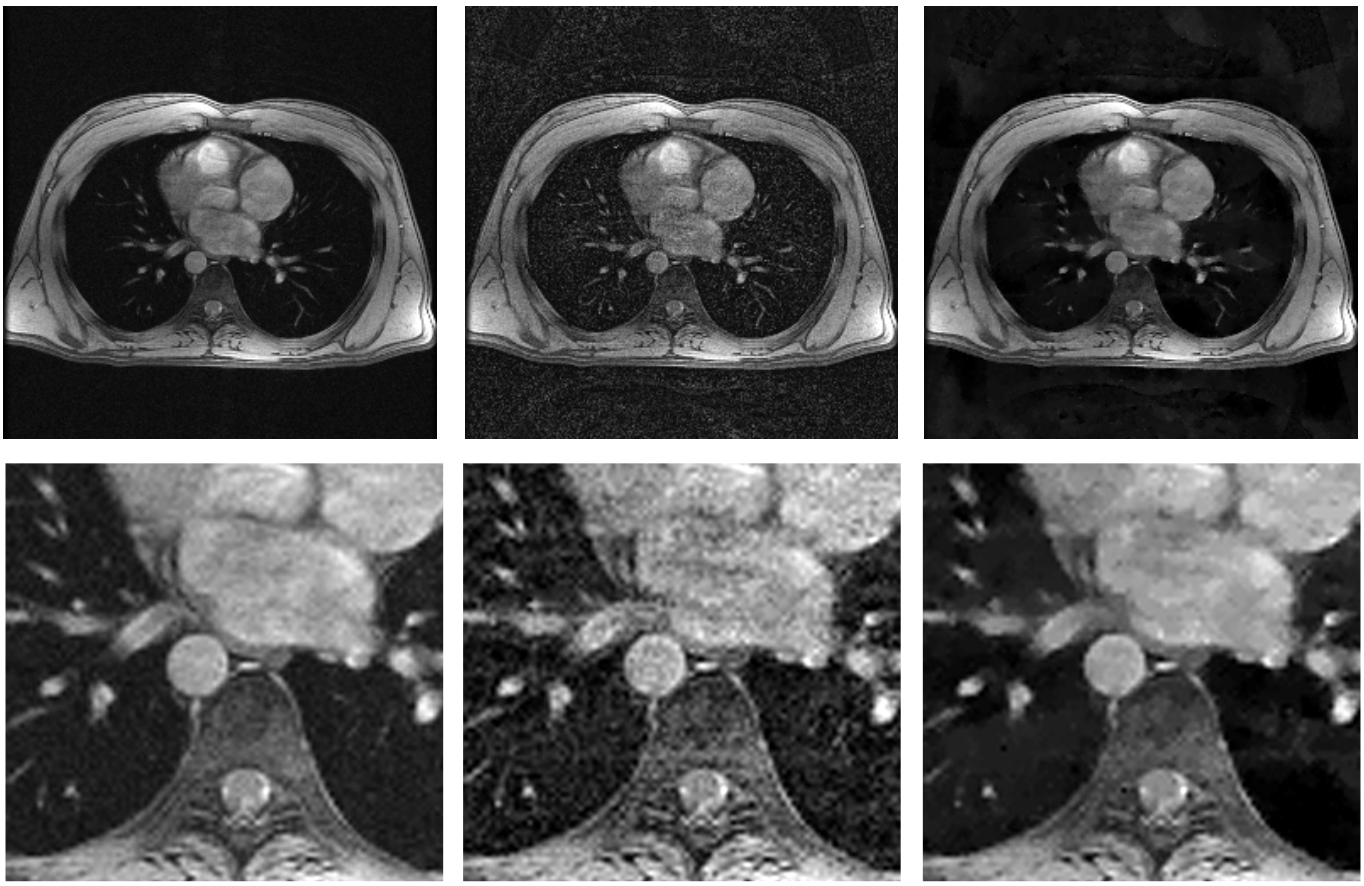

Non-accelerated reconstruction | Regularized SENSE ($\mu = 0.1$) | Our reconstruction

Figure 2. In vivo cardiac results, 3× acceleration, 8 coils. Zoomed versions shown below. The unaccelerated scans took 30 seconds.

Theorem 3 *For a given α-expansion move, a pixel p can be inconsistent in the minimum $\tilde{\mathcal{E}}$ solution $\chi_{\min}$ only if it satisfies the following inequalities:*

$$\begin{aligned} x_p \geq \gamma'_R \quad \textit{OR} \quad x_p \leq \gamma_R, \\ \gamma_S \leq x_p \leq \gamma'_S, \\ 0 \leq x_p \leq x_{max}. \end{aligned}$$

PROOF: The last inequality is self-evident. Denote by $\tilde{\mathcal{E}}_p(\chi_p, \bar{\chi}_p)$ the terms of $\tilde{\mathcal{E}}$ involving the (inconsistent) assignment at p. Assume this assignment, without loss of generality, to be $\chi_p = x_p$ and $\bar{\chi}_p = \alpha$. We impose the local minimum conditions

$$\tilde{\mathcal{E}}_p(x_p, \alpha) \leq \tilde{\mathcal{E}}_p(x_p, x_p), \quad \tilde{\mathcal{E}}_p(x_p, \alpha) \leq \tilde{\mathcal{E}}_p(\alpha, \alpha).$$

Evaluating these conditions produces, after some algebra, two quadratic inequalities, which simplify, neglecting ϵ, to the above result. If ϵ cannot be neglected the bounds become more complicated, but it is easy to show by expanding in a Taylor series and omitting high order terms, that this merely perturbs the γ's by small quantities of $\mathcal{O}(\epsilon)$. This analysis is omitted here due to space limitations. ■

Although the bounds on x_p provided by this theorem are not particularly intuitive, it is obvious that the intersection of all six inequalities above severely restricts possible inconsistent assignments. It is also possible to give a statistical argument that under reasonable assumptions about the distributions of intensities very few pixels can be inconsistent.

5.1.3 Experimental values of our bounds

We also performed some experiments to determine the number of inconsistent pixels and the value of our bound C. Over multiple runs on real MR images with different parameters, we observed a small number of inconsistent pixels (typically no more than 2%, and often nearly 0). We also computed the bound C corresponding to a given $|\mathcal{I}|$, for various values of parameters α, λ, etc. For 2% inconsistency (which is a conservative value), we found C to be between 1% and 10% of the lowest energy we obtained.

6. Experimental results

We experimented with structural T1-weighted brain and cardiac images, acquired both at 4T and 1.5T High field strength (4T) structural MRI brain data was obtained using a Siemens whole body scanner equipped with a standard birdcage transmit, 8 channel phased-array receiver headcoil. Volumetric T1-weighted im-

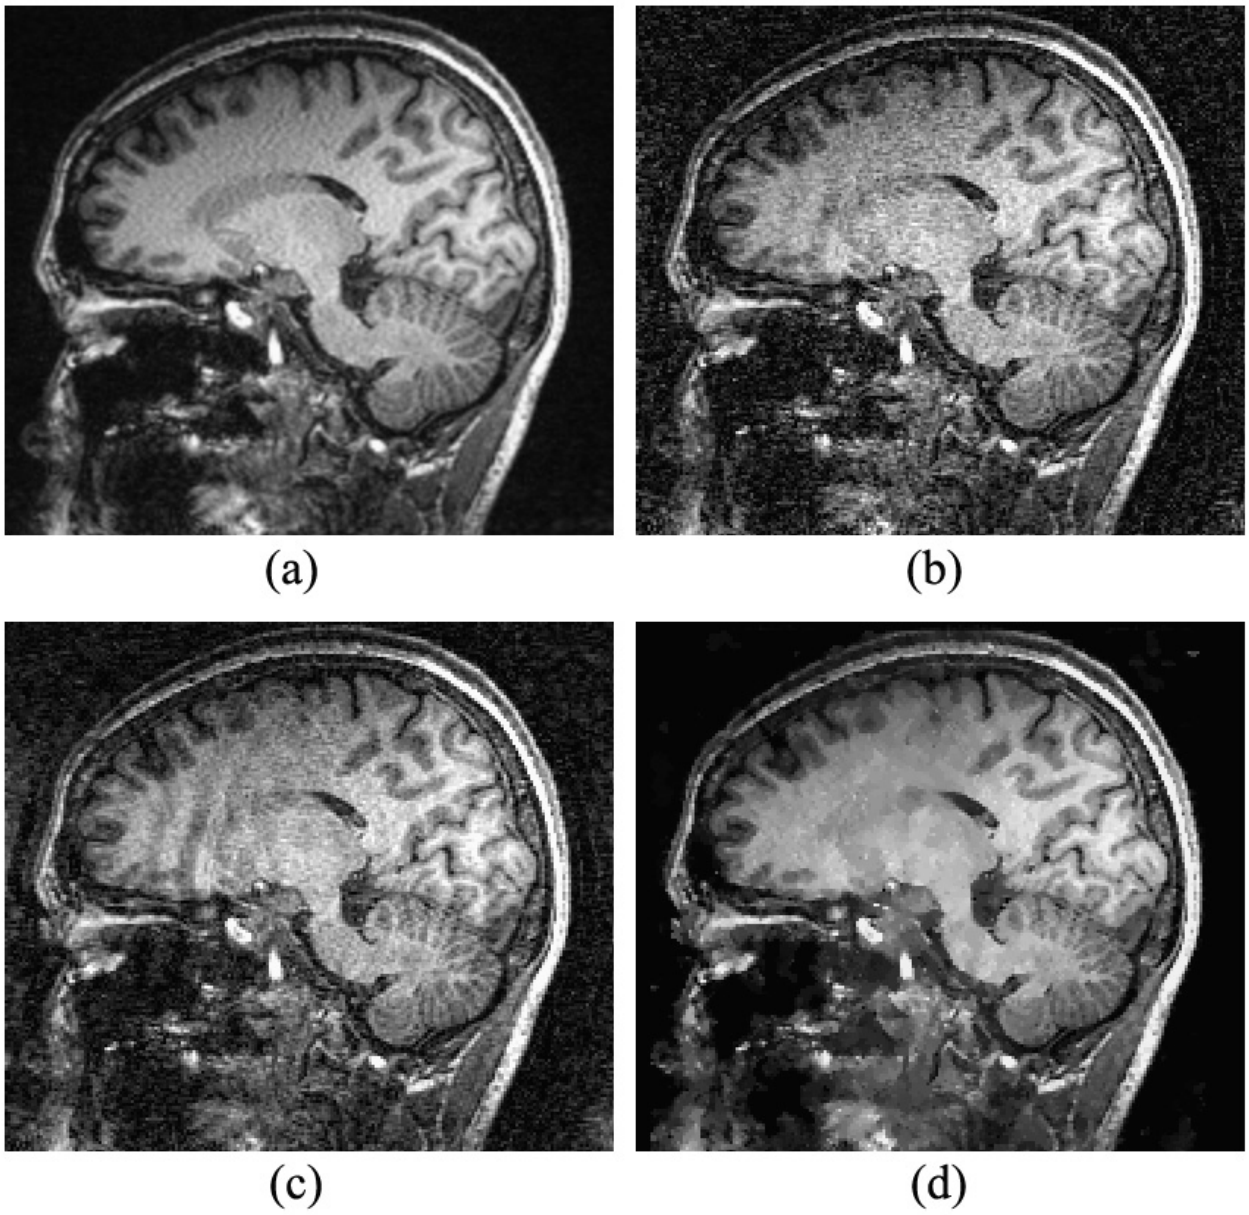

Figure 3. In vivo brain results, $4\times$ acceleration, 8 coils. The unaccelerated scan at top left (a) took 8:00, the accelerated scans took 2:00. (b) SENSE, $\mu = 0.08$, (c) SENSE, $\mu = 0.16$, (d) Our reconstruction.

ages were acquired using a MPRAGE sequence with TI/TR = 950/2300 ms timing and a flip angle of 8°. We also acquired several parallel imaging data sets of axial and oblique slices of the torso region using a GE 1.5T Excite-11 system. Additional details concerning the acquisition parameters are given in [10].

Reconstruction results for 3x acceleration of a cardiac slice are shown in figure 2, and 4x acceleration of a sagittal brain slice in figure 3. Two reconstruction using different regularization factors μ are shown for the SENSE method in order to demonstrate that its output cannot be improved by mere regularization. Higher regularization in SENSE does reduce noise, but at the cost of introducing aliasing artifacts which make the images diagnostically useless. In contrast our reconstruction has very little noise or aliasing, and exhibits good sharpness of edges.

These results also point to some important features and shortcomings of our approach. First, there is visible loss of dynamic range in our reconstructions compared to unaccelerated data. This is because in our experiments we discretized the intensities to 8 bits. The original unaccelerated reconstruction used a larger range of intensities. This is a limitation of our implementation, and not

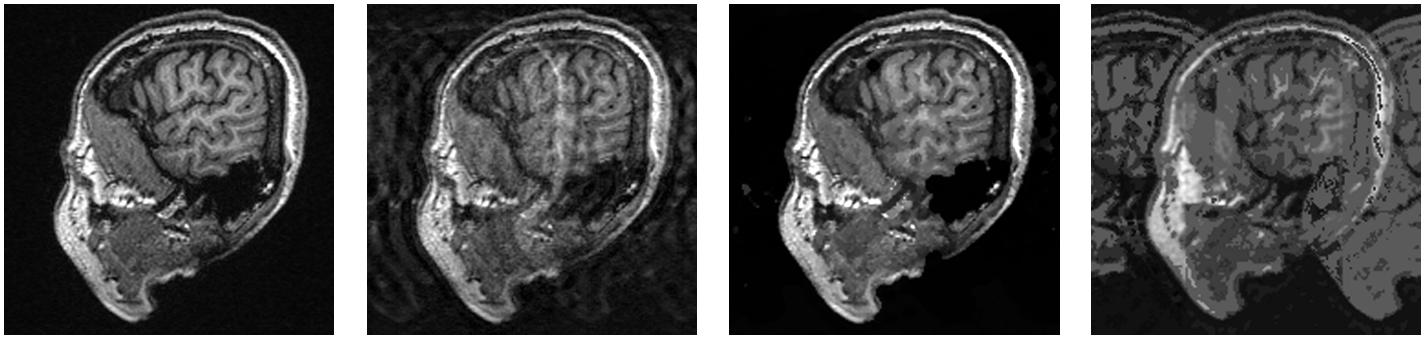

Non-accelerated reconstruction | Regularized SENSE ($\mu = 0.3$) | Our reconstruction | [11] reconstruction

Figure 4. In vivo brain results, $5\times$ acceleration, 8 coils. The unaccelerated scan at left took 8:00, the accelerated scans took 1:36.

of our algorithm, since the running time is linear in the number of labels.

Comparing a parallel image reconstruction to the unaccelerated reconstruction using the signal to noise (SNR) metric is widely done in parallel imaging [9]. We computed SNR by obtaining two reconstructions from the same original data set after adding random noise. The sum of the two images gives a measure of the signal and the difference image a measure of noise, as described in [12]. Details of our evaluation procedure are given in [10].

	Acceleration	SNR of SENSE	Our SNR
Fig. 2	3	15	36
Fig. 3	4	10	23
Fig. 4	5	10	17

The SNR comparison against SENSE is show in the table above. For SENSE, the regularization parameter μ was tuned to achieve the best SNR. Our method typically results in a 100% improvement in SNR, as well as a significant increase in visual quality.

6.1. Comparison with [11]

Our work clearly builds upon [11], who first addressed the parallel imaging problem using graph cuts. However, the algorithm we propose has better theoretical properties than the earlier method of [11], and works much better in practice.

Theoretically, the earlier algorithm has quite weak convergence properties, since [11] only proved that their α-expansion does not increase the energy.[2] Obviously this guarantee could be provided by never changing any pixels. Unlike the guarantee that they provide, we can show that the expansion move we compute is close to optimal, in two senses. First, the persistence results of [3] mean that consistent pixels (which are in practice the vast majority) have their optimal labels. Second, the results given in the previous section bound the error in the expansion move that we compute.

Besides the theoretical difference between our method and the earlier algorithm, there is a significant difference in practice. While it is possible to obtain good results using the previous method, the method is sensitive to its starting point, and in our experience it is not uncommon for the algorithm to simply return its input (i.e., to produce the expansion move that never changes any pixels). An example of the poor performance we sometimes observed in practice is given in figure 4. Although it is possible to reduce the aliasing artifacts at the cost of introducing a great deal of local noise, the experimental results of our new method are vastly better.

[2] A similar result is provided in [13] for a very broad class of energy functions.

Acknowledgements

We thank Vladimir Kolmogorov for spending a considerable amount of time explaining the pseudo-Boolean optimization algorithm in [3]; his help was invaluable in clarifying this important work. We also thank Bryan Kressler and Yi Wang for helping provide the data.

References

[1] R. Bammer and S. O. Schoenberg. Current concepts and advances in clinical parallel Magnetic Resonance imaging. *Top Magn Reson Imaging*, 15(3):129–58, 2004.

[2] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, 2001.

[3] P. L. Hammer, P. Hansen, and B. Simeone. Roof duality, complementation and persistency in quadratic 0-1 optimization. *Mathematical Programming*, 28:121–155, 1984.

[4] V. Kolmogorov and C. Rother. Minimizing non-submodular functions with graph cuts - a review. Technical Report MSR-TR-2006-100, Microsoft Research, 2006.

[5] V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE Trans Pattern Anal Mach Intell*, 26(2):147–59, 2004.

[6] S. Li. *Markov Random Field Modeling in Computer Vision*. Springer-Verlag, 1995.

[7] Z.-P. Liang and P. C. Lauterbur. *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. IEEE press, 1999.

[8] F. H. Lin, K. K. Kwong, J. W. Belliveau, and L. L. Wald. Parallel imaging reconstruction using automatic regularization. *Magn Reson Med*, 51(3):559–567, 2004.

[9] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and B. Peter. SENSE: sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42(5):952–962, 2001.

[10] A. Raj, G. Singh, R. Zabih, B. Kressler, Y. Wang, N. Schuff, and M. Weiner. Bayesian parallel imaging with edge-preserving priors. *Magnetic Resonance in Medicine*, 2006. to appear.

[11] A. Raj and R. Zabih. A graph cut algorithm for generalized image deconvolution. In *International Conference on Computer Vision (ICCV)*, 2005.

[12] S. Reeder, B. Wintersperger, O. Dietrich, T. Lanz, A. Greiser, M. Reiser, G. Glazer, and S. Schoenberg. Practical approaches to the evaluation of signal-to-noise ratio performance with parallel imaging: Application with cardiac imaging and a 32-channel cardiac coil. *Magnetic Resonance in Medicine*, 54:748–754, 2005.

[13] C. Rother, S. Kumar, V. Kolmogorov, and A. Blake. Digital tapestry. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005.

[14] D. Scharstein and R. Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision* , 47:7–42, 2002.