INPAINTING WITH SPARSE LINEAR COMBINATIONS OF EXEMPLARS

Brendt Wohlberg*

T-7 Mathematical Modeling and Analysis Los Alamos National Laboratory Los Alamos, NM 87545, USA

ABSTRACT

We introduce a new exemplar-based inpainting algorithm that represents the region to be inpainted as a sparse linear combination of example blocks, extracted from the image being inpainted or an external training image set. This method is conceptually simple, being computed by minimization of a simple functional, and avoids the complexity of correctly ordering the filling in of missing regions of other exemplar-based methods. Initial performance comparisons on small inpainting regions indicate that this method provides similar or better performance than other recent methods.

Index Terms— Image restoration, Inpainting, Exemplar

1. INTRODUCTION

Exemplar based methods are becoming increasingly popular for problems such as denoising [1, 2], superresolution [3, 4, 5], texture synthesis [6], and inpainting [7, 8]. The common theme of these methods is the use of a set of actual image blocks, extracted either from the image being restored, or from a separate training set of representative images, as an image model. In the case of inpainting, the approach is usually to progressively replace missing regions with the best matching parts of the same image, carefully choosing the order in which the missing region is filled to minimize artifacts [8]. Instead, we propose an inpainting method that represents missing regions as sparse linear combinations of other regions in the same image (in contrast to [9], in which sparse representations on standard dictionaries, such as wavelets, are employed), computed by minimizing a simple functional.

2. SPARSE LINEAR COMBINATIONS OF EXEMPLARS

The image model of the proposed approach is that each block of the restored image should be a sparse linear combination of other image blocks, either from known regions (i.e. not intersecting the inpainting region) of the image being restored, or from a separate training image set. These image blocks are overlapped to reduce blocking artifacts, and, more importantly, to enable information from the exterior of the inpainting region to propagate to blocks entirely within the interior, as illustrated in Fig. 1.



Fig. 1. Image blocks which cross the boundary of the inpainting region can be chosen as the best match to the image model subject to the constraint that they match the known image pixels outside the inpainting region, but in a tiled block structure, blocks interior to the inpainting region are unconstrained. In an overlapping block structure, an additional constraint on the mismatch between overlapping blocks allows the constraint from the known pixels to propagate to interior blocks.

Within this framework, the solution is computed by minimizing a global functional which

- 1. penalizes the l^1 norm of the linear combination coefficients to encourage a sparse, low complexity, solution,
- 2. constrains (or penalizes) the mismatch between solution blocks and known pixels, and
- 3. constrains (or penalizes) the mismatch between overlapping parts of different solution blocks.

In order to discuss this approach in more detail, we need to establish some notation. Denote the image to be inpainted by vector s, the inpainting region mask by r, and the inpainted result by u. To reduce the complexity of computing the mismatch between overlapping parts of different blocks, the blocks are arranged in indexed grids, with the overlap being produced by an offset of the entire grid, as indicated in

^{*}This research was supported by the U.S. Department of Energy through the LANL/LDRD Program.

Fig. 2. This structure allows the total block overlap mismatch to be computed as the mismatch between the grids, without having to track overlapping parts of individual blocks. Each block is indexed by the number of its grid and its number within that grid, block $s_{k,l}$ being the l^{th} block in grid k.



Fig. 2. Structure of overlapping block grids.

Denote the operator that extracts block k, l from an image as $B_{k,l}$ and the operator that "inserts" the same block back into the image as $B_{k,l}^T$, and make the definitions

$$R = \operatorname{diag}(\mathbf{r}) \qquad \mathbf{r}_{k,l} = B_{k,l}\mathbf{r}$$
$$R_{k,l} = \operatorname{diag}(\mathbf{r}_{k,l}) \qquad \mathbf{s}_{k,l} = R_{k,l}B_{k,l}\mathbf{s} .$$

We define $\Phi_{k,l}$ and $\alpha_{k,l}$ as the dictionary and coefficients respectively for block k, l. The dictionary is indexed because it is reasonable (and consistent with other exemplar-based image restoration algorithms) to choose a distinct dictionary for each block by finding similar blocks within the same image, or within an external training set of blocks extracted from similar images. For the initial experiments reported here, however, a simpler approach is adopted in which all blocks share a common dictionary, constructed from all complete (i.e. not intersecting the inpainting region) blocks within a specified training region in the image to be inpainted. This approach remains effective since finding the linear combination with minimum ℓ^1 norm has the effect of selecting a representation in terms of similar blocks, when possible.

Now, for each block $\mathbf{u}_{k,l} = \Phi_{k,l} \boldsymbol{\alpha}_{k,l}$, we wish to minimize or constrain (either equal to zero, or less than some upper bound) the following terms:

- 1. Solution sparsity $\|\boldsymbol{\alpha}_{k,l}\|_1$
- 2. Mismatch with known pixels in $s_{k,l}$

$$\|R_{k,l}\Phi_{k,l}\boldsymbol{\alpha}_{k,l}-R_{k,l}B_{k,l}\mathbf{s}\|_2$$

3. Mismatch in overlap with block grid n

$$\|\Phi_{k,l}\boldsymbol{\alpha}_{k,l} - B_{k,l}\sum_{m} B_{n,m}^{T}\Phi_{n,m}\boldsymbol{\alpha}_{n,m}\|_{2} \quad n \neq k$$

Indexed symbols are combined over all blocks to give the definitions

$$\begin{split} \boldsymbol{\alpha} &= \begin{pmatrix} \alpha_{0,0} \\ \alpha_{0,1} \\ \vdots \\ \alpha_{1,0} \\ \alpha_{1,1} \\ \vdots \end{pmatrix} \quad \tilde{\mathbf{s}} = \begin{pmatrix} \mathbf{s}_{0,0} \\ \mathbf{s}_{0,1} \\ \vdots \\ \mathbf{s}_{1,0} \\ \mathbf{s}_{1,1} \\ \vdots \end{pmatrix} \\ R &= \begin{pmatrix} R_{0,0} & 0 & 0 & 0 & 0 & \cdots \\ 0 & R_{0,1} & 0 & 0 & 0 & \cdots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \cdots \\ 0 & 0 & 0 & R_{1,0} & 0 & \cdots \\ 0 & 0 & 0 & 0 & R_{1,1} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \\ C &= \begin{pmatrix} 0 & 0 & \cdots & B_{1,0}^T & B_{1,1}^T & \cdots \\ 0 & 0 & \cdots & B_{1,0}^T & B_{1,1}^T & \cdots \\ 0 & 0 & \cdots & B_{1,0}^T & B_{1,1}^T & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ B_{0,0}^T & B_{0,1}^T & \cdots & 0 & 0 & \cdots \\ B_{0,0}^T & B_{0,1}^T & \cdots & 0 & 0 & \cdots \\ \vdots & \vdots \end{pmatrix} . \end{split}$$

Matrices Φ and *B* have the same diagonal block structure (in terms of their indexed components) as *R*. Note that these definitions have been simplified by assuming there are only two block grids, but in principle a larger number is possible, and may well enhance performance at the expense of greater computational cost.

Our global penalties/constraints may then be written as

- 1. Solution sparsity $\|\boldsymbol{\alpha}\|_1$
- 2. Mismatch with known pixels in s

 $\|R\Phi \boldsymbol{\alpha} - \tilde{\mathbf{s}}\|_2$

3. Mismatch between overlapping block grids

 $\|(I - BC)\Phi\alpha\|_2$

The two most obvious ways of combining these distinct goals are the unconstrained problem

$$\min \left\{ \|\boldsymbol{\alpha}\|_1 + \gamma_0 \| R \Phi \boldsymbol{\alpha} - \tilde{\mathbf{s}} \|_2 + \gamma_1 \| (I - BC) \Phi \boldsymbol{\alpha} \|_2 \right\}$$

with weights γ_0 and γ_1 , and the constrained problem

$$\min \|\boldsymbol{\alpha}\|_1 \text{ s.t. } \|R\Phi\boldsymbol{\alpha} - \tilde{\mathbf{s}}\|_2 \le \sigma_0 \text{ and } \|(I - BC)\Phi\boldsymbol{\alpha}\|_2 \le \sigma_1$$

with upper bounds σ_0 and σ_1 . For the preliminary experiments reported here we concentrate on the latter problem, but in the absence of an efficient solver for this exact problem, we define

$$A = \begin{pmatrix} \gamma_0 R\Phi \\ \gamma_1 (I - BC)\Phi \end{pmatrix} \quad \mathbf{b} = \begin{pmatrix} \gamma_0 \tilde{\mathbf{s}} \\ \mathbf{0} \end{pmatrix}$$

and solve the related problem

$$\min \|\boldsymbol{\alpha}\|_1$$
 such that $\|A\boldsymbol{\alpha} - \mathbf{b}\| \leq \sigma$

using SPGL1 [10, 11]. The parameters γ_0 and γ_1 balance the influences of the two "mismatch" components of the constraint, and σ determines the maximum of this combined constraint. The final restored image (or image region, since the computational domain can be restricted to a small region around the inpainting region) is obtained by averaging the overlapping block grids to obtain a single value for each image pixel.

3. RESULTS

We illustrate the performance of the proposed method using two small test regions extracted from the well-known Lena and Barbara test images, and compare performance with the Field of Experts method [12] and the method of Criminisi et al. [8], computed using the publicly available software [13] and [14] respectively. Results for the first region, extracted from the hat in the Lena image, are displayed in Figure 3, and the results for the second, extracted from the textured pants in the Barbara image, are displayed in Figure 4, and a comparison for SNR values is displayed in Table 1. For the results presented here, we chose parameters empirically to be $\gamma_0 = 1.00, \, \gamma_1 = 0.01$, and $\sigma = 40.0$. In each case, a subregion of the image was extracted for training (i.e. as a source for the image blocks populating the dictionary), and the smaller subregion actually displayed was tiled by the overlapping block grids and used for minimization of the functional. When evaluating the performance of the competing algorithms, these algorithms were presented with the larger (training) subregion for a fair comparison.

Image	FoE [12]	Criminisi et al. [8]	Proposed
Lena	15.4dB	15.6dB	17.8dB
Barbara	6.2dB	2.3dB	8.1dB

Table 1. Comparison of inpainting SNR values on test images displayed in Figures 3 and 4.

In terms of subjective image quality, the FoE reconstruction (which, in its defense, was perhaps designed for narrower inpainting regions) is noticeably inferior to the other two on the Lena image, and very much inferior on the Barbara image. (Despite the significant subjective quality difference on the second example, the FoE SNR is significantly higher than that of the method of Criminisi *et al.*, illustrating the limited utility of SNR as an objective quality measure in this context.) The proposed method provides slightly better subjective image quality than that of Criminisi *et al.* for the Lena image (more clearly visible on an image display than in the printed examples), and a more significant performance advantage for the Barbara image. For these comparisons, runs times were similar to those of the FoE method, and 2-3 orders of magnitude longer than the method of Criminisi *et al.*.

4. CONCLUSIONS

The simple initial implementation of the proposed inpainting methods exhibits high computational cost, but delivers comparable or better results than the significantly more conceptually complex methods with which it was compared, providing evidence of the potential of the proposed approach. A number of significant improvements remain to be considered, including application of a more appropriate solver for the minimization problem (particularly important given the computational expense), and selection of optimized dictionaries for each solution block.

5. REFERENCES

- A. Buades, B. Coll, and J.-M. Morel, "Image denoising by non-local averaging," in *Proc. ICASSP*, March 2005, vol. 2, pp. 25–28.
- [2] K. Dabov, A. Foi, V. Katkovnik, and K. O. Egiazarian, "Image denoising by sparse 3-d transform-domain collaborative filtering.," *IEEE Trans. Image Proc.*, vol. 16, no. 8, pp. 2080–2095, 2007.
- [3] W. T. Freeman, T. R. Jones, and E. C. Pasztor, "Example-based super-resolution," *IEEE Comput. Graph. and Appl.*, vol. 22, no. 3, March 2002.
- [4] S. Baker and T. Kanade, "Limits on super-resolution and how to break them," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 9, pp. 1167–1183, 2002.
- [5] M. Elad and D. Datsenko, "Example-based regularization deployed to super-resolution reconstruction of a single image," *Comput. J.*, 2007, DOI: 10.1093/comjnl/bxm008.
- [6] A. A. Efros and T. K. Leung, "Texture synthesis by nonparametric sampling," in *Proc. ICCV*, Corfu, Greece, September 1999, pp. 1033–1038.
- [7] I. Drori, D. Cohen-Or, and H. Yeshurun, "Fragmentbased image completion," *ACM Trans. on Graph.*, vol. 22, no. 3, pp. 303–312, 2003.
- [8] A. Criminisi, P. Pérez, and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *IEEE Trans. Image Proc.*, vol. 13, no. 9, pp. 1200– 1212, 2004.
- [9] M. J. Fadili, J. L. Starck, and F. Murtagh, "Inpainting and zooming using sparse representations," *Comput. J.*, 2007, DOI: 10.1093/comjnl/bxm055.



Fig. 3. Inpainting comparison on a subimage cropped from the Lena image. The region to be inpainted is 20×8 pixels in size.



Fig. 4. Inpainting comparison on a subimage cropped from the Barbara image. The region to be inpainted is 8×30 pixels in size. Images have been rotated by 90° to save space.

- [10] E. van den Berg and M. P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," Tech. Rep. TR-2008-01, Department of Computer Science, University of British Columbia, Jan. 2008.
- [11] E. van den Berg and M. P. Friedlander, "SPGL1: A solver for large-scale sparse reconstruction," Matlab code available from http://www.cs.ubc.ca/ labs/scl/spgl1, June 2007.
- [12] S. Roth and M. J. Black, "Fields of experts: a framework

for learning image priors," *Proc. CVPR*, vol. 2, pp. 860–867, June 2005.

- [13] S. Roth, "Fields of experts demo," Matlab code available from http://www.cs.brown.edu/~roth/ research/foe/downloads.html, 2005.
- [14] S. Bhat, "Object removal by exemplar-based inpainting," Matlab code available from http://www.cc. gatech.edu/~sooraj/inpainting/, 2004, Implementation of [8].