



Published in final edited form as:

Inf Process Med Imaging. 2017 June ; 10265: 659–671. doi:10.1007/978-3-319-59050-9_52.

Population Based Image Imputation

Adrian V. Dalca^{1,2}, Katherine L. Bouman¹, William T. Freeman^{1,3}, Natalia S. Rost⁴, Mert R. Sabuncu^{2,5}, and Polina Golland¹

¹Computer Science and Artificial Intelligence Lab, MIT, Cambridge, USA

²Martinos Center for Biomedical Imaging, Massachusetts General Hospital, HMS, Charlestown, MA, USA

³Google Research, Cambridge, MA, USA

⁴Department of Neurology, Massachusetts General Hospital, HMS, Boston, USA

⁵School of Electrical and Computer Engineering, Cornell, Ithaca, USA

Abstract

We present an algorithm for creating high resolution anatomically plausible images consistent with acquired clinical brain MRI scans with large inter-slice spacing. Although large databases of clinical images contain a wealth of information, medical acquisition constraints result in sparse scans that miss much of the anatomy. These characteristics often render computational analysis impractical as standard processing algorithms tend to fail when applied to such images. Highly specialized or application-specific algorithms that explicitly handle sparse slice spacing do not generalize well across problem domains. In contrast, our goal is to enable application of existing algorithms that were originally developed for high resolution research scans to significantly undersampled scans. We introduce a model that captures fine-scale anatomical similarity across subjects in clinical image collections and use it to fill in the missing data in scans with large slice spacing. Our experimental results demonstrate that the proposed method outperforms current upsampling methods and promises to facilitate subsequent analysis not previously possible with scans of this quality.

1 Introduction

Increasingly open acquisition efforts in clinical practice are driving dramatic increases in the number and size of patient cohorts in clinical archives. Unfortunately, clinical volumes are typically of dramatically lower resolution than the research scans that motivate most methodological development. Specifically, while individual slices in the scan can be of high resolution, slice spacing is often significantly larger, resulting in missing voxels, as illustrated in Fig. 1. This presents significant challenges for even basic tasks, such as skull stripping and registration, which are often necessary for downstream analysis [4,9]. We present a novel method for constructing high resolution anatomically plausible volumetric images consistent with the available slices in sparsely sampled clinical scans. The restored

images promise to enable computational analysis of clinical scans with existing techniques originally developed for isotropic research scans. Importantly, our method does not require any high resolution scans or expert annotations, but instead imputes the missing structure by learning from the available collection of clinical scans.

Our work is motivated by a study that includes brain MRI scans of thousands of stroke patients acquired within 48 h of stroke onset. The study aims to quantify white matter disease burden, necessitating skull stripping and deformable registration into a common coordinate frame [16]. The volumes are severely under-sampled ($0.85 \text{ mm} \times 0.85 \text{ mm} \times 6 \text{ mm}$) due to clinical constraints of acute stroke care, as illustrated in Fig. 1. Such undersampling is typical of modalities that aim to characterize tissue properties such as T2-FLAIR, even in research studies like ADNI2 [9].

Since clinically acquired scans violate many algorithms' underlying assumptions, even basic tasks present significant challenges [6,20,21]. In undersampled scans, the image is no longer smooth, and the anatomical structure may change substantially between consecutive slices (Fig. 1). Application-specific algorithms promise to address these problems for certain clinical scans but do not generalize well across applications and imaging modalities. In contrast, we harness the data available in a given clinical image collection to reconstruct the high resolution images that represent plausible anatomy from the low resolution scans (Fig. 2). The resulting images can then be analyzed by widely used algorithms that require nearly isotropic high resolution input. The imputed data acts as a medium for improving analysis tasks. For example, although imputed data should not be used in the clinical evaluation, the brain mask obtained through skull stripping of the restored scans can be applied to the original clinical scan for further analysis.

Prior Work

Traditional image restoration, or superresolution, techniques depend on having enough information in a single scan to synthesize data. Unfortunately, clinical slices are often sampled too sparsely for functional interpolation, such as linear, cubic or spline [19], to succeed. Similarly, patch-based methods that rely on redundancy within a single scan to “hallucinate” missing fine scale structure [13–15] fail to produce anatomically plausible reconstructions at this level of sparsity. Superresolution algorithms that use multiple images of the same subject to improve a single scan [2,10,15] are unsuitable for clinical data where multiple similar acquisitions are not commonly available.

Nonparametric upsampling methods proposed to tackle the problem of superresolution often rely on an external dataset of high resolution data or cannot handle extreme undersampling present in clinical scans. For example, some methods fill in missing data by matching a low resolution patch from the scan with a high resolution patch from the training dataset [3,7,10,11,18]. A recent approach to improve resolution from a collection of scans with sparse slices jointly upsamples all images using non-local means [17]. However this method has only been demonstrated on slice spacing of roughly three times the in-plane resolution, and in our experience similar non-parametric methods fail to upsample clinical scans with more significant undersampling.

Parametric methods and low dimensional embeddings have been used to model the common structure of image patches from full resolution images, but are typically not designed to handle missing data. Specifically, priors [22] and Gaussian Mixture Models [23] have been used in both medical and natural images for classification [1] and denoising [5,23]. The procedures used for training these models rely on having full resolution patches with no missing data. Unfortunately, high (full) resolution training datasets are not readily available for many image contrasts and scanners, and may not adequately represent pathology or other properties of clinical populations. Acquiring the appropriate high resolution data is often infeasible, and here we explicitly focus on the realistic medical scenario where only low resolution image sets are available.

Overview

Our method takes advantage of the fact that local fine scale structure is intrinsically shared in a population of medical images, and each scan with sparse slices captures some partial aspect of this structure. We borrow ideas from Gaussian Mixture Model (GMM) patch priors [23], low dimensional Gaussian embeddings [8], and missing data models [8,12] to develop a probabilistic model that describes sparse 3D patches from all volumetric images in a collection around a particular location using a low-dimensional GMM with partial observations. We derive an iterative algorithm to learn the model parameters. We demonstrate our algorithm using scans from the ADNI cohort as well as the motivating stroke study, and also illustrate a preliminary illustration of potential improvements in the down-stream analysis using an example task of skull stripping. Finally, we discuss initialization tradeoffs and modelling choices.

2 Method

We employ a Gaussian Mixture Model (GMM) to capture local structure in sparse 3D patches in the vicinity of a particular location across the entire collection. We treat a patch as a high dimensional manifestation of a low dimensional representation, with the intuition that the covariation within image patches has small intrinsic dimensionality relative to the number of voxels in the patch. In this section, we describe the model, the resulting learning algorithm, and our image restoration procedure.

2.1 Generative Model

Let $\{Y_1, \dots, Y_N\}$ be an image collection of scans with large slice separation, roughly aligned into a common atlas space (we use affine transformations in our experiments). For each image Y_i in the collection, only a few slices are known, and we seek to restore the anatomically plausible high resolution volume by imputing the missing voxel values.

We capture local structure using image patches. We assume a constant patch shape, and use y_i to denote a D -length vector that contains voxels of the image patch centered at a certain location of image Y_i . We perform inference at each location independently. We model the set of patches $\mathcal{Y} = \{y_i\}$ at a common location as drawn from a K -component multivariate GMM. If generated by cluster k , patch y_i is a high dimensional observation of a low dimensional patch representation x_i of length d .

$$y_i = \mu_k + W_k x_i + \varepsilon_i, \text{ where } (1)$$

$$x_i \sim \mathcal{N}(0, I_{D \times D}),$$

$$\varepsilon_i \sim \mathcal{N}(0, \sigma_k^2 I_{D \times D}), \text{ and } \varepsilon_i \perp x_i.$$

Here patch μ_k is the mean of cluster k , matrix W_k shapes the covariance structure of the cluster, and σ_k^2 is the variance of image noise. The likelihood of all patches at this location under the mixture model is

$$p(\mathcal{Y}; \{\mu_k\}, \{W_k\}, \{\sigma_k^2\}, \pi) = \prod_i \sum_k \pi_k \mathcal{N}(y_i; \mu_k; C_k), \quad (2)$$

where $C_k = W_k W_k^T + \sigma_k^2 I_{D \times D}$ and π is a vector of cluster proportions.

Let \mathcal{O}_i be the set of observed voxels, i.e., from the known slices, in patch y_i and $y_i^{\mathcal{O}_i}$ be the corresponding vector of their intensity values:

$$y_i^{\mathcal{O}_i} = \mu_k^{\mathcal{O}_i} + W_k^{\mathcal{O}_i} x_i + \varepsilon_i, \quad (3)$$

where $W_k^{\mathcal{O}_i}$ comprises rows of W_k that correspond to the observed voxel set \mathcal{O}_i . The likelihood of the *observed* data is therefore

$$p(\mathcal{Y}^{\mathcal{O}}; \{W_k\}, \{\mu_k\}, \{\sigma_k^2\}, \pi) = \prod_i \sum_k \pi_k \mathcal{N}(y_i^{\mathcal{O}_i}; \mu_k^{\mathcal{O}_i}, C_k^{\mathcal{O}_i \mathcal{O}_i}), \quad (4)$$

where $\mathcal{Y}^{\mathcal{O}} = \{y_i^{\mathcal{O}_i}\}$, and matrix $C_k^{\mathcal{O}_i \mathcal{O}_i}$ extracts the rows and columns of C_k that correspond to the observed voxel subset \mathcal{O}_i .

2.2 Learning

We learn the maximum likelihood estimates of the parameters $\{\mu_k\}$, $\{W_k\}$, $\{\sigma_k^2\}$ and π under the likelihood (4). As traditional Expectation Maximization for our model does not lead to closed form update equations, we employ the Expectation Conditional Maximization (ECM) [8,12] variant of the Generalized Expectation Maximization, where parameter updates depend on the previous parameter estimates. Due to space limitations, we omit the derivations and provide the resulting updates along with their interpretations.

The **expectation step** updates the class memberships based on the observed voxels of the patches at this location:

$$\gamma_{ik} \triangleq p(k|y_i^{\mathcal{O}_i}; \{W_k\}, \{\mu_k\}, \{\sigma_k^2\}, \pi) \leftarrow \frac{\pi_k \mathcal{N}(y_i^{\mathcal{O}_i}; \mu_k^{\mathcal{O}_i}; C_k^{\mathcal{O}_i \mathcal{O}_i})}{\sum_k \pi_k \mathcal{N}(y_i^{\mathcal{O}_i}; \mu_k^{\mathcal{O}_i}; C_k^{\mathcal{O}_i \mathcal{O}_i})}. \quad (5)$$

Next, we compute the statistics of the low dimensional representation x for each patch as “explained” by each cluster:

$$\hat{x}_{ik} \triangleq E[x_i|k] = ((W_k^{\mathcal{O}_i})^T W_k^{\mathcal{O}_i} + \sigma_k^2 I_{D \times D})^{-1} (W_k^{\mathcal{O}_i})^T (y_i^{\mathcal{O}_i} - \mu_k^{\mathcal{O}_i}), \quad (6)$$

$$S_{ik} \triangleq E[x_i x_i^T | k] - \hat{x}_{ik} \hat{x}_{ik}^T = \sigma_k^2 ((W_k^{\mathcal{O}_i})(W_k^{\mathcal{O}_i})^T + \sigma_k^2 I_{D \times D})^{-1}. \quad (7)$$

The **maximization step** uses the observed voxels to update the model parameters. We let be the set of patches in which voxel j is observed, y_i^j be the j^{th} element of vector y_i , and W_k^j be the j^{th} row of matrix W_k . We update the mean as the average residual of the predicted patch voxels $W_k^j \hat{x}_{ik}$ and the observed values y_i^j :

$$\mu_k^j \leftarrow \frac{\sum_{i \in \mathcal{P}_j} \gamma_{ik} (y_i^j - W_k^j \hat{x}_{ik})}{\sum_{i' \in \mathcal{P}_j} \gamma_{i'k}}. \quad (8)$$

The covariance factors and residual variance are updated based on the statistics of low dimensional representation from (6) and (7):

$$W_k^j \leftarrow \left[\sum_{i \in \mathcal{P}_j} \gamma_{ik} (\hat{x}_{ik} \hat{x}_{ik}^T + S_{ik}) \right]^{-1} \sum_{i \in \mathcal{P}_j} \gamma_{ik} (y_i^j - \mu_k^j) \hat{x}_{ik}^T \quad (9)$$

$$\sigma_k^2 \leftarrow \frac{1}{\sum_j \sum_{i \in \mathcal{P}_j} \gamma_{ik}} \sum_j \sum_{i \in \mathcal{P}_j} \gamma_{ik} [(y_i^j - W_k^j \hat{x}_{ik} - \mu_k^j)^2 + W_k^j S_{ik} (W_k^j)^T]. \quad (10)$$

Finally, we update the cluster proportions:

$$\pi_k = \frac{1}{N} \sum_i \gamma_{ik}. \quad (11)$$

Throughout learning, we work in the atlas space, and approximate voxels as either observed or missing in this space by thresholding interpolation weights and ignoring interpolation effects due to affine alignment. Intuitively, learning such a model with sparse data is possible because each image patch provides a slightly different subset of voxel observations that contribute to the parameter estimation (Fig. 3). The estimation can be extended to carry out the learning by appropriately transforming model parameters into the subject-specific space in order to optimally use the observed voxels, but this leads to computationally prohibitive updates.

In our experiments, all subject scans have the same acquisition direction. Despite different affine transformations to the atlas space for each subject, some voxel *pairs* are still never observed in the same patch, resulting in missing entries of the covariance. Representing covariance using a low-rank approximation helps to alleviate this lack of observations.

2.3 Restoration

To restore an individual patch y_i , we first estimate the most likely cluster \hat{k} for patch y_i by selecting the cluster with the highest membership γ_{ik} . We then estimate the low dimensional representation $\hat{x}_{i\hat{k}}$ given the observed voxels $y_i^{\mathcal{O}_i}$ using (6). Finally, we reconstruct the high resolution patch:

$$z_i = \mu_{\hat{k}} + W_{\hat{k}} \hat{x}_{i\hat{k}} \quad (12)$$

using the estimates of the model parameters $\mu_{\hat{k}}$ and $W_{\hat{k}}$. We average overlapping restored patches using standard techniques and form the restored volume Z_i [13].

2.4 Implementation

We stack together the affinely registered sparse images from the entire collection and split the stack into overlapping subvolumes of $18 \times 18 \times 18$ voxels in the isotropically sampled common atlas space. Subvolumes are centered 9 voxels apart in each direction. Within each subvolume, we learn the mixture model parameters. Instead of selecting just one patch from each volume at a given location, we collect all overlapping patches within the subvolume centered at that location. This aggregation provides more data for each model, which is crucial when working with severely undersampled volumes. Moreover, it offers robustness in the face of image misalignment. Given the learned parameters at each location, we restore all overlapping patches within a subvolume. We use a cubic patch of size $9 \times 9 \times 9$ voxels, and found $K=5$ clusters and $d=21$ to be sufficient. We use a hierarchical implementation, where the subvolume parameters are trained at three iterative scales.

While learning is performed in the common atlas space, we restore each volume in its original subject space to limit effects of interpolation. We apply the inverse of the estimated subject-specific affine transformation to the cluster statistics, and use them to restore patches via (12) in the original subject space.

Our implementation is freely available at <https://github.com/adalca/papago>.

3 Experiments

We demonstrate the proposed imputation algorithm on two datasets and evaluate the results both visually and quantitatively. We also include a preliminary example of how imputation can aid in a skull stripping task.

3.1 Data

ADNI Dataset—We evaluate our algorithm using 326 T1-weighted brain MR images from ADNI [9]. We downsample the isotropic 1 mm³ images to slice separation of 5 mm (1 mm × 1 mm in-plane), and use these low resolution images as input. All subjects are affinely registered to a T1 atlas. The original images serve as the ground truth for quantitative comparisons.

Clinical Dataset—We also demonstrate our algorithm on a clinical set of 127 T2-FLAIR brain MR scans in a stroke patient cohort. These scans are severely anisotropic (0.85 × 0.85 mm in-plane, slice separation of 6 mm). All subjects are affinely registered to a T2-FLAIR atlas and intensity normalized. The slices are resampled to 1.2 mm × 1.2 mm resolution.

3.2 Evaluation

We compare our algorithm to three upsampling methods: nearest neighbour (NN) interpolation, linear interpolation, and non-local means (NLM) upsampling [13]. For ADNI images, we found the hierarchical implementation was unnecessary and only ran the final scale of our algorithm. We compare the reconstructions to the original isotropic volumes both visually and quantitatively (ADNI images only). We use mean squared error

$MSE(Z, Z_o) = \frac{1}{N} \sum (Z - Z_o)^2$ of the reconstructed image Z relative to the original high

resolution scan Z_o , and peak signal to noise ratio $PSNR = \log_{10} \frac{\max(Z_o)}{MSE(Z, Z_o)}$. Both metrics are commonly used in measuring the quality of reconstruction of compressed or noisy signals. Additionally, we illustrate a preliminary example application where skull stripping fails using the original scan and improves dramatically if an imputed image is used.

3.3 Results

Figure 4 illustrates representative restored images for typical subjects in the ADNI dataset. Our method produces more plausible structure both in coronal and sagittal slices. The method restores anatomical structures that are almost entirely missing in the other reconstructions, such as the dura or the sulci of the temporal lobe. Figure 5 reports the error statistics in the ADNI data. Due to high variability among subject scans, we report

improvements of each method over the nearest neighbor interpolation baseline in the same scan. Our algorithm offers significant improvement compared to the linear interpolation and NLM.

We also show a preliminary result where imputed data facilitates downstream image analysis. Specifically, the first step in many analysis pipelines is brain extraction – isolating the brain from the rest of the anatomy. Typical algorithms assume that the brain consists of a single connected component separated from the skull and dura by cerebral spinal fluid [20], and often fail on sparsely sampled scans that no longer have clear contrast between these regions. Figure 6 illustrates an example result where the brain extraction fails on the original subject scan but succeeds on our reconstructed image.

Figure 7 illustrates representative restoration improvements in the clinical population. Our method produces more plausible structure, as can be especially seen in the close-up panels focusing on anatomical details.

4 Discussion and Conclusions

We propose an image imputation method that employs a large collection of low-resolution images to infer fine-scale anatomy of a particular subject. We introduce a model that captures structural similarity across subjects in large clinical image collections, and fills in the missing data in low resolution scans. The method produces anatomically plausible volumetric images consistent with sparsely sampled input scans.

Latent Structure

In this paper, we explicitly model and estimate the latent low-dimensional embedding for each patch. Additionally, we also explored an alternative choice that instead models each missing voxel as a latent variable. The resulting ECM algorithm estimates the expected missing voxel statistics directly, and then updates the cluster parameters. The most notable difference between this formulation and simpler algorithms that iteratively fill in missing voxels and then estimate GMM model parameters is in the estimation of the expected data covariance, which captures the covariance of the missing and observed data (c.f. [12], Chap. 8). We found that this variant often got stuck in local minima and had difficulty moving away from the initial missing voxel estimates, and was an order of magnitude slower than the presented method. We provide both implementations in our code.

Initialization

In contrast to the classical EM algorithm, the M-step of the ECM algorithm employs previous parameter estimates to perform parameter updates. This makes the initialization more challenging compared to the classical GMM learning, where initializing cluster memberships is sufficient, and also leads to slower convergence than simpler GMMs. We experimented with several initialization schemes, and provide them in our implementation. The experimental results are initialized by first learning a simple GMM from the linearly interpolated volumes, and using the resulting parameter as initializations for our method. This leads to results that improve on the linear interpolation but still maintain slightly blocky effects caused by interpolation. More agnostic initializations, such as random parameter

values, lead to more realistic anatomy but noisier final estimates. Different methods perform well in different regions of the brain. Future research will further investigate the effects of initialization on the resulting reconstruction.

Restoration Method

Our restoration method, assumes that the observed voxels are noisy manifestations of low dimensional patch representations, and reconstructs the entire patch, including the observed voxels, leading to smoother images. We also explored an alternative reconstruction method of filling in just the missing voxels given the observed voxels (not shown). This formulation imputes the most likely missing voxels assuming the observed voxels are true observations, leading to sharper but noisier patches. The two restoration methods therefore yield images with different characteristics. This tradeoff is a function of the noise in the original acquisition: higher noise in the clinical acquisition leads to noisier reconstructions using the alternative method, whereas in the ADNI dataset the two methods perform similarly.

Varying Resolution

The proposed method can be used for general image imputation using datasets of various resolutions. For example, while acquiring a large high resolution dataset for a clinical study is often infeasible, our algorithm will naturally make use of any additional image data available. Even a small number of acquisitions in different directions or higher resolution than the study scans promise to improve accuracy of the resulting reconstruction.

Slice Thickness

In many clinical datasets the slice spacing is unknown or varies by site, scanner or acquisition. Therefore, throughout our model we simply treat the original data as high resolution planes. Explicitly modeling varying slice thickness is an interesting direction of future research.

Our method does not require high volumetric resolution scans or expert annotations, but can instead build the missing structure by learning from collections of clinical scans of similar quality to that of the input image. This enables the use of untapped clinical data for large scale scientific studies, promising to facilitate novel clinical analyses.

Acknowledgments

We acknowledge the following funding sources: NIH NINDS R01NS086905, NIH NICHD U01HD087211, NIH NIBIB NAC P41EB015902, NIH R41AG052246-01, 1K25EB013649-01, 1R21AG050122-01, and Wistron Corporation.

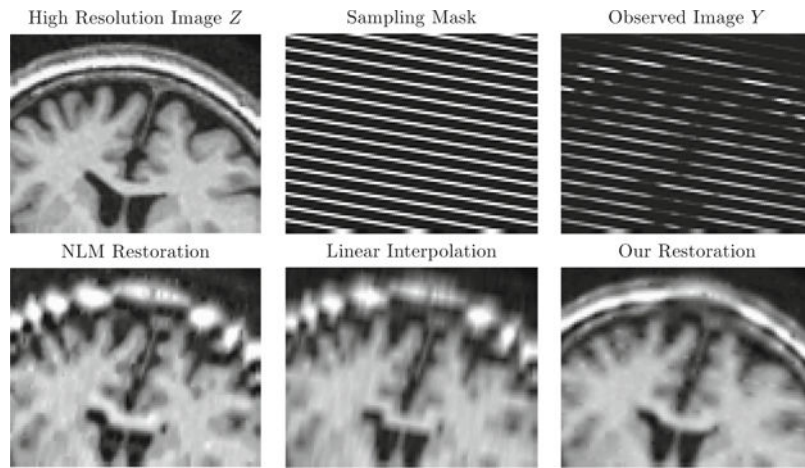
References

1. Bhatia KK, Rao A, Price AN, Wolz R, Hajnal JV, Rueckert D. Hierarchical manifold learning for regional image analysis. *TMI*. 2014; 33(2):444–461.
2. Carmi E, Liu S, Alon N, Fiat A, Fiat D. Resolution enhancement in MRI. *Magn Reson Imaging*. 2006; 24(2):133–154. [PubMed: 16455402]
3. Coupé P, Manjón JV, Fonov V, Pruessner J, Robles M, Collins LD. Patch-based segmentation using expert priors: application to hippocampus and ventricle segmentation. *NeuroImage*. 2011; 54(2): 940–954. [PubMed: 20851199]

4. Di Martino A, Yan CG, Li Q, Denio E, Castellanos FX, Alaerts K, et al. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Mol Psychiatry*. 2014; 19(6):659–667. [PubMed: 23774715]
5. Elad M, Aharon M. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE TMI*. 2006; 15(12):3736–3745.
6. Hill DL, Batchelor PG, Holden M, Hawkes DJ. Medical image registration. *Phys Med Biol*. 2001; 46(3):R1. [PubMed: 11277237]
7. Iglesias, JE., Konukoglu, E., Zikic, D., Glocker, B., Leemput, K., Fischl, B. Is synthesizing MRI contrast useful for inter-modality analysis?. In: Mori, K.Sakuma, I.Sato, Y.Barillot, C., Navab, N., editors. MICCAI 2013. Vol. 8149. Springer; Heidelberg: 2013. p. 631-638.LNCS
8. Ilin A, Raiko T. Practical approaches to principal component analysis in the presence of missing values. *J Mach Learn Res*. 2010; 11:1957–2000.
9. Jack CR, Bernstein MA, Fox NC, Thompson P, Alexander G, Harvey D, et al. The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods. *J Magn Reson Imaging*. 2008; 27(4):685–691. [PubMed: 18302232]
10. Jog, A., Carass, A., Prince, JL. ISBI. IEEE; 2014. Improving magnetic resonance resolution with supervised learning; p. 987-990.
11. Konukoglu, E., van der Kouwe, A., Sabuncu, MR., Fischl, B. Example-based restoration of high-resolution magnetic resonance image acquisitions. In: Mori, K.Sakuma, I.Sato, Y.Barillot, C., Navab, N., editors. MICCAI 2013. Vol. 8149. Springer; Heidelberg: 2013. p. 131-138.LNCS
12. Little, RJ., Rubin, DB. Statistical Analysis with Missing Data. Wiley; Hoboken: 2014.
13. Manjon JV, Coupé P, Buades A, Fonov V, Louis Collins D, Robles M. Non-local MRI upsampling. *Med Image Anal*. 2010; 14(6):784–792. [PubMed: 20566298]
14. Manjón JV, Coupé P, Buades A, Louis Collins D, Robles M. New methods for MRI denoising based on sparseness and self-similarity. *Med Image Anal*. 2012; 16(1):18–27. [PubMed: 21570894]
15. Plenge, E., Poot, DHJ., Niessen, WJ., Meijering, E. Super-resolution reconstruction using cross-scale self-similarity in multi-slice MRI. In: Mori, K.Sakuma, I.Sato, Y.Barillot, C., Navab, N., editors. MICCAI 2013. Vol. 8151. Springer; Heidelberg: 2013. p. 123-130.LNCS
16. Rost NS, Fitzpatrick K, Biffi A, Kanakis A, Devan W, Anderson CD, Cortellini L, Furie KL, Rosand J. White matter hyperintensity burden and susceptibility to cerebral ischemia. *Stroke*. 2010; 41(12):2807–2811. [PubMed: 20947843]
17. Rousseau, F., Kim, K., Studholme, C. ISBI. IEEE; 2010. A groupwise super-resolution approach: application to brain MRI; p. 860-863.
18. Rousseau F, Habas PA, Studholme C. A supervised patch-based approach for human brain labeling. *IEEE Tran Med Imag*. 2011; 30(10):1852–1862.
19. Schoenberg, IJ., Schoenberg, IJ. Cardinal Spline Interpolation. Vol. 12. SIAM; Philadelphia: 1973.
20. Ségonne F, Dale AM, Busa E, Glessner M, Salat D, Hahn HK, Fischl B. A hybrid approach to the skull stripping problem in MRI. *Neuroimage*. 2004; 22(3):1060–1075. [PubMed: 15219578]
21. Sridharan, R., et al. Quantification and analysis of large multimodal clinical image studies: application to stroke. In: Shen, L.Liu, T.Yap, P-T.Huang, H.Shen, D., Westin, C-F., editors. MBIA 2013. Vol. 8159. Springer; Cham: 2013. p. 18-30.LNCS
22. Yang J, Wright J, Huang TS, Ma Y. Image super-resolution via sparse representation. *IEEE Trans Image Process*. 2010; 19(11):2861–2873. [PubMed: 20483687]
23. Zoran D, Weiss Y. Natural images, Gaussian mixtures and dead leaves. *Advances in Neural Information Processing Systems*. 2012:1736–1744.



Fig. 1.
Axial, sagittal and coronal slices from an example scan in our clinical dataset.

**Fig. 2.**

Problem setup and preview of the results. Top row: (unobserved) isotropic image Z that we seek to recover is sampled according to the sampling mask which produces observed image Y . Bottom row: restoration results for non-local means (NLM) upsampling, linear interpolation, and our method, respectively. The most dramatic improvement can be seen in restorations of the skull, dura matter, and ventricles.

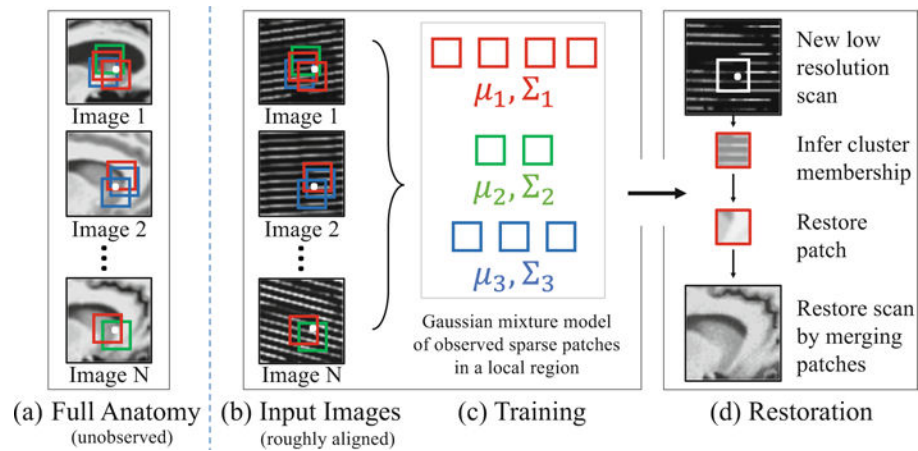
**Fig. 3.**

Image imputation. (a) Full resolution images, for illustration only. These are unobserved by the algorithm. (b) Sparse planes acquired in clinical scans. (c) During learning, we train a GMM that captures the low dimensional nature of patch variability at a particular location (white dot). (d) Given an image from the collection, or a new image, we infer the most likely cluster for each 3D patch, and restore the missing data using the learned model and the observed voxels. We quilt the final volume from overlapping restored patches. 2D images are shown for illustration only, the algorithms operate fully in 3D.




Fig. 4. Representative ADNI restorations. Representative reconstruction by NLM, linear interpolation, and our method, and the original high resolution images for two representative subjects in the study. Our method reconstructs more anatomically plausible substructures as can be especially seen in the close-up panels, for example in the skull or temporal lobe.




Fig. 5. Reconstruction accuracy statistics. Accuracy improvement over nearest neighbor interpolation for different image restoration methods. Left: MSE (lower is better), Right: PSNR (higher is better). Our method performs significantly better. For Nearest neighbor interpolation, $\text{MSE} = 0.004 \pm 0.001$ and $\text{PSNR} = 23.7 \pm 1.1$. All statistics were computed over 40 scans randomly chosen from the ADNI dataset. Image intensities are scaled to a $[0, 1]$ range.




Fig. 6. Skull stripping example. Example of a skull stripping failure for linear and NLM interpolation. Skull stripping dramatically improves when applied to the imputed image for this example.




Fig. 7. Representative clinical restorations. Reconstruction using NLM, linear interpolation and our method for two representative subjects. Our method reconstructs more plausible substructures, as can be especially seen in the close-up panels of the skull and periventricular regions.