

Low-rank Atlas Image Analyses in the Presence of Pathologies

Xiaoxiao Liu, Marc Niethammer, Roland Kwitt,
Nikhil Singh, Matt McCormick, Stephen Aylward

Abstract—We present a common framework, for registering images to an atlas and for forming an unbiased atlas, that tolerates the presence of pathologies such as tumors and traumatic brain injury lesions. This common framework is particularly useful when a sufficient number of protocol-matched scans from healthy subjects cannot be easily acquired for atlas formation and when the pathologies in a patient cause large appearance changes.

Our framework combines a low-rank-plus-sparse image decomposition technique with an iterative, diffeomorphic, group-wise image registration method. At each iteration of image registration, the decomposition technique estimates a “healthy” version of each image as its low-rank component and estimates the pathologies in each image as its sparse component. The healthy version of each image is used for the next iteration of image registration. The low-rank and sparse estimates are refined as the image registrations iteratively improve.

When that framework is applied to image-to-atlas registration, the low-rank image is registered to a pre-defined atlas, to establish correspondence that is independent of the pathologies in the sparse component of each image. Ultimately, image-to-atlas registrations can be used to define spatial priors for tissue segmentation and to map information across subjects.

When that framework is applied to unbiased atlas formation, at each iteration, the average of the low-rank images from the patients is used as the atlas image for the next iteration, until convergence. Since each iteration’s atlas is comprised of low-rank components, it provides a population-consistent, pathology-free appearance.

Evaluations of the proposed methodology are presented using synthetic data as well as simulated and clinical tumor MRI images from the brain tumor segmentation (BRATS) challenge from MICCAI 2012.

Index Terms—Low-Rank and Sparse Decomposition, Sparse Images, Unbiased Atlas, Atlas-based Segmentation

I. INTRODUCTION

HEREIN we present a common framework for (1) registering images to atlases and (2) forming unbiased atlases from a collection of images. This framework specifically addresses the challenging situation in which the images contain large, deformation inducing, unsegmented pathologies.

Image-to-atlas registration is used to assess and plan the treatment of patients suffering from traumatic brain injuries (TBI), brain tumors, or stroke [1]. For these cases, image-to-atlas registration is used to estimate tissue priors and to map adjunct information, such as functional site locations, from that atlas into the patient. However, when those patient

images contain large lesions, then lesion-induced appearance changes may inhibit image-to-atlas registrations and confound the tissue priors and mappings.

Atlas image formation typically involves registering images from a sample of healthy subjects, however, in clinical practice, acquiring such a collection of images can be time consuming, expensive, and possibly pose a risk to the healthy subjects. For research projects having limited time and financial resources or involving new imaging protocols or children, it can be problematic to obtain a sufficient number of protocol-matched scans from healthy subjects for atlas formation.

Several excellent approaches to image-to-atlas registration involving images containing pathologies have been previously proposed. One of the most straightforward methods to eliminate a lesion’s influence during registration is “masking.” Masking prevents a lesion’s content from being considered during the computation of the image similarity metric. Other methods attempt to address this problem by joint registration and segmentation which tolerates missing correspondences [2], geometric metamorphosis that separates estimating healthy tissue deformation from modeling tumor change [3], or personalized atlas construction that accounts for diffeomorphic and non-diffeomorphic changes [4]. While effective, these methods require explicit lesion or anatomical structures segmentations or initial tumor localizations which, in many cases, are actually the goal of the process. In [5], a patient-specific, piecewise, most-similar atlas was proposed by combining subgroups of images selected for each local region. The selection criterion was based on the degree of contraction and dilation of the structures, thus it can tolerate anatomical variabilities. The patient-specific atlas was then used for anatomical structure segmentations for head and neck CT images.

On the other hand, forming an unbiased atlas from images containing pathologies has not been the focus of much research. However, unbiased atlas formation when lesions are not present has been well studied in the context of estimating within-population variability and cross-population differences [6]. During most of those atlas formation processes, individual images are iteratively mapped into a common coordinate system. This construction process can be formulated as a Fréchet mean estimation via diffeomorphic image registrations. Typically, lesions in the input images would not only degrade registration accuracy but also propagate into the resulting atlas and corrupt its variability estimates.

We propose to eliminate the effects of lesions during image-to-atlas registration and atlas formation by leveraging the inherent *low-rank* structure of the input data. This is

Copyright (c) 2010 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

combined with separating out the *sparse* components in the data which are not consistent with the low-rank structure. This technique, commonly referred to as *low-rank plus sparse* matrix decomposition [7] can mitigate the confounding effects of the pathologies. In a preliminary study, cf. [8], we have shown how low-rank plus sparse decomposition can be integrated into an iterative image registration framework to match individual images containing pathologies to a normal atlas. In this article, we briefly review that work and present a common framework that encompasses that work and extends it to include estimating an unbiased low-rank atlas from data containing pathologies. The common framework represents a new theoretic focus and results in implementation improvements to the original work, e.g., introduces a non-greedy optimization strategy. Correspondingly, the experiments in this paper are new and primarily focus on the parameters and robustness of the common framework and its application to unbiased atlas formation, rather than focusing on individual-to-atlas registration evaluation as in [8].

The advantages of unbiased, low-rank atlas formation are two-fold: (1) population information is exploited to assess which parts of an image are likely lesions (i.e., they are inconsistent with the population) and which parts of an image should be considered normal anatomy. This is done *without* explicit individual pathology segmentations; and (2) the recovery of the low-rank structures replaces the visual appearance of the lesion regions (in each image) with population-consistent normal appearance. As a result, the effect of pathologies on the estimated atlas is greatly reduced, if not eliminated.

The paper is structured as follows: Section II introduces the proposed methodology. Section III presents experimental results on both synthetic and clinical data, evaluated in comparison to conventional image-to-atlas registration and unbiased atlas building methods. Advantages and limitations of the method are discussed in Section IV.

II. METHODOLOGY

We first introduce the low-rank and sparse decomposition technique (Section II-A) and the classic unbiased atlas formation method (Section II-B), which are key components of our proposed framework (Section II-C).

A. Low-rank and sparse decomposition

In [7] Peng et al., the authors propose to decompose a matrix of vectorized images into the sum of a low-rank and a sparse component in the context of simultaneous rigid image alignment. The intuition is that the portion of each image that *cannot* be explained by the low-rank model is allocated to the sparse part. Hence, the low-rank component could be interpreted as a blending of recorded values and values inferred from the population; the sparse component then contains each subject's *anomalous* values. Technically, the allocation of image intensities to each of those components is driven by the amount of linear-correlation across the images. Given a collection of n images with m voxels, we have:

D a $M \times N$ matrix in which each image I_i is a column vector that contains the m spatially-ordered voxel intensities in I_i .

L a $M \times N$ matrix that contains the low-rank representations L_i for each of the images in the collection D .

S a $M \times N$ matrix that is the sparse component such that $S_i = D_i - L_i$.

The low-rank representation of D is then defined as

$$\begin{aligned} \{L^*, S^*\} &= \arg \min_{L, S} (\|L\|_* + \lambda \|S\|_1) \\ &\text{subject to } D = L + S \end{aligned} \quad (1)$$

where $\|\cdot\|_*$ denotes the nuclear norm (surrogate for the rank) and $\|\cdot\|_1$ denotes the 1-norm (surrogate for sparsity). The optimal $\{L^*, S^*\}$ can be efficiently computed using an augmented Lagrangian multiplier approach [9]. The optimization problem is convex as both the nuclear and the 1-norm are convex, the constraint is linear and the problem is defined over a convex domain. Hence, a globally optimal solution can be obtained.

By defining

$$\begin{aligned} X &= \{L, S\}, \\ f(X) &= \|L\|_* + \lambda \|S\|_1, \\ h(X) &= D - L - S \end{aligned} \quad (2)$$

we can solve the low-rank representation using an augmented Lagrangian function, which is defined as

$$\begin{aligned} g(L, S, Y, \mu) &= \|L\|_* + \lambda \|S\|_1 + \langle Y, D - L - S \rangle \\ &\quad + \frac{\mu}{2} \|D - L - S\|_F^2 \end{aligned} \quad (3)$$

where Y is the Lagrange multiplier and μ is a positive scalar of the ALM (Augmented Lagrangian Multiplier Method). More implementation details can be found at [9].

B. Unbiased low-rank atlas formation

An unbiased atlas of a sample population is defined as the representative image \hat{I} that requires the least amount of energy to deform into each individual image I_i from the population [6]. The energy can be formulated as:

$$\begin{aligned} E(\{\Phi_i\}, \hat{I}) &= \sum_{i=1}^N \text{Reg}[\Phi_i] + \frac{1}{\sigma^2} \text{Sim}[I_i \circ \Phi_i, \hat{I}], \\ &= \sum_{i=1}^N \int_0^1 \|\kappa v_i(t)\|^2 dt + \frac{1}{\sigma^2} \|I_i \circ \Phi_i - \hat{I}\|_2^2, \\ &\text{s.t. } \frac{d\Phi_i(t)}{dt} = v_i(\Phi_i, t), t \in [0, 1] \end{aligned} \quad (4)$$

where the regularity term $\text{Reg}[\Phi_i]$ measures the smoothness of the deformation, and the similarity term $\text{Sim}[I_i \circ \Phi_i, \hat{I}]$ measures the differences between the estimated atlas and the deformed individual images. The deformation fields $\Phi_i(t)$ are defined as the flow of smooth time-indexed velocity fields $v_i(t)$ and are generated by integrating $v_i(t)$, forward in time. The penalty on the diffeomorphic deformation between the pair of images are formulated as the Sobolev norm via a partial differential operator κ on the time-dependent velocity vector fields $v(t)$. This differential operator κ also governs the smoothness of the resulting deformation fields. Image

similarity is measured by the the sum-of-squared intensity differences (SSD): $\|L_i \circ \Phi_i - \hat{I}\|_2^2$.

The minimization of the energy in Eq. (4) can then be simplified by an alternating optimization scheme in which, for fixed Φ_i , the image \hat{I} that minimizes the energy is the voxel-wise arithmetic mean of the deformed images:

$$\hat{I} = \frac{1}{N} \sum_{i=1}^N L_i \circ \Phi_i. \quad (5)$$

The stability convergence study in [6] suggests that roughly ten images are needed to create a stable atlas representing neuroanatomy.

C. Iterative low-rank image registration framework

In our preliminary study [8], we have proposed to integrate low-rank plus sparse decomposition into an iterative registration framework in which a group of input images, potentially containing large pathologies and deformations, are registered to a normal-control atlas. Our premise is that by identifying the low-rank plus sparse components of each input image, its low-rank component, which contains reduced or eliminated pathologies, can be more accurately registered with a normal-control atlas, compared to the *direct* registration of an image containing a pathology to an atlas.

The low-rank plus sparse decomposition exploits the fact that lesions generally do not manifest in consistent locations or with consistent appearance in populations. These inconsistencies result in lesions being reduced in the low-rank component and allocated to the sparse component. In this iterative registration framework, we seek to minimize the following energy:

$$E(\{\Phi_i\}, \hat{I}) = \sum_{i=1}^N \text{Reg}[\Phi_i] + \frac{1}{\sigma^2} \text{Sim}[L_i \circ \Phi_i, \hat{I}], \quad (6)$$

where i denotes the identifier of input image L_i , $\text{Reg}[\Phi_i]$ is the regularity measure for deformation Φ_i , penalizing spatially non-smooth transformations; L_i is the low-rank image from the low-rank plus sparse matrix decomposition; \hat{I} is the target reference image to which each L_i is mapped onto during each registration step and $\text{Sim}[L_i \circ \Phi_i, \hat{I}]$ is the image similarity measurement between the deformed low-rank image $L_i \circ \Phi_i$ and the atlas \hat{I} .

When applying this framework to atlas-based segmentation applications, \hat{I} is fixed to be a provided atlas which is often a population atlas formed by a group of selected images from healthy individuals from unrelated imaging sources.

When extending this framework to forming an unbiased atlas, on the other hand, the goal is to estimate an atlas such that it is central with respect to the data population. In this sense, we replace \hat{I} in the above framework with an average low-rank image \hat{L} which will converge to an unbiased atlas of the low-rank population, namely the *unbiased low-rank atlas*. The advantage of the low-rank atlas is its ability to recover a “normal” atlas from “corrupted” input images.

Our alternating optimization strategy to solve the minimization problem with energy Eq (6) is as follows: we first keep

\hat{L} fixed while solving for $\{\Phi_i\}$ and subsequently keep the transformations $\{\Phi_i\}$ fixed while solving for \hat{L} . The first part performs independent pairwise registrations between L_i and the fixed image \hat{L} . The second part requires minimizing the dissimilarity measurement

$$E(\hat{L}) = \sum_{i=1}^N \|L_i - \hat{L}\|_2^2, \quad (7)$$

which is achieved by the arithmetic mean of the images, i.e.,

$$\hat{L} = \frac{1}{N} \sum_{i=1}^N L_i \circ \Phi_i. \quad (8)$$

Note that when fixing the atlas \hat{L} , unbiased atlas construction simplifies to a group-wise registration. The group-wise approach is essential because it allows for improved decomposition of the images into low-rank/sparse components.

The general unbiased low-rank atlas framework for our method is shown in Fig. 1. The algorithm steps are listed as follows:

- (0) Initialization: Affinely align input images, with respect to (w.r.t.) a common reference atlas image, to prepare the initial iteration data I_i^0 , set $k = 1$;
- (1) For iteration k , compute the low-rank image L_i^k from $\{I_1^k, \dots, I_N^k\}$ via low-rank and sparse decomposition, solving Eq. (1);
- (2) Compute the arithmetic mean image \hat{L}^k from $\{L_1^k, \dots, L_N^k\}$;
- (3) Solve for diffeomorphic transforms $\{\Phi_1^k, \dots, \Phi_N^k\}$, each Φ_i^k maps the low-rank image $L_i^k \circ (\Phi_i^{(k-1)})^{-1}$ to the current low-rank atlas \hat{L}^k ;
- (4) Apply the transforms from the previous step to update the input images: $I_i^{(k+1)} = I_i^0 \circ \Phi_i^k$;
- (5) Set $k \leftarrow k + 1$ and continue with step (1) until convergence.

Note that in step (3) we first deform the low-rank image L_i^k back to the original input image’s physical space using the *inverse* transform of the previous iteration’s deformation ($\Phi_i^{(k-1)}$) before calculating the mapping. The purpose is to avoid accumulated errors when composing the deformations over the iterations. This non-greedy strategy is theoretically more rigorous and is considered to be a major improvement comparing to the greedy version proposed in [8]. Convergence is reached when the total change in deformation is small enough (a numerically small tolerance threshold can be set).

Setting the parameters of our equations is straightforward. Convergence typically occurred within 5~10 iterations. The weight λ in Eq. (1) was adjusted for each data type, e.g., T2 MRI or FLAIR MRI. The range was roughly 0.5~1.0. The weight could be determined by any of several factors, including the type of imaging modality, sample size, image dimensions, and lesion intensity contrast. At first, we heuristically determined the weight for each experiment by evaluating several different weights on the initial input images for a low-rank and sparse decomposition, and we picked the one which seemed to produce the best separation between the low-rank/normal and the sparse/lesion variations within

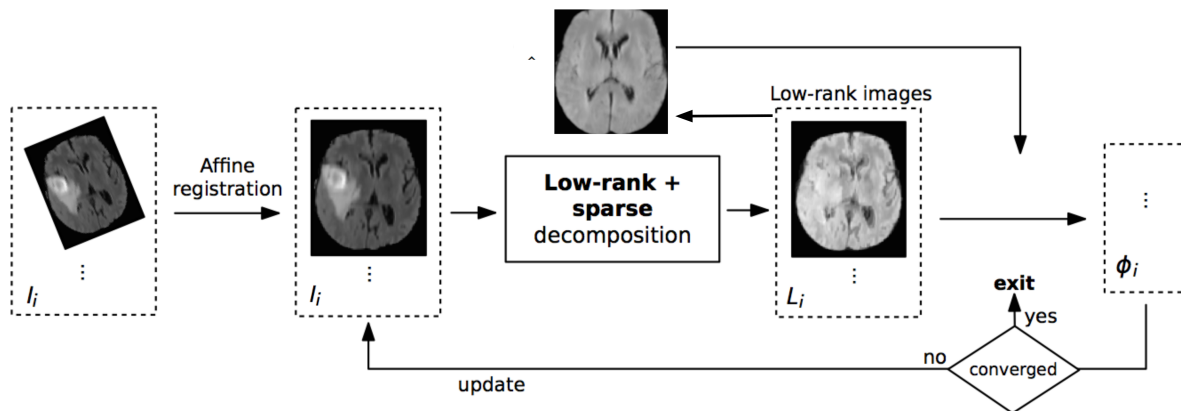


Fig. 1: An illustration of the proposed low-rank atlas formation framework, where I_i refers to the i -th input image, D_i is the i -th vector of the input matrix D , L_i is low-rank component of the input image and ϕ_i refers to the registration map generated from diffeomorphic image registration at each iteration.

the populations. Subsequently, we conducted a simple weight sensitivity study using the synthetic example. That study is given in Section III-B2, and it revealed that our method converges to similar results for a relatively wide range of weight values.

The precise control of the quality of the decomposition would require the precise quantification of the efficacy of the sparse components, and such considerations are highly application specific and beyond the scope of this initial theoretic and exploratory paper.

Implementation. The proposed framework is largely implemented in Python¹ and uses ANTS [10] for image registrations. Running on 8-input MRI volumes (of a voxel size $140 \times 190 \times 155$) on a dual-core 16GB RAM PC takes about 3 to 5 hours. Each low-rank plus sparse decomposition merely takes a minute and most of the computation time is spent during dense diffeomorphic registrations. With recent advances in adapting image registration algorithms to the GPU, total runtime could potentially be reduced to minutes.

III. VALIDATION

In this section, we first introduce the evaluation metrics and then present a series of experimental results.

A. Evaluation metrics

The following metrics are used to evaluate our iterative atlas-based image registration framework:

1) *TCS*D for atlas-based segmentation accuracy: The premise of atlas-based segmentation is that by registering an atlas with a target image, the tissue labels in the atlas provide spatial priors for the tissues in the target image. When image-to-atlas registration is successful, the tissue labels of the atlas should align with the corresponding tissue labels in the target image. Therefore we compute the standard deviation of the target image intensities under each tissue label in the atlas. Smaller *tissue-class standard deviation (TCS*D) values indicate

more accurate image-to-atlas registrations. It is important to note that we exclude the tumor region when calculating TCS

2) *Dissimilarity Metrics for low-rank atlas estimation accuracy:* For simulated ground-truth (pathology-free) data to which corruptions are artificially added, we can easily measure the accuracy of the estimated low-rank atlas by comparing it with an unbiased atlas formed from the ground-truth.

Two metrics are used to compare the atlases. First, we use the same similarity metric, i.e., *SSD*, that we used to drive the optimization in Eq. (4) to measure the dissimilarity between the estimated low-rank atlas and the unbiased ground-truth atlas. Second, we use *geodesic distance (GD)* which measures the amount of diffeomorphic deformation it takes to match one atlas to the other. Geodesic distance is calculated by integrating the deformable registration's velocity field over time. In most cases these two metrics are consistent with each other (e.g., larger *SSD* values typically indicate larger amounts of deformation).

3) *Visual inspection and entropy measurement for atlas image quality:* We can qualitatively assess the iterative registration process by visual examinations of the low-rank and sparse components of each image, after each iteration.

In image-to-atlas registrations, a particular image's sparse component at the final iteration (i.e., after reaching convergence) should be specific to the lesion(s). Conversely, at the final iteration the low-rank component should contain only healthy-looking tissue, that is well aligned with the healthy atlas.

To visually evaluate the quality of the unbiased low-rank atlas, we can inspect the appearance of the atlas image in terms of the *sharpness* and *shape* of high-contrast structures. Similar to conventional unbiased atlas formation, the low-rank atlas obtained in the first iteration of the registration is typically quite blurry and then sharpens up over the iterations. Residues of pathology structures or missing normal structures indicate unexpected low-rank plus sparse decomposition results, typically due to the fact that pathologies do not constitute random corruptions of the image, but are rather structured

¹available online at <https://github.com/KitwareMedical/pyLAR>

| Patient | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|---------|------|------|------|------|------|------|------|------|
| Ours | 46.7 | 42.2 | 44.9 | 48.0 | 46.3 | 43.0 | 47.3 | 62.3 |
| BSpline | 68.0 | 48.3 | 50.8 | 53.4 | 40.5 | 53.7 | 58.2 | 39.4 |

TABLE I: TCSD comparison to traditional BSpline image-to-atlas registration for the GM class.

outliers. Quantitatively, we use image intensity entropy as a goodness metric of the atlas image quality [6], [11]. The discrete image entropy is defined as an expected uncertainty in the random variable associated with the intensities of the given image, yielding entropy measured in bits. Sharp images have relatively low entropy, while blurry or noisy images (relatively flat histograms) tend to have higher entropy.

B. Case studies

The following three data sources are used for validation: a) simulated “bullseye” data, used as a toy example for testing and illustration; b) synthetic tumor MRI images, generated using TumorSim [12] and available as part of the BRATS’12 [13] challenge; c) clinical MRI from tumor patients, also available as part of the BRATS’12 challenge.

1) *Atlas-based tissue segmentation:* Here we summarize a study that we previously conducted to assess the utility of our framework when a healthy atlas is used as the target atlas \hat{L} [8]. In this study, the SRI24 atlas [14] was used as the target healthy atlas for image-to-atlas registrations and to provide gray-matter (GM), white-matter (WM), and cerebrospinal fluid (CSF) tissue labels. These labels enabled us to compute, after registration, the TCSD metric of Section III-A1 for each tissue class. A subset of 8 FLAIR images from the BRATS ’12 challenge data were used.

Quantitatively, the box plots in Fig. 2 show that the overall TCSD for the GM class improves over the iterations. Each box summarizes the TCSD values from all 8 patients at each iteration. The numbers in Table I confirm that our framework outperforms BSpline registration in most cases in this study. However, our method performs worse on two cases, patients 5 and 8. On inspection, we realized that these images are quite different from the six other images in the sample. Both have much narrower and more distorted ventricles. Due to their distinctive appearance the ventricular areas are allocated to the sparse component instead of the low-rank component, and thereby the ventricles are excluded from consideration during registration and overall alignment quality is degraded. This is a failing of our sample size; eight subjects are too few to fully capture normal variability using our framework. If more patient images with similar ventricular variations would have been included, it is very likely that the low-rank plus sparse decomposition would have been more effective and the ventricles would have appropriately driven the registrations for patients 5 and 8.

2) *Low-rank atlas estimation results on bullseye data:* We simulated a group of bullseye images, see Fig. 3(a), for testing and illustration purposes. Eight bullseye patterns (*i.e.*, a 2D image) are generated by composing three concentric disks of different radii and intensity. The radii of the middle ring varies in these images. These radii changes are meant to represent

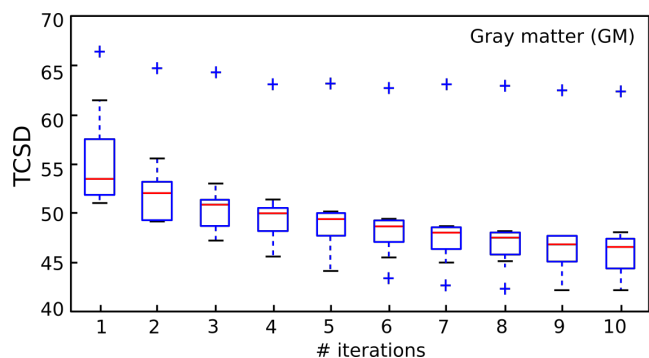


Fig. 2: Change in TCSD values for GM labels transcribed from the atlas after each iteration of our method.

normal anatomic variations. Furthermore, we inserted a bright disk with random radius at a random location into each image, to simulate pathologies.

Influence of the weight parameter on atlas generation:

In theory, the influence of the weight parameter could be dependent on the type and size of the pathologies present, sample size, amount of normal variation in the data, and numerous other factors. We used the bullseye data to study its influence in practice. Fig. 3(a), shows the first iteration of the low-rank plus sparse decomposition results using two selected weights. With the weight being 0.6, the initial decomposition allocates not only the “tumor” to the sparse component, but also some of the true structural variation, as illustrated by the ring artifacts in Fig. 3(a). When the estimation converges after seven iterations, the sparse components contains much less non-tumor structures, while the low-rank component visually appears much sharper and closer to the true mean geometry. When the weight is set to be 1.0, after a single iteration the separation between the background geometry and the inserted “tumor” blobs are much cleaner as shown on the lower half of the Fig. 3(a). After convergence, it too produced an atlas that is appears to be similar to the true mean geometry. Most importantly, while neither weight produces optimal decompositions for all cases in the first iteration, both weights do converge after a similar number of iterations and do produce similar atlas images that are very close to the ground truth.

A quantitative analysis of the influence of the weight parameters on the final atlas generated is shown in Fig. 3(b). The relatively flat curve between the weights 0.4 and 1.6, and their small SSD values, shows that the proposed framework converges to a stable low-rank atlas for a wide range of weight values.


Influence of the portion of data containing pathologies:

In order to evaluate the robustness of the proposed method to the presence of pathologies, we studied how the final atlas varies when the portion of patients that contain pathologies in the input population is changed. Specifically, we generated various mixtures of the “tumor-free” and the “tumor” injected data from the BRATS challenge and compared low-rank atlas results with traditional unbiased atlas results using those mixtures. The results are given in Fig. 3(c). Low-rank atlas formation exhibits stable performance when the input sample


has patients with no pathologies or when every patient has a pathology, while traditional atlas building results degrade dramatically when the portion of tumor data is increased in the input sample.

3) *Low-rank atlas estimation results from brain MRI with simulated tumors:* The synthetic brain tumor dataset in BRATS'12 was created by injecting tumors into the MRI

data from five normal BrainWeb [15] subjects. Tumors were injected into the tumor-free images via a series of physical and statistical modeling techniques [12]. Local image deformations were induced to simulate inter-subject variations. While the use of simulated inter-subject differences may induce bias and degrade generalizability of results, the utility of this data arises from being a public standard dataset used in the



(a) Low-rank atlas estimation results, after a single iteration of our framework, on a group of 8 simulated bullseye images. The upper half shows the first iteration low-rank and sparse decomposition result with a weight of 0.6: the top row shows the original input images, the second row shows the low-rank components and the third row shows the sparse components; The lower half visualizes the first iteration decomposition results with a weight of 1.0. Note that the sparse images are rescaled to highlight the differences for illustration.



(b) SSD (to the ground-truth atlas) of the estimated low-rank atlas with varying weights on the sparse component. The estimated atlases are displayed at four weight values from the left to the right: 0.4, 0.6, 1.0 and 1.9. The ground truth atlas, for comparison, is displayed on the rightmost.

(c) SSD comparison between the estimated low-rank atlases and the traditional unbiased atlases as the portion of patients containing pathologies in the training sample is increased.

Fig. 3: Evaluation of low-rank atlas building on the bullseye dataset.




Fig. 4: Low-rank atlas estimation results on synthetic FLAIR images: a) The axial slices of the eight input FLAIR images. Our estimated low-rank FLAIR atlas (c) is compared with the conventional unbiased atlas (b) and the sample mean T1 atlas (d) formed from healthy BrainWeb images.




Fig. 5: Low-rank atlas estimation results on synthetic T1 images : a) The axial slices of the eight input T1 images. Our estimated low-rank atlas (c) is compared with the conventional unbiased atlas (b) and the sample mean T1 atlas (d) formed from healthy BrainWeb images.

BRATS MICCAI challenge and from the fact that the expected unbiased healthy atlas for these data can be computed as the unbiased atlas formed from the five healthy MRI images (without tumor injection), which were used to generate the tumor cases. We refer to this tumor-free atlas as the tumor-free

| | w.r.t. sample mean | | w.r.t. SRI24 | |
|-------------------------|--------------------|---------|--------------|---------|
| | SSD | GD | SSD | GD |
| Unbiased atlas | 4.08e8 | 1.88e-2 | 19.42e+8 | 1.91e-2 |
| Low-rank unbiased atlas | 4.29e8 | 1.64e-2 | 18.64e+8 | 1.87e-2 |

TABLE II: Comparison of the T1 low-rank atlas with conventional unbiased atlas using both SSD and geodesic distance (GD), with respect to the sample mean (BrainWeb) and a population mean (SRI24).

sample mean. A challenge with this data, however, is that the tumor insertion software changes the intensities of all tissues in the images, as can be seen from Fig. 5. Therefore, when we compute the distance between estimated atlases and the tumor-free sample mean, the *mutual information (MI)* metric must be used instead of SSD.

Comparison with the tumor-free sample mean: Using 8 synthetic FLAIR images that have tumor(s) injected at different locations in the brain, as shown in Fig. 4, we compared our estimated low-rank FLAIR atlas (middle) with the traditional unbiased atlas building results from that tumor data (left) and the tumor-free sample mean (right). The geodesic distances of the traditional unbiased atlas and the low-rank atlas to the tumor-free sample mean are $1.29e-2$ and $1.10e-2$, respectively. Our low-rank atlas is also slightly sharper (measured by image entropy).

Influence of imaging modality: We repeated the above experiment using the corresponding T1 images, as seen in Fig. 5(a). This alternative MRI protocol was chosen because in it the lesion intensity nearly matches the intensity range of healthy tissues; potentially confounding the low-rank decompositions. Fig. 5(b)-(d) shows low-rank atlas results (middle) in comparison with traditional unbiased atlas results (left) and the tumor-free sample mean (right). We also computed the SSD metric and geodesic distances between the computed atlases, the tumor-free sample mean, and a published population atlas, i.e., the SRI24 normal T1 atlas. Those results are given in Table. II. The low-rank atlas has a slightly lower entropy (3.5058) than the unbiased atlas (3.5071). The low-rank and traditional atlases perform nearly equally well with respect to the tumor-free sample mean in terms of SSD and GD, and the the low-rank atlas was a much closer match to the published population atlas than the traditional unbiased atlas. As predicted, compared to the FLAIR images, the benefits of our low-rank framework are reduced due to the reduced conspicuity of the tumors, but the benefits of our framework do persist.

4) *Low-rank atlas estimation results from clinical brain tumor data:* The BRATS'12 challenge also provided clinical FLAIR data from patients with brain tumors, see Fig. 6(a). We used that data to illustrate the expected clinical utility of low-rank atlas formation.

Large and high-contrast lesions in the clinical tumor data makes the advantage of the low-rank strategy even pronounced compared to the results from conventional atlas building, which clearly fails as shown in Fig. 6(b). Our low-rank atlas calculated from 16 input FLAIR images, see Fig. 6(c), recovers major structures regardless of the presence of large tumors in

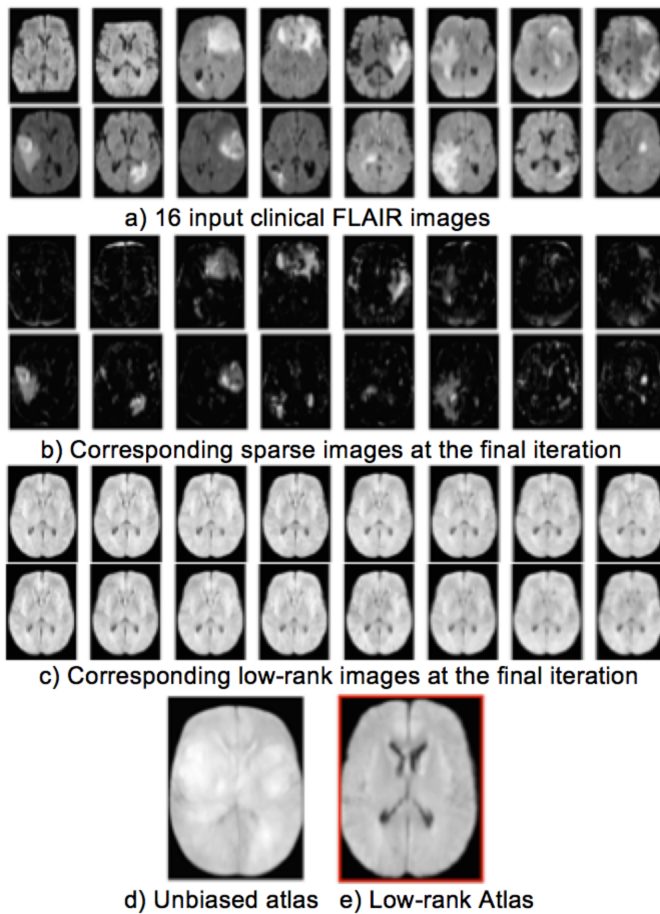


Fig. 6: Low-rank atlas estimation results on real patients' FLAIR images: a) The axial slices of the 16 input FLAIR images. The sparse images b) and low-rank images c) are in correspondence with the input images. Our estimated low-rank atlas e) with an entropy of 3.13, is compared with the conventional unbiased atlas d) with an entropy of 3.76).




Fig. 7: Intensity entropy measurement over the iterations during the atlas formation of 16 FLAIR images (Fig. 6) : unbiased atlas v.s. low-rank atlas.

these clinical data. Fig. 7 shows the image entropy values for each iteration during the atlas formation for both low-rank atlas and unbiased atlas. As the low-rank atlas gets sharper over the iterations, its entropy values decreases gradually. On the other hand, the traditional unbiased atlas suffers from the corrupted data from the beginning and can not recover with more iterations.




Fig. 8: A low-rank atlas will be degraded when the training images have tumors in consistent locations (indicated by the red and blue arrows). The low-rank atlas (image at the far right and outlined in red) contains tumor intensities at locations that repeatedly had tumors in the training data.

IV. DISCUSSION AND CONCLUSION

The proposed low-rank image registration framework is able to handle large pathologies / lesions when registering input images to a normal atlas or when computing an unbiased atlas. Its results are insensitive to the value of its main parameter, the weight parameter which influences the allocation of image data into low-rank and sparse components based on normal variation observed in the training sample.

Our experiments show that the low-rank atlases produced by our framework match well with the atlas that would have been formed if the patients data did not contain tumors. The atlases it produces are insensitive to the portion of training patients that contain tumors. The atlases it produces are similar to healthy population atlases.

The low-rank and sparse decompositions of each individual's data can also serve as a spatial prior for tissue segmentation. However, as discussed above, each image's sparse component may contain some "normal" anatomic variation that may arise due to (a) normal variation that isn't well represented by the sample, (b) limited lesion conspicuity relative to normal anatomic variation, or (c) cross-patient lesion correlations that cause those correlations to be assumed as normal anatomy.

The method does require that the lesions are randomized spatially across the training sample. This is especially critical when the sample size is small, as normal anatomic variations which are weakly represented may be incorrectly allocated to the sparse component and common parts of the tumors may be interpreted as normal anatomic variation and thereby allocated to the low-rank component. For example, as shown in Fig. 8, when lesions repeatedly appear at nearly the same locations across the training sample, the decomposition is not able to statistically distinguish those lesions from "normal" tissue and therefore will not allocate them to the sparse component. Instead, the resulting low-rank atlas will contain the "averaged" appearance of the repeated lesions at those locations.

The method performs better when the lesions are more conspicuous in the data. As indicated in our comparison in applying our method to T1 and FLAIR synthetic brains, our method works better for those imaging modalities where the lesions have high contrast with respect to normal tissues. When the intensity ranges of the lesion and the normal tissues overlap, it is mathematically challenging to separate the intensity differences caused by geometric variation and the pathology in the current framework.

The novel contributions of this paper are 1) the use of low-

rank plus sparse image decomposition in image-to-atlas registration, 2) the integrated formulation of this decomposition into atlas formation, and 3) the use of sparse components as a prior for lesion identification and segmentation. These contributions allow images containing pathologies to drive atlas formation, and they allow images containing pathologies (large lesions and deformations) to nevertheless be well registered with normal-control atlases. The robustness of proposed methods is validated on both synthetic and clinical MRI datasets.

For future work, adding spatial coherence constraints during the low-rank and sparse decomposition optimization may mitigate some of the limitations seen when the sample size is small and the lesions are less conspicuous. The definition of the constraints could, for example, be customized according to expected lesion patterns. For example, it could be useful to penalize thin structures or surfaces allocated to the sparse part thereby discouraging normal variation to be allocated to the sparse component. Conversely, one could encourage spatially contiguous regions to be allocated to the sparse part. This could likely be achieved by some form of total variation penalty on the sparse component.

ACKNOWLEDGMENT

This work was supported, in part, by the National Institutes of Health under grants R44CA143234, R41EB015775, R43EB016621, R41NS081792, R43CA165621, and by NSF under grants EECS-1148870 and EECS-0925875.

REFERENCES

- [1] A. Irimia, B. Wang, S. Aylward, M. Prastawa, D. Pace, G. Gerig, D. Hovda, R. Kikinis, P. Vespa, and J. Van Horn, "Neuroimaging of structural pathology and connectomics in traumatic brain injury: Toward personalized outcome prediction." *NeuroImage: Clinical*, vol. 1, pp. 1–17, 2012.
- [2] N. Chitphakdithai and J. Duncan, "Non-rigid registration with missing correspondences in preoperative and postresection brain images," in *Med Image Comput Assist Interv.*, ser. LNCS, J. Tianzi, N. Navab, J. Pluim, and M. Viergever, Eds., vol. 6361, 2010, pp. 367–374.
- [3] M. Niethammer, G. Hart, D. Pace, P. Vespa, A. Irimia, J. Van Horn, and S. Aylward, "Geometric metamorphosis," in *MICCAI 2011*, ser. LNCS, G. Fichtinger, A. Martel, and T. Peters, Eds., vol. 6892, 2011, pp. 639–646.
- [4] B. Wang, M. Prastawa, S. Awate, A. Irimia, M. Chambers, P. Vespa, J. Van Horn, and G. Gerig, "Segmentation of serial MRI of TBI patients using personalized atlas construction and topological change estimation," in *ISBI*, 2012.
- [5] L. Ramus, O. Commowick, and G. Malandain, "Construction of patient specific atlases from locally most similar anatomical pieces," in *Med Image Comput Assist Interv.*, vol. 13, 2010, pp. 155–162.
- [6] S. Joshi, B. Davis, M. Jomier, and G. Gerig, "Unbiased diffeomorphic atlas construction for computational anatomy," *NeuroImage*, vol. 23, pp. 151–160, 2004.
- [7] Y. Peng, A. Ganesh, J. Wright, W. Xu, and Y. Ma, "RASL: Robust alignment by sparse and low-rank decomposition for linearly correlated images," *TPAMI*, vol. 34, no. 11, pp. 2233–2246, 2012.
- [8] X. Liu, M. Niethammer, R. Kwitt, M. McCormick, and S. Aylward, "Low-rank to the rescue – atlas-based analyses in the presence of pathologies," *MICCAI*, vol. LNCS 8675, pp. 97–104, 2014.
- [9] Z. Lin, M. Chen, and Y. Ma, "The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices," *arXiv preprint arXiv:1009.5055*, 2010.
- [10] B. Avants, C. Epstein, M. Grossman, and J. Gee, "Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain." *Med Image Anal.*, vol. 12, pp. 26–41, 2008.
- [11] D. Atkinson, D. H. DL, P. Stoylo, P. Summers, and S. Keevil, "Automatic correction of motion artifacts in magnetic resonance images using an entropy focus criterion," *IEEE Trans Med Imaging*, vol. 16(6), pp. 903–910, 1997.
- [12] M. Prastawa, E. Bullitt, and G. Gerig, "Simulation of brain tumors in MR images for evaluation of segmentation efficacy," *Med. Image Anal.*, vol. 13, no. 2, pp. 297–311, 2009.
- [13] B. Menze, M. Reyes, and K. V. Leemput, "The multimodal brain tumor image segmentation benchmark (brats)," pp. 33, 2014.
- [14] T. Rohlfing, N. M. Zahr, E. V. Sullivan, and A. Pfefferbaum, "The SRI24 multichannel atlas of normal adult human brain structure," *Hum. Brain Mapp.*, vol. 31, no. 5, pp. 798–819, 2010.
- [15] D. Collins, A. Zijdenbos, V. Kollokian, J. Sled, N. Kabani, C. Holmes, and A. Evans, "Design and construction of a realistic digital brain phantom," *IEEE Transactions on Medical Imaging*, vol. 17, pp. 463–468, 1998.