# Low-Rank to the Rescue – Atlas-based Analyses in the Presence of Pathologies

Xiaoxiao Liu<sup>1</sup>, Marc Niethammer<sup>2</sup>, Roland Kwitt<sup>3</sup>, Matthew McCormick<sup>1</sup>, and Stephen Aylward<sup>1</sup>

<sup>1</sup> Kitware Inc., USA
 <sup>2</sup> University of North Carolina at Chapel Hill, USA
 <sup>3</sup> Department of Computer Science, University of Salzburg, Austria

Abstract. Low-rank image decomposition has the potential to address a broad range of challenges that routinely occur in clinical practice. Its novelty and utility in the context of atlas-based analysis stems from its ability to handle images containing large pathologies and large deformations. Potential applications include atlas-based tissue segmentation and unbiased atlas building from data containing pathologies. In this paper we present atlas-based tissue segmentation of MRI from patients with large pathologies. Specifically, a healthy brain atlas is registered with the low-rank components from the input MRIs, the low-rank components are then re-computed based on those registrations, and the process is then iteratively repeated. Preliminary evaluations are conducted using the brain tumor segmentation challenge data (BRATS '12).

## 1 Introduction

Image-based lesion detection and segmentation are needed to assess and plan the treatment of patients suffering from traumatic brain injuries (TBI), brain tumors, or stroke [2]. One popular method for such image analysis involves registering an atlas to the patient's images to estimate tissue priors. However, if the patient's images contain large pathologies, then lesion-induced deformations may inhibit atlas registration and confound the tissue priors. Furthermore, even forming an appropriate, unbiased atlas for segmentation may be problematic.

In unbiased atlas building [3], when images with lesions are used to form the atlas, the lesions propagate into and corrupt the atlas. However, particularly for research projects with limited time and financial resources or involving a new imaging protocol or children, it can be problematic to obtain a sufficient number of protocol-matched scans from healthy subjects for atlas formation. Hence, registration methods tolerant to such image corruptions are desirable.

The iterative, low-rank image registration framework presented in this paper tolerates the presence of large lesions during image registration, and it can thereby aid in atlas-based segmentation and unbiased atlas formation by mitigating the effects described above. While our approach is general, in this paper we focus on registration in the presence of pathologies for the purpose of atlas-based tissue segmentation for illustration.

### 2 Liu et al.

The most straightforward method to eliminate a lesion's influence during registration is to "mask" it so that the lesion's voxels are not 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 [1], geometric metamorphosis that separates estimating healthy tissue deformation from modeling tumor change [5], or personalized atlas construction that accounts for diffeomorphic and non-diffeomorphic changes [9]. While effective, these methods require explicit lesion segmentations or initial lesion localizations, which, in this case, is actually the goal of the process.

**Contribution.** We propose to exploit *population information* to assess which parts of an image are likely lesions (they are inconsistent with the population) and which parts of an image should be considered *normal*. We adopt a recent machine learning technique, i.e., the decomposition of matrices into a low-rank and sparse components [6], in an iterative registration process to achieve this objective.

### 2 Low-rank plus sparse decomposition

In [6], Peng et al. propose to decompose a matrix of vectorized images into the sum of a low-rank and a sparse component (containing residuals) in the context of simultaneous 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 linearcorrelation across the images. Given a collection of n images having 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, s.t.  $S_i = D_i L_i$ .

The low-rank representation of D is then defined as

$$\{L^*, S^*\} = \underset{L,S}{\arg\min}(\|L\|_* + \lambda \|S\|_1) \text{ s.t. } D = L + S , \qquad (1)$$

where  $||L||_*$  is the nuclear norm of L and  $||S||_1$  is the 1-norm of S. Since the problem is convex, a globally optimal solution  $\{L^*, S^*\}$  can be obtained using, e.g., an augmented Lagrangian approach [4].

# 3 Integrate low-rank decomposition into an iterative registration framework

We have integrated the 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 and 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 lowrank component and allocated to the sparse component. Thereby, the sparse component can be used to inform spatial and intensity priors for localizing and segmenting lesions.

Our method also supports unbiased atlas formation using data containing pathologies. Specifically, in the above framework the normal-control atlas  $I_A$ can be replaced by the mean low-rank image, at each iteration. In unbiased atlas-building the goal is to estimate an atlas image such that it is central with respect to the data population. This is achieved by minimizing

$$E(\{\Phi_i^{-1}\}, I_A) = \sum_{i=1,\dots,N} Reg[\Phi_i^{-1}] + \sigma^{-2} Sim[I_i \circ \Phi_i^{-1}, I_A],$$
(2)

with respect to the unknown atlas image  $I_A$  and the unknown transformations  $\{\Phi^{-1}\}$ . Here,  $Reg[\Phi]$  denotes a regularity measure for the transformation  $\Phi$ , typically penalizing spatially non-smooth transformations and Sim[I, J] is a chosen similarity measure between the images I and J. This could simply be the sum-of-squared intensity differences (SSD).

To optimize this energy using alternating optimization, we first keep  $I_A$  fixed while solving for  $\{\Phi_i^{-1}\}$  and subsequently keep the transformations  $\{\Phi_i^{-1}\}$  fixed while solving for  $I_A$ . The first part performs independent pairwise registrations between  $\{I_i\}$  and the fixed image  $I_A$ . The second part requires, for SSD, minimizing

$$E(I_A) = \sum_{i=1,\dots,N} \|I_i - I_A\|^2$$
(3)

which is achieved by the mean image  $I_A = 1/N \sum_i I_i$ . However, when lesions are present in  $\{I_i\}$ , the mean image is degraded by the undesired involvement of the lesions in the average. We can instead minimize Eq. (1) to obtain the low-rank approximations  $\{L_i\}$  of the warped images  $\{D_i = I_i \circ \Phi^{-1}\}$ , and then minimize

$$E(I_A) = \sum_{i=1,\dots,N} Sim[L_i, I_A]$$
(4)





Fig. 1. An illustration of the proposed low-rank iterative image registration 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 *i*-th input image and  $\phi_i$  refers to the *i*-th registration map generated from BSpline image registration at each iteration.

with respect to the unknown atlas  $I_A$ . Again, for SSD, the solution will be the mean over  $\{L_i\}^4$ . Other similarity measures, such as normalized cross correlation, could be used for registration, but they may require more challenging optimizations and may not be meaningful for atlas construction as they may make the atlas-image non-unique. Also note that when fixing the atlas,  $I_A$ , unbiased atlas construction simplifies to group-wise registration. The group-wise approach is essential because it allows for the population-based decomposition of the images into low-rank/sparse components (cf.  $\S^2$ ). A general framework for our method is shown in Fig. 1. The algorithm proceeds as follows:

- (1) Solve for affine transform  $(\phi_i^0)^{-1}$  registering each  $I_i$  to the atlas image  $I_A$ .
- (2) For each iteration j, compute the low-rank image  $L_i^j$  by solving Eq. (1).
- (2) For each normality, compare the first end of  $(\phi_i^j)^{-1}$  registering low-rank images  $L_i^j$  to  $I_A$ . (4) Compose and apply transforms to  $I_i$ , s.t.  $I_i^{j+1} = I_i \circ (\phi_i^0)^{-1} \cdots \circ (\phi_i^j)^{-1}$ .
- (5) Set  $j \leftarrow j + 1$  and continue with step (2) until convergence.

Given a low-rank plus sparse decomposition, the registration step can be based on any standard deformable registration algorithm and its associated convergence characteristics apply. In our experiments, BSpline transforms and the Mattes mutual information (MMI) metric are used to register the low-rank images with the atlas, cf. step (3). The number of BSpline control points is increased gradually over the iterations to effect a coarse-to-fine optimization strategy. At each iteration we are maximizing the mutual information between the atlas image and each individual low-rank image, cf. Eq. (4). As our algorithm alternates between low-rank decomposition and registration, it can be considered a greedy strategy. Convergence is reached when the total change in deformation is small. In our experiments with the BRATS '12 dataset (8 inputs), results converge within 10 iterations. For the TumorSim data [7] (20 simulated T1 images),  $\lambda$  in Eq. (1) is set to 0.5 and for the patient data (8 FLAIR images),  $\lambda$  is set to 0.8.

<sup>&</sup>lt;sup>4</sup> In case of SSD it is advisable to initially histogram-normalize each  $I_i$ .

5

### 4 Experimental study

We have conducted initial assessments of our method for atlas-to-image registration using two evaluation metrics with simulated and patient data.

**Quantitative assessment of atlas-to-image registration**: 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 atlas-to-image registration is successful, the atlas' tissue labels should align with the corresponding tissues 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 (TCSD)* values indicate more accurate atlas-to-image registrations.

We calculate TCDS at each iteration to evaluate the convergence of the iterative framework. We also use it to compare our method with traditional BSpline atlas-to-image registration.

**Qualitative assessment for lesion segmentation**: The iterative atlas-toimage registration process can be examined by inspecting the parts of each image allocated to its sparse component in each iteration. This *sparse image* at the final iteration (after reaching convergence) should be sensitive and specific to the lesion. By reviewing the sparse image's evolution over the iterations, we can qualitatively assess the effectiveness of our method in matching each patient's image to the healthy atlas, while not burdened by lesions.

Note that only a qualitative assessment of lesions is made. Sparse images will contain some normal anatomic variation as well as the lesions. Over the iterations, the variations between the individual patients and the healthy atlas are minimized via the deformable registrations between the low-rank images and the healthy atlas. The sparse images after convergence could then serve as a strong prior for subsequent tumor segmentation algorithms, but lesion's heterogeneity as well as "normal" small-scale anatomic variations must be appropriately handled by subsequent lesion segmentation algorithms.

### 4.1 Case studies and results

For the following case studies we used the SRI24 atlas [8] as the healthy atlas for registration and to provide gray-matter (GM), white-matter (WM), and cerebrospinal fluid (CSF) tissue labels after registration to compute the TCSD metric.

**Case 1 (simulation data)**: The training data in the BRATS '12 challenge included MRI scans into which simulated high-grade and low-grade glioma tumors were injected using *TumorSim* [7]. We selected 20 cases containing large tumors and large deformations to form a set of challenging image-to-atlas registration tasks. Fig. 2(a) shows the first 4 subjects and their corresponding low-rank and sparse components during the 1<sup>st</sup> iteration. Fig. 2(b) shows the TCSD values for the GM and WM classes after each iteration. The tightening of the statistics in Fig. 2(b) illustrates the convergence of the registration between each subject and the healthy atlas over time.



(a) 4/20 simulated T1 brain volumes with the corresponding low-rank and sparse components (during the 1<sup>st</sup> iteration).



Fig. 2. Results on TumorSim dataset from BRATS '12.

Case 2 (clinical data): A subset of 8 FLAIR images from BRATS '12 challenge are tested using the same experimental setup as the simulated data. Fig. 3(a)illustrates the input data and the results. We show the low-rank and sparse components of the first three input FLAIR images at the initial iteration in Fig. 3(a) (left). The low-rank plus sparse decomposition results at the  $2^{nd}$  and the  $10^{\text{th}}$  (final) iteration of the first input patient are shown in Fig. 3(a) (middle). On the right-hand side of Fig. 3(a), the top image is the normal-control SRI24 atlas T1 image that we used as the fixed normal-control image during each registration. As we can see from the box plots in Fig. 3(b), the overall TCSD for the GM class improves over the iterations ( each box contains the TCDS values of all 8 patients at each iteration).

To compare with a direct BSpline registration, we used the same BSpline parameters, e.g., number of control points, that are used in the final iteration of our method. We exclude the tumor region when calculating TCDS for each tissue class. The middle image of Fig. 3(a) is the atlas CSF label contour overlaid on the direct BSpline-registered (i.e., deform the original input FLAIR image to match the healthy atlas) image; the bottom image is the atlas CSF label contour



(a) Left: first three FLAIR images and their initial low-rank and sparse components. Middle: low-rank and sparse components at the 2<sup>nd</sup> and 10<sup>th</sup> iteration of the first input. Right: SRI24 T1 atlas (top), atlas CSF label contour overlaid on the direct BSpline registered image (middle), atlas CSF label contour overlaid on the deformed input image at final iteration using our approach (bottom).



(b) *Left*: Change in the TCSD for GM labels transcribed from the atlas after each iteration of our method. *Right*: TCSD comparison to traditional BSpline atlas-to-image registration for the GM class.

Fig. 3. Results on TBI patient dataset from BRATS '12.

overlaid on top of the converged deformed input image. Much better alignments are seen on the edges of the ventricles than it is the case for the middle image. The direct BSpline registration produces larger TCSD values for the GM class in most cases in this study, as shown in the table of Fig. 3(b). Our method performs worse on two cases (patient 5 and 8). Different from others they both have much narrower and distorted ventricles, which are high-contrast landmarks for guiding the registration optimization. Due to their distinctive appearance, the decomposed low-rank images contains very little truth geometries (mostly assigned to the sparse images), therefore the registration based on the low-rank component is not reliable. If more similar type datasets are included, the decomposition would be more effective and registration would have been resolved. Furthermore, only eight cases is not sufficient to represent a population, further 8 Liu et al.

work is needed to determine if and how many additional cases would be needed to represent a healthy population given subjects with pathologies.

### 5 Discussion and future work

The novel contributions of this paper are 1) the integrated formulation of low-rank image decomposition into atlas formation, 2) the use of low-rank image decomposition in atlas-to-image registration, and 3) the use of low-rank image decomposition as a prior for lesion identification and segmentation. These contributions are significant, because they 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. However, our current iterative registration framework needs to be better evaluated on TBI data sets with ground truth tissue labels. A near-term extension of this work is to form an unbiased atlas without a reference healthy atlas image, which is useful when only data containing pathologies is available. Future work will also focus on the development of a lesion segmentation pipeline using the sparse image as a spatial and intensity prior and a non-greedy implementation. **Acknowledgements.** This work was supported, in-part, by the NIBIB (R41EB015775), the NINDS (R41NS081792) and the NSF (EECS-1148870).

### References

- Chitphakdithai, N., Duncan, J.: Non-rigid registration with missing correspondences in preoperative and postresection brain images. In: Tianzi, J., Navab, N., Pluim, J., Viergever, M. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 367–374 (2010)
- Irimia, A., Wang, B., Aylward, S., Prastawa, M., Pace, D., Gerig, G., Hovda, D., Kikinis, R., Vespa, P., Van Horn, J.: Neuroimaging of structural pathology and connectomics in traumatic brain injury: Toward personalized outcome prediction. NeuroImage: Clinical 1, 1–17 (2012)
- Joshi, S., Davis, B., Jomier, M., Gerig, G.: Unbiased diffeomorphic atlas construction for computational anatomy. NeuroImage 23, 151–160 (2004)
- 4. Lin, Z., Chen, M., Ma, Y.: The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. arXiv preprint arXiv:1009.5055 (2010)
- Niethammer, M., Hart, G., Pace, D., Vespa, P., Irimia, A., Van Horn, J., Aylward, S.: Geometric metamorphosis. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011. LNCS, vol. 6892, pp. 639–646 (2011)
- Peng, Y., Ganesh, A., Wright, J., Xu, W., Ma, Y.: RASL: Robust alignment by sparse and low-rank decomposition for linearly correlated images. TPAMI 34(11), 2233–2246 (2012)
- 7. Prastawa, M., Bullitt, E., Gerig, G.: Simulation of brain tumors in MR images for evaluation of segmentation efficacy. Med. Image Anal. 13(2), 297–311 (2009)
- Rohlfing, T., Zahr, N.M., Sullivan, E.V., Pfefferbaum, A.: The SRI24 multichannel atlas of normal adult human brain structure. Hum. Brain Mapp. 31(5), 798–819 (2010)
- Wang, B., Prastawa, M., Awate, S., Irimia, A., Chambers, M., Vespa, P., Van Horn, J., Gerig, G.: Segmentation of serial MRI of TBI patients using personalized atlas construction and topological change estimation. In: ISBI (2012)