Published in final edited form as:

IEEE Trans Med Imaging. 2016 August ; 35(8): 1937–1951. doi:10.1109/TMI.2016.2538998.

Abnormality Detection via Iterative Deformable Registration and Basis-Pursuit Decomposition

Ke Zeng,

Center for Biomedical Image Computing and Analytics, Department of Radiology, University of Pennsylvania, Philadelphia, PA 19104 USA

Guray Erus,

Center for Biomedical Image Computing and Analytics, Department of Radiology, University of Pennsylvania, Philadelphia, PA 19104 USA

Aristeidis Sotiras,

Center for Biomedical Image Computing and Analytics, Department of Radiology, University of Pennsylvania, Philadelphia, PA 19104 USA

Russell T. Shinohara, and

Department of Biostatistics and Epidemiology, University of Pennsylvania, Philadelphia, PA 19104 USA

Christos Davatzikos

Center for Biomedical Image Computing and Analytics, Department of Radiology, University of Pennsylvania, Philadelphia, PA 19104 USA

Abstract

We present a generic method for automatic detection of abnormal regions in medical images as deviations from a normative data base. The algorithm decomposes an image, or more broadly a function defined on the image grid, into the superposition of a normal part and a residual term. A statistical model is constructed with regional sparse learning to represent normative anatomical variations among a reference population (e.g., healthy controls), in conjunction with a Markov random field regularization that ensures mutual consistency of the regional learning among partially overlapping image blocks. The decomposition is performed in a principled way so that the normal part fits well with the learned normative model, while the residual term absorbs pathological patterns, which may then be detected through a statistical significance test. The decomposition is applied to multiple image features from an individual scan, detecting abnormalities using both intensity and shape information. We form an iterative scheme that interleaves abnormality detection with deformable registration, gradually improving robustness of the spatial normalization and precision of the detection. The algorithm is evaluated with simulated images and clinical data of brain lesions, and is shown to achieve robust deformable registration

Personal use is permitted, but republication/redistribution requires IEEE permission.

Correspondence to: Ke Zeng.

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

and localize pathological regions simultaneously. The algorithm is also applied on images from Alzheimer's disease patients to demonstrate the generality of the method.

Index Terms

Abnormality detection; basis pursuit; brain pathology; convex optimization; medical image registration; sparse representation

I. Introduction

THE task of detecting pathological regions plays a central role in medical image analysis, especially in brain imaging. As manual delineation of the pathological regions is time-consuming and suffers from large intra- and inter-rater variability, extensive efforts have been devoted to the development of fully automatic methods that could reduce both processing time and rater variability. The ultimate goal is to obtain a reproducible algorithm that could automatically process hundreds or thousands of images in research studies and clinical trials.

The literature on pathological region detection is abundant, therefore, a comprehensive summary is beyond the scope of this introduction. For multiple sclerosis (MS) brain lesions alone, a recent survey [1] listed 80 papers that describe automatic segmentation procedures. These methods focus exclusively on the delineation of MS lesions, thus relying heavily on MS lesion-specific characteristics, the most distinct of which is that MS lesions appear brighter than normal white matter on T2-weighted Magnetic Resonance (MR), Proton Density (PD) and Fluid Attenuation Inversion Recovery (FLAIR) images [1]. Although a wide range of techniques have been applied to this particular detection task where informative prior knowledge is available, the authors of [1] still conclude that "a robust, accurate, fully-automated lesion segmentation method suitable for use in clinical trials is still not available".

The discussion on MS lesion detection serves as an illustrative example of the challenges and uncertainties associated with abnormality detection, even when well-tuned methods are applied to a pathology with distinct features. In reality, however, there are hundreds of pathologies causing imaging abnormalities with diverse characteristics in different imaging modalities. For example, in stroke lesions, a hemorrhage appears as a bright region and ischemic stroke appears as a dark region in Computed Tomography scans [2], while in T₁-weighted MR images an infarct lesion is shown with intensities similar to cerebrospinal fluid (CSF) or grey matter (GM) [3].

Researchers have devoted extensive efforts to developing tailored algorithms for a targeted pathology, where they analyze the characteristics and construct specialized models for their objective. Those algorithms can be divided into two categories: supervised and unsupervised. In supervised methods, a classification rule is inferred from a training set that consists of annotated images from patients with the targeted disease. Pathology in the test image is then identified according to the learned classification rule (e.g., [4]). Due to the potential heterogeneity of the disease, the training set and the learning method must be

carefully chosen to produce accurate and reproducible results. Furthermore, the training set has to be manually segmented, a time-consuming and subjective procedure as already noted. Unsupervised methods, on the other hand, do not rely on an annotated training set and can be directly applied to a test image. The pathological regions may be modeled either as additional tissue classes, or simply as outliers (e.g., [5]). The detection relies on experts' knowledge of the imaging appearance of the targeted pathology and on how such prior knowledge is incorporated into the mathematical model.

Encoding the characteristics of a target pathology can be a challenging task under both supervised and unsupervised settings. The pathology of interest may vary greatly in shape and location, while its intensity profile may be indistinguishable from the intensity profile of a normal tissue type. In cases where such difficulties can be addressed by appropriate modeling, the resulting algorithms often lack the ability to extend to new domains. A framework that has encoded features dedicated to one type of pathology is difficult to generalize across other abnormalities with different characteristics. This lack of generalization ability suggests a need for separate detectors for each of the existing pathologies, which is a daunting task.

In this work, we take an alternate view to abnormality detection, which complements the individualized approaches, emphasizing generality over specificity. This alternative approach is built on the statistical perspective that pathological patterns follow probabilistic laws that are very distinct from the normative anatomical variation and is similar in spirit to the path taken in [6]. We therefore focus exclusively on the normative variation, with the premise that if one could capture the normative variation among the healthy population, then it would be possible to spot not just one specific pathology, but a broader class of pathologies as deviations from the normative model.

A generically formulated framework that identifies pathologies as deviations from normal variations can be useful in various clinical scenarios. For example, its outputs can serve as initial detectors, statistical priors or features that help subsequent, or concurrent specialized detectors. Moreover, as the use of imaging becomes increasingly widespread, automated screening of imaging data, by flagging scans with abnormalities that need further expert attention, can dramatically improve efficiency and throughput. Along the same lines, and for more subtle abnormalities, a tool for directing the attention of expert readers towards regions displaying relatively large deviation from normality would be quite helpful.

Relatively little attention is given to pathology detection approaches that do not target a particular abnormality, and capturing the normative variation among images from a healthy population features challenges of its own. The dimensionality of a typical brain scan is prohibitively large when compared to the typical number of available samples, which makes it impractical to accurately model the entire image directly. Moreover, spatial normalization, or co-registration, is essential for this task since locational information is highly relevant in any abnormality detection problem [1], [7]. However, pathological regions are by definition atypical parts in the patient's brain that lack correspondence with normal anatomies, and therefore the registration step should be robust to such topological changes and must be performed carefully.

In [8]–[10], an image is decomposed into a normal part (also referred to as the projection/normal projection) plus a residual, and abnormalities are detected as statistically significant parts of the residual. Regional modeling is employed in [8]–[10] to address the dimensionality challenge. In [9], image patches are randomly sampled from the brain and processed iteratively, with the local model built from principal component analysis in the wavelet domain. However, the decomposition is not robust enough; many of the pathological patterns are falsely absorbed into the normal part, since often the underlying anatomy does not follow a Gaussian distribution. In [8], an image is treated as a collection of minimally overlapping local patches, and consensus is enforced among overlapping parts. The final projection is cleaner from pathologies but looks overly smooth, suffering from significant loss of high frequency information. Further, in [8]–[10], the registration step is performed straightforwardly, completely neglecting the presence of the pathological region, a problematic procedure known to produce false deformation fields [11], [12].

Deformable registration between a normal template image and a patient image with pathologies and topological changes is, by itself, an active research area. A widely used method to isolate the impact of the abnormal region is Cost Function Masking (CFM) [11], where the pathological regions are excluded during the optimization. However, applying CFM requires a binary mask of the abnormal region which is not known a priori in our case. In some cases, such as tumors, registration can be combined with a dedicated model, which describes the growth of the pathology (see for example [13]–[15]), into a unified framework. Such an approach is by nature restricted to one specific pathology, and thus does not fit into a generic framework. More sophisticated generic approaches, that account for topological changes, have also been developed, for example in [16]–[18]. In the theory of metamorphosis [16], a Riemannian metric is defined on the image manifold, accounting for both geometric deformation and intensity changes. A geodesic path may be computed between two images, yielding a geometric diffeomorphism between the source and the target that is optimal up to some residual intensity error. However, the intensity part of the Riemannian metric is assumed to be spatially homogeneous, a property not well suited for more locally supported pathologies. Further, it is not clear how to balance the geometric deformation and intensity change in the Riemannian metric. In [17], images are embedded as surfaces in $\mathbb{R}^4$ Riemannian space, and registration is performed as a surface revolution that matches one embedded image to another, again accounting for both shape and intensity changes necessary to match them. As in CFM, an a priori estimate of the pathological regions is required in order to attain a robust deformation field. In [18], an expectation maximization (EM) type algorithm is used to perform registration with missing data, where one alternates between segmenting missing data and estimating the deformation field. As [18] focuses on handling corrupted data with drastically different image appearances, the detection step is built on simple uni-variate statistics, which could have difficulty recognizing multi-variate patterns.

Towards addressing both abnormality detection and registration, we follow our previous work [10] and treat a 3-D image as a collection of overlapping local regions that are mutually constrained. We propose a multi-domain learning framework, where we use a standard template domain for learning shape-based information while centering the

intensity-based learning around the patient's image domain by warping normal training samples to it. The proposed framework differs from [8]–[10], where the learning is performed only on a template domain, an approach that we experimentally found to induce template bias. The patient-specific warping approach is customary in multiatlas segmentation method (e.g., [19], [20]) and has the advantage that the pathological regions will be kept intact during the spatial normalization. Each local patch $\mathbf{y}$ is written as the superposition of $\mathbf{D}_1\alpha_1^*$ and $\mathbf{D}_2\alpha_2^*$ by solving for the minimal ℓ_1 -norm solution of an under-determined linear system. Here, $\mathbf{D}_1$ is a location-specific dictionary that accounts for normative anatomical variation, and $\mathbf{D}_2$ is a generic dictionary that accounts for the unknown pathological patterns. The decomposition of $\mathbf{y}$ naturally reads $\mathbf{y} = \mathbf{D}_1\alpha_1^* + \mathbf{D}_2\alpha_2^*$, with $\mathbf{D}_1\alpha_1^*$ being the normal projection and $\mathbf{D}_2\alpha_2^*$ being the residual. Ultimately, all partially overlapping local patches are processed jointly in a consistent way by introducing appropriate linking terms to the local decomposition. An iterative registration and detection scheme is proposed to achieve robust registration and accurate detection simultaneously. At each iteration, the robustness and accuracy of the registration improve by leveraging the detection results, either through CFM or image inpainting. The updated registration output in turn leads to refined abnormality detection.

The remainder of the paper is organized as follows. The methodology is presented in Section II. In Section III we validate the method on both simulated abnormalities and clinical brain MR Image data. Section IV reviews the key components of the methodology. It also discusses the limitations, possible extensions and clinical applications of the current work. We conclude the paper in Section V.

II. Method

A. Model for Normal Test Sample

Let N be the number of training samples, $\{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_N\}$ be the set of normative training images, (i, j, k) be a fixed location in the 3-D image grid, and $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\}$ be the collection of vectorized 3-D training blocks of size $p_x \times p_y \times p_z$ centered at (i, j, k) . Given a normal vectorized test block $\mathbf{y} \in \mathbb{R}^p$, $p = p_x p_y p_z$, extracted from the same spatial location, we assume that $\mathbf{y}$ may be approximated by a parsimonious linear combination of training samples, where parsimony is measured through the ℓ_1 -norm. In other words, if we let $\mathbf{A} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N]$, then

$$\mathbf{y} \approx \mathbf{A}\mathbf{x} = \sum_{n=1}^N \mathbf{x}_n \mathbf{y}_n \quad (1)$$

for some $\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]^T \in \mathbb{R}^N$ with $\|\mathbf{x}\|_1$ small. The example-based dictionary $\mathbf{A}$ is effective in the sense that every training example could be represented by just one atom. Furthermore, the resulting highly correlated columns in $\mathbf{A}$ provide an additional structure that is useful for our purpose, as detailed below.

B. Model for the Pathologies

Frame the setting as in the previous subsection, we now assume that the test patch $\mathbf{y} \in \mathbb{R}^p$ contains abnormalities, in which case (1) no longer holds. Instead, we assume that the abnormal patch $\mathbf{y}$ can be decomposed as the superposition of a normal part $\tilde{\mathbf{y}}$ and an abnormal part $\mathbf{r}$. Together with (1), we may write

$$\mathbf{y} = \tilde{\mathbf{y}} + \mathbf{r} \approx \mathbf{A}\mathbf{x} + \mathbf{r} \quad (2)$$

with a coefficient vector $\mathbf{x}$ that has small ℓ_1 energy. The only requirement we impose on the abnormal part $\mathbf{r}$ is that the number of non-trivial elements in $\mathbf{r}$, or the number of elements that significantly differ from zero, is strictly smaller than the dimension p . This hypothesis will hold in a wide range of cases, as long as the block size is sufficiently large and the abnormalities are focal in space.

The columns in the dictionary $\mathbf{A}$ are highly correlated with each other since they are image blocks drawn from the same spatial location. With this additional structure, it has been both empirically and theoretically verified that, after normalizing the columns in $\mathbf{A}$ to unit ℓ_2 -length, $\mathbf{A}\mathbf{x}$ and $\mathbf{r}$ can be recovered via the following ℓ_1 -norm problem

$$(P1) \quad \min \|\mathbf{x}\|_1 + \|\mathbf{r}\|_1 \text{ subject to } \mathbf{A}\mathbf{x} + \mathbf{r} = \mathbf{y}, \quad (3)$$

even when the abnormal part $\mathbf{r}$ is not sparse [21], [22]. Interested readers are referred to the “cross-and-bouquet” model discussed in [21], [22] for more technical insights.

Note that (P1) could also be written as

$$(P1-BP) \quad \min \|\boldsymbol{\alpha}\|_1 \text{ subject to } \mathbf{D}\boldsymbol{\alpha} = \mathbf{y} \quad (4)$$

with $\boldsymbol{\alpha} = [\mathbf{x}; \mathbf{r}]$ and $\mathbf{D} = [\mathbf{A}, \mathbf{I}]$, which is a classical procedure known as the basis pursuit [23]. Here $\mathbf{A}$ is our choice of the location-specific dictionary that accounts for normative anatomical variation, and the identity matrix $\mathbf{I}$ is used as the generic dictionary that accounts for the unknown pathological patterns, though $\mathbf{I}$ may be replaced by more specific dictionaries that can better represent a target pathology. By seeking the minimal ℓ_1 -norm solution to the under-determined linear system $\mathbf{D}\boldsymbol{\alpha} = \mathbf{y}$, we achieve a natural decomposition of $\mathbf{y}$ into the normal part and the residual.

C. Geometric and Probabilistic Interpretation

In this section, we present geometric and probabilistic interpretations for the optimization problem (P1).

Under the additional assumption that (P1) admits a unique solution (see [24] for a necessary and sufficient condition), we may rewrite it in the following two constrained forms,

$$\min \|y - Ax\|_1 \text{ subject to } \|x\|_1 = \|x^*\|_1$$

or

$$\min \|y - Ax\|_1 \text{ subject to } \|x\|_1 \leq \|x^*\|_1,$$

where x^* is the unique optimal solution for (P1). The constrained forms suggest that 1) we are approximating the manifold on which y resides with the facets of a polytope, that is $\{z = Ax : \|x\|_1 = \|x^*\|_1\}$, and 2) we are searching for the point closest to y in terms of ℓ_1 -norm within the polytope $\{z = Ax : \|x\|_1 \leq \|x^*\|_1\}$, justifying the terminology “normal projection”. A toy example is presented in appendix B in the supplementary file to demonstrate the geometric intuition.

From a stochastic point of view, we are modeling y under the Bayesian linear regression framework, or $y = Ax + r$, with a Laplacian prior $p(x) \propto \exp(-\|x\|_1)$ and a Laplacian likelihood $p(r) \propto \exp(-\|r\|_1)$. It is clear that x^* corresponds to the maximum a posteriori (MAP) estimate for x .

D. Joint Projection for all Local Blocks

One may envision the block-level decomposition process discussed in Section II-B as a local system that outputs the partial estimate $\tilde{y} = Ax^*$ through (P1). Those partial results need to be fused together to obtain a global estimation. A straightforward way to reconstruct the whole image from local blocks is through simple averaging, customary in the literature [25]. However, estimating the overlapping image blocks independently is a suboptimal approach, as pointed out in [26].

To address the foregoing issue, a distributed estimation algorithm is proposed in [26] and adopted in [8] for abnormality detection on 2-D image slices, where consensus between minimally overlapped subsystems is enforced through hard constraints. However, in our case the estimation task can not be easily “distributed”, as we work on image blocks with significant overlaps from a large 3-D image, rendering the method in [26] unsuitable for our purpose. Instead, we add extra penalty terms, which link the individual blocks, to the objective function, an alternative approach to achieve the goal of consistency that is equivalent to hard constraints in the limit.

We now present the full formulation of the joint optimization problem. Let $\mathcal{S}$ be the index set for the location of the blocks within the image $\mathcal{I}$, A_s be the example-based dictionary for location $s \in \mathcal{S}$, and y_s be the test block centered around this location. The proposed optimization problem may be written as

$$(P2) \quad \min_{x_s, s \in \mathcal{S}} \sum_{s \in \mathcal{S}} (\|x_s\|_1 + \|y_s - A_s x_s\|_1) + \frac{w}{2} \sum_{s_1, s_2 \in \mathcal{S}, s_1 \neq s_2} \Gamma(s_1, s_2),$$

where the $\Gamma(s_1, s_2)$'s are the penalty terms that enforce consistency across partially overlapping patches. Specifically,

$$\Gamma(s_1, s_2) = \begin{cases} 0 & \text{if } \mathbf{y}_{s_1} \cap \mathbf{y}_{s_2} = \emptyset \\ \|\mathbf{P}_{s_1}^{s_2} \mathbf{A}_{s_1} \mathbf{x}_{s_1} - \mathbf{P}_{s_2}^{s_1} \mathbf{A}_{s_2} \mathbf{x}_{s_2}\|_2^2 & \text{otherwise,} \end{cases}$$

where $\mathbf{P}_{s_1}^{s_2}$ ($\mathbf{P}_{s_2}^{s_1}$) is the operator that extracts from $\mathbf{y}_{s_1}$ ($\mathbf{y}_{s_2}$) the overlap with $\mathbf{y}_{s_2}$ ($\mathbf{y}_{s_1}$), respectively (see appendix C in the supplementary file for a graphical illustration of $\mathbf{P}_{s_1}^{s_2}$ and $\mathbf{P}_{s_2}^{s_1}$). Note that sending the weight parameter w to infinity is equivalent to requiring $\mathbf{P}_{s_1}^{s_2} \mathbf{A}_{s_1} \mathbf{x}_{s_1} = \mathbf{P}_{s_2}^{s_1} \mathbf{A}_{s_2} \mathbf{x}_{s_2}$.

From a different perspective, we are effectively modeling $\{\mathbf{y}_s\}_{s \in \mathcal{S}}$ with a Markov random field in the exponential family, where the nodes are $\{\mathbf{x}_s\}_{s \in \mathcal{S}}$ and $\{\mathbf{y}_s\}_{s \in \mathcal{S}}$. The graph structure reads 1) $\mathbf{x}_s$ and $\mathbf{y}_s$ are connected for all $s \in \mathcal{S}$, and 2) $\mathbf{y}_{s_1}$ and $\mathbf{y}_{s_2}$ are connected if and only if they overlap. The joint likelihood for $\{\{\mathbf{x}_s\}_{s \in \mathcal{S}}, \mathcal{S}\}$ admits the form

$$p(\{\mathbf{x}_s\}_{s \in \mathcal{S}}, \{\mathbf{y}_s\}_{s \in \mathcal{S}}) \propto \exp \left\{ - \sum_{s \in \mathcal{S}} (\|\mathbf{x}_s\|_1 + \|\mathbf{y}_s - \mathbf{A}_s \mathbf{x}_s\|_1) - \frac{w}{2} \sum_{s_1, s_2 \in \mathcal{S}, s_1 \neq s_2} \Gamma(s_1, s_2) \right\},$$

from which it is clear that (P2) solves for precisely the MAP estimate of $\{\mathbf{x}_s\}_{s \in \mathcal{S}}$.

The individual blocks now “speak” to one another through the $\Gamma(s_1, s_2)$ terms, but without complicating much the inference task. Indeed, though non-separable in the $\mathbf{x}_s$, (P2) is still convex and may be solved efficiently in a Gauss-Seidel fashion. If we assume that the index set $\mathcal{S}$ is ordered, the updating rule may be stated as

$$\begin{aligned} \mathbf{x}_s^{(k+1)} &:= \arg \min_{\mathbf{x}_s} (\|\mathbf{x}_s\|_1 + \|\mathbf{y}_s - \mathbf{A}_s \mathbf{x}_s\|_1 \\ &+ \sum_{r < s, \mathbf{y}_r \cap \mathbf{y}_s \neq \emptyset} \frac{w}{2} \|\mathbf{P}_s^r \mathbf{A}_s \mathbf{x}_s - \mathbf{P}_r^s \mathbf{A}_r \mathbf{x}_r^{(k+1)}\|_2^2 \\ &+ \sum_{u > s, \mathbf{y}_u \cap \mathbf{y}_s \neq \emptyset} \frac{w}{2} \|\mathbf{P}_s^u \mathbf{A}_s \mathbf{x}_s - \mathbf{P}_u^s \mathbf{A}_u \mathbf{x}_u^{(k)}\|_2^2). \end{aligned} \quad (5)$$

In our implementation, each Gauss-Seidel update is solved using the Alternating Direction Method of Multipliers [27]. Let $\{\mathbf{x}_s^*, s \in \mathcal{S}\}$ be the optimal solution of (P2), the projection $\tilde{\mathcal{S}}$ of $\mathcal{S}$ is reconstructed from $\{\tilde{\mathbf{y}}_s := \mathbf{A}_s \mathbf{x}_s^*, s \in \mathcal{S}\}$ through average pooling. Specifically, for each voxel v , we denote by $\{\mathbf{y}_v^1, \mathbf{y}_v^2, \dots, \mathbf{y}_v^M\}$ the blocks in which v is contained. The final value at v is set as the average of $\{\mathbf{y}_v^1(v), \mathbf{y}_v^2(v), \dots, \mathbf{y}_v^M(v)\}$. We refer interested readers to appendix A in the supplementary file for details on the optimization.

E. Block Size Selection

The size of the local blocks determines the scale at which we study the image content. Therefore, it is one of the most influential, if not the most crucial parameter in the proposed model. On one hand, the representation ability of the model decreases with increasing block size. In particular, if we extract tiny image blocks, we would be able to capture fine details of the image, thus expect a faithful reconstruction $\hat{\mathcal{J}}$ when $\mathcal{J}$ is a normal image; whereas, in the other extreme, when we consider the entire image as a single block, one would expect to lose significant detailed information. In the quantitative results presented in Section III-B, we will verify that smaller block size does lead to lower projection error when $\mathcal{J}$ is normal.

However, the ability to represent normative anatomical variation is not the only factor that demands consideration. An overly localized model might fail to recognize abnormal patterns as outliers. The abnormal regions may be well represented by a normal dictionary when the block size is inadequate, as the dictionary lacks the ability to capture more global patterns. In particular, it has been demonstrated in [22] (see [22, Fig. 12] for example) that (P1) starts to break down when approximately 40% of $\mathbf{y}$ is corrupted. In our case, this implies that abnormalities have a much higher chance of being represented by $\mathbf{Ax}$ than by $\mathbf{r}$ when the fraction of abnormalities in a block exceeds 40%.

The effect of block size on robustness of the model is best illustrated with a concrete example. Fig. 1 shows a patient's image with lesions. The red box in the left panel of Fig. 1 is a uniform area of bright lesions, and the corresponding box from a normative image is most likely a uniform area of normal-appearing white matter. Those two uniform areas are indistinguishable after being normalized to have unit ℓ_{norm} , therefore, the lesion block will be reconstructed perfectly through (P1). This will not be the case at the scale of the green block in the central panel, where the contrast between the lesions and surrounding normative tissues is preserved and the lesion fraction does not exceed 40%. The lesions will most likely be correctly absorbed into the residual term $\mathbf{r}$, as suggested by the theoretical results in [21], [22].

It is clear from the foregoing discussion that one must carefully balance between model accuracy and robustness when deciding the block size, to which the correctness of (P1) is highly sensitive. Fortunately, treating the blocks jointly through (P2) turns out to be helpful in alleviating the situation. One instance where jointly processing all blocks is beneficial is shown in the right panel of Fig. 1. In this case, (P1) will correctly decompose the four green blocks, and the consistency penalties will “pull” the decomposition of the red block towards the correct direction.

It is also clear that much may be gained by selecting the block size adaptively, provided that the location of the abnormalities is given. Although we do not have the precise location, an estimation from the previous round is available in the proposed iterative framework, and it can be utilized for block size selection in future iterations. In the numerical experiment the default block size is set to ensure that image content is viewed at an intermediate scale. The block size is adaptively increased around the (estimated) abnormal region.

F. Generation of the Abnormality Score

We now describe how to transform the decomposition results to a statistical abnormality map. Let $\mathcal{D}$ be a measure that quantifies the difference between two given images at each voxel. Considering that the brain contains highly variable, complex structure such as the cortex, we expect that the proposed model is not able to fully capture normal variations within healthy subjects, resulting in a non-trivial $\mathcal{D}(\mathcal{I}, \tilde{\mathcal{I}})$ even when $\mathcal{I}$ is considered normal. Therefore, it is important to standardize the difference $\mathcal{D}(\mathcal{I}, \tilde{\mathcal{I}})$ through permutation tests, rather than use it directly, as we quantify the abnormality level of each voxel. More precisely, for every normal image $\mathcal{I}_n$ in $\{\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_N\}$ we compute its normal projection $\tilde{\mathcal{I}}_n$ using the remaining subjects $\{\mathcal{I}_1, \dots, \mathcal{I}_{n-1}, \mathcal{I}_{n+1}, \dots, \mathcal{I}_N\}$ as training samples. This step defines a null distribution for the difference between $\mathcal{I}_n$ and $\tilde{\mathcal{I}}_n$, based on which statistical tests may be performed for any test pair $(\mathcal{I}, \tilde{\mathcal{I}})$.

The difference measure may be simply set as the voxel-wise intensity difference, though some prior knowledge of the abnormality pattern would be helpful when making decisions for $\mathcal{D}$. Loosely speaking, an ideal $\mathcal{D}$ would emphasize the abnormal regions in $\mathcal{D}(\mathcal{I}, \tilde{\mathcal{I}})$. In the numerical implementations, we have experimented solely with the voxel-wise intensity difference, and the Crawford-Howell t-test [28] is employed to test $\mathcal{D}(\mathcal{I}, \tilde{\mathcal{I}})$ against $\{\mathcal{D}(\mathcal{I}_1, \tilde{\mathcal{I}}_1), \dots, \mathcal{D}(\mathcal{I}_N, \tilde{\mathcal{I}}_N)\}$. The resulting test score is used to quantify the abnormality level for each voxel. The abnormality score may also be thresholded to create a binary mask.

G. Abnormality Detection for Generic Functions

With minor modifications, the proposed abnormality detection framework can be applied to generic functions defined on an image grid. All the model ingredients remain valid given a suitable choice of the location-specific dictionary $\mathbf{A}_s$. As the location-specific training samples are not necessarily highly correlated with each other for generic functions, we do not expect an unobserved sample to be reconstructed as a sparse combination of raw training samples. A general class of alternatives includes the output of matrix decomposition algorithms that aim to extract underlying basic elements (atoms) from raw observations. Principal component analysis, non-negative matrix factorization [29]–[31] and dictionary learning methods [32] are all instances of such algorithms. Furthermore, as we lose the robustness provided by the “cross-and-bouquet” structure, it may become necessary to constrain the ℓ_1 -norm of the coefficient $\mathbf{x}$ to prevent over-fitting.

In this paper, we employ the following constrained version of (P2) when decomposing generic functions on the image grid. One example for such generic functions that could benefit the proposed method is the log Jacobian determinant of a deformation field, which captures shape-based abnormalities that are not reflected in the intensity domain. The constrained version (P2-C) has the form

$$\begin{aligned}
\text{(P2-C)} \quad & \min_{\mathbf{x}_s, s \in \mathcal{S}} \sum_{s \in \mathcal{S}} (\|\mathbf{x}_s\|_1 + \|\mathbf{y}_s - \mathbf{A}_s \mathbf{x}_s\|_1) \\
& + \frac{w}{2} \sum_{s_1, s_2 \in \mathcal{S}, s_1 \neq s_2} \Gamma(s_1, s_2) \text{ s.t. } \|\mathbf{x}_s\|_1 \leq u_s, \forall s \in \mathcal{S}.
\end{aligned}$$

We fix the location-specific dictionary $\mathbf{A}_s$ to be the top eigenvectors of the sample covariance matrix that explain 95 percent of the total variance. We empirically set a location-specific upper bound u_s for the ℓ_1 -norm of $\mathbf{x}$. Specifically, we first run the unconstrained version (P2) within the normative database. Each normal logJacDet $\mathcal{J}_n$ is decomposed using the remaining samples $\{\mathcal{J}_1, \dots, \mathcal{J}_{n-1}, \mathcal{J}_{n+1}, \dots, \mathcal{J}_N\}$. We record the ℓ_1 -norm ρ_s^n of the optimal solution $\mathbf{x}_s^n$ at location s . The location-specific upper bound u_s is then

determined by the 80th percentile of $\{\rho_s^1, \rho_s^2, \dots, \rho_s^N\}$.

H. Iterative Registration-Detection Procedure

We now address two assumptions that we have made implicitly so far: 1) all images have been non-linearly aligned to an unspecified coordinate system, and 2) such a registration step can be robust in the presence of abnormalities, i.e., the registration should accurately align the normal regions, while relaxing the deformation around the abnormalities instead of aggressively matching normal regions to abnormal ones (or vice versa).

Regarding the first assumption, there are two competing needs to be met. On the one hand, as mentioned in the previous section, including shape information is useful as it provides complementary information to image intensities. However, shape information can only be captured in a fixed reference space by warping both training samples and the “query” image to this reference. On the other hand, we aim to detect abnormalities in the patient’s brain, a task that can be best accomplished in the patient’s native space.

Regarding the second assumption, the registration step is a challenging task by itself. Regions of the patient’s brain are affected by the disease and the topological property of these regions has been altered. Such topological changes will impact the registration, resulting in heavily distorted training samples, which may in turn compromise the detection task.

To tackle the aforementioned difficulties, we propose an iterative scheme that incorporates both template domain and subject domain learning, while interleaving registration with abnormality detection. The iterative scheme needs to be carefully initialized as an insufficient initial detection will misguide the subsequent registration, which in turn will lead to inaccurate detection. We propose to initialize by learning in a fixed reference space so that we can take into account both intensity-based and shape-based information. This fixed reference space is constructed as an unbiased template $\mathcal{T}$ that is representative of the training set [33]. The training samples $\mathcal{J}_1, \dots, \mathcal{J}_N$ are then non-linearly registered to the template $\mathcal{T}$. We record the log Jacobian determinant of the corresponding deformation fields as $\mathcal{J}_1, \dots, \mathcal{J}_N$. We also denote by $\mathcal{J}_1^T, \dots, \mathcal{J}_N^T$ the warped images. The subject image $\mathcal{S}$ is also warped to $\mathcal{T}$, giving a corresponding diffeomorphic deformation field Φ , log

Jacobian determinant $\mathcal{J}$ and warped image $\mathcal{I}^T$. Let $\mathcal{M}\mathcal{S}_0$ (S for shape) be the abnormality score for $\mathcal{I}$ with respect to the training population $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_N$, and let $\mathcal{M}\mathcal{I}_0$ (I for intensity) be the abnormality score for $\mathcal{I}^T$ with respect to the training population $\mathcal{I}_1^T, \mathcal{I}_2^T, \dots, \mathcal{I}_N^T$ then the initial abnormality map $\mathcal{A}\mathcal{M}_0$ is set as

$$\mathcal{A}\mathcal{M}_0 = \max(|\mathcal{M}\mathcal{S}_0|, |\mathcal{M}\mathcal{I}_0|) \circ \Phi^{-1}.$$

Here Φ^{-1} is well defined as we have ensured that Φ is a diffeomorphism by limiting the maximum allowed displacement in each DRAMMS iteration [34] and combining the per-iteration deformation fields through composition. Lastly, the max operator is used to combine $\mathcal{M}\mathcal{S}_0$ and $\mathcal{M}\mathcal{I}_0$ based on the premise that an abnormality is detectable in either $\mathcal{M}\mathcal{S}_0$, if it has been aggressively matched to a normal region by Φ , or in $\mathcal{M}\mathcal{I}_0$, if it has been preserved by a smoother deformation field.

Based on $\mathcal{A}\mathcal{M}_0$ and its thresholded binary version $\mathcal{B}\mathcal{M}_0$, we refine the registration, the decomposition (P2) and after that the detection result during each subsequent iteration. Within each iteration, robust patient-specific registration is performed by leveraging the current estimate of the pathological regions. Using the updated registration results, an intensity-based abnormality map $\mathcal{M}\mathcal{I}_{k+1}$ is then generated through the decomposition (P2) and subsequent significance test. The update rule for $\mathcal{A}\mathcal{M}$ reads

$\mathcal{A}\mathcal{M}_{k+1} = \max\{|\mathcal{M}\mathcal{I}_{k+1}|, \mathcal{A}\mathcal{M}_0\}$, after which $\mathcal{B}\mathcal{M}_{k+1}$ is updated as

$\mathcal{B}\mathcal{M}_{k+1} := H_t(\mathcal{A}\mathcal{M}_{k+1})$. Here max represents the voxel-wise max operator and H_t represents the hard thresholding operator with threshold t , that is,

$$H_t(x) = \begin{cases} 0 & \text{if } |x| \leq t, \\ 1 & \text{if } |x| > t. \end{cases}$$

The procedure is summarized in Algorithm 1, where all the operators are overloaded in a voxel-wise fashion.

Algorithm 1

Iterative registration and abnormality detection procedure

Input: $\{\mathcal{I}_1, \dots, \mathcal{I}_N\}$, $\mathcal{I}$ distance measure $\mathcal{D}$, initial abnormality map $\mathcal{A}\mathcal{M}_0$ and binary mask $\mathcal{B}\mathcal{M}_0$, total number of iterations K , DRAMMS regularization weights $\{g_1, \dots, g_K\}$ thresholds $\{t_1, \dots, t_K\}$

Output: Projection $\tilde{\mathcal{I}}$ and abnormality map $\mathcal{A}\mathcal{M}$ for $\mathcal{I}$ initialization

for $k = 1$ to K **do**

Registration Part:

if $k = 1$ **then**

 Register $\{\mathcal{I}_1, \dots, \mathcal{I}_N\}$ to $\mathcal{I}$ with g_k , neglecting the regions within $\mathcal{B}\mathcal{M}_{k-1}$ using CFM.

else

Register $\{\mathcal{I}_1, \dots, \mathcal{I}_N\}$ to $\mathcal{I}^{k-1}$ with g_k
end if
Detection Part:
 Step 1: Decompose $\mathcal{I}$ into $\tilde{\mathcal{I}}^k + \mathcal{R}^k$ through (P2)
 Step 2: Compute $\mathcal{M}\mathcal{I}_k$ through Crawford-Howell t-test
 Step 3: Update $\mathcal{A}\mathcal{M}_k$ as $\mathcal{A}\mathcal{M}_k := \max\{|\mathcal{M}\mathcal{I}_k|, \mathcal{A}\mathcal{M}_0\}$ and $\mathcal{B}\mathcal{M}_k$ as $\mathcal{B}\mathcal{M}_k := H_{t_k}(\mathcal{A}\mathcal{M}_k)$
 Step 4: Estimate a pathology-free version $\mathcal{I}^k$ for $\mathcal{I}$ via

$$\mathcal{I}_k := (1 - \mathcal{B}\mathcal{M}_k) * \mathcal{I} + \mathcal{B}\mathcal{M}_k * \tilde{\mathcal{I}}^k \quad (6)$$

end for

Output $\tilde{\mathcal{I}} := \tilde{\mathcal{I}}^K$ and $\mathcal{A}\mathcal{M} := \mathcal{A}\mathcal{M}_K$

The registration part and step 4 in the detection part merit further explanation. Through equation (6) in Algorithm 1 we form a new image $\mathcal{I}^k$ by inpainting the original image $\mathcal{I}$ with its current normal projection $\tilde{\mathcal{I}}^k$. In particular, $\mathcal{I}^k$ takes the value of in the estimated normal regions ($\mathcal{B}\mathcal{M}_k=0$) and the value of $\tilde{\mathcal{I}}^k$ in the estimated abnormal regions ($\mathcal{B}\mathcal{M}_k=1$). Therefore, $\mathcal{I}^k$ is an estimate of what a patient's brain looks like before the development of the pathologies, and a substitute for $\mathcal{I}$ as the reference image in the subsequent registration. This procedure is adopted in Algorithm 1, except for the first iteration where $\mathcal{I}^0$ is not available and one instead resorts to CFM.

The max operator that is used to combine $|\mathcal{M}\mathcal{I}_k|$ and $\mathcal{A}\mathcal{M}_0$ also demands some clarification. The initial abnormality scores capture shape-based abnormalities that can not be revealed by the future intensity-based abnormality maps. Therefore, merging the initial abnormality scores with $\mathcal{M}\mathcal{I}_k$ at each subsequent iteration allows the transfer of important information regarding pertaining shape abnormalities. It could be possible that along with this useful information, falsely identified abnormalities propagate along iterations. However, we have empirically found that the initial abnormality scores tend to under-segment the lesions rather than giving artificially high abnormality scores to normal regions. This behavior is demonstrated in Fig. 5 in the experiments section.

III. Numerical Experiments

In this section, we present the experimental validations of the proposed method. We first experiment on brain images from healthy subjects to examine how well normative variation can be captured by the model. We then apply the proposed abnormality detection framework on both simulated and clinical brain lesion data to analyze its performance on registration and abnormality detection. The proposed framework is also evaluated on an Alzheimer's

Disease database to demonstrate the generality of the method. Unless otherwise specified, DRAMMS [35], [36] is set as the default choice for deformable registration throughout the experiment.

A. Database: Brain Lesions

The normative training set (NC) consists of MRI scans from 72 healthy subjects, with voxel size $0.93 \text{ mm} \times 0.93 \text{ mm} \times 1.5 \text{ mm}$. All scans are manually inspected to ensure they do not include any visible pathologies, beyond those common within the subjects' age range (e.g., peri-ventricular hyper-intensities). The first test set (TestSet 1), involves images from 14 subjects with brain lesions, with manual masks of the lesion created by a radiologist. The second test set (TestSet2), involves images from 20 subjects from an aging population with high prevalence of abnormalities. On TestSet2, manual masks for white matter lesions and infarcts have been created by one expert. The average age of the training subjects is 59.3 with a standard deviation of 4. The mean ($\pm$ standard deviation) age is 60.8 ± 5.8 for TestSet1 and for TestSet2. We mainly use FLAIR images for validation, but a multi-modality extension is straight-forward and will be discussed in Section IV.

B. Performance on Normative Images

In the first experiment we examine the performance of the model on capturing normative brain variation. 20 subjects are randomly selected from NC and treated as a test subject in a leave-one-out fashion. For each left out subject, we register the remaining samples to it using a moderate regularization parameter ($g=0.3$) for DRAMMS. We do not follow the iterative scheme, as pathological regions are absent in the test image. The focus here is to quantify the distance between $\tilde{\mathcal{J}}$ and $\mathcal{J}$ for a normal $\mathcal{J}$, given good spatial normalization.

Ideally, $\tilde{\mathcal{J}}$ should be exactly equal to $\mathcal{J}$ as the latter is a normal image. In practice, however, it is unrealistic to expect a perfect $\tilde{\mathcal{J}}$ through the numerical procedure (P2), and one should be content when $\tilde{\mathcal{J}}$ and $\mathcal{J}$ are reasonably close. Here, the distance measure is chosen as the root-mean-square error (RMSE). We look into the effect of block size on RMSE to verify the hypothesis made in Section II-E, that is, the representation ability of the model decreases with increasing block size. We evaluate the RMSE using 3 different block sizes: $7.44 \text{ mm} \times 7.44 \text{ mm} \times 6 \text{ mm}$, $14.88 \text{ mm} \times 14.88 \text{ mm} \times 12 \text{ mm}$, and $29.76 \text{ mm} \times 29.76 \text{ mm} \times 24 \text{ mm}$. The step size for the sliding windows is set as half of the block size, and the consensus weight as $w=1$, a value used throughout the experiments as it leads to sufficient consistency. The mean $\pm$ std RMSE with the 3 incremental block sizes are 8.64 ± 1.13 , 9.70 ± 1.19 , 10.22 ± 1.21 , respectively, which confirms our foregoing hypothesis.

For a more concrete understanding about the quality of $\tilde{\mathcal{J}}$, we also compute the RMSE between the original images and their smoothed versions. The images are smoothed with isotropic Gaussian kernels of standard deviations 1 mm, 1.5 mm and 2 mm. The mean $\pm$ std RMSE for the 3 incremental bandwidths is 9.04 ± 0.82 , 10.69 ± 1.03 , and 13.61 ± 1.17 , respectively. For the selected 20 normal images, the quality of the projection $\tilde{\mathcal{J}}$, using the medium block size, is slightly superior to that of the smoothed original image with a standard deviation 1.5 mm in terms of RMSE. Representative examples of $\mathcal{J}$ and $\tilde{\mathcal{J}}$ are shown in appendix D in the supplementary file.

C. Simulated Experiment: Brain Lesions

We next evaluate the framework on simulated abnormalities. The advantages of a simulated study are threefold. Firstly, ground truth is available to evaluate the method for such data. Secondly, one may freely vary the location and size of the abnormalities and examine how these factors affect the results. Lastly, one may quantify the quality of any registration that involves simulated images as we know the normal image $\mathcal{I}^G$ on which the abnormality is simulated. Such knowledge provides us the “ground truth” for any training sample $\mathcal{I}_n$, which is the registration result when warping $\mathcal{I}_n$ to $\mathcal{I}^G$.

The simulated abnormality is ellipsoidal in shape and it consists of a necrotic interior with surrounding bright lesions. The intensity profile of the pathological region is set according to its typical appearance on FLAIR images. More specifically, the intensity of the necrotic interior is matched to the intensity profile of the cerebrospinal fluid (CSF), while the intensity of the peripheral region is given a 60% signal increase when compared to normal white matter. While it is relatively easy to spot hyper-intense lesions, the necrotic parts are much more difficult to detect as they may share location and intensity profile with cortical CSF. We simulate lesions on 4 different anatomical locations (Zone 1, Zone 2, Zone 3 and Zone 4) with 5 progressively larger sizes on the same image. The semi-principal axes length ranges from $10 \text{ mm} \times 10 \text{ mm} \times 8 \text{ mm}$ to $19 \text{ mm} \times 19 \text{ mm} \times 15 \text{ mm}$. Representative examples of simulated lesions with intermediate size are shown in Fig. 2. Note that Zone 1, Zone 2 and Zone 3 lesions are all simulated in the cortex, while Zone 4 lesions are positioned in deep white matter.

Next we explain the implementation details, starting with how the initial abnormality map is generated. Recall from Section II-H that we have constructed an unbiased template $\mathcal{T}$ for the normative set NC. The normative samples are warped to $\mathcal{T}$. The log Jacobian determinant maps (logJacDet) of the associating deformation fields are recorded as

$\mathcal{J}_1, \mathcal{J}_2, \dots, \mathcal{J}_N$ and the warped images as $\mathcal{I}_1^T, \mathcal{I}_2^T, \dots, \mathcal{I}_N^T$. We register each simulated test image to $\mathcal{T}$ and save the corresponding logJacDet $\mathcal{J}$ together with the warped image $\mathcal{I}^T$.

Each warped image is decomposed directly with (P2) using the example-based dictionary.

The intensity-based abnormality map $\mathcal{M}\mathcal{I}_0$ is the Crawford-Howell t-statistics of $\mathcal{I}^T - \tilde{\mathcal{I}}^T$ against $\{\mathcal{I}_1^T - \tilde{\mathcal{I}}_1^T, \mathcal{I}_2^T - \tilde{\mathcal{I}}_2^T, \dots, \mathcal{I}_N^T - \tilde{\mathcal{I}}_N^T\}$. The block size is set to $14.88 \text{ mm} \times 14.88 \text{ mm} \times 12 \text{ mm}$. We solve the modified version (P2-C) when decomposing the logjacDet $\mathcal{J}$'s, with the block size set to $7.44 \text{ mm} \times 7.44 \text{ mm} \times 6 \text{ mm}$. The shape-based abnormality map, or $\mathcal{M}\mathcal{S}_0$, is also calculated using the Crawford-Howell t-test. Representative examples of $(\mathcal{M}\mathcal{I}_0, \mathcal{M}\mathcal{S}_0)$ pairs are shown in Fig. 3. $\mathcal{M}\mathcal{S}_0$ better captures the necrotic part of the lesion, while the bright surroundings are more salient in $\mathcal{M}\mathcal{I}_0$.

The initial abnormality map $\mathcal{A}\mathcal{M}_0$ is thresholded at $t = 3$ when producing the binary $\mathcal{B}\mathcal{M}_0$, after which we proceed to the iterative registration-detection scheme. The specific parameter settings are detailed as follows. The total number of iterations is set to $K = 5$. We fix the DRAMMS regularization weights to $g_k = 0.3$ and the threshold values to t_k for all k . The

threshold approximately sets a significance level $p < 0.005$ (if the values were truly normally distributed). The block size is determined in an adaptive way. The default block size is $14.88 \text{ mm} \times 14.88 \text{ mm} \times 12 \text{ mm}$, and the block size is doubled until at least 60% of the volume is classified as normal in the previous iteration. The 60% threshold is set based on the 40% breakdown point of (P1), as detailed in Section II-E.

We now demonstrate that the discriminative power and the registration accuracy of the outputs simultaneously improve along the iterations in Algorithm 1. We quantify the discriminative power of the abnormality map $\mathcal{A}\mathcal{M}_k$ with two measures: Area Under Curve (AUC) and Hellinger Distance (HD) between the abnormality score distribution of the True Positives (TPs) and that of the True Negatives (TNs). The HD measure is defined in what follows. Let $\mathbf{P}_0$ denote the abnormality score distribution of the TNs

($P_0 = \{\mathcal{A}\mathcal{M}_k(v) | \mathcal{G}\mathcal{T}(v)=0\}$) and $\mathbf{P}_1$ denote that of the TPs ($P_1 = \{\mathcal{A}\mathcal{M}_k(v) | \mathcal{G}\mathcal{T}(v)=1\}$), where $\mathcal{G}\mathcal{T}$ is the ground truth labelling. $\mathbf{P}_0$ and $\mathbf{P}_1$ are indicative of the quality of the abnormality maps, as a good abnormality map will give low values to $\mathbf{P}_0$ and high values to $\mathbf{P}_1$. Therefore, higher values of HD imply better separation between normal and abnormal regions in the abnormality map.

The registration accuracy at each iteration is quantified in the following manner. Let $\mathcal{J}^G$ be the normal image on which the abnormalities are simulated. Since no manual landmarks or parcellation are available for any training sample $\mathcal{J}_n$, the gold standard is chosen as the warped image $\mathcal{J}_n^G$ when registering $\mathcal{J}_n$ to $\mathcal{J}^G$. For each simulated image $\mathcal{J}^S$, we use $\mathcal{J}_n^k$ to represent $\mathcal{J}_n$ warped to $\mathcal{J}^S$ at iteration k . The quality of $\mathcal{J}_n^k$ is measured by the relative difference Γ_n^k between $\mathcal{J}_n^k$ and $\mathcal{J}_n^G$ within the lesion mask, that is,

$\Gamma_n^k = \|M(\mathcal{J}_n^k - \mathcal{J}_n^G)\|_2 / \|M(\mathcal{J}_n^G)\|_2$, where M is the operator that extracts the lesion region from the image. Twenty images from NC are randomly chosen for validation. Therefore, registration accuracy at each iteration is quantified through $20 \times 20 = 400$ counts of relative differences. We also examine the quality of the normal projection $\tilde{\mathcal{J}}_k$, by measuring $\Lambda_k = \|M(\tilde{\mathcal{J}}_k - \mathcal{J}^G)\|_2 / \|M(\mathcal{J}^G)\|_2$, which is the relative distance to the “ground truth” $\mathcal{J}^G$ within the lesion region. Smaller values of Λ_k imply that $\tilde{\mathcal{J}}_k$ better approximates what the patient’s image looks like before pathology development.

The statistics for the four aforementioned measures are summarized in separate panels in Fig. 4. We also include AUC and HD for $\mathcal{A}\mathcal{M}_0$, the initial abnormality map derived using both intensity-based and shape-based information. From Fig. 4 we observe a converging trend for all four measurements as the number of iterations increases. In particular, their medians stabilize after the second iteration. The box plots show clear improvement for the HD, as well as a decreasing pattern for both Γ_n^k and Λ_k along the iterations. The AUCs of the abnormality maps peak at the second iteration, and degrade slightly onwards. However, the AUC medians differ by less than 0.0001 between the second and the final iteration and this difference is not statistically significant.

A visual inspection reveals greater details on how the outputs progress with the number of iterations. We show the abnormality maps qualitatively by presenting three axial slices in

Fig. 5. The initial abnormality maps, together with abnormality maps at iteration 1 and 2, are plotted alongside FLAIR image slices. In all three cases, the lesion regions are better delineated as the number of iterations increases. In particular, the initial abnormality maps fail to cover the entire lesion region and under-segment the abnormalities. The missing parts are successfully filled in by the following iterations, leading to more accurate lesion detection. This behavior is also observed in cases not shown here.

Next we compare the output of Algorithm 1 with that of three registration schemes: direct affine registration (AR), direct deformable registration (DR) and deformable registration with cost function masking (DRM). AR with CFM is not included as in our experiments the scale of the abnormalities is not sufficiently large to influence affine registration. To ensure a fair comparison between the proposed method (DRI, I for inpainting) and DRM, we use $\mathcal{B}\mathcal{M}_4$, the abnormality mask estimated at the penultimate iteration, as the input for CFM. As before, we start by examining the registration accuracy. The warped images by the proposed scheme are denoted by $\mathcal{J}_n^{DRI}$. The warped images by the other three candidates are denoted by $\mathcal{J}_n^{AR}$, $\mathcal{J}_n^{DR}$ and $\mathcal{J}_n^{DRM}$, respectively. The same image batch from NC is recycled from the previous experiment for validation. We continue to measure registration accuracy by the relative difference $\Gamma^{(\bullet)}$ between $\mathcal{J}_n^{(\bullet)}$ and $\mathcal{J}_n^G$ within the lesion regions.

Fig. 6 summarizes the results on the twenty simulated images. The results are grouped by lesion zones. We observe the best performance for all candidate methods in cases where lesions are simulated in the deep white matter (Zone 4). This indicates that DRAMMS does not try to aggressively match normal anatomy to the simulated infarcts in Zone 4. All methods achieve comparable median differences, with DR being the least consistent, as suggested by its larger upper quartile values.

Not surprisingly, AR performs stably with respect to the scale of the abnormalities. In fact, the relative differences tend to decrease as lesion size grows. This is expected since the underlying matching criterion is global and the transformation model has only 12 degrees of freedom. In contrast, DR is highly sensitive to the location and size of the abnormalities. Its performance drops drastically as we increase the size of the abnormalities that are simulated in cortical regions (Zone 1, 2 and 3). DRM does not share the same issue; its performance does not degrade with increasing lesion size, even when the provided mask is merely an estimate of the ground truth. Incorporating CFM with deformable registration also tightens the inter-quartile range of the relative differences. However, one only observes the benefit of CFM in terms of median difference when the lesion size surpasses a certain threshold. For example in Zone 3, DRM starts to outperform DR from the intermediate size onwards. Meanwhile in Zone 2, it compares unfavorably to its vanilla counter part until the lesion size reaches the largest value.

DRI produces competitive results throughout this experiment. In particular, it achieves the lowest median difference in 13 out of the 15 cases with cortical lesions, except for SIM2-1 and SIM3-3, where its median is slightly higher than that of DR and AR, respectively. DRI clearly outperforms DRM ($p < 10^{-40}$ under the Wilcoxon signed-rank test) in cases with

cortical lesions, which demonstrates the premise of using an inpainting step such as equation (6) in Algorithm 1 for matching normative images to images with abnormalities.

We revisit AUC and HD for evaluating detection accuracy, but this time we compare the proposed method to those of the three other competitors. All methods achieve AUCs that are higher than 0.999 for the five samples in Zone 4. Fig. 7 showcases the results for the other cases.

DRI and DRM achieve comparable performance in AUC and HD. The difference between these two methods is not statistically significant under the Wilcoxon signed-rank test ($p = 0.82$ and $p = 0.86$ for AUC and HD, respectively). DRI and DRM substantially outperform AR and DR. The differences between $\{\text{DRI}, \text{DRM}\}$ and $\{\text{AR}, \text{DR}\}$ are statistically significant, with $p < 10^{-4}$ for all pairwise comparisons. The detection performance of affine and deformable registration both degrade with increasing lesion size, with deformable registration decaying at a faster speed. One may conclude that AR is insufficient for the proposed regional learning scheme in this simulated experiment, and DR needs to couple with CFM or inpainting for reliable performance.

As the last synthetic experiment, we evaluate the performance of the proposed multi-variate method against a standard univariate statistical test. We fix the registration as the “ground truth” used in previous experiments, that is, the training database $\{\mathcal{I}_1, \dots, \mathcal{I}_N\}$ warped to the subject space with the reference image being $\mathcal{I}^G$ (the base image on which the abnormalities are simulated). We again denote by $\{\mathcal{I}_1^G, \dots, \mathcal{I}_N^G\}$ the warped images. The uni-variate statistics are computed as the Crawford-Howell t-statistics of $\mathcal{I}^S$ against $\{\mathcal{I}_1^G, \dots, \mathcal{I}_N^G\}$. The uni-variate statistical parametric maps are then compared to the abnormality maps generated by DRI in terms of AUC and HD. Both methods achieve AUCs that are consistently higher than 0.995 for the 5 images where infarcts are simulated in the deep white matter (Zone 4). The results for the remaining 15 cases are summarized in the box plots shown in Fig. 8. We observe a statistically difference for both AUC and HD ($p < 0.05$ and $p < 0.01$, respectively).

D. Experiment on Clinical Data: Brain Lesions

We now turn to the experiments on clinical data. As detailed in Section III-A, TestSet1 involves images from 14 patients with white matter lesions and/or cortical infarcts. Periventricular hyper-intensities are the dominant abnormality type in 9 out of the 14 patients, a pathology that is also present in NC. TestSet2 involves elder subjects that are 15 years older, on average, than the subjects in the training database. In other words, NC is not an age-appropriate training database for TestSet2, and we expect that age-related differences, such as cortical atrophy and ventricular enlargement, would be identified as abnormal by our statistical framework.

There are two inherent difficulties when experimenting with the given clinical database. First of all, assessing registration quality is no longer possible due to the lack of a ground truth. Furthermore, the presence of periventricular lesions in NC and the age-difference between NC and TestSet2 indicate that deviations from NC will not entirely match the lesions that are

present in TestSet1 and TestSet2. As a consequence, we expect these factors to act as confounders and compromise the ability of the abnormality detectors to separate lesions from normal tissues. However, this behavior is desirable as the proposed method constitutes a generic abnormality detection framework that does not specifically target lesion delineation.

We use parameters identical to the ones used in the simulated experiments, except that we now terminate Algorithm 1 after two iterations. Different from the simulated scenario, we use T1 images for registration and FLAIR images for detection. T1 images are routinely collected, characterized by high tissue contrast and high resolution, making them very suitable for fine registration. Similar to FLAIR, lesions appear distinct from normal tissues in T1 images and pose similar challenges to registration [11]. We continue using AUC and HD to quantify discriminative power. The proposed method is compared to AR and DR. As the simulated experiment shows that DRI outperforms DRM in registration but behaves similarly in detection, we exclude DRM in this experiment as registration quality can not be measured for the clinical database.

The box plots shown in Fig. 9 summarize the statistics for the candidate methods. DRI achieves the highest median both in AUC and HD. AR follows closely as the second best method, while DR performs worse than its competitors. Under the Wilcoxon signed-rank test, AR and DRI both perform significantly better than DR ($p < 10^{-4}$). The difference between AR and DRI is not as substantial but remains statistically significant, with $p < 0.05$ and $p < 0.01$ for AUC and HD, respectively.

We qualitatively evaluate the methods by visually inspecting the outputs. Fig. 10 shows axial slices from four representative cases. For each method, we show the normal projection and abnormality map. Overall, the projection images based on deformable registrations are substantially sharper than the ones based on affine registration. This is hardly surprising as finer registration leads to better localized dictionary A_g . Fig. 10 also shows that the projection images computed from DR and DRI have highly similar appearance outside the lesion regions, implying that the inpainting steps primarily affect the abnormal parts of the brain.

Let us focus our visual inspection on lesion areas in the brain. One immediately notices that most periventricular bright lesions are (at least partially) preserved in the projections, which is the result of their presence in the normative dataset. Examples of this behavior are marked by the green circles in row 1. When viewing the abnormality maps, one notices that bright lesions are more salient in the AR and DRI maps, and they appear less notable in the DR maps. We mark instances of such differences by red circles in row 1.

We next inspect the cortical infarcts that are present in rows 2 and 3. In row 2, the necrosis is largely preserved in the DR projection, because the training samples have been “distorted” to match the patient’s image. Both AF and DRI are able to partially recover the lost tissues, with DRI being able to generate a sharper projection. A similar observation can be made for the subject shown in row 3. Both AF and DRI successfully recover the entire necrotic part of the infarct with DRI producing a shaper image, while DR only fills in part of the necrosis.

The abnormality maps highly correlate with the normal projections. In particular, necroses appear more salient in the abnormality maps when tissues are better recovered in the projections.

Fig. 10 also sheds light on why AR outperforms DR in this experiment. Although DR better matches normal regions in the training subject to normal regions in the test subject, thus producing sharper projections than AR, it also matches normal regions in the training subject to lesions in the test subject in an aggressive way. The aggressive matching “abnormalizes” the dictionaries used in the subsequent detection step, making DR insensitive to lesions.

Aside from lesions, sulcus is one of the dominant regions that may be identified as abnormal. One may observe instances of sulci flagged as abnormal by all three detectors, with AR being the most aggressive and DR being the most conservative. However, it is difficult to judge which method makes better decisions. On the one hand, the detector might flag sulci as abnormal because the underlying model fails to capture the complex normative variation in the cortex. On the other hand, the flagged sulci might indicate actual cerebral atrophy. In the second case, the atrophy might be caused by normal aging, injuries or other diseases.

Ventricles may also be identified as abnormal, which can be seen in row 1 and row 3. DR is insensitive to ventricular differences, as direct deformable registration accurately matches ventricle shapes between the normal samples and the patient’s image, even when the latter’s ventricle is substantially enlarged. AR and DRI signify enlarged ventricles in distinct ways. Affine registration fails to match the shape of the ventricles, thus flags the misalignment as pathological. In DRI, the high abnormality scores in the ventricles come from the shape-based maps that are computed using the log Jacobian determinant. As a result, they appear much more diffused, spreading throughout the ventricles rather than clustering around the borders.

E. Case Example: Alzheimer’s Disease

In this section, we present additional experimental results on Alzheimer’s Disease (AD) patients to demonstrate the generality of the proposed framework. In particular, T1 scans from a population of healthy older adults along with a cohort of AD patients are pre-processed using a previously validated and published pipeline [37], producing regional tissue volumetric (RAVENS) maps that allow us to quantify the amount of brain tissues in the vicinity of each voxel. The RAVENS maps are normalized by individual intracranial volume to adjust for global differences in intracranial size, and are smoothed to incorporate neighborhood information using a 5-mm Full Width at Half Maximum Gaussian filter.

The proposed framework is applied to the log Gray Matter-RAVENS maps for abnormality detection. As all RAVENS maps reside in the same standardized template space (atlas) by definition, we only need to perform the detection part of the framework. Specifically, the log GM-RAVENS maps are decomposed through the constraint version (P2-C), following the same parameters as the ones used to decompose the logJacDet $\mathcal{J}^s$, except that the block size is set to $8 \text{ mm} \times 8 \text{ mm} \times 8 \text{ mm}$.

Since this experiment was merely designed for demonstrating the generality of our approach beyond lesion-like abnormalities, we only performed a qualitative evaluation of the method based on Fig. 11, in which representative examples of the abnormality maps were displayed. The abnormality maps have been thresholded at $t = -2.5$ before being overlaid on the atlas. A highly negative abnormality score at a given voxel indicates substantial loss of gray matter around that voxel. Fig. 11 shows abnormality map slices from three AD subjects. The hippocampus regions of subject 1 and 2 (shown in row 1 and 2, respectively) are salient in the abnormality maps with extremely negative scores, which is consistent with the typical pattern of atrophy found in AD. Interestingly, this is not the case for the subject shown in row 3, where the hippocampus is not at all salient in the abnormality map. Instead, a pronounced frontal and cerebellar atrophy pattern is detected by the proposed method. It is conceivable that either this patient had a different underlying pathology that was clinically manifested as an AD phenotype, or some co-existing pathology to AD pushed this individual beyond the threshold of clinical AD for smaller levels of hippocampal atrophy. Regardless of the (unknown) underlying ground truth, this example demonstrates the potential clinical utility of the proposed approach as a means for constructing visual representations of abnormalities to be further evaluated by clinicians along with other variables.

IV. Discussion

We have presented an abnormality detection framework that is based on a generic formulation. The framework makes minimal assumptions on the characteristics of the abnormalities and potentially targets a wide range of pathologies. The method utilizes both intensity and shape information from the image to capture abnormalities from multiple perspectives. All image features on the 3-d grid, treated as a collection of overlapping local blocks, are processed in a unified way. They are decomposed into a normal part and a residual part using joint basis-pursuit, after which abnormalities are identified as statistically significant areas in the residual. The method also incorporates an iterative registration and detection scheme aiming to tackle registration and detection simultaneously. At each iteration, the robustness and accuracy of the registration improve by leveraging the detection outcome. The registration improvement in turn benefits subsequent detection steps.

The proposed framework may be modified and generalized in many ways. For instance, only the Jacobian determinant is used as a feature for detecting abnormalities in the deformation field. One may perform statistical learning, using the same method, on other aspects of the deformation field (divergence, curl etc.), or on the deformation field itself. One may also perform a sample selection procedure prior to registering the training database to the subject, similar to the atlas selection process in multi-atlas segmentation [19]. The extension to a multi-modal setting is straightforward, at the expense of higher computational load. For example, one could apply the method to each modality separately and fuse the results. The proposed framework is also flexible. In fact, the decomposition (P2) may be tailored to the specific applications, as is done in (P2-C). Moreover, one could replace DRAMMS by any other registration method of his/her choice. Last but not least, the abnormality map could potentially be used as prior information (and/or features) by specific detectors to boost the their performance.

The proposed framework can be useful in various clinical scenarios. For instance, the abnormality maps can be used for automated screening of scans by flagging the ones with abnormalities that need further expert attention. Along the same lines, and for more subtle abnormalities, a tool for directing the attention of expert readers towards regions displaying relatively large deviation from normality would be quite helpful. For example, the abnormalities present in Alzheimer's Disease patients, such as hippocampal atrophy and posterior cingulate hypometabolism, would not qualify as a "lesion". Nonetheless, we have shown that a 3-D abnormality map produced by the proposed approach could help as a computationally derived image of brain atrophy, both for visual assessment and as a quantitative assay.

Despite its many advantages, the current work is still limited in several aspects. The proposed framework is computationally intensive, as one needs to register a entire database to a test image. The computational burden may be alleviated by the aforementioned sample selection procedure, or by using faster, but less accurate, registration methods. These alternatives may have a negative impact on detection accuracy, the details of which are left to a future study. Secondly, the method is sensitive to age and scanner differences, as are most statistical learning approaches. For lesion detection, the training database should, ideally, be age appropriate for the test image and should be acquired using the same imaging protocol. If such requirements fail to be met, confounding factors (e.g., age and scanner differences) will compromise the performance of the method on detecting lesions. Thirdly, the registration part of the framework relies on cost function masking and image inpainting, which can not account for substantial tissue displacements around certain types of lesions (e.g., tumors). As a result, simultaneous detection and registration, especially the latter component, can not be tackled by the framework in its current form for those lesions. In a future work we plan to extent CFM by exploiting more advanced growth models (e.g., the ones proposed in [13]–[15]). Lastly, while the proposed method is built upon a generic formulation that is not limited by a specific application, it is validated mainly on brain lesion MRI data in this paper. Future work will leverage its potential in applications involving other anatomical regions, image modalities and abnormality types.

It is important to note that a generic approach as such is not meant as a substitute, but rather as a complement to the abundant specialized methods. In particular, a generically formulated method may not reach the accuracy of a specifically trained model when applied to the pathology the latter targets. Its main advantage lies in the potential to produce solid results on a wider class of abnormalities. In practice, a choice could be made based on how one wishes to balance generality and specificity in a particular application.

V. Conclusion

We presented a framework that automatically detects abnormal regions in medical images. The framework is built upon a generic formulation, where pathological regions are detected as deviations from a statistical model that encompasses both shape and intensity information learned from a database of a healthy population with the aim to capture normative variation. The framework interleaves registration and abnormality detection in an iterative fashion, gradually improving the robustness of the spatial normalization and the precision of the

detection. We validated the framework on both simulated and clinical brain lesion data, achieving competitive registration accuracy and discriminatory power in both scenarios. The iterative scheme was compared to direct affine and deformable registration, while the multi-variate basis-pursuit decomposition was compared to a uni-variate statistical procedure. The iterative scheme is able to produce the most accurate registration results in the presence of cortical infarcts, and the multi-variate model is superior to the uni-variate method at separating lesions from normal regions. The framework was also applied on images from AD patients for detecting atrophy patterns, so as to demonstrate its potential to generalize across different abnormality detection problems.

References

1. García-Lorenzo D, Francis S, Narayanan S, Arnold DL, Collins DL. Review of automatic segmentation methods of multiple sclerosis white matter lesions on conventional magnetic resonance imaging. *Med Image Anal.* 2013; 17(1):1–18. [PubMed: 23084503]
2. Gillebert CR, Humphreys GW, Mantini D. Automated delineation of stroke lesions using brain {CT} images. *NeuroImage, Clin.* 2014; 4(0):540–548. [PubMed: 24818079]
3. Shen S, Szameitat A, Sterr A. Detection of infarct lesions from single MRI modality using inconsistency between voxel intensity and spatial location; A 3-d automatic approach. *IEEE Trans Inf Technol Biomed.* Jul; 2008 12(4):532–540. [PubMed: 18632333]
4. Harmouche R, Subbanna N, Collins D, Arnold D, Arbel T. Probabilistic multiple sclerosis lesion classification based on modeling regional intensity variability and local neighborhood information. *IEEE Trans Biomed Eng.* May; 2015 62(5):1281–1292. [PubMed: 25546852]
5. Prastawa M, Bullitt E, Ho S, Gerig G. A brain tumor segmentation framework based on outlier detection. *Med Image Anal.* 2004; 8(3):275–283. [PubMed: 15450222]
6. Mourão-Miranda J, et al. Patient classification as an outlier detection problem: An application of the one-class support vector machine. *NeuroImage.* 2011; 58(3):793–804. [PubMed: 21723950]
7. Gordillo N, Montseny E, Sobrevilla P. State of the art survey on MRI brain tumor segmentation. *Magn Reson Imag.* 2013; 31(8):1426–1438.
8. Zacharaki E, Bezerianos A. Abnormality segmentation in brain images via distributed estimation. *IEEE Trans Inf Technol Biomed.* Mar; 2012 16(3):330–338. [PubMed: 22157062]
9. Erus G, Zacharaki EI, Davatzikos C. Individualized statistical learning from medical image databases: Application to identification of brain lesions. *Med Image Anal.* 2014; 18(3):542–554. [PubMed: 24607564]
10. Zeng K, Erus G, Tanwar M, Davatzikos C. Brain abnormality segmentation based on l_1 -norm minimization. *Proc SPIE.* 2014; 9034:903 409–903 409–7.
11. Brett M, Leff AP, Rorden C, Ashburner J. Spatial normalization of brain images with focal lesions using cost function masking. *NeuroImage.* 2001; 14(2):486–500. [PubMed: 11467921]
12. Andersen SM, Rapcsak SZ, Beeson PM. Cost function masking during normalization of brains with focal lesions: Still a necessity? *NeuroImage.* 2010; 53(1):78–84. [PubMed: 20542122]
13. Konukoglu E, et al. Image guided personalization of reaction-diffusion type tumor growth models using modified anisotropic eikonal equations. *IEEE Trans Med Imag.* Jan; 2010 29(1):77–95.
14. Gooya A, et al. Glistr: Glioma image segmentation and registration. *IEEE Trans Med Imag.* Oct; 2012 31(10):1941–1954.
15. Swanson KR, Bridge C, Murray J, A EC Jr. Virtual and real brain tumors: Using mathematical modeling to quantify glioma growth and invasion. *J Neurol Sci.* 2003; 216(1):1–10. [PubMed: 14607296]
16. Trounev A, Younes L. Metamorphoses through lie group action. *Foundat Comput Math.* 2005; 5(2): 173–198.
17. Li X, Long X, Laurienti P, Wyatt C. Registration of images with varying topology using embedded maps. *IEEE Trans Med Imag.* Mar; 2012 31(3):749–765.

18. Periaswamy S, Farid H. Medical image registration with partial data. *Med Image Anal.* 2006; 10(3):452–464. [PubMed: 15979375]
19. Aljabar P, Heckemann R, Hammers A, Hajnal J, Rueckert D. Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy. *Neuroimage.* 2009; 46(3):726–738. [PubMed: 19245840]
20. van der Lijn F, den Heijer T, Breteler MM, Niessen WJ. Hippocampus segmentation in MR images using atlas registration, voxel classification, and graph cuts. *Neuroimage.* 2008; 43(4):708–720. [PubMed: 18761411]
21. Wright J, Ma Y. Dense error correction via l1-minimization. *Proc IEEE Int Conf Acoust., Speech Signal Process.* 2009:3033–3036.
22. Wright J, Yang A, Ganesh A, Sastry S, Ma Y. Robust face recognition via sparse representation. *IEEE Trans Pattern Anal Mach Intell.* Feb; 2009 31(2):210–227. [PubMed: 19110489]
23. Chen S, Donoho D, Saunders M. Atomic decomposition by basis pursuit. *SIAM J Sci Comput.* 1998; 20(1):33–61.
24. Zhang H, Yin W, Cheng L. Necessary and sufficient conditions of solution uniqueness in l1-norm minimization. *J Optimizat Theory Appl.* 2015; 164(1):109–122.
25. Elad, M. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing* Springer. 2010. [Online]. Available: <http://books.google.com/books?id=d5b6lJl9BvAC>
26. Samar S, Boyd S, Gorinevsky D. Distributed estimation via dual decomposition. *Proc Eur Control Conf.* 2007:1511–1519.
27. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundat Trends Mach Learn.* 2011; 3(1):1–122.
28. Crawford JR, Garthwaite PH. Testing for suspected impairments and dissociations in single-case studies in neuropsychology: Evaluation of alternatives using Monte Carlo simulations and revised tests for dissociations. *Neuropsychology.* 2005; 19(3):318. [PubMed: 15910118]
29. Paatero P, Tapper U. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics.* 1994; 5(2):111–126.
30. Lee DD, Seung HS. Learning the parts of objects by non-negative matrix factorization. *Nature.* 1999; 401(6755):788–791. [PubMed: 10548103]
31. Sotiras A, Resnick SM, Davatzikos C. Finding imaging patterns of structural covariance via non-negative matrix factorization. *NeuroImage.* 2015; 108:1–16. [PubMed: 25497684]
32. Elad M, Figueiredo MAT, Ma Y. On the role of sparse and redundant representations in image processing. *Proc IEEE.* Jun; 2010 98(6):972–982.
33. Fonov V, et al. Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage.* 2011; 54(1):313–327. [PubMed: 20656036]
34. Choi Y, Lee S. Injectivity conditions of 2d and 3d uniform cubic b-spline functions. *Graph Models.* 2000; 62(6):411–427.
35. Ou Y, Sotiras A, Paragios N, Davatzikos C. DRAMMS: Deformable registration via attribute matching and mutual-saliency weighting. *Med Image Anal.* 2011; 15(4):622–639. [PubMed: 20688559]
36. Ou Y, Akbari H, Bilello M, Da X, Davatzikos C. Comparative evaluation of registration algorithms in different brain databases with varying difficulty: Results and insights. *IEEE Trans Med Imag.* Oct; 2014 33(10):2039–2065.
37. Goldszal A, et al. An image processing protocol for the analysis of MR images from an elderly population. *J Comput Assist Tomogr.* 1998; 22:827–837. [PubMed: 9754125]
38. Efron B, Hastie T, Johnstone I, Tibshirani R. Least angle regression. *Ann Statist.* 2004; 32(2):407–499.
39. Beck A, Teboulle M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J Imag Sci.* 2009; 2(1):183–202.

Fig. 1.
Illustrative examples showing the effect of block size on the decomposition results.

Fig. 2.
Representative examples of the simulated abnormal images. Each column shows lesion region simulated at one selected location and the region is of intermediate size.

Fig. 3.

Representative examples of $(\mathcal{M} \mathcal{I}_0, \mathcal{M} \mathcal{I}_0)$ pairs. Top row: simulated images warped to $\mathcal{T}$; middle row: $\mathcal{M} \mathcal{I}_0$; bottom row: $\mathcal{M} \mathcal{I}_0$

Fig. 4.

Box-and-Whisker plots showing the discriminative power and registration accuracy of the output along the iterations. The top left panel shows the AUC of the abnormality map at each iteration. The top right panel show the HD between P_0 and P_1 . The bottom two panels show the relative distance Γ_n^k and Λ_k , respectively.

Fig. 5. Representative axial slices along with the corresponding abnormality maps at various iterations; from left to right: FLAIR image slices, initial abnormality maps, abnormality maps at iteration 1 and abnormality maps at iteration 2.

Fig. 6.

Box-plot showing the relative differences Γ_n^{Method} for the competing methods. SIMi-j denotes the case where lesions are simulated at Zone i with size j , where larger i signifies larger abnormalities. The Zone indexing is consistent with what is shown in Fig. 2.

Fig. 7.

Bar-plot comparing the AUC and HD between the four competing methods.

Fig. 8. Bar-plot comparing the AUC and HD between univariate statistical maps and the abnormality maps computed using the proposed method.

Fig. 9.
Bar-plot comparing the AUC and HD between AR, DR and DRI for the clinical dataset.

Fig. 10.

Visual depiction of the outputs of each method. Each row showcases axial slices from one representative subject in the clinical database. From left to right: FLAIR image slice, AR projection, DR projection, DRI projection, AR abnormality map, DR abnormality map and DRI abnormality map.

Fig. 11.

Visual depiction of the results on AD patients. Each row showcases three representative slices of the abnormality map from one subject. The abnormality maps are overlaid on the atlas. From left to right: a sagittal slice (left hemisphere), a coronal slice, and another sagittal slice (right hemisphere).