

EFFICIENT REGISTRATION OF PATHOLOGICAL IMAGES: A JOINT PCA/IMAGE-RECONSTRUCTION APPROACH

Xu Han¹ Xiao Yang¹ Stephen Aylward² Roland Kwitt³ Marc Niethammer¹

¹ The University of North Carolina at Chapel Hill, USA

²Kitware Inc., USA

³The University of Salzburg, Austria

ABSTRACT

Registration involving one or more images containing pathologies is challenging, as standard image similarity measures and spatial transforms cannot account for common changes due to pathologies. Low-rank/Sparse (LRS) decomposition removes pathologies prior to registration; however, LRS is memory-demanding and slow, which limits its use on larger data sets. Additionally, LRS blurs normal tissue regions, which may degrade registration performance. This paper proposes an efficient alternative to LRS: (1) normal tissue appearance is captured by principal component analysis (PCA) and (2) blurring is avoided by an integrated model for pathology removal and image reconstruction. Results on synthetic and BRATS 2015 data demonstrate its utility.

1. INTRODUCTION

Image registration in the presence of pathologies is challenging as standard image similarity measures (e.g., sum of squared differences, mutual information, and normalized cross-correlation (NCC)) and standard spatial transforms (e.g., B-Spline and deformation fields) do not account for common changes arising from pathologies and cannot establish reliable spatial correspondences. Pathological image registration is needed, for example, to support (a) disease diagnosis and treatment planning using atlas-based tissue segmentation to identify traumatic brain injuries, tumors, or strokes [1]; and (b) treatment monitoring using longitudinal images for brain tumor recurrence assessment [2].

A variety of approaches have been proposed to address pathological image registration. For example, cost function masking [3] and geometric metamorphosis [4] exclude pathological regions from measurements of image similarity. However, these approaches require *prior segmentations* of the pathological regions, which is non-trivial and/or labor-intensive. Joint registration and segmentation approaches have also been proposed, which include estimating a latent label field to indicate missing correspondences [5, 2].

A conceptually different approach is to *learn* normal image appearance from population data and to use it to esti-

mate a quasi-normal image from an image with pathologies. This quasi-normal image can then be used for registration. Quasi-normal images can, for example, be estimated via a low-rank/sparse (LRS) decomposition [6] or by learning a direct mapping from a pathological to a quasi-normal image via an autoencoder [7]. LRS suffers from three shortcomings: *First*, the ideal LRS decomposition is computed based on already aligned images. Hence, in practice, registration and LRS decomposition steps need to be alternated making the algorithm costly. *Second*, LRS decomposes the full population sample, causing high memory demand. Jointly with the first shortcoming, this severely limits the number of subjects that can be used for the decompositions to capture population variation. *Third*, while LRS reconstructs pathological image areas, making them appear quasi-normal, it also blurs normal tissue and hence may impair registrations in areas unaffected by the pathology. While the autoencoder approach by Yang et al. [7] does not blur normal tissue and does not require alternating registrations for a full population of images, it requires a large number of training images and has so far not been extended to 3D. *This paper proposes an approach inspired by LRS, which overcomes all three of its shortcomings.*

Contributions. *First*, we use *normal* images as our population. This is different from the original LRS framework [6] which iteratively estimates quasi-normal images from groups of pathological images (interleaved with registration to a normal atlas). Instead, we can register the normal images to the atlas only *once*. Additional registrations are performed only for the pathological image, greatly reducing computational cost. *Second*, when LRS is applied to a population of normal images and one pathological image, the most desirable decomposition would be to allocate all normal images to the low-rank part and to decompose only the pathological image into its low-rank and sparse components¹. Instead, we completely replace the LRS decomposition. Specifically, we mimic the low-rank component via a PCA basis obtained from the normal images in atlas space. We decompose the pathological image into (i) a quasi-normal part which is *close*

¹While desirable, this will not happen in practice, because part of the normal images will also be allocated to the sparse part, causing image blurring.

to the PCA space and (ii) an abnormal part which has low total variation (TV) and replaces the sparse component of the LRS decomposition. This new decomposition is highly beneficial as it avoids image blurring (by only requiring *closeness* to the PCA space) and captures large pathologies (via TV) while avoiding attributing image detail and misalignments to the pathology as in LRS. Similar to [6], our approach does not require prior knowledge of the location of the pathology.

Organization. Sec. 2 discusses the LRS registration model and our proposed approach. Sec. 3 presents experimental results on synthetic and real data. The paper concludes with a discussion in Sec. 4 and an outline of ideas for future work.

2. METHODOLOGY

Review of Low-Rank/Sparse (LRS). The standard LRS decomposition requires minimization of

$$E(L, S) = \|L\|_* + \lambda \|S\|_1, \quad \text{s.t. } D = L + S, \quad (1)$$

where D is a given data matrix, $\|\cdot\|_*$ is the nuclear norm (i.e., a convex approx. for the matrix rank), and $\lambda > 0$ weighs the contribution of the sparse part, S , in relation to the low-rank part L . In imaging applications, each image is represented as a column of D . The low-rank term then captures common information across columns. The sparse term, on the other hand, captures uncommon/unusual information. As Eq. (1) is convex, minimization results in a global optimum, e.g., computed via an augmented Lagrangian approach [8].

To use LRS for the registration of pathological images requires joint registration and decomposition, as the decomposition relies on good spatial alignment, while good spatial alignment requires a good decomposition. This can be accomplished via alternating optimization [6]. Upon convergence, the low-rank matrix contains the normal parts of all images, while the sparse matrix contains the estimated pathologies. Since LRS does not consider spatial image information, small misalignments, unavoidable in image registration, may be considered abnormal and allocated to the sparse part. Also, image details may be allocated to the sparse part and cause blurring in the estimated normal image parts. Furthermore, solving the LRS decomposition iteratively [8] requires a singular value decomposition (SVD) at each iteration with a complexity of $\mathcal{O}(\min\{mn^2, m^2n\})$ [9] for an $m \times n$ matrix. For large images $m \gg n$ and hence the computational complexity will grow quadratically with the number of images, n , making LRS costly for large sample sizes, which are beneficial to capture data variation.

Proposed PCA-based model. Our proposed model assumes that we have a collection of normal images available. In fact, our goal is to register *one* pathological image to a normal-control atlas. Hence, we can first register all the normal images to the atlas using a standard image similarity measure. These normal images do not need to be re-registered during

the iterative solution approach, resulting in a dramatic reduction of computational cost, which then allows using large image populations to capture normal data variation. Since we know *a priori* which images are normal, we can mimic the low-rank part of LRS by a PCA decomposition of the atlas-aligned normal images; we obtain PCA basis images $\{\beta_i\}$ and the mean image M . We are now only concerned with a *single* pathological image I . Let \hat{I} denote the pathological image after subtracting M , B the PCA basis matrix, and L and S are images of the same size as I^2 . Our *first model* minimizes

$$E(S, \hat{L}, \alpha) = \gamma \|\hat{L} - B\alpha\|_1 + \|\nabla S\|_{2,1}, \quad \text{s.t. } \hat{I} = \hat{L} + S \quad (2)$$

akin to the TV-L1 model [10], where $\|\nabla S\|_{2,1} = \sum_i \|\nabla S_i\|_2$ and i denotes spatial location. The *second model* minimizes

$$E(S, \hat{L}, \alpha) = \frac{\gamma}{2} \|\hat{L} - B\alpha\|_2^2 + \|\nabla S\|_{2,1}, \quad \text{s.t. } \hat{I} = \hat{L} + S \quad (3)$$

and is akin to the Rudin-Osher-Fatemi (ROF) model [11]. Both models result in a TV term, S , which explains the parts of \hat{I} which are (i) spatially contiguous, (ii) relatively large, and (iii) cannot be explained by the PCA basis, e.g., a tumor region. The quasi-low-rank part \hat{L} remains close to the PCA space but retains fine image detail. Adding M to \hat{L} results in the reconstructed quasi-normal image L . In principle, model (2) would be preferred, because of the attractive geometric scale-space properties of the TV-L1 model [10]. However, we use model (3) in our experiments as it is simpler to optimize. Unfortunately, just as the ROF model [10], it suffers from an intensity loss of the reconstruction, which we counteract by adapting the iterative regularization approach proposed by Osher et al. [12] for the ROF model, which iteratively adds “noise” back to the reconstructed images. Specifically, we first solve (3) (obtaining $\tilde{L}_0 = \hat{L}$ and α_0), followed by a small number of regularization steps. For each iteration $k \geq 1$, we minimize

$$E(S_k, \tilde{L}_k, \alpha_k) = \frac{\gamma}{2} \|\tilde{L}_k - B\alpha_k\|_2^2 + \|\nabla S_k\|_{2,1}, \quad (4)$$

$$\text{s.t. } \hat{I}_k = \tilde{L}_k + S_k,$$

where $\hat{I}_k = \hat{I} + \tilde{L}_{k-1} - B\alpha_{k-1}$. After n iterations, the TV part, S_n , will contain an approximation of the pathology, from which we obtain the quasi-normal image $\hat{L}_n = \hat{I} - S_n$. The quasi-normal image reconstructs pathological areas while retaining detailed image information in normal image areas.

Implementation details. We solve our PCA model via a primal-dual hybrid gradient method [13]. Compared to LRS *memory requirements are lower and runtime is faster*.

3. EXPERIMENTS

We use the ICBM atlas [14] as our normal atlas image.

Pseudo-tumor data (2D). We evaluate the performance of our model in 2D. We pick 250 images from the OASIS

²Images are vectorized; the spatial gradient ∇ is defined correspondingly.

cross-sectional MRI dataset [15] as the normal population. We simulate 50 test cases by picking another set of 50 OA-SIS images and registering them to the BRATS 2015 T1c images [16], followed by injecting the BRATS tumors into these warped images. The registrations simulate tumor mass effects. Each image is of size 196×232 with 1mm isotropic pixels. We select 50 fixed normal images as the population for LRS, to test a scenario which would still be computable in 3D given the high computational demand of LRS. We select 250 normal images for our PCA model and choose the top 150 PCA modes as the PCA basis. We test *without* regularization and *with* at most two regularization steps.

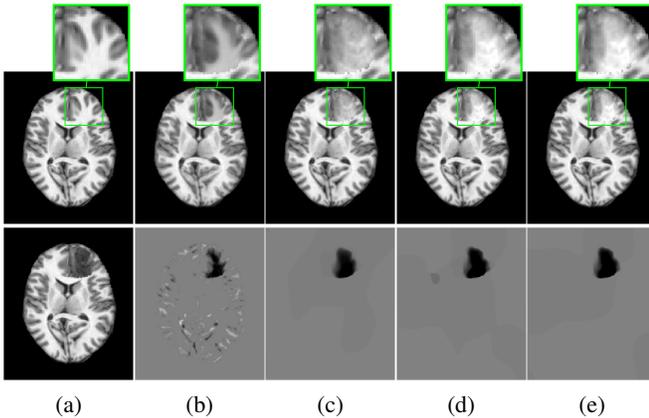


Fig. 1: Example quasi-normal reconstructions. (a) ground truth (*top*) and tumor (*bottom*); (b)-(e) Reconstruction result (*top*) and tumor (*bottom*): (b) LRS; (c) PCA model w/o regularization; (d) PCA model w/ one and (e) w/ two regularization steps.

Fig. 1 shows reconstruction results for LRS and for our PCA-based models. For each model, we perform cross-validation, partitioning the 50 test cases into 10 folds, with 9 folds for training and 1 fold for testing. We train each model with $\lambda = \{0.005, 0.0067, 0.0084, 0.01, 0.0117, 0.0133, 0.015\}$, for LRS, and $\gamma = \{0.5, 1, 1.5, 2, 2.5, 3\}$, for our PCA models. We evaluate the *mean registration error* compared to the ground truth registration result. This is done in three areas: the tumor area, the normal areas near the tumor (within 10mm) and the normal areas far from the tumor (>10 mm). We weigh the deformation errors in these areas 4:1:1 and, for each model, pick the parameter that gives the smallest errors.

Fig. 1 shows a good but blurry LRS reconstruction as the sparse part captures the tumor *and* misalignments. Also, the small and round left posterior ventricle in the ground truth image is not reconstructed faithfully by LRS. Our PCA models capture only the tumor in S , resulting in a sharper and more precise reconstruction. Furthermore, regularization yields an even better tumor separation.

Fig. 2 shows atlas-to-image registration results for images with and without tumor, LRS reconstruction and our PCA-based models with and without regularization. Fig. 3 shows the spatial error distributions, compared to the ground truth

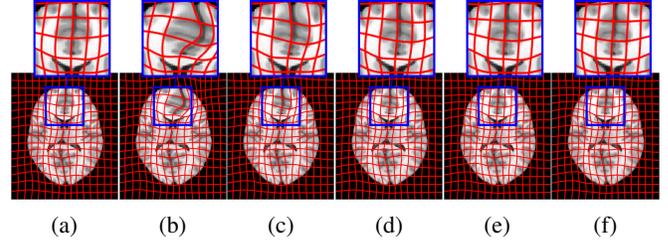


Fig. 2: Example atlas-to-image registrations: (a) ground truth; (b) tumor; (c) LRS; (d) PCA model w/o regularization; (e) PCA model w/ one step and (f) w/ two regularization steps.

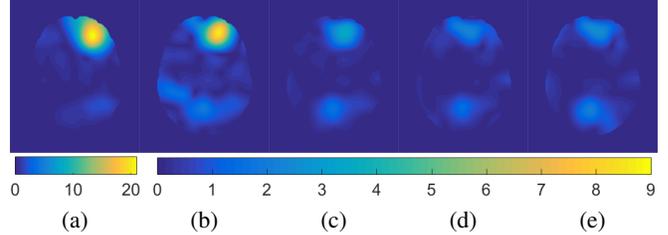


Fig. 3: Example atlas-to-image registration errors [mm]: (a) tumor; (b) LRS; (c) PCA model w/o regularization; (d) PCA model w/ one step of regularization; (e) PCA model w/ two steps of regularization.

registration. We use `NiftyReg` [17] (with standard settings) and NCC for registrations. Errors are computed using Euclidean distance. Direct registration of the tumor image results in large registration errors. Registration to the low-rank reconstruction greatly reduces the error in the tumor areas but retains errors near the cortex, mainly due to its blurry reconstruction. Our PCA models further reduce registration errors in the tumor areas *and* keep errors near the cortex low.

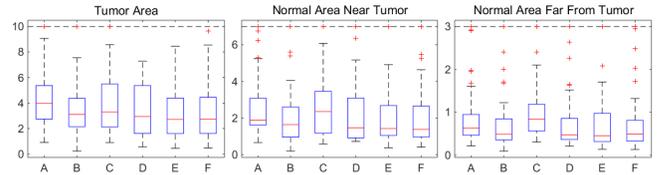


Fig. 4: Mean deformation errors [mm] for test cases. A: Tumor image; B: cost function masking; C: LRS; D: PCA model w/o regularization; E: PCA w/ one and F: w/ two regularization steps.

Fig. 4 shows mean deformation errors over all test cases in the 3 areas. We also add cost function masking for comparison. Note that the tumors selected from BRATS to generate our 2D test cases are relatively mild resulting in relatively small deformation errors even when using tumor images for registration. LRS (C) reduces errors in the tumor areas but has higher errors in the normal areas. Our PCA models (D, E, F) show better results in both the tumor and the normal areas. Paired *t*-tests between LRS and our PCA models show statistically significant differences in all areas for the PCA models with regularization, and in the normal areas for the PCA model without regularization. Moreover, the PCA models with regularization show similar performance to cost function masking but do not require a tumor segmentation.

Pseudo-tumor Dataset (3D). We also generate 3D pseudo-tumor data for evaluation. We pick 100 OASIS images and select the first 50 PCA modes as the basis. We also simulate 20 test images with tumor. Each image is of size $196 \times 232 \times 189$. Different from the 2D experiment, the tumors for our 3D test cases are picked randomly from BRATS, including cases with large tumors and deformations. For cross validation, we sep-

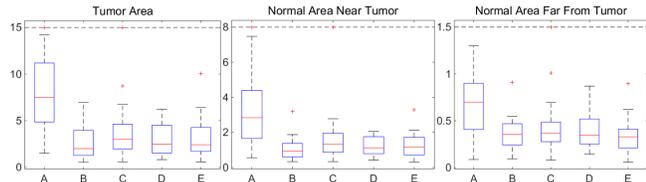


Fig. 5: Mean deformation errors [mm] for 3D test cases. A: Tumor image; B: cost function masking; C: PCA model w/o regularization; D: PCA model w/ one and E: w/ two regularization steps.

arate the twenty test cases into ten 9:1 folds. The training parameters for our models are $\gamma = \{1, 1.5, 2, 2.5, 3\}$. Registration errors in the three different areas are weighted as before, i.e., 4:1:1. Fig. 5 shows box plots of the mean deformation errors. Directly registering to tumor images results in large errors. The quasi-normal images reconstructed by our PCA models greatly reduce the deformation errors in all the areas. As in 2D, our PCA models show similar performance to cost function masking but do not require a tumor segmentation.

BRATS Dataset (3D). Finally, we also apply our model to the BRATS 2015 data [16]. As the BRATS data was acquired at different institutions and on different scanners, we pick 80 BRATS T1c images as the population which show consistent image appearance and contain the full brain. To obtain our PCA model we locally impute image intensities in the tumor areas, prior to computing the PCA basis, using the mean intensity over all images that do not contain a tumor at that location. We also pick the first 50 PCA modes as our basis.

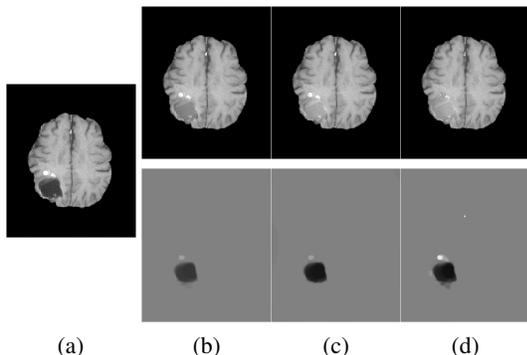


Fig. 6: Example BRATS reconstructions: (a) tumor image; (b)-(d) reconstructions (top) and tumors (bottom); (b) PCA model w/o regularization; (c) PCA model w/ one and (d) w/ two regularization steps.

Fig. 6 shows decomposition results for our PCA models. We pick $\gamma = 5.0$ for the model without and $\gamma = 2.0$ for mod-

els with regularization. The goal is to allocate as much of the tumor as possible to the abnormal part, S , while keeping the normal tissue in the quasi-normal part of the decomposition. Qualitatively, our models identify tumor/normal areas, while retaining image details in normal tissue areas. Finally, Fig. 7

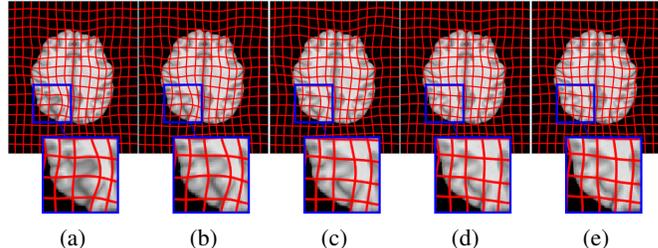


Fig. 7: Example BRATS atlas-to-image registration results: (a) tumor image; (b) cost function masking; (c) PCA model w/o regularization; (d) PCA model w/ one and (e) w/ two regularization steps.

shows atlas-to-image registration results for the PCA models, the tumor image, and cost function masking. The results show the significant impact of the tumor on the registration, which is mitigated by cost function masking and our PCA models, in particular, with regularization.

Memory use. For LRS, $D \in \mathbb{R}^{m \times n}$, where m is the number of pixels/voxels and n the number of images. Each $196 \times 232 \times 189$ 3D image (stored as double) consumes about 65MB of memory. Hence, 3GB of memory is required to store D for $n = 50$. As the LRS algorithm [8] requires storing several variables of the size of D , memory use quickly becomes prohibitive, in particular for GPU implementations. Our model only stores *one* copy of the pre-computed PCA basis thereby substantially reducing memory use ($\approx 4\text{GB}/8\text{GB}$ for $n = 50$ in single/double precision) and consequentially facilitating larger sample sizes even on the GPU.

Runtime. For the 3D cases, with $n = 50$, an LRS decomposition takes one hour to run and uses up to 40GB of memory thereby precluding a GPU implementation. Due to the low memory requirements of our PCA models, a GPU implementation is possible resulting in a runtime of ≈ 3 minutes / decomposition. The 3D image registrations are computed on the CPU (≈ 3 minutes). Therefore, with 6 registration iterations, our algorithm requires ≈ 40 minutes / test case and takes about 1 hour if extra regularization steps are computed, whereas the LRS approach takes >6 hours.

4. DISCUSSION

To conclude, our experiments show that the proposed PCA-based model (i) improves image reconstructions and registrations over the LRS model, while (ii) requiring less memory, at (iii) substantially reduced computational cost. On the tested pseudo-tumor data, our models achieve performance close to cost function masking, without requiring tumor segmentations. Future work should include a quantitative assessment of the registration results on 3D BRATS data via landmarks.

5. REFERENCES

- [1] A. Irimia, B. Wang, S. Aylward, M. Prastawa, D. Pace, G. Gerig, D. Hovda, R. Kikinis, P. Vespa, and J. Van Horn, "Neuroimaging of structural pathology and connectomics in traumatic brain injury: Toward personalized outcome prediction," *NeuroImage: Clinical*, vol. 1, no. 1, pp. 1–17, 2012.
- [2] D. Kwon, M. Niethammer, H. Akbari, M. Bilello, C. Davatzikos, and K. M. Pohl, "PORTR: Pre-operative and post-recurrence brain tumor registration," *TMI*, vol. 33, no. 3, pp. 651–667, 2014.
- [3] M. Brett, A. Leff, C. Rorden, and J. Ashburner, "Spatial normalization of brain images with focal lesions using cost function masking," *NeuroImage*, vol. 14, no. 2, pp. 486–500, 2001.
- [4] M. Niethammer, G. Hart, D. Pace, P. Vespa, A. Irimia, J. Van Horn, and S. Aylward, "Geometric metamorphosis," in *MICCAI*, 2011.
- [5] N. Chitphakdithai and J. Duncan, "Non-rigid registration with missing correspondences in preoperative and postresection brain images," in *MICCAI*, 2010.
- [6] X. Liu, M. Niethammer, R. Kwitt, N. Singh, M. McCormick, and S. Aylward, "Low-rank atlas image analyses in the presence of pathologies," *TMI*, vol. 34, no. 12, pp. 2583–2591, 2015.
- [7] X. Yang, X. Han, E. Park, S. Aylward, R. Kwitt, and M. Niethammer, "Registration of pathological images," in *MICCAI SASHIMI workshop*. Springer, 2016, pp. 97–107.
- [8] Z. Lin, M. Chen, and Y. Ma, "The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices," *arXiv preprint <https://arxiv.org/abs/1009.5055>*, 2010.
- [9] M. Holmes, A. Gray, and C. Isbell, "Fast SVD for large-scale matrices," in *Workshop on Efficient Machine Learning at NIPS*, 2007, vol. 58, pp. 249–252.
- [10] T. Chan and S. Esedoglu, "Aspects of total variation regularized L1 function approximation," *SIAM J. Appl. Math.*, vol. 65, no. 5, pp. 1817–1837, 2005.
- [11] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1, pp. 259–268, 1992.
- [12] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin, "An iterative regularization method for total variation-based image restoration," *Multiscale Model. Simul.*, vol. 4, no. 2, pp. 460–489, 2005.
- [13] T. Goldstein, M. Li, X. Yuan, E. Esser, and R. Baraniuk, "Adaptive primal-dual hybrid gradient methods for saddle-point problems," *arXiv preprint <https://arxiv.org/abs/1305.0546>*, 2013.
- [14] J. C. Mazziotta, A. W. Toga, A. Evans, P. Fox, and J. Lancaster, "A probabilistic atlas of the human brain: Theory and rationale for its development: The international consortium for brain mapping (ICBM)," *Neuroimage*, vol. 2, no. 2, pp. 89–101, 1995.
- [15] D. Marcus, A. Fotenos, J. Csernansky, J. Morris, and R. Buckner, "Open access series of imaging studies: longitudinal MRI data in nondemented and demented older adults," *J. Cognitive Neurosci.*, vol. 22, no. 12, pp. 2677–2684, 2010.
- [16] B. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, et al., "The multimodal brain tumor image segmentation benchmark (BRATS)," *TMI*, vol. 34, no. 10, pp. 1993–2024, 2015.
- [17] M. Modat, G. Ridgway, Z. Taylor, M. Lehmann, J. Barnes, D. Hawkes, N. Fox, and S. Ourselin, "Fast free-form deformation using graphics processing units," *Comput. Meth. Prog. Bio.*, vol. 98, no. 3, pp. 278–284, 2010.