SIBIA-GlS: Scalable Biophysics-Based Image Analysis for Glioma Segmentation

Andreas Mang[2], Sameer Tharakan[1], Amir Gholami[1], Naveen Himthani[1], Shashank Subramanian[1], James Levitt[1], Muneeza Azmat[1], Klaudius Scheufele[3], Miriam Mehl[3], Christos Davatzikos[4] Bill Barth[4], and George Biros[1]

[1] Institute for Computational Engineering and Science, U Texas, Austin
[2] Department of Mathematics, University of Houston
[3] Department of Computer Science, University of Stuttgart
[4] Department of Radiology, University of Pennsylvania
[5] Texas Advanced Computing Center, U Texas, Austin

1 Introduction

SIBIA-GlS is baaed on Gooya et al [8] and on SIBIA (Scalable Integrated Biophysics-based Image Analysis) [6], a set of algorithms and software for biophysics-based image analysis. The SIBIA-GlS pipeline comprises two main steps: the ***ML-step*** in which we use supervised machine learning to create probability maps for the target classes ("whole tumor", "edema", "tumor core", and "enhancing tumor"); and the ***SIBIA-step*** in which we combine these probabilities with a biophysical model of tumor growth coupled with large-scale diffeomorphic registration to implicitly impose spatial correlations Besides [8], our work borrows from [10] and the algorithms and workflows in [22] that summarizes the BRATS 2012 and 2013 competitions. In this work, we used the training, validation, and testing datasets of the BRATS 2017 competition [1–3]. The BRATS 17 competition has several metrics to assess the quality of a segmentation. In this paper we just focus on the *"whole tumor" Dice score*. We achieve 0.87 median Dice score (0.84 mean) and 0.84 median Dice score (0.83 mean) for the training and validation images respectively. The histograms are shown below.

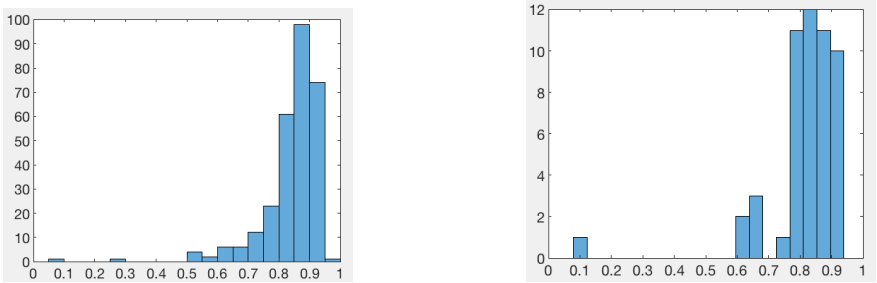

Fig. 1 *Dice histograms for "whole-tumor" label. Left: training (285 images); Right: validation (46 images)*

2 Methods

In this section, we discuss the methodology and the overall formulation for SIBIA-GlS. First, a few words regarding preprocessing. We normalize the intensities by centering their mean after removing the bottom and top 1% outliers (for each modality separately). Then, we affinely register all images to a normal (segmented) probabilistic atlas image using an L^2 similarity measure and the T1 patient image. All the classification steps take in the atlas space.

Notation: With boldface, we denote vector fields; with normal face fonts we denote scalar fields. We define the following probability fields: π_G (gray matter), π_W (white matter), π_F (cerebrospinal fluid and ventricles), π_{ED} (edema), π_{TC} (tumor core), π_{EN} (enhancing tumor), and the π_{WT} (whole tumor) probability. We also define $\boldsymbol{\pi}^A$ to be the vector probability for a reference brain (that may or may not have a tumor, depending on the context), so that $\boldsymbol{\pi}^A = \{\pi_G, \pi_W, \pi_F, \pi_{WT}\}$. We define $\boldsymbol{\pi}^S$ to be the patient vector probability map.

We use SIBIA-GlS for the "whole-tumor" label and then binary classification for the other labels. SIBIA-GlS consists of four main components. First, the ***inverse tumor growth model*** is used to biophysically constrain the ML classifier. Second, the ***registration problem*** used for atlas-based segmentation (with or whithout tumor). Third, ***supervised machine learning*** framework that provides the initial "whole-tumor" probability. Fourth, the ***overall coupling*** that combines the three first components to produce a final "whole-tumor" probability. We briefly describe these components below.

Tumor model: Given $\boldsymbol{\pi}^A(0) := \boldsymbol{\pi}^A(x, 0)$, the probability map of a healthy brain (i.e., the *"atlas"*), the *forward* tumor operator $\boldsymbol{\mathcal{T}}$ is given by

$$\boldsymbol{\pi}^A(1) := \boldsymbol{\pi}^A(\boldsymbol{x}, 1) = \boldsymbol{\mathcal{T}}(\boldsymbol{g}, \boldsymbol{\pi}^A(0)). \tag{1}$$

Here $\boldsymbol{g}$ are ***tumor growth model parameters*** that control the tumor growth. $\boldsymbol{\pi}^A(0)$ comprises π_W, π_G, and π_F, where $\boldsymbol{\pi}^A(1)$ comprises π_W, π_G, π_F, and π_{WT}. To simplify the notation, we suppress the dependence on the normal atlas $\boldsymbol{\pi}^A(0)$; we simply write $\boldsymbol{\pi}^A = \boldsymbol{\mathcal{T}}(\boldsymbol{g})$, where $\boldsymbol{g}$ is the vector of parameters that control the tumor growth dynamics. In this work we used a simple reaction-diffusion model and g is the initial condition for the tumor parameterized by 125 Gaussians.

In the ***inverse tumor problem***, given $\boldsymbol{\pi}^A_*$ (data with tumor) and $\boldsymbol{\pi}^A(0)$ (a normal brain), we solve an optimization problem for $\boldsymbol{g}$: $\min_{\boldsymbol{g}}(\boldsymbol{\pi}^A(1) - \boldsymbol{\pi}^A_*)^2$, were $\boldsymbol{\pi}^A(1)$ is given by (1). We have omitted (due to space limitations) an additional regularization term [5–7, 11, 16] that controls the reconstruction of g.

Registration: We use a velocity-based formulation for diffeomorphic registration [4, 9, 12–15]. Given a vector field $\boldsymbol{\pi}^S(0)$ and a velocity field $\boldsymbol{v}$, the *forward* image registration problem computes a deformation of $\boldsymbol{\pi}^S(0) := \boldsymbol{\pi}^S(x, 0)$, let's call it $\boldsymbol{\pi}^S(1) := \boldsymbol{\pi}^S(x, 0)$. We abstract this operation using $\boldsymbol{\mathcal{R}}$. That is,

$$\boldsymbol{\pi}^S(1) = \boldsymbol{\mathcal{R}}(\boldsymbol{v}, \boldsymbol{\pi}^S(0)). \tag{2}$$

In the ***inverse registration problem***, we're given two vector fields $\boldsymbol{\pi}^S(0)$ and $\boldsymbol{\pi}^S_*$ and we seek to compute $\boldsymbol{v}$ such that the difference between $\boldsymbol{\pi}^S(1)$ and

$\boldsymbol{\pi}_*^S$ is a small as possible. Formally, $\min_{\boldsymbol{v}}(\boldsymbol{\pi}^S(1) - \boldsymbol{\pi}_*^S)^2$ such that $\boldsymbol{\pi}^S(1)$ is given by (2). We have omitted the necessary regularization for the velocity [13].

Supervised classification: We first identify "whole-tumor" voxels using binary classification. Then, we classify the "whole-tumor" voxels to "tumor core" and "edema". Lastly, we classify the "tumor-core" voxels to "enhancing" and "non-enhancing". All these are binary classifications. As featured, we use 288 2D Gabor features per voxel, and we use 50,000,000 training voxels. For the machine learning step we use nearest neighbor classification (using an in-house code [20,21]) for the "whole tumor" label. For the binary classifications to distinguish edema, and enhancing tumor we used 25,000,000 points and used LightGBM, an open-source, random-forest classifier. Both classifiers return probability maps. The whole-tumor probability maps are passed to the next step to introduce spatial correlation. Then, the updated π_{WT}, along with the π_{TC} and π_{EN} from LightGBM are combined and threshold to produce the final labels.

Coupled formulation (SIBIA): The inputs to our problem are $\boldsymbol{\pi}^S(0)$ (initial patient probability map) and $\boldsymbol{\pi}^A(0)$ (normal atlas without tumor). The outputs are $\boldsymbol{g}$ and $\boldsymbol{v}$, and $\boldsymbol{\pi}^S(1)$ and $\boldsymbol{\pi}^A(1)$, which contains π_{WT}-the main output of the SIBIA-GlS part in the pipeline. Formally, the optimization problem (omitting, for notational simplicity, regularization terms for g and u) is given by

$$\min_{\boldsymbol{v},\,\boldsymbol{g}} \; (\boldsymbol{\pi}^A(1) - \boldsymbol{\pi}^S(1))^2 \qquad \text{such that} \quad \begin{aligned} \boldsymbol{\pi}^A(1) &= \boldsymbol{\mathcal{T}}(\boldsymbol{g}, \boldsymbol{\pi}^A(0)), \\ \boldsymbol{\pi}^S(1) &= \boldsymbol{\mathcal{R}}(\boldsymbol{v}, \boldsymbol{\pi}^S(0)). \end{aligned} \tag{3}$$

Here we assume we have probability maps $\boldsymbol{\pi}^A(0)$ and $\boldsymbol{\pi}^S(0)$ for both atlas and patient images and then we try to match them. That means that we need to segment gray matter, white matter, and CSF and ventricles, in addition to tumor probabilities. We need this information to be able to calibrate the tumor growth model through the determination of $\boldsymbol{g}$. Tumor doesn't grow in ventricles and material properties differ in white and grey matter. We obtained these tissue-type probability maps using probabilistic atlas segmentation averaging registration with 10 normal brains and two large-deformation diffeomorphic registration algorithms, DEMONS [17], and our own CLAIRE [14].

3 Results

3.1 SIBIA results

First, we report results for the SIBIA component, assuming a correct and known π_{WT} and shows how we can grow a tumor in atlas space and then map it to a given tumor in the patient space, to test our joint inversion / segmentation approach. By adjusting the regularization parameters, we can adjust the regularization provided by SIBIA. We report representative results in Figure 2. SIBIA can effectively match the given segmentation. All SIBIA runs are in reduced 128^3 resolution. The total run time for SIBIA is under 1 minute using a 10-node configuration (2-socket Xeon E5-2690 v3 (Haswell) with 12 cores/socket).

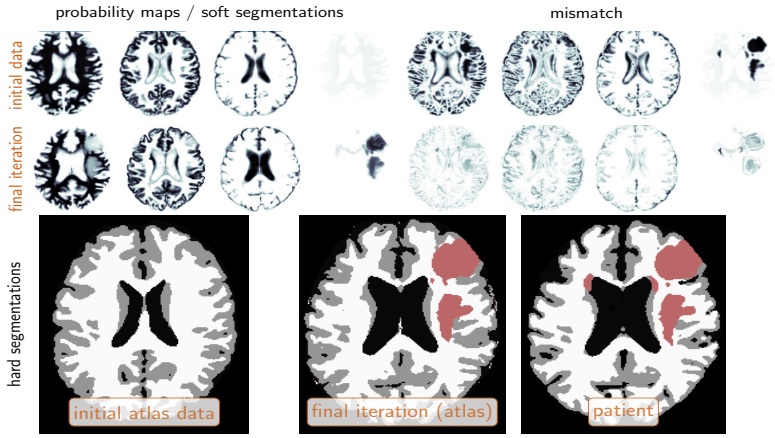

Fig. 2 *The top row shows the initial configuration (4 images to the left: probability maps* $\boldsymbol{\pi}^A(0)$ *(in particular, from left to right,* π_G*,* π_W*,* π_F *and* π_{WT}*) at iteration zero; 4 images to the right: mismatch (pointwise residual) between* $\boldsymbol{\pi}^A(1)$ *and* $\boldsymbol{\pi}^S(1)$ *at iteration zero. The second row shows the same configuration at the final iteration of our coupled tumor inversion and registration scheme. The three images on the bottom show the corresponding hard segmentation (left: initial label maps given for the atlas image; middle: deformed configuration of the atlas image (registered to the patient data) with simulated tumor; right: synthetic patient image. The obtained atlas based segmentation (middle image) and the ground truth segmentation for the patient are very similar.*

3.2 BRATS Validation results

In figure 3, we present segmentation for seven different BRATS17 cases (from the training set) for which we get different Dice scores. We can recover multifocal tumors and quite complex shapes, but we do have cases with quite bad dice scores both because of false negatives and positives. Overall, there are several technical reasons that create problems in our algorithm and we will discuss them in a longer version of this paper.

4 Discussion

We presented preliminary results for our joint formulation for combining atlas-based and machine learning-based segmentation. Below we list some observations on our efforts.

- The processing time per patient is about one hour using 10 dual-socket x86 nodes. The most expensive parts is the nearest neighbor classifier and the 20 diffeomorphic registrations for the gray matter, white matter, and CSF. Both of them are in full resolution ($204^2 \times 155$). SIBIA is much faster because we use 128^3 resolution.
- The nearest-neighbor classifier requires 8 nodes and it uses distributed memory parallelism based on MPI so that it can handle the problem of finding nearest neighbors. Notice for each brain we have to find the neighbors of 1.5M voxels in a dataset of 50M voxels. SIBIA and CLAIRE use MPI. The

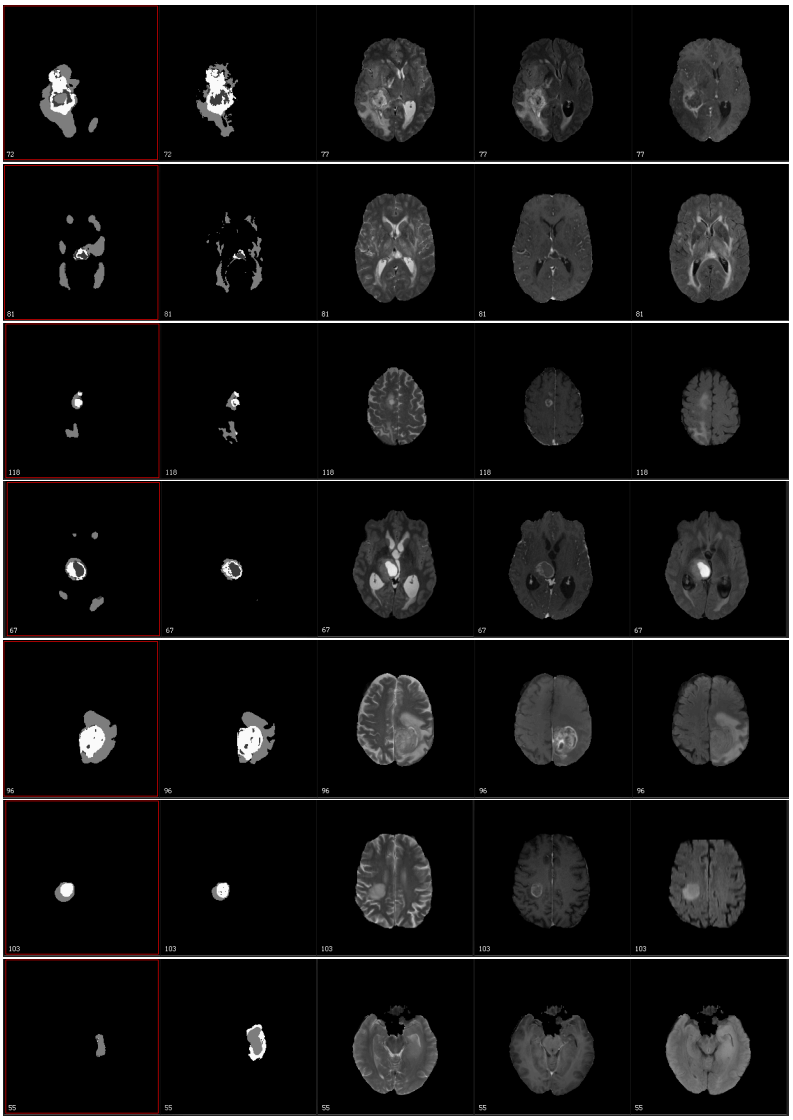

Fig. 3 *Images from the training set:* ***Top-to-bottom:*** *different BRATS brains: CBICA_ABO (0.87), CBICA_AWH (0.81), CBICA_ATB (0.81), CBICA_ATX (0.53), TCIA_242 (0.93), 2013_26 (0.53), TCIA_177 (0.23). In parenthesis, we report the "whole-tumor" Dice score.* ***Left-to-right****: In the first two columns we show our SIBIA-GlS segmentation (outlined with a red box) followed by the ground truth segmentation (provided by the BRATS17 organizers). In the last three columns, we show the T2, T1ce and FLAIR MRI images for each case. In the segmentation images (the first two columns), white is enhancing tumor, light gray is edema, and dark gray is non-enhancing tumor.*

nearest neighbors were run on the Stampede 2 system, CLAIRE and SIBIA on Lonestar 5 system, and combining the segmentations on the Maverick system, all at TACC.

- The primary classification is key for SIBIA. We need an acceptable initial guess for the segmentation based on machine learning for good results. SIBIA cannot correct really bad initial segmentations.
- Our approach is fully automatic. We only need to specify algorithm parameters (for the inverse solvers, iterative solvers, etc) . The results are not overly sensitive to these parameters. The only important parameter is the thresholding parameter for π_{WT}, which is used to produce a hard segmentation. Although 0.5 seems obvious, using 0.3 produces much better results (meaning that a voxel is labeled as "whole-tumor" if $\pi_{WT} > 0.3$).
- Our approach integrates biophysical simulations with machine learning, optimization, and image analysis. It can provide bio-physical parameters (growth rate, mass effect, and others) that might be critical for clinical studies to assess the current and/or future state of an individual patient.
- The solvers for the individual building blocks are based on state-of-the-art technology in scientific computing [6, 12, 14, 18–20].
- Unlike [8], we do not iterate to update the SIBIA segmentation. That's a limitation of our scheme and will address in future work.
- High-grade and low-grade gliomas have different characteristics but we have not taken this into account. We simple merged all the training data.

5 Conclusion

We have presented preliminary results for SIBIA-GlS—a new framework for biophysics-based image analysis for glioma segmentation. We demonstrated that our approach yields promising results. However, several issues remain open. We obtain excellent results for SIBIA-GlS if the initial proposal for the segmentation does not contain significant noise. Improving on this initial segmentation will be key for our future work. In addition to that we will extend our biophysical model to, e.g., include edema and mass effect. We expect that this will significantly improve our current results.

Acknowledgments

This material is based upon work supported by NSF grant CCF-1337393; by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under Award Numbers DE-SC0010518 and DE-SC0009286; and by NIH grant 10042242. Any opinions, findings, and conclusions or recommendations expressed herein are those of the authors and do not necessarily reflect the views of the DOE, NIH, and NSF. Computing time on the Texas Advanced Computing Centers Stampede system was provided by an allocation from TACC and the NSF. We're particularly grateful to the stuff of TACC for providing the resources that enabled us to participate in this competition.

References

1. Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J., Freymann, J., Farahani, K., Davatzikos, C.: Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. Nature Scientific Data (2017), in press
2. Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J., Freymann, J., Farahani, K., Davatzikos, C.: Segmentation labels for the pre-operative scans of the TCGA-GBM collection (2017), http://doi.org/10.7937/k9/tcia.2017.klxwjj1q
3. Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J., Freymann, J., Farahani, K., Davatzikos, C.: Segmentation labels for the pre-operative scans of the TCGA-LGG collection (2017), http://doi.org/10.7937/k9/tcia.2017.gjq7r0ef
4. Chen, K., Lorenz, D.A.: Image sequence interpolation using optimal control. Journal of Mathematical Imaging and Vision 41, 222–238 (2011)
5. Gholami, A., Mang, A., Biros, G.: An inverse problem formulation for parameter estimation of a reaction-diffusion model of low grade gliomas. Journal of Mathematical Biology 72(1), 409–433 (2016)
6. Gholami, A., Mang, A., Scheufele, K., Davatzikos, C., Mehl, M., Biros, G.: A framework for scalable biophysics-based image analysis. In: Proc ACM/IEEE Conference on Supercomputing (2017), (accepted)
7. Gholami, A., Scheufele, K., Davatzikos, C., Mang, A., Mehl, M., Biros, G.: A framework for scalable biophysics-based image analysis. In: Proceedings of SC17. The SCxy Conference series, ACM/IEEE, Denver, Colorado (November 2017)
8. Gooya, A., Pohl, K.M., Bilello, M., Cirillo, L., Biros, G., Melhem, E.R., Davatzikos, C.: GLISTR: Glioma image segmentation and registration. Medical Imaging, IEEE Transactions on 31(10), 1941–1954 (2013)
9. Hart, G.L., Zach, C., Niethammer, M.: An optimal control approach for deformable registration. In: Proc IEEE Conference on Computer Vision and Pattern Recognition. pp. 9–16 (2009)
10. Hogea, C., Davatzikos, C., Biros, G.: Brain-tumor interaction biophysical models for medical image registration. SIAM Journal on Scientific Computing 30(6), 3050–3072 (2008)
11. Hogea, C., Davatzikos, C., Biros, G.: An image-driven parameter estimation problem for a reaction-diffusion glioma growth model with mass effects. Journal of Mathematical Biology 56(6), 793–825 (2008)
12. Mang, A., Biros, G.: An inexact Newton–Krylov algorithm for constrained diffeomorphic image registration. SIAM Journal on Imaging Sciences 8(2), 1030–1069 (2015)
13. Mang, A., Biros, G.: Constrained H^1-regularization schemes for diffeomorphic image registration. SIAM Journal on Imaging Sciences 9(3), 1154–1194 (2016)
14. Mang, A., Gholami, A., Biros, G.: Distributed-memory large-deformation diffeomorphic 3D image registration. In: Proc ACM/IEEE Conference on Supercomputing. No. 72 (2016)
15. Mang, A., Ruthotto, L.: A Lagrangian Gauss–Newton–Krylov solver for mass- and intensity- preserving diffeomorphic image registration. SIAM Journal on Scientific Computing (2017), in press
16. Mang, A., Toma, A., Schuetz, T.A., Becker, S., Eckey, T., Mohr, C., Petersen, D., Buzug, T.M.: Biophysical modeling of brain tumor progression: from unconditionally stable explicit time integration to an inverse problem with parabolic PDE constraints for model calibration. Medical Physics 39(7), 4444–4459 (2012)

17. Mansi, T., Pennec, X., Sermesant, M., Delingette, H., Ayache, N.: ilogdemons: A demons-based registration algorithm for tracking incompressible elastic biological tissues. International Journal of Computer Vision 92(1), 92–111 (2011)
18. March, W.B., Xiao, B., Biros, G.: Askit: Approximate skeletonization kernel-independent treecode in high dimensions. SIAM Journal on Scientific Computing 37(2), A1089–A1110 (2015)
19. March, W.B., Xiao, B., Yu, C., Biros, G.: An algebraic parallel treecode in arbitrary dimensions. In: Proceedings of IPDPS 2015. 29th IEEE International Parallel and Distributed Computing Symposium, Hyderabad, India (May 2015), http://dx.doi.org/10.1109/IPDPS.2015.86
20. March, W.B., Xiao, B., Yu, C.D., Biros, G.: Askit: An efficient, parallel library for high-dimensional kernel summations. SIAM Journal on Scientific Computing 38(5), S720–S749 (2016), http://dx.doi.org/10.1137/15M1026468
21. March, W.B., Yu, C., Xiao, B., Biros, G.: LIBASKIT home page (2015), http://padas.ices.utexas.edu/libaskit
22. Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., Weber, M.A., Arbel, T., Avants, B.B., Ayache, N., Buendia, P., Collins, D.L., Cordier, N., Corso, J.J., Criminisi, A., Das, T., Delingette, H., Demiralp, Ç., Durst, C.R., Dojat, M., Doyle, S., Festa, J., Forbes, F., Geremia, E., Glocker, B., Golland, P., Guo, X., Hamamci, A., Iftekharuddin, K.M., Jena, R., John, N.M., Konukoglu, E., Lashkari, D., Mariz, J.A., Meier, R., Pereira, S., Precup, D., Price, S.J., Raviv, T.R., Reza, S.M.S., Ryan, M., Sarikaya, D., Schwartz, L., Shin, H.C., Shotton, J., Silva, C.A., Sousa, N., Subbanna, N.K., Szekely, G., Taylor, T.J., Thomas, O.M., Tustison, N.J., Unal, G., Vasseur, F., Wintermark, M., Ye, D.H., Zhao, L., Zhao, B., Zikic, D., Prastawa, M., Reyes, M., Leemput, K.V.: The multimodal brain tumor image segmentation benchmark (BRATS). Medical Imaging, IEEE Transactions on 34(10), 1993–2024 (2015)