2021 43rd Annual International Conference of the
IEEE Engineering in Medicine & Biology Society (EMBC)
Oct 31 - Nov 4, 2021. Virtual Conference

Improving Localization of Brain Tumors through 3D GAN Inpainting

Leon Weninger[1], Andre Gilerson[1] and Dorit Merhof[1]

Abstract— **For survival prediction of brain tumor patients based on MRI scans, radiomic features have been a major research focus in the last years. However, radiomic features do not take the location of the lesion into account, which, in relation to the functional regions of the brain, could be a significant factor in predicting survival. An automatic and exact localization of the tumor in relation to specific functional areas is not straightforward, as typical brain parcellation methods fail in presence of large lesions. Here, we propose a model that replaces the tumorous region in 3D brain MRI scans with healthy tissue in order to improve the registration process towards a brain template. Further, we assemble a set of features for quantitative description of brain tumor location. On an openly available dataset, registration is strongly improved. The extracted location features also have better predictive performance when used after the proposed registration step and reach accuracies in survival prediction comparable to radiomic features.**

Clinical relevance— **This work improves the quantification of the location of brain tumors in the human brain and proposes an extension of radiomic features to include the location, resulting in a refined prediction of patient survival.**

I. INTRODUCTION

In recent years, algorithms have improved radiological diagnosis of brain tumors by automatically assessing the risks posed by the tumor. These algorithms mostly employ radiomic features or use convolutional neural networks (CNNs). However, tumor location is largely ignored by these approaches, as radiomic features are location-independent, and CNNs focus more on texture than on location.

This focus on Radiomics or CNNs can also be seen in the Multimodal Brain Tumor Segmentation (BraTS) challenge, a public challenge which allows to perform an unbiased comparison of different brain tumor segmentation and survival prediction algorithms [1], [2]. The top competitors in survival prediction for both 2018 and 2019 focus on different combinations of mostly phenotypic features. However, although the location of the tumor plays an important role in survival [3], a rigorous automated approach in quantitative location estimation is still lacking.

We propose a set of multiple location features to uniquely define tumor location in the brain based on tumor segmentation maps and an anatomical atlas. Similar to radiomic features, these location features can then be used to assess the impact of the lesion with machine learning models.

[1]Imaging and Computer Vision, RWTH Aachen University, Germany leon.weninger@lfb.rwth-aachen.de

*This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – grant 2699533 72/GRK2150 and the DFG grant ME 3737/19-1.

One problem of an automated location quantification is that typical image registration tools are perturbed by brain lesions, inducing noise in the location estimation. Thus, before registration, the tumorous tissue should be replaced by matching healthy tissue. Previous publications on brain tumor image inpainting can be found for tumor segmentation [4] or detection [5], as well as for image perturbation removal [6], and were mostly applied to 2D images. Here, we present a novel model that replaces tumorous tissue by healthy tissue in a 3D MRI scan.

The performance of the tumor-replacing model was first evaluated on a synthetic dataset. The obtained model was then applied to the BraTS challenge dataset, where the proposed location features were subsequently extracted and used to predict patient survival. The lesion inpainting improves registration of brain tumor MRI scans, and the obtained location features have a predictive power similar to top-performing previous approaches for patient survival prediction.

II. MATERIALS AND METHODS

A. Data

Two different datasets were employed, a brain tumor dataset, comprising data from the BraTS 2019 dataset, as well as an artificially created brain tumor dataset based on MRI scans from healthy subjects. All data used in this work were publicly available, thus no separate ethics approval was necessary. All analyses were conducted in accordance with the standards of Good Clinical Practice and the Declaration of Helsinki.

The BraTS training dataset contained T1, T1 contrast-enhanced, T2, and FLAIR MRI scans, groundtruth segmentation labels, the patient age, and partly tumor grade, resection status and survival information. Here, only the datasets comprising survival information were utilized, reducing the number of datasets to 208 patients. The data was collected from multiple institutions, and the scans are of variable quality and resolution [7]. The three segmentation labels comprised contrast-enhancing tissue, the non-enhancing tumor core, as well as peritumoral edema.

Healthy MRI scans were obtained from the HCP [8] and the IXI dataset [9]. The HCP data contained 1113 T1 weighted MRI scans captured on Connectome Skyra MRI scanners. The IXI dataset contained 581 T1 weighted MRI Scans captured in 3 different London hospitals using 3T and 1.5T scanner systems. In total, this joint dataset was composed 1694 3D MRI scans of healthy patients.

Before usage, a bounding box was applied to every brain scan, discarding all empty planes. The scans were further

downsampled to a resolution of $128 \times 128 \times 128$ and $64 \times 64 \times 64$ using trilinear interpolation. A histogram equalization was applied, and the value range of the brain region mapped to [0,1].

For the artificial brain tumor dataset, malign areas were synthetically inserted into the scans of healthy subject by extracting exemplary tumors from the complete BraTS dataset, i.e., also from patients with missing survival information. This process is visualized in Fig. 1. Next to inserting artificial tumors in the healthy subjects, a second synthetic dataset was created, where designated tumorous areas were filled with zeros instead of tumorous tissue.

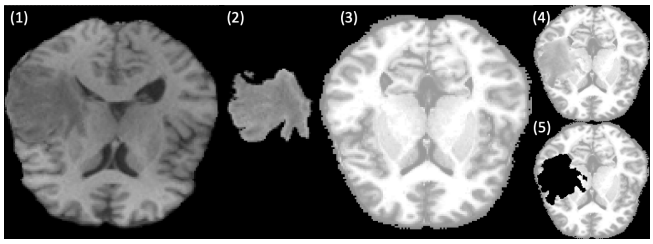

Fig. 1. Creation of artificial brain tumor dataset: Brain tumor patient (1), extracted tumor (2), healthy subject (3), same subject with artificially inserted tumor (4), or designated tumor area replaced with zeros (5).

B. Tumor inpainting

The task of the employed neural network was to replace the tumor data with healthy-looking tissue that is consistent with, and therefore based on, the rest of the brain. For this, the network must find an internal encoding representing the healthy brain state while discarding noise as well as the tumor data. By choosing an encoding smaller than the input, this behavior is forced during network training. Since the input and output will be of the same size, commonly chosen architectures are based on either variational autoencoders or U-Nets [10]. Variational autoencoders are effective in unsupervised tumor segmentation. However, multiple papers suggest that their output is blurry in actual image inpainting tasks [11]. Therefore, a state-of-the-art variation of the U-Net was chosen, the Attention U-Net (AUNET) [12]. The architecture of the AUNET closely follows the architecture of the U-Net, except for additional attention gates at each skip connection. These attention gates are designed to allow the network to prune undesirable feature responses, forcing the network to disregard information obtained from the tumor region. The utilized architecture of the attention gate is an adaption from [13] to volumetric inputs. In Fig. 2, the architecture of the net is shown. The initial feature map size was set to 8, and for every downsampling step the number of feature maps doubles. Training was performed on a Nvidia RTX 2080 Ti GPU with a batch size of 4 and an Adam optimizer with default parameters.

Following the idea of enantiomorphic replacement of tumor tissue [14], it was evaluated whether a second input image, containing the normal input image flipped at the sagittal plane, would improve the results. As brains have a symmetry axis at the sagittal plane, it was hypothesized that this information could help the neural network, as exemplary tissue would be at the correct location.

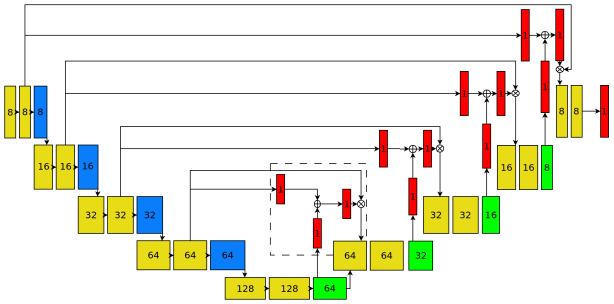

Fig. 2. The AUnet architecture. Yellow blocks designate $3 \times 3 \times 3$ convolutions followed by instance norm, blue blocks convolutions with stride 2 for downsampling, green blocks $2 \times 2 \times 2$ transposed convolutions for upsampling, and red blocks $1 \times 1 \times 1$ convolutions. The number inside each block is the number of feature maps. All blocks use LeakyReLU activation functions. The striped rectangle marks the first attention gate.

For the loss function, different combinations were considered: A simple L1 loss, the structural similarity loss (SSIM) [15], a combination of the two, as well as a generative adversarial network (GAN) discriminator loss in combination with the L1 loss. For the GAN loss, the discriminator is based on a simple image classification architecture. Local information is captured in progressively downsampled feature maps using $3 \times 3 \times 3$ convolutions with stride 2. Each convolution is followed by batch normalization and a LeakyReLU activation function. After 4 downsampling steps, the final feature map is fed into a fully connected layer with 64 hidden neurons and LeakyReLU activation, which in turn leads into a fully connected output layer with a single output neuron. The discriminator was trained using either a final sigmoid activation function in combination with a binary cross entropy loss, or with the approximated Wasserstein distance, further called Wasserstein GAN (WGAN).

C. Location feature extraction

The individual scans of the BraTS dataset, with and without previous tumor inpainting, were registered to the MNI-ICBM 152 6th generation template [16] using the symmetric diffeomorphic registration of the ANTs toolbox [17] with default parameters in order to use the MNI structural atlas [18], which contained 9 different regions. Relatively to the atlas regions, the locations of the 3 tumor regions were quantified using the following features: *One-hot*, the binary location of the center-of-mass of the tumor region in question in an atlas region, *Multi-hot*, a binary encoding whether the tumor region overlaps with a certain atlas region, *Center-distance*, the distance between the center-of-mass of the tumor region and the atlas region, *Covered-volume-ratio*, the percentage of template area covered by tumor volume, the *Hausdorff-distance*, the *Minimal-distance* between any two voxels of tumor region and atlas region in question, and finally the *Mean-min-distance*, the mean of the minimal distances of each voxel in the tumor region to the specific atlas region. All location features were extracted for the three

tumor classes and the nine atlas regions, resulting in a total number of 189 features.

D. Survival prediction

First, the extracted 189 features were reduced using the variance inflation factor (VIF) method [19], with the recommended default VIF threshold of 10. The remaining features were regressed against the patient survival in days, using a simple linear regressor as well as Lasso and Ridge regression. For Lasso and Ridge regression, the regularization parameters were set to default.

III. RESULTS

As groundtruth data for registration of the patient dataset does not exist, we relied on two indirect measurements of success: First, the registration error on the artificial dataset, using the average L1 error of the warp field. Second, the prediction error of the location features with and without tumor inpainting.

A. Registration error

10% of the artificial dataset with known groundtruth was set apart as test set, containing only healthy brains and extracted tumor areas that were not present in the training dataset. As groundtruth, the warp fields of the registrations of the healthy brain scans using the symmetric diffeomorphic registration of the ANTs toolbox were considered. A baseline was established by taking the mean registration error of the scans with artificially inserted tumor before inpainting. Using the same registration algorithm, the warp field error of the reconstructed brains was evaluated for neural networks trained on either the dataset with tumorous tissue or on the dataset with left-out areas. The warp field errors were evaluated for a low-resolution ($64 \times 64 \times 64$), as well as for a higher-resolution ($128 \times 128 \times 128$) setting. The differences between baseline and registration after inpainting using different inputs and loss functions for the low-resolution case can be seen in Table I.

TABLE I

BEST WARP ERRORS AFTER 30 EPOCHS FOR A $64 \times 64 \times 64$ INPUT SIZE, BY LOSS FUNCTION AND INPUT TYPE. FLIP: THE INPUT IMAGE FLIPPED AT THE SAGITTAL PLANE IS ADDED AS A SECOND INPUT. HIGHLIGHTED IN RED ARE WARP ERRORS <0.2, HIGHLIGHTED AS BOLD IS THE OVERALL BEST.

	Loss function				
Input	L1	SSIM	L1-SSIM	GAN	WGAN
Tumor tissue	0.432	0.281	0.219	0.247	0.168
Tumor + flip	0.409	0.266	0.214	0.230	**0.161**
Zeros	0.468	0.205	0.165	0.174	0.187
Zeros + flip	0.516	0.249	0.178	0.168	0.167
Baseline	0.693				

The best results were achieved using the combined L1 and SSIM loss or the WGAN loss function. Thus, these two loss functions were further evaluated on the higher-resolution data. Between replacement with tumor-tissue or zeros, no clear performance difference could be established. Both approaches were further evaluated at the higher-resolution (Table II). The lowest error was achieved using the combined L1-SSIM loss on the dataset where tumor areas were cut out. Exemplary output for the best performing network can be seen in Fig. 3.

TABLE II

BEST WARP ERRORS AFTER 30 EPOCHS WITH $128 \times 128 \times 128$ INPUT SIZE, BY LOSS FUNCTION AND INPUT TYPE. FLIP.: THE INPUT IMAGE FLIPPED AT THE SAGITTAL PLANE IS ADDED AS A SECOND INPUT. THE LOWEST ERROR IS HIGHLIGHTED IN BOLD.

	Loss function	
Input	L1-SSIM	WGAN
Tumor tissue	0.142	0.370
Tumor + flip	0.146	0.382
Zeros	**0.135**	0.368
Zeros + flip	0.140	0.382
Baseline	0.989	

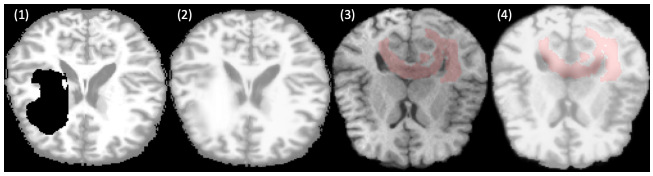

Fig. 3. Examples for the output generated by the best performing AUNET: Input (1) and output (2) image of the synthetic validation dataset, and input (3) and output image (4) of an exemplary scan of the BraTS dataset, with tumorous area marked in red.

B. Location features

The primary goal behind training the tumor inpainting networks was to improve the registration for the extraction of location features. Thus, the best performing model was used to inpaint the tumors of the BraTS challenge training set. Location features were extracted before and after inpainting, and their predictiveness is shown in Fig. 4. The evaluation was performed using 20 iterations of an 80-20 shuffle train-test split, using the negative median absolute error. As baseline benchmark, previously published methods for the BraTS challenge could be employed. In 2018, Feng et al. [20], won the first place in the BraTS survival prediction task, and the code of the implementation was readily available. We thus used it to compare the performance of our proposed location-based survival prediction. Using the same data splits for evaluation, the approach of Feng et al. achieved an error of -183, which is further used as baseline benchmark.

Survival prediction using location features was improved using the tumor inpainting. With inpainting, it even achieves results comparable to top-performing BraTS submissions, without taking tumor texture or the age of the patient into account.

IV. DISCUSSION

An integral part of location feature extraction is the registration of the brain scan to the target atlas. The regis-

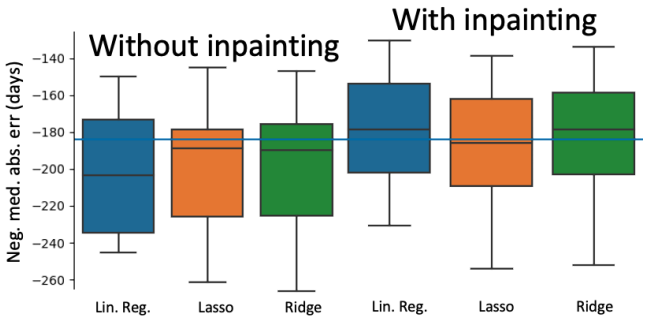

Fig. 4. Negative median absolute error (higher is better) for survival prediction using the proposed location features before and after tumor inpainting. Lin. Reg.: Linear regression, Lasso: Lasso regression, Ridge: Ridge regression

tration algorithm matches local brain structures from the individual subject to the target. We showed that using a deep learning image inpainting algorithm replacing voxels containing tumor tissue with plausible healthy tissue improves the registration process. To that end, a state-of-the-art architecture and different loss functions were explored. Following the idea of enantiomorphic filling, it was further assessed if adding a second input channel with the input image mirrored around the sagittal plane would improve results. However, in general, the network seemed to be able to infer sufficient information without the flipped channel. A surprising result was that the GAN based losses achieved similar validation errors during the training using $64\times64\times64$ images, but underperformed when using higher-resolution $128 \times 128 \times 128$ images. Nonetheless, this shows that GAN based approaches are suitable for 3D tumor inpainting and need to be explored further.

The best performing model was then used to improve the registration process for the scans in the BraTS challenge, and thus extract less noisy location information of the tumors. There was a clear improvement in the survival prediction using the proposed location features extracted from the inpainted images compared to original images, achieving results similar to current top-performing algorithms of the BraTS challenge.

One limitation of our methodology is the necessity of the tumor segmentation map for survival prediction. These segmentation maps either need to be manually annotated or inferred using a segmentation algorithm. Furthermore, it is unclear how stable the location features will be when used in a clinical pipeline, where segmentation errors of the tumor will lead to errors in localization.

V. CONCLUSION

An integral part of location feature extraction is the registration of the brain scan to the target atlas. We showed that a deep learning approach to inpainting improves the registration of lesioned brains to an atlas. Subsequently, extracted location features were better suited for survival prediction of brain tumor patients, and similarly predictive as radiomic features. However, it is still necessary to further explore tumor location for survival prediction, e.g., by including location features in radiomic approaches.

REFERENCES

[1] Spyridon Bakas, Mauricio Reyes, et Int, and Bjoern Menze, "Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the BRATS challenge," *CoRR*, vol. arXiv:1811.02629, 2018.

[2] Bjoern Menze, Andras Jakab, Stefan Bauer, et al., "The multimodal brain tumor image segmentation benchmark (BRATS)," *IEEE Transactions on Medical Imaging*, vol. 34, no. 10, pp. 1993–2024, 2015.

[3] Mohammad Sami Walid, "Prognostic factors for long-term survival after glioblastoma," *The Permanente journal*, vol. 12, no. 4, pp. 45–48, 2008.

[4] Christoph Baur, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab, "Fusing unsupervised and supervised deep learning for white matter lesion segmentation," in *Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning*, 2019, vol. 102, pp. 63–72.

[5] Changhee Han, Leonardo Rundo, Ryosuke Araki et al., "Infinite brain MR images: Pggan-based data augmentation for tumor detection," *CoRR*, vol. arXiv:1903.12564, 2019.

[6] Karim Armanious, Youssef Mecky, Sergios Gatidis, and Bin Yang, "Adversarial inpainting of medical image modalities," *CoRR*, vol. arXiv:1810.06621, 2018.

[7] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, et al., "Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features," *Nature Scientific Data*, vol. 4, no. 1, 2017.

[8] D.C. Van Essen, K. Ugurbil, E. Auerbach, et al., "The human connectome project: A data acquisition perspective," *NeuroImage*, vol. 62, no. 4, pp. 2222–2231, 2012.

[9] Biomedical Image Analysis Group, Imperial College London, "IXI dataset," http://brain-development.org/ixi-dataset/, 2020.

[10] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, 2015, pp. 234–241.

[11] Tony F. Chan and Jianhong (Jackie) Shen, "Variational image inpainting," *Communications on Pure and Applied Mathematics*, vol. 58, no. 5, pp. 579–619, 2005.

[12] Saumya Jetley, Nicholas A. Lord, Namhoon Lee, and Philip H. S. Torr, "Learn to pay attention," *CoRR*, vol. arXiv:1804.02391, 2018.

[13] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, et al., "Attention u-net: Learning where to look for the pancreas," *CoRR*, vol. arXiv:1804.03999v3, 2018.

[14] Parashkev Nachev, Elizabeth Coulthard, H. Rolf Jäger, Christopher Kennard, and Masud Husain, "Enantiomorphic normalization of focally lesioned brains," *NeuroImage*, vol. 39, no. 3, pp. 1215–1226, 2008.

[15] Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, apr 2004.

[16] Günther Grabner, Andrew L. Janke, Marc M. Budge, David Smith, Jens Pruessner, and D. Louis Collins, "Symmetric atlasing and model based segmentation: An application to the hippocampus in older adults," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*, pp. 58–66. 2006.

[17] Brian B. Avants, Nicholas J. Tustison, Michael Stauffer, Gang Song, Baohua Wu, and James C. Gee, "The insight ToolKit image registration framework," *Frontiers in Neuroinformatics*, vol. 8, apr 2014.

[18] John Mazziotta, Arthur Toga, Alan Evans, et al., "A probabilistic atlas and reference system for the human brain: International consortium for brain mapping (icbm)," *Philos Trans R Soc Lond B Biol Sci*, vol. 356, no. 1412, pp. 1293–1322, 2001.

[19] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani, *An Introduction to Statistical Learning: With Applications in R*, Springer Publishing Company, Incorporated, 2014.

[20] Xue Feng, Nicholas J. Tustison, Sohil H. Patel, and Craig H. Meyer, "Brain tumor segmentation using an ensemble of 3d u-nets and overall survival prediction using radiomic features," *Frontiers in Computational Neuroscience*, vol. 14, pp. 25, 2020.