

# Are we using appropriate segmentation metrics? Identifying correlates of human expert perception for CNN training beyond rolling the DICE coefficient

Florian Kofler<sup>1,2,3,12</sup>, Ivan Ezhov<sup>1,2</sup>, Fabian Isensee<sup>4,5</sup>, Fabian Balsiger<sup>6</sup>, Christoph Berger<sup>1</sup>, Maximilian Koerner<sup>1</sup>, Beatrice Demiray<sup>7</sup>, Julia Rackerseder<sup>7,8</sup>, Johannes Paetzold<sup>1,9,10</sup>, Hongwei Li<sup>1,11</sup>, Suprosanna Shit<sup>1,2</sup>, Richard McKinley<sup>6</sup>, Marie Piraud<sup>12</sup>, Spyridon Bakas<sup>13,14,15</sup>, Claus Zimmer<sup>3</sup>, Nassir Navab<sup>7</sup>, Jan Kirschke<sup>3</sup>, Benedikt Wiestler<sup>2,3,16</sup>, Bjoern Menze<sup>1,11,16</sup>

<sup>1</sup>*Department of Informatics, Technical University Munich, Germany*

<sup>2</sup>*TranslaTUM - Central Institute for Translational Cancer Research, Technical University of Munich, Germany*

<sup>3</sup>*Department of Diagnostic and Interventional Neuroradiology, School of Medicine, Klinikum rechts der Isar, Technical University of Munich, Germany*

<sup>4</sup>*Applied Computer Vision Lab, Helmholtz Imaging, Germany*

<sup>5</sup>*Division of Medical Image Computing, German Cancer Research Center (DKFZ), Germany*

<sup>6</sup>*Support Center for Advanced Neuroimaging (SCAN), Institute for Diagnostic and Interventional Neuroradiology, Inselspital, Bern University Hospital, University of Bern, Bern, Switzerland*

<sup>7</sup>*Computer Aided Medical Procedures (CAMP), Technical University of Munich, Germany*

<sup>8</sup>*ImFusion GmbH, Munich, Germany*

<sup>9</sup>*Helmholtz Zentrum München, Germany*

<sup>10</sup>*Imperial College London, Exhibition Rd, South Kensington, London SW7 2BX, United Kingdom*

<sup>11</sup>*Department of Quantitative Biomedicine, University of Zurich, Switzerland*

<sup>12</sup>*Helmholtz AI, Helmholtz Zentrum München, Germany*

<sup>13</sup>*Center for Biomedical Image Computing and Analytics (CBICA), University of Pennsylvania, Philadelphia, Pennsylvania, USA*

<sup>14</sup>*Department of Pathology and Laboratory Medicine, Perelman School of Medicine, University of Pennsylvania, Philadelphia, Pennsylvania, USA*

<sup>15</sup>*Department of Radiology, Perelman School of Medicine, University of Pennsylvania, Philadelphia, Pennsylvania, USA*

<sup>16</sup>*contributed equally as senior authors*

## Abstract

Metrics optimized in complex machine learning tasks are often selected in an ad-hoc manner. It is unknown how they align with human expert perception. We explore the correlations between established quantitative segmentation quality metrics and qualitative evaluations by professionally trained human raters. Therefore, we conduct psychophysical experiments for two complex biomedical semantic segmentation problems. We discover that current standard metrics and loss functions correlate only moderately with the segmentation quality assessment of experts. Importantly, this effect is particularly pronounced for clinically relevant structures, such as the enhancing tumor compartment of glioma in brain magnetic resonance and grey matter in ultrasound imaging. It is often unclear how to optimize abstract metrics, such as human expert perception, in convolutional neural network (CNN) training. To cope with this challenge, we propose a novel strategy employing techniques of classical statistics to create complementary compound loss functions to better approximate human expert perception. Across all rating experiments, human experts consistently scored computer-generated segmentations better than the human-curated reference labels. Our results, therefore, strongly question many current practices in medical image segmentation and provide meaningful cues for future research.

**Keywords:** machine learning, deep learning, interpretation, metrics, segmentation, glioma, MR, biomedical image analysis

## 1. Introduction

In the medical imaging community, challenges have become a prominent forum to benchmark the latest methodological advances for complex machine learning problems (Maier-Hein et al., 2018). Across many semantic segmentation challenges (Menze et al., 2014; Sekuboyina et al., 2020; Bilic et al., 2019; Bakas et al., 2018), convolutional neural networks (CNNs), and in particular the U-Net architecture (Ronneberger et al., 2015), gained increasing popularity over the years. Most challenges rely upon a seemingly ad-hoc combination of Dice coefficient (DICE) with other metrics for scoring (Menze et al., 2014; Maier-Hein et al., 2018).

How these metrics and the resulting rankings reflect clinical relevance and expert opinion regarding segmentation quality is poorly understood. For example, segmentation models in the BraTS challenge (Bakas et al., 2018) are evaluated on three label channels: *enhancing tumor (ET)*, *tumor core (TC)* and *whole tumor (WT)* compartments. Volumetric Dice coefficients and Hausdorff distances are aggregated equally across the three challenges to obtain the final challenge ranking. In contrast, the ET channel is of higher importance from a medical perspective. An increase in ET volume defines tumor progression (Wen et al., 2010), and glioma surgery aims for a gross total resection of the ET (Weller et al., 2014).

Taha and Hanbury (2015) investigated how different metrics capture certain aspects of segmentation quality. However, there is a clear gap in knowledge of approximating expert assessment using established metrics and achieving a clinically meaningful representation of segmentation performance. This problem is of central importance when defining loss functions for CNN training; consequently, the optimized metrics are often selected in an ad-hoc fashion. In contrast to a plethora of volumetric loss functions (Hashemi et al., 2018; Milletari et al., 2016; Sudre et al., 2017; Rahman and Wang, 2016; Brosch et al., 2015; Salehi et al., 2017), only a few non-volumetric losses are established for CNN training such as Karimi and Salcudean (2019)<sup>1</sup>.

We aim to identify the quantitative correlates of human expert perception concerning semantic segmentation quality assessment. Therefore, we conduct psychophysical segmentation rating experiments on multi-modal 3D magnetic resonance (MR) glioma data and 2D ultrasound (US) data for grey- and white matter segmentation with expert radiologists.

Building upon these insights, we propose a method exploiting techniques of classical statistics and experimental psychology to form new compound losses for convolutional neural network(CNN) training. Even though our CNN training experiments do not achieve the initially intended goal, we derive meaningful insights for future research in semantic segmentation. Across our quality assessment experiments, human experts consistently score computer-generated segmentations better than human-curated reference labels. This finding raises many questions about the status quo of semantic segmentation in medical images.

---

1. This is especially true when searching for multi-class segmentation losses, which support GPU-based training. For instance, according to the authors, there is no implementation of Karimi and Salcudean (2019) with GPU support.

## 2. Methodology

### 2.1 Collection of expert ratings

To better understand how experts judge the quality of medical segmentations, we conduct experiments where participants rate segmentation quality on a six-degree Likert scale (see Figure 1). The selection options range from one star for *strong rejection* to six stars for *perfect* segmentations. Participants assign star ratings by pressing the corresponding keys on the keyboard. The hotkeys are selected, accounting for several international keyboard layouts to minimize and equalize reaction times. All stimuli appear on a neutral-grey background. The segmentation labels and other colorful items are presented employing a color-blind-friendly color palette (Wong, 2011). Participants could toggle the display of the semi-opaque segmentation labels by pressing the space bar. A fixation cross appears in a quasi-random inter-trial interval (ITI) of  $125 / 250 / 500 / 750$  milliseconds to ensure participants cannot anticipate the onset of subsequent trials. We recorded how often participants toggled the display of the segmentation labels and reaction times for all tasks. By comparing the individual ratings against the mean, the methodology enables analyzing the participants’ individual biases (see Appendix). To allow the recording of accurate reaction times, the experiment was realized using *JS Psych v6.1* (De Leeuw, 2015) embedded into a *Vue JS v2.6* web application.




Figure 1: Chronological sequence of the experiments from left to right. Participants conduct the evaluations in a suitable environment for reading medical exams. The experiments start with participants declaring consent. Afterward, participants answer a survey regarding their age, gender, and various items to measure their expertise. Subsequently, the stimulus trials are presented in random sequence to account for order effects. Following the rating assignment, the experiment automatically progresses to the presentation of the next trial. At each experiment’s end, experts have the opportunity to provide feedback during a post-trial survey.

### 2.2 Computation of metrics

We calculate a comprehensive set of segmentation quality metrics for each experimental condition using *pymia* (Jungo et al., 2021), summarized in Table 1. The metrics compare the segmentations to the expert-curated reference images.

Table 1: Metrics considered in the analysis. The metrics are categorized according to Taha and Hanbury (2015).

Metric	Abbreviation	Reference	Category
Dice coefficient	DICE	(Dice, 1945)	overlap
Jaccard coefficient	JACRD	(Jaccard, 1912)	overlap
Sensitivity	SNSVTY	(Tharwat, 2020)	overlap
Specificity	SPCFTY	(Tharwat, 2020)	overlap
Fallout	FALLOUT	(Tharwat, 2020)	overlap
False negative rate	FNR	(Tharwat, 2020)	overlap
Accuracy	ACURCY	(Tharwat, 2020)	overlap
Precision	PRCISON	(Tharwat, 2020)	overlap
False positive	FP	(Pearson, 1904)	overlap
True positive	TP	(Pearson, 1904)	overlap
False negative	FN	(Pearson, 1904)	overlap
True negative	TN	(Pearson, 1904)	overlap
F-measure	FMEASR	(Tharwat, 2020)	overlap
Volume similarity	VOLSMTY	(Cárdenes et al., 2009)	volume
Prediction volume	PREDVOL	(Kohlberger et al., 2012)	volume
Reference volume	REVOL	(Kohlberger et al., 2012)	volume
Global consistency error	GCOERR	(Martin et al., 2001)	misc
Rand index	RNDIND	(Rand, 1971)	pair-counting
Adjusted rand index	ADJRIND	(Hubert and Arabie, 1985)	pair-counting
Mutual information	MUTINF	(Cover, 1999)	information-theoretical
Variation of information	VARINFO	(Meilă, 2003)	information-theoretical
Interclass correlation	ICCORR	(Shrout and Fleiss, 1979)	probabilistic
Probabilistic distance	PROBDST	(Gerig et al., 2001)	probabilistic
Cohen Kappa coefficient	KAPPA	(Cohen, 1960)	probabilistic
Area under curve	AUC	(Powers, 2020)	probabilistic
Hausdorff distance	HDRFDST	(Huttenlocher et al., 1993)	distance
Average distance	AVGDIST	(Fujita, 2013)	distance
Mahalanobis distance	MAHLNBS	(Mahalanobis, 1936)	distance
Surface Dice overlap	SURFDICE	(Nikolov et al., 2018)	distance
Surface overlap	SURFOVLP	(Nikolov et al., 2018)	distance

### 2.3 Loss function implementations

Table 2 summarizes the loss functions implementations employed in our experiments. For many losses, multiple implementations exist. We discovered that multiple implementations of the same losses could produce different loss values. Implementations might arise from different means of approximation, such as the choice of  $\epsilon$ . To investigate whether the choice of implementation plays a role, we included multiple implementations in the analysis. For Tversky loss, we choose combinations of hyperparameters that have proven successful for

similar segmentation problems; in the following analysis, the parameters for  $\alpha$  and  $\beta$  are written behind the abbreviation.

Table 2: Losses considered in our analysis. Implementations are Github links.

Loss name	Abbreviation	Implementation	Reference
Asymmetric	ASYM	junma	Hashemi et al. (2018)
Binary cross-entropy	BCE	pytorch	n/a
(Soft) Dice	DICE	monai	Milletari et al. (2016)
Soft Dice	SOFTD	nnUNet	Drozdal et al. (2016)
Generalized Dice	GDICE.L	liviaets	Sudre et al. (2017)
Generalized Dice	GDICE.W	wolny	Sudre et al. (2017)
Generalized Dice	GDICE.M	monai	Sudre et al. (2017)
Hausdorff DT	HDDT	patryg	Karimi and Salcudean (2019)
Hausdorff ER	HDER	patryg	Karimi and Salcudean (2019)
Jaccard	IOU	junma	Rahman and Wang (2016)
Jaccard	JAC	monai	Milletari et al. (2016)
Sensitivity-Specifity	SS	junma	Brosch et al. (2015)
Tversky	TVERSKY	monai	Salehi et al. (2017)

## 2.4 Statistical analysis

All statistical analyses are computed with *R*. We implement linear mixed models via *lme4* (Bates et al., 2015) and evaluate their performance using the package *performance* (Lüdtke et al., 2021). The principal component analysis is illustrated with *factorminer* (Lê et al., 2008).

## 2.5 Construction of compound losses

To achieve a better correlation with human segmentation quality assessment, we aim to design new loss functions by combining established loss functions in a complementary fashion. The resulting compound loss functions are constructed as a weighted linear combination of established loss functions per label channel:

$$L = \sum_i \alpha_i \sum_j w_{ij} loss_{ij} \quad (1)$$

where  $\alpha_i$  denotes weights for channel  $i$ , and  $w_{ij}$  denotes weights for loss  $j$  in channel  $i$ . Here, the model coefficients obtained from the linear mixed models serve as weights.

Further information regarding the mixed model experiments can be found in Section 3 and Appendix D.

### 3. Experiments

#### Experiment 1: MR segmentation rating

**Motivation:** To understand how human experts evaluate segmentation quality, we conduct a perception study. We select the BraTS 2019 test set (Menze et al., 2014) as a platform for our experiments for several reasons. First, glioma segmentation represents a complex multi-class segmentation problem. Second, the annotations have been curated by multiple domain experts over the years. Consequently, unlike in other available segmentation data sets, the annotations not only represent the interpretation of a single individual and are of high quality. Last, the dataset is non-public, so we can be sure that the evaluated algorithms are not trained on the annotations.

**Procedure:** We randomly select three exams from eight institutions that provided data to the BraTS 2019 test set (Menze et al., 2014). Additionally, we present one exam with apparently erroneous segmentation labels to check whether participants follow the instructions. We display stimulus material according to four experimental conditions, i.e. four different segmentation labels for each exam: the human-annotated *reference* segmentations, segmentations from team *zyx* (Zhao et al., 2019), a *SIMPLE* (Langerak et al., 2010) fusion of seven segmentations (Zhao et al., 2019; Feng et al., 2019; Isensee et al., 2018; McKinley et al., 2019; Weninger et al., 2019; Marcinkiewicz, 2019; Nuechterlein and Mehta, 2019) obtained from BraTS-Toolkit (Kofler et al., 2020), and a segmentation randomly selected without replacement from all the teams participating in BraTS 2019. In order to have an unbiased view, we select the non-public BraTS test set for this experiment, as we can be sure that none of the evaluated algorithms is trained on it. For all **25** exams, the respective center of mass is displayed on **3** axes, namely axial, coronal, and sagittal view for each of the **4** experimental conditions. This results in a total of  $25 * 3 * 4 = 300$  trials presented to each expert radiologist. Figure 2 depicts an example trial.

Furthermore, we calculate metrics comparing to the official reference labels for all BraTS evaluation criteria: *WT*, *TC*, and *ET*, as well as the individual labels: *necrosis* and *edema*. In addition, we compute mean aggregates for BraTS and individual labels.

**Results:** A total of  $n=15$  radiologists (*two* female, *13* male) from *six* institutions participated in the experiment. Participants had an average experience of  $10.0 \pm 5.1$  working years as radiologists and a mean age of  $37.7 \pm 4.8$ . Uninterrupted, participants needed approximately 50 minutes to complete the experiment. Figure 3 and Figure 4 depict how experts evaluated the segmentation quality. Figure 5 depicts the Pearson correlation matrix between segmentation quality metrics and expert assessment. We find only low to moderate correlations across all metrics; correlations are especially low for the clinically important *enhancing tumor* label. It is important to note that the frequently used DICE coefficient is outperformed by other less established metrics.



Figure 2: Stimulus material and rating controls presented to the participants in the BraTS segmentation rating experiment. One trial consisted of the presentation of an MR exam in either axial, sagittal, or coronal view, along with the controls for the quality rating and a legend for the segmentation labels. We presented the glioma’s center of mass according to the TC, defined by the *nerosis and enhancing tumor* voxels in 2D slices of T1, T1c, T2, T2 FLAIR in a horizontal row arrangement. The stimulus presentation is conducted in line with best practices in experimental psychology; further details are outlined in Figure 1 and in Section 2.1




Figure 4: Mean expert assessment in star rating across exams for the different experimental conditions. Ratings *without* (left) vs. ratings *with* bias correction (right). Diamonds indicate mean scores. Expert radiologists rated the *simple* fusion with a mean of 4.79 stars. This is slightly higher than the best individual docker algorithm *zyx* with 4.71 stars. The fusions’ mean improvement is mainly driven by more robust segmentation performance with fewer outliers towards the bottom. We observe that both of these algorithmic segmentations receive slightly higher scores than the human expert-annotated *reference* labels at 4.47 stars. The quality of the randomly selected BraTS submissions (*rnd*) still lags behind with 4.04 stars. Three more qualitative segmentation rating examples can be found in the Appendix.



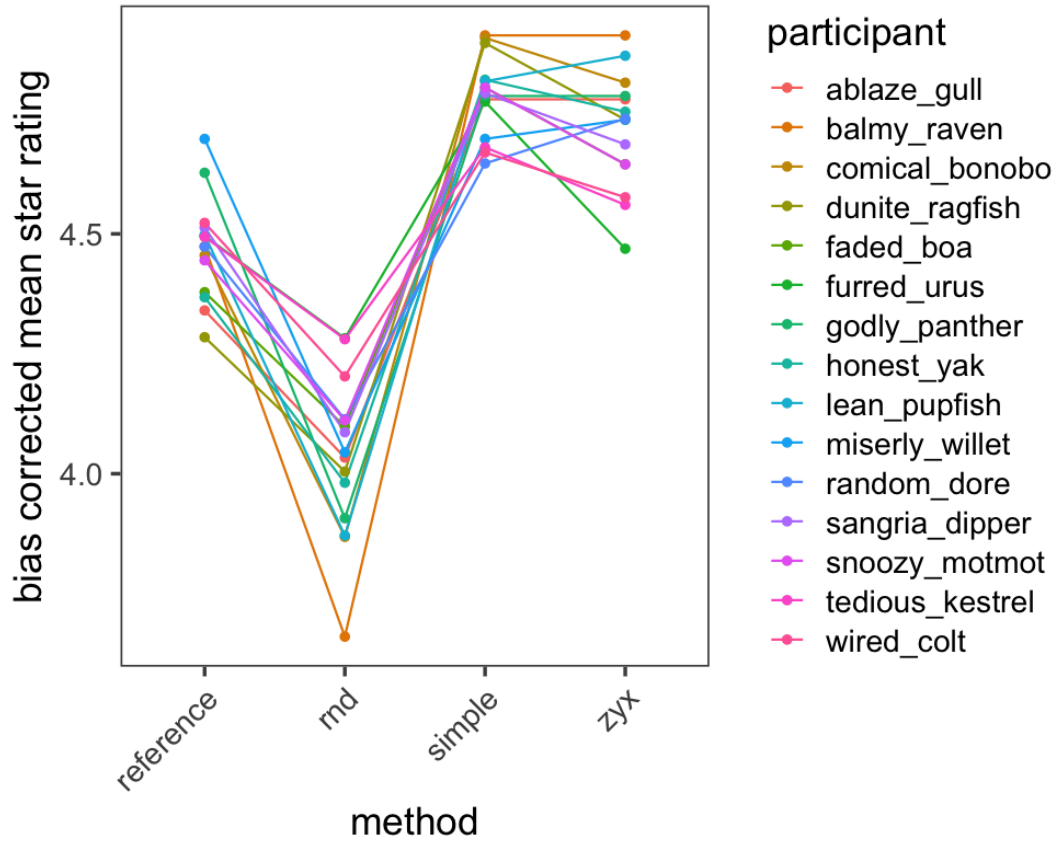


Figure 3: Bias-corrected mean star ratings per expert and method for the MR segmentation rating experiment. Raw ratings and the bias correction methodology are detailed in Appendix. Notably, the human-annotated reference labels (*gt*) are predominantly rated worse than computer-generated annotations of *simple* and *zyx*.




Figure 5: Pearson correlation matrix between expert assessment and segmentation quality metrics. The rows show the correlations for the individual label channels. In addition, we present the mean aggregates for the BraTS labels (*enhancing tumor*, *tumor core*, and *whole tumor*), as well as the single channels (*enhancing tumor*, *necrosis*, and *edema*). For the abbreviations in the figure refer to (Taha and Hanbury, 2015; Jungo et al., 2021).

## Experiment 2: US segmentation rating

**Motivation:** To understand how the findings from Experiment 1 translate into other segmentation problems, we conduct another expert perception study. We investigate the task of tissue segmentation in ultrasound (US) imaging in a dataset from Demiray et al. (2019). The dataset has the interesting property that the reference annotations are created in a multi-stage collaborative human-machine interaction: Demiray et al. (2019) base their work on the RESECT dataset (Xiao et al., 2017). In addition to co-registered pre-operative T1-weighted and T2-weighted MR images and the intra-operative US images from different stages of the intervention, RESECT provides corresponding anatomic landmarks for both modalities. First, the co-registered pre-operative MR and pre-durotomy US are re-registered with an affine registration with LC2 metric (Fuerst et al., 2014) in ImFusion Suite<sup>2</sup>. Second, two human experts selected the registration result with smaller registration errors according to the RESECT landmarks as a starting point for manual corrections of the registration. Third, the tissues in the MR volumes are segmented with FreeSurfer (Fischl, 2012) and manually corrected by two medical imaging experts (work experience >3 years). Fourth, the tissue segmentation maps are transformed from the MR space into the US space. Last, the human experts again perform quality control on the resulting annotations.

**Procedure:** For the perception study, human experts evaluate the performance of different tissue segmentation algorithms against the reference. Analogue to the first experiment, the experts are blinded regarding the source of the segmentation. In addition to the reference annotation, we present two segmentation algorithms from Demiray et al. (in preparation). First, a segmentation algorithm based on DeepMedic (Kamnitsas et al., 2017) and another based on nnUnet (Isensee et al., 2021). Further, a control condition with broken annotations is presented to the participants. Figure 6 illustrates one example of the total 57 trials. In line

2. ImFusion GmbH, Munich, Germany, <https://www.imfusion.com>

with Experiment 1, we compute segmentation quality metrics between the two candidate algorithms, the control condition, and the reference labels and correlate them with expert assessment.




Figure 6: Stimulus material presented to the participants during the US segmentation rating experiment. One trial consisted of the presentation of a US image (left) with the source area delineated in a corresponding MR image (right). In addition, the star rating assessment control elements were displayed analog to the BraTS experiment, see Figure 2. As in the previous experiment, participants could toggle the display of segmentation maps for grey- and white matter with the space bar. Further details regarding the stimulus presentation are outlined in Section Section 2.1

**Results:** A total of  $n=8$  radiologists (*three* female, *five* male), from *two* institutions participated in the experiment. Participants had an average work experience of  $5.25 \pm 5.63$  years and a mean age of  $35.0 \pm 5.42$ . Uninterrupted, participants needed around twenty minutes to complete the experiment. Figure 7 depicts how experts evaluated the segmentation quality across experimental conditions.




Figure 7: Bias-corrected mean star ratings per expert and condition for the ultrasound segmentation rating experiment. Expert radiologists rated the *nnu* (inspired by (Isensee et al., 2021)) and *dpm* (inspired by (Kamnitsas et al., 2017, 2015)) candidate algorithms much higher than the control condition (*con*) consisting out of purposely wrongly manufactured segmentations. Notably, again the U-Net segmentation achieves the highest ratings.

Furthermore, Figure 8 depicts the Pearson correlation matrix between segmentation quality metrics and the expert assessment. Similar to the first Experiment, we find only low correlations across all metrics, especially for the grey matter label.

### Experiment 3: Evaluation of established loss functions

**Motivation:** To evaluate how these findings could translate into CNN training. A differentiable loss function is required to train a CNN with stochastic gradient descent (SGD). As all metrics do not fulfill this criterion, we investigate how established loss functions correlate with expert quality assessment.

**Procedure:** To achieve this, we use the binary label maps to compute established segmentation losses (defined in Table 2) between the candidate segmentations and the *reference* labels for the BraTS segmentation task. To test the validity of this approach, we conduct a supportive experiment, comparing *binary* segmentations to *non-binary* network outputs and

## ARE WE USING APPROPRIATE SEGMENTATION METRICS?




Figure 8: Pearson correlation matrix between expert assessment and segmentation quality metrics for the US segmentation rating experiment. The rows show the correlations for grey matter (*gm*) and white matter (*wm*).

achieve comparable results across all loss functions. In line with the previous experiments, we then correlate the resulting losses with expert assessment. Furthermore, we conduct a principal component analysis (PCA) to better understand loss functions’ complementarity.

**Results:** Figure 9 illustrates the correlations with expert assessment; again, we only find low to moderate correlations. Furthermore, the results of the PCA are depicted in the Appendix.




Figure 9: Pearson correlation matrix between expert assessment and segmentation losses. With the exception of one Generalized Dice Score implementation, we observe only low correlations for the enhancing tumor and tumor core channel. In contrast, multiple losses are moderately correlated for the whole tumor channel. We observe huge variance across the implementations of *Generalized Dice*. However, the implementations abbreviated as IOU and JAC for *Jaccard index*, as well as DICE and SOFTD for *Dice coefficient*, provide very similar signals, as expected.

### Experiment 4: Construction of compound loss functions

**Motivation:** As all investigated established loss functions achieve, at best moderate correlation with expert assessment, we set out to explore whether we can achieve a better fit with expert assessment when combining these.

**Procedure:** We approach this in a two-step process. First, we identify promising loss combinations: Losses produce different signals (gradients) when they react to input signals (network outputs). With hierarchical clustering on a Euclidean distance matrix, we analyze similarity in loss response patterns, see Section 3.




Figure 10: Phylogenetic tree illustrating hierarchical clustering on a Euclidean distance matrix. Losses are colored according to ten cluster groups. We included the expert assessment for reference; one can observe how it resides somewhere between the distance-based losses and the group of generalized Dice losses and SS loss. These findings are in line with the *PCA*, see Appendix.

Applying Equation (1), we can now build complementary loss combinations by selecting established loss functions from the different and hopefully complementary cluster groups.

In a second step, we evaluate our loss combinations' predictive performance for the expert assessment using linear mixed models (LMM) as detailed in section Section 2.4.

Therefore, we average the human expert rating across views to obtain a *quasi-metric* variable allowing us to apply mixed models. We deem this approach valid, as the distribution is consistent across views; see Appendix. We then model the human expert assessment as a dependent variable and predict it by plugging in the loss values of our candidate combinations as fixed-effect predictor variables. Mixed models allow accounting for the *non-independence* of data points by additionally modeling random factors. Therefore, we include *exam* as a random factor, as some glioma are more difficult to segment. This increased difficulty is expressed in less segmentation performance and greater variability across algorithmic and human segmentations. Additionally we include the *segmentation method* as a random factor, as some algorithms are better than others and human segmentations are worse than good algorithms according to our experts, compare Section 3. To identify promising compound loss candidates for CNN training, we evaluate the predictive power of our models by computing Pseudo  $R^2$  (Nakagawa et al., 2017) while monitoring typical mixed regression model criteria, such as multi-collinearity, non-normality of residuals and outliers, homoscedasticity, homogeneity of variance and normality of random effects, see Appendix D. Once a promising loss combination is detected, we obtain the weights for our loss formulation (see Section 2.5) from the corresponding LMM.

**Results:** Our methodology allows for generating a variety of loss candidates. To evaluate the performance of our approach, we obtain four promising loss candidates. The first two candidates use the same combination across all label channels. Candidate *gdice\_bce* uses *GDICE\_W* in combination with *BCE*. Candidate *gdice\_ss\_bce* iterates upon this by adding *SS* loss to the equation. In contrast the candidates *channel\_wise\_weighted* and *channel\_wise* use different losses per channel: Channel *whole tumor* uses the *gdice\_ss\_bce* candidate loss, while channel *tumor core* is computed via the *gdice\_bce* variant. In contrast, the *enhancing tumor* channel relies solely on *GDICE\_W*, as this is the only loss candidate with at least moderate correlation with expert assessment, compare Figure 9. While candidate *channel\_wise* treats all channels equally, the weighted variant prioritizes the more clinically relevant *tumor core and enhancing tumor channels* by a factor of *five*. Detailed analysis for all models is provided in the Appendix.

### Experiment 5: Quantitative evaluation of compound losses

**Motivation:** To explore how the four identified compound losses from Experiment 4 perform with regard to established metrics, we perform another experiment.

**Procedure:** Therefore, we train nnU-Net (Isensee et al., 2021) using a standard BraTS trainer (Isensee et al., 2020) with a moderate augmentation pipeline on the *fold 0 split* of the *BraTS 2020* training set, using the rest of the training set for validation. The official *BraTS 2020* validation set is used for testing. To make our experiment fully reproducible, we select the BraTS 2020 data for this experiment as the training data is publicly available, and metric evaluation on the validation set can be obtained through the CBICA imaging portal. Apart from replacing the default *DICE+BCE* loss with our custom-tailored loss candidates and, in some cases, necessary learning rate adjustments, we keep all training

parameters constant.<sup>3</sup> To shed some light on the training variance, we also train another five times with dice\_bce baseline.

**Results:** We observe no significant quantitative improvement; the resulting volumetric dice scores are detailed in Figure 11.

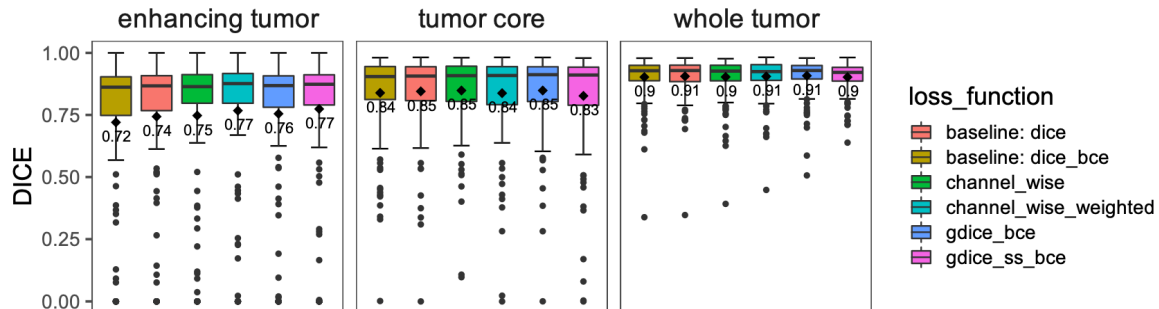


Figure 11: Dice comparison of loss candidates vs. baselines across BraTS label channels. Diamonds indicate mean scores. P-values, without adjustment for repeated tests, for paired-samples T-tests comparing our candidates with the dice\_bce baseline from left to right:  $0.1547$ ,  $0.04642$ ,  $0.05958$ ,  $0.01881$  for the 95% significance level. While we consider the volumetric DICE score improvement as non-significant, interestingly, it is larger than the training variance, see Appendix.

## Experiment 6: Qualitative evaluation of compound losses

**Motivation:** Even though we did not achieve a systematic improvement according to established segmentation quality metrics, we conduct another expert perception study to evaluate the generated compound losses with expert radiologists and fully close the loop. As we learned from previous experiments that established quantitative metrics are not fully representative of expert assessment, we still consider this experiment worthy.

**Procedure:** Experts evaluate only the axial views of the gliomas' center of mass on a random sample of 50 patients from our test set. We present our four candidate losses vs. the *reference* and *DICE+BCE* baseline as conditions. This way, the experiment incorporates 300 trials.

**Results:** Three male senior radiologists from two institutions participated in the experiment. Participants reported three, five, and eleven years of work experience and were aged 31, 37, and 40 years. Figure 12 illustrates how experts can hardly detect performance differences between the evaluated computer-generated segmentations. Remarkably, again the human-annotated *reference* is rated worse.

3. Friendly note to beloved reviewers: Naturally, the code for our nnU-net training will be publicly available on GitHub for full reproducibility.



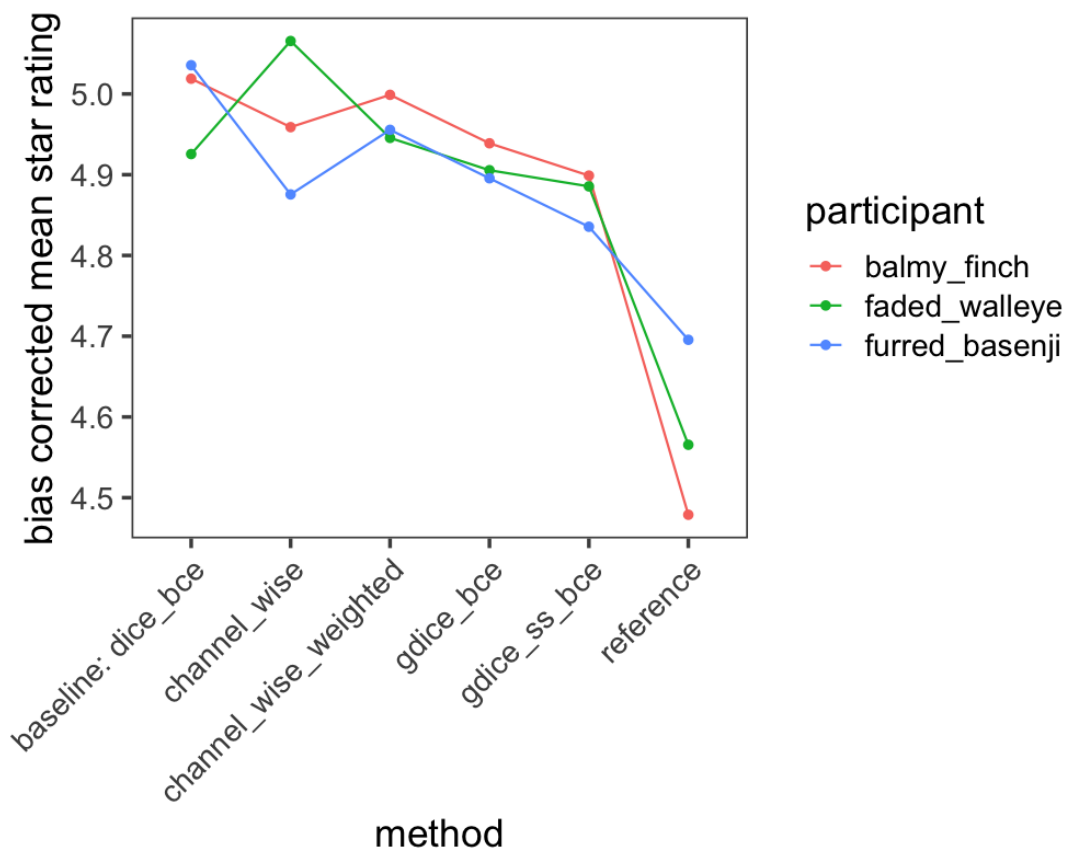


Figure 12: Bias-corrected mean star ratings per expert and condition for the second MR segmentation rating experiment. Raw ratings and the bias correction methodology are detailed in the Appendix. Expert assessment of the four loss candidates vs. *reference* and the *DICE+BCE* baseline. We observe only subtle differences in expert rating between the *DICE+BCE* baseline and our candidates. Notably, as in the first perception study, the human-annotated *reference* is rated worse, compare Figure 3 & Figure 4.

## 4. Discussion

We conducted multi-center expert perception studies with 15 expert radiologists for 3D MR- and eight expert radiologists for 2D ultrasound imaging segmentations. Furthermore, to achieve a better fit with human expert assessment, we developed a method to construct compound loss functions for CNN training. These compound loss functions represent an effort to better approximate abstract metrics that cannot be optimized directly, such as human expert assessment. For this purpose, we select components based on the cluster - and principal component analysis and obtain weights from linear mixed models. Even though we do not manage to train more expert-pleasing segmentation networks, our findings provide meaningful cues for future research.

Existing literature (Reinke et al., 2018; Maier-Hein et al., 2018) reported on the discrepancy between expert segmentation quality assessment and established segmentation quality metrics. Our approach, combining psychophysical experiments and classical statistics, identifies the quantitative correlates of qualitative human expert perception to shed more light on this complex phenomenon. While both entities, *segmentation quality metrics* and *expert radiologists*, try to measure segmentation quality, our analysis reveals that their measurements correlate only moderately at best, especially for clinically relevant features such as the *enhancing tumor* label. Therefore, we highlight the need to develop new segmentation metrics and losses that better correlate with expert assessment.

One limitation of our study is that similarity metrics are always dependent on the employed reference annotations. It is evident that *DSC* fails to represent expert opinion when the segmentation of small structures matters, such as for multiple sclerosis lesion segmentation (Kofler et al., 2022b) or metastasis segmentation (Buchner et al., 2022). However, in the absence of actual ground truth, it is impossible to disentangle to which extent the observed low correlations are due to the incapability of metrics to capture expert assessment or a result of poor reference annotations (Kofler et al., 2022c). Nevertheless, it should be noted that the BraTS dataset arguably features above-average annotation quality, as, unlike other datasets in the biomedical field, it was curated by several experts over multiple years.

Another limitation of our study is that in the BraTS segmentation rating experiments, radiologists judge the complex 3D segmentations on three 2D views. However, while axial views are rated slightly better, the rating effects are quite consistent across views (see Appendix B).

When training a modern convolutional neural network with our expert-inspired compound loss functions, we cannot produce a significant improvement over the well-established *Dice+BCE* baseline. As we find *BCE* to be one of few loss functions complementary to confusion-matrix-based losses and the *Dice+BCE* baseline is mathematically similar to our identified loss candidates, this is perhaps not surprising (see Section 3). Hence our findings might explain why the empirically found *Dice+BCE* provides a solid baseline across many segmentation tasks.

As algorithms and humans can usually generate satisfactory pre-operative glioma segmentations, poor segmentations are under-represented in our sample. We hypothesize that if we replace our random sample from the BraTS test set with an equally distributed sample covering the full spectrum of segmentation quality, we might find higher correlations

between established similarity metrics and human expert perception. Up to the level of inter-rater agreement, it seems possible to train visually appealing segmentation networks using a plain soft Dice loss (Kofler et al., 2022c,a).<sup>4</sup>

Interestingly, across all our experiments, experts consistently scored CNN segmentations better than the human-curated *reference* labels. This is not only true for sophisticated ensembles (see Figure 3) but also for segmentations generated by single CNNs (compare Figure 7, Figure 12, Appendix B). This highlights the need to develop better annotation procedures.

We can only speculate why computers might produce better segmentations than human annotators. In the annotation process, human annotators produce systematic and random errors. With enough training data, a CNN can abstract and will only learn systematic errors represented in the training data. However, training a CNN can eliminate random, non-systematic human annotation errors by averaging these out.

When interviewed in a single-blind setting, radiologists repeatedly attributed their judgments to the better tracing of contours<sup>5</sup>. This seems plausible, as, unlike computers, humans do not possess unlimited diligence to pixel/voxel-perfectly trace contours. Further, this is supported by the fact that the difference between algorithmic and human-annotated reference segmentations is least pronounced in the axial view Appendix B that is the basis for most expert annotations (see Appendix B). However, how these complex decisions are formed remains an open research question.

**Conclusion:** Even though the terms are often used synonymously in the literature, our experiments highlight that it is imperative to reflect upon whether *reference* labels actually qualify as *ground truth*. In our experiments, a hypothetical volumetric DICE score of 1, so 100 percent overlap with the *reference*, would translate to a worse segmentation quality, according to expert radiologists. Therefore, our findings question the status of the Dice coefficient as the *de facto gold standard* for measuring segmentation quality beyond expert agreement. Furthermore, they question whether segmentation performance in general and challenge rankings, in particular, should be determined based on similarity metrics with potentially erroneous human-curated *reference* annotations (Kofler et al., 2022c).

Future research should investigate whether humans are still able to outperform CNNs when applying exhaustive quality assurance and control procedures. While more participants would certainly add to the credibility of our experiments, we found the reported effects to be consistent across participants. As senior radiologists are hard to come by, one should mention that this, to the best of our knowledge, represents the largest study featuring expert segmentation quality ratings. As primate perception follows a fixed, partially understood set of rules, for instance, (Wagemans et al., 2012; Roelfsema, 2006), it might be that our findings also generalize to other segmentation problems, even outside the medical domain. However, further research is required to explore this.

---

4. This might also explain why surrogate models for human star rating reached a higher correlation with volumetric *DSC* (Kofler et al., 2022a) on a dataset with a broader spectrum of segmentation quality.

5. We find inconclusive multifaceted, moderate correlations between shape features of specific labels and expert ratings.

## Acknowledgments

Even though we cannot reveal their names to maintain anonymity, we want to thank the participating radiologists who enabled these studies.

Bjoern Menze, Benedikt Wiestler and Florian Kofler are supported through the SFB 824, subproject B12.

Supported by Deutsche Forschungsgemeinschaft (DFG) through TUM International Graduate School of Science and Engineering (IGSSE), GSC 81.

Suprosanna Shit and Ivan Ezhov are supported by the Translational Brain Imaging Training Network (TRABIT) under the European Union’s ‘Horizon 2020’ research & innovation program (Grant agreement ID: 765148).

Ivan Ezhov, Suprosanna Shit are funded by DComEX (Grant agreement ID: 956201).

With the support of the Technical University of Munich – Institute for Advanced Study, funded by the German Excellence Initiative.

Supported by Anna Valentina Lioba Eleonora Claire Javid Mamasani.

Research reported in this publication was supported in part by the National Cancer Institute (NCI) of the National Institutes of Health, under award number U01CA242871. The content of this publication is solely the responsibility of the authors and does not represent the official views of the NIH.

Suprosanna Shit is supported by the Graduate School of Bioengineering, Technical University of Munich.

Jan Kirschke has received Grants from the ERC, DFG, BMBF and is Co-Founder of Bone-screen GmbH.

Bjoern Menze acknowledges support by the Helmut Horten Foundation.

Research reported in this publication was partly supported by AIME GPU cloud services.

## Ethical Standards

The work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding treatment of animals or human subjects.

## Conflicts of Interest

We declare we don’t have conflicts of interest.

## References

- Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018.
- Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48, 2015. doi: 10.18637/jss.v067.i01.
- Patrick Bilic, Patrick Ferdinand Christ, Eugene Vorontsov, Grzegorz Chlebus, Hao Chen, Qi Dou, Chi-Wing Fu, Xiao Han, Pheng-Ann Heng, Jürgen Hesser, et al. The liver tumor segmentation benchmark (lits). *arXiv preprint arXiv:1901.04056*, 2019.
- Tom Brosch, Y Yoo, Lisa YW Tang, David KB Li, Anthony Traboulsee, and R Tam. Deep convolutional encoder networks for multiple sclerosis lesion segmentation. vol. 9556, 2015.
- Josef A Buchner, Florian Kofler, Lucas Etzel, Michael Mayinger, Sebastian M Christ, Thomas B Brunner, Andrea Wittig, Björn Menze, Claus Zimmer, Bernhard Meyer, et al. Development and external validation of an mri-based neural network for brain metastasis segmentation in the aurora multicenter study. *Radiotherapy and Oncology*, 2022.
- Rubén Cárdenes, Rodrigo de Luis-Garcia, and Meritxell Bach-Cuadra. A multidimensional segmentation evaluation for medical image data. *Computer methods and programs in biomedicine*, 96(2):108–124, 2009.
- Jacob Cohen. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46, 1960.
- Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- Joshua R De Leeuw. jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior research methods*, 47(1):1–12, 2015.
- Beatrice Demiray, Julia Rackerseder, Stevica Bozhinoski, and Nassir Navab. Weakly-supervised white and grey matter segmentation in 3d brain ultrasound. *arXiv*, 2019.
- Lee R Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3):297–302, 1945.
- Michal Drozdal, Eugene Vorontsov, Gabriel Chartrand, Samuel Kadoury, and Chris Pal. The importance of skip connections in biomedical image segmentation. In *Deep learning and data labeling for medical applications*, pages 179–187. Springer, 2016.
- Xue Feng, Nicholas Tustison, and Craig Meyer. Brain tumor segmentation using an ensemble of 3d u-nets and overall survival prediction using radiomic features. In Alessandro Crimi, Spyridon Bakas, Hugo Kuijf, Farahani Keyvan, Mauricio Reyes, and Theo van Walsum, editors, *International MICCAI Brainlesion Workshop*, pages 279–288, Cham, 2019. Springer. ISBN 978-3-030-11726-9.

- Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.
- Bernhard Fuerst, Wolfgang Wein, Markus Müller, and Nassir Navab. Automatic ultrasound–mri registration for neurosurgery using the 2d and 3d lc2 metric. *Medical image analysis*, 18(8):1312–1319, 2014.
- Osamu Fujita. Metrics based on average distance between sets. *Japan Journal of Industrial and Applied Mathematics*, 30(1):1–19, 2013.
- Guido Gerig, Matthieu Jomier, and Miranda Chakos. Valmet: A new validation tool for assessing and improving 3d object segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 516–523. Springer, 2001.
- Seyed Raein Hashemi, Seyed Sadegh Mohseni Salehi, Deniz Erdogmus, Sanjay P Prabhu, Simon K Warfield, and Ali Gholipour. Asymmetric loss functions and deep densely-connected networks for highly-imbalanced medical image segmentation: Application to multiple sclerosis lesion detection. *IEEE Access*, 7:1721–1735, 2018.
- Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of classification*, 2(1):193–218, 1985.
- Daniel P Huttenlocher, Gregory A. Klanderman, and William J Rucklidge. Comparing images using the hausdorff distance. *IEEE Transactions on pattern analysis and machine intelligence*, 15(9):850–863, 1993.
- Fabian Isensee, Philipp Kickingereder, Wolfgang Wick, Martin Bendszus, and Klaus H Maier-Hein. No new-net. In *International MICCAI Brainlesion Workshop*, pages 234–244. Springer, 2018.
- Fabian Isensee, Paul F Jäger, Peter M Full, Philipp Vollmuth, and Klaus H Maier-Hein. nnu-net for brain tumor segmentation. In *International MICCAI Brainlesion Workshop*, pages 118–132. Springer, 2020.
- Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021.
- Paul Jaccard. The distribution of the flora in the alpine zone. 1. *New phytologist*, 11(2):37–50, 1912.
- Alain Jungo, Olivier Scheidegger, Mauricio Reyes, and Fabian Balsiger. pymia: A python package for data handling and evaluation in deep learning-based medical image analysis. *Computer methods and programs in biomedicine*, 198:105796, 2021.
- Konstantinos Kamnitsas, Liang Chen, Christian Ledig, Daniel Rueckert, and Ben Glocker. Multi-scale 3d cnns for segmentation of brain lesions in multi-modal mri. *proceeding of ISLES challenge, MICCAI*, 2015.

- Konstantinos Kamnitsas, Christian Ledig, Virginia FJ Newcombe, Joanna P Simpson, Andrew D Kane, David K Menon, Daniel Rueckert, and Ben Glocker. Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation. *Medical image analysis*, 36:61–78, 2017.
- Davood Karimi and Septimiu E Salcudean. Reducing the hausdorff distance in medical image segmentation with convolutional neural networks. *IEEE Transactions on medical imaging*, 39(2):499–513, 2019.
- Florian Kofler, Christoph Berger, Diana Waldmannstetter, Jana Lipkova, Ivan Ezhov, Giles Tetteh, Jan Kirschke, Claus Zimmer, Benedikt Wiestler, and Bjoern H Menze. Brats toolkit: translating brats brain tumor segmentation algorithms into clinical and scientific practice. *Frontiers in neuroscience*, page 125, 2020.
- Florian Kofler, Ivan Ezhov, Lucas Fidon, Izabela Horvath, Ezequiel de la Rosa, John LaMaster, Hongwei Li, Tom Finck, Suprosanna Shit, Johannes Paetzold, Spyridon Bakas, Marie Piraud, Jan Kirschke, Tom Vercauteren, Claus Zimmer, Benedikt Wiestler, and Bjoern Menze. Deep quality estimation: Creating surrogate models for human quality ratings, 2022a. URL <https://arxiv.org/abs/2205.10355>.
- Florian Kofler, Suprosanna Shit, Ivan Ezhov, Lucas Fidon, Izabela Horvath, Rami Al-Maskari, Hongwei Li, Harsharan Bhatia, Timo Loehr, Marie Piraud, et al. blob loss: instance imbalance aware loss functions for semantic segmentation. *arXiv preprint arXiv:2205.08209*, 2022b.
- Florian Kofler, Johannes Wahle, Ivan Ezhov, Sophia Wagner, Rami Al-Maskari, Emilia Gryska, Mihail Todorov, Christina Bukas, Felix Meissen, Tingying Peng, et al. Approaching peak ground truth. *arXiv preprint arXiv:2301.00243*, 2022c.
- Timo Kohlberger, Vivek Singh, Chris Alvino, Claus Bahlmann, and Leo Grady. Evaluating segmentation error without ground truth. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 528–536. Springer, 2012.
- Thomas Robin Langerak, Uulke A van der Heide, Alexis NTJ Kotte, Max A Viergever, Marco Van Vulpen, and Josien PW Pluim. Label fusion in atlas-based segmentation using a selective and iterative method for performance level estimation (simple). *IEEE transactions on medical imaging*, 29(12):2000–2008, 2010.
- Sébastien Lê, Julie Josse, and François Husson. FactoMineR: A package for multivariate analysis. *Journal of Statistical Software*, 25(1):1–18, 2008. doi: 10.18637/jss.v025.i01.
- Daniel Lüdecke, Mattan S Ben-Shachar, Indrajeet Patil, Philip Waggoner, and Dominique Makowski. performance: An r package for assessment, comparison and testing of statistical models. *Journal of Open Source Software*, 6(60), 2021.
- Daniel Lüdecke, Mattan S. Ben-Shachar, Indrajeet Patil, Philip Waggoner, and Dominique Makowski. performance: An R package for assessment, comparison and testing of statistical models. *Journal of Open Source Software*, 6(60):3139, 2021. doi: 10.21105/joss.03139.

- Prasanta Chandra Mahalanobis. On the generalized distance in statistics. National Institute of Science of India, 1936.
- Lena Maier-Hein, Matthias Eisenmann, Annika Reinke, Sinan Onogur, Marko Stankovic, Patrick Scholz, Tal Arbel, Hrvoje Bogunovic, Andrew P Bradley, Aaron Carass, et al. Why rankings of biomedical image analysis competitions should be interpreted with care. *Nature communications*, 9(1):1–13, 2018.
- Michal and et al Marcinkiewicz. Segmenting brain tumors from mri using cascaded multi-modal u-nets. In *International MICCAI Brainlesion Workshop*, pages 13–24, Cham, 2019. Springer. ISBN 978-3-030-11726-9.
- David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 416–423. IEEE, 2001.
- Richard McKinley, Raphael Meier, and Roland Wiest. Ensembles of densely-connected cnns with label-uncertainty for brain tumor segmentation. In Alessandro Crimi, Spyridon Bakas, Hugo Kuijf, Farahani Keyvan, Mauricio Reyes, and Theo van Walsum, editors, *International MICCAI Brainlesion Workshop*, pages 456–465, Cham, 2019. Springer. ISBN 978-3-030-11726-9.
- Marina Meilă. Comparing clusterings by the variation of information. In *Learning theory and kernel machines*, pages 173–187. Springer, 2003.
- Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014.
- Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016.
- Shinichi Nakagawa, Paul CD Johnson, and Holger Schielzeth. The coefficient of determination  $r^2$  and intra-class correlation coefficient from generalized linear mixed-effects models revisited and expanded. *Journal of the Royal Society Interface*, 14(134):20170213, 2017.
- Stanislav Nikolov, Sam Blackwell, Alexei Zverovitch, Ruheena Mendes, Michelle Livne, Jeffrey De Fauw, Yojan Patel, Clemens Meyer, Harry Askham, Bernardino Romera-Paredes, et al. Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy. *arXiv preprint arXiv:1809.04430*, 2018.
- Nicholas Nuechterlein and Sachin Mehta. 3d-espnet with pyramidal refinement for volumetric brain tumor image segmentation. In Alessandro Crimi, Spyridon Bakas, Hugo Kuijf, Farahani Keyvan, Mauricio Reyes, and Theo van Walsum, editors, *International MICCAI Brainlesion Workshop*, pages 245–253, Cham, 2019. Springer. ISBN 978-3-030-11726-9.



- Karl Pearson. *On the theory of contingency and its relation to association and normal correlation*, volume 1. Dulau and Company London, UK, 1904.
- David MW Powers. Evaluation: from precision, recall and f-measure to roc, informedness, markedness and correlation. *arXiv preprint arXiv:2010.16061*, 2020.
- Md Atiqur Rahman and Yang Wang. Optimizing intersection-over-union in deep neural networks for image segmentation. In *International symposium on visual computing*, pages 234–244. Springer, 2016.
- William M Rand. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical association*, 66(336):846–850, 1971.
- Annika Reinke, Matthias Eisenmann, Sinan Onogur, Marko Stankovic, Patrick Scholz, Peter M Full, Hrvoje Bogunovic, Bennett A Landman, Oskar Maier, Bjoern Menze, et al. How to exploit weaknesses in biomedical challenge design and organization. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 388–395. Springer, 2018.
- Pieter R Roelfsema. Cortical algorithms for perceptual grouping. *Annu. Rev. Neurosci.*, 29:203–227, 2006.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- Seyed Sadegh Mohseni Salehi, Deniz Erdogmus, and Ali Gholipour. Tversky loss function for image segmentation using 3d fully convolutional deep networks. In *International workshop on machine learning in medical imaging*, pages 379–387. Springer, 2017.
- Anjany Sekuboyina, Amirhossein Bayat, Malek E Husseini, Maximilian Löffler, Markus Rempfler, Jan Kukačka, Giles Tetteh, Alexander Valentinitzsch, Christian Payer, Martin Urschler, et al. Verse: a vertebrae labelling and segmentation benchmark. *arXiv.org e-Print archive*, 2020.
- Patrick E Shrout and Joseph L Fleiss. Intraclass correlations: uses in assessing rater reliability. *Psychological bulletin*, 86(2):420, 1979.
- Carole H Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pages 240–248. Springer, 2017.
- Abdel Aziz Taha and Allan Hanbury. Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. *BMC medical imaging*, 15(1):1–28, 2015.
- Alaa Tharwat. Classification assessment methods. *Applied Computing and Informatics*, 2020.

- Johan Wagemans, James H Elder, Michael Kubovy, Stephen E Palmer, Mary A Peterson, Manish Singh, and Rüdiger von der Heydt. A century of gestalt psychology in visual perception: I. perceptual grouping and figure-ground organization. *Psychological bulletin*, 138(6):1172, 2012.
- Michael Weller, Martin Van Den Bent, Kirsten Hopkins, Jörg C Tonn, Roger Stupp, Andrea Falini, Elizabeth Cohen-Jonathan-Moyal, Didier Frappaz, Roger Henriksson, Carmen Balana, et al. Eano guideline for the diagnosis and treatment of anaplastic gliomas and glioblastoma. *The lancet oncology*, 15(9):e395–e403, 2014.
- Patrick Y Wen, David R Macdonald, David A Reardon, Timothy F Cloughesy, A Gregory Sorensen, Evanthia Galanis, John DeGroot, Wolfgang Wick, Mark R Gilbert, Andrew B Lassman, et al. Updated response assessment criteria for high-grade gliomas: response assessment in neuro-oncology working group. *Journal of clinical oncology*, 28(11):1963–1972, 2010.
- Leon Weninger, Oliver Rippel, Simon Koppers, and Dorit Merhof. Segmentation of brain tumors and patient survival prediction: Methods for the brats 2018 challenge. In Alessandro Crimi, Spyridon Bakas, Hugo Kuijf, Farahani Keyvan, Mauricio Reyes, and Theo van Walsum, editors, *International MICCAI Brainlesion Workshop*, pages 3–12, Cham, 2019. Springer. ISBN 978-3-030-11726-9.
- Bang Wong. Points of view: Color blindness. *Nature Methods*, 8(6):441, 2011. ISSN 1548-7105.
- Yiming Xiao, Maryse Fortin, Geirmund Unsgård, Hassan Rivaz, and Ingerid Reinertsen. Retrospective evaluation of cerebral tumors (resect): A clinical database of pre-operative mri and intra-operative ultrasound in low-grade glioma surgeries. *Medical physics*, 44(7): 3875–3882, 2017.
- Yuan-Xing Zhao, Yan-Ming Zhang, Ming Song, and Cheng-Lin Liu. Multi-view semi-supervised 3d whole brain segmentation with a self-ensemble network. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 256–265. Springer, 2019.

## Appendix A. Terminology definitions

Here we define some basic concepts of Neuroradiology and experimental Psychology used in the manuscript.

**Exam:** A radiology exam consists of a set of image recordings recorded in one session. For instance, a glioma exam in the BraTS challenge consists of a T1, T1c, T2, and FLAIR image.

**Stimulus (material):** Unlike the control elements, which are constant over all trials, the stimulus material is varied across experimental conditions. The participants' responses are recorded and analyzed with regard to different stimuli.

**Fixation cross:** In the inter-trial interval (ITI), we display fixation crosses to avoid confounding the recording of reaction times with *unnecessary* eye movements such as (re-)focusing on the screen.

## Appendix B. Supplemental figures




Figure 13: PCA with two factors. The scree plot and multiple nongraphical tests show that two factors are sufficient to describe the loss landscape. PCA analysis with the full set of loss functions on the *left*, for the sake of clarity, we repeated the PCA analysis with a subset of commonly used loss functions on the *right*. Similar to the hierarchical clustering, one can see how expert assessment can be described as a combination of two factors. While the volumetric losses load mainly on dimension one, only *BCE* and the two Hausdorff losses cover dimension two. This might explain why the empirically found baseline *DICE+BCE* performs well across many segmentation tasks.

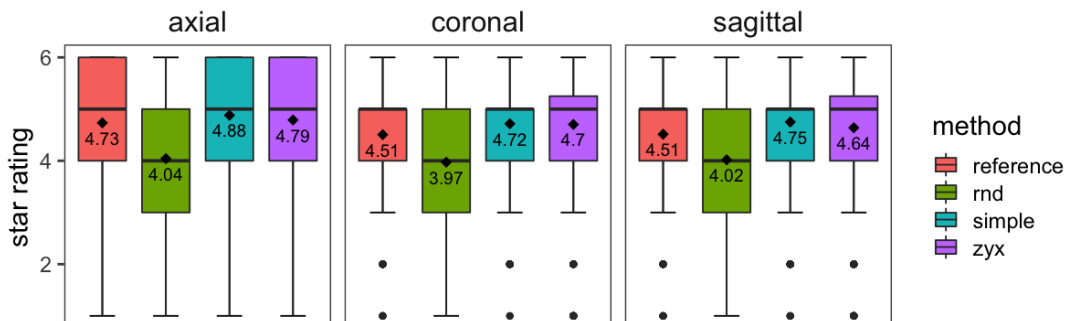


Figure 14: Star rating across views. Diamonds indicate mean scores. The distribution is constant per *axial*, *coronal* and *sagittal* view. In general, segmentation quality ratings seem slightly higher on the axial view. We speculate this might occur because doctors usually annotate on the axial view, so the resulting annotation quality should be slightly higher. In addition, the resolution is often higher in the axial plane leading to less ambiguity when judging the quality of segmentation images.

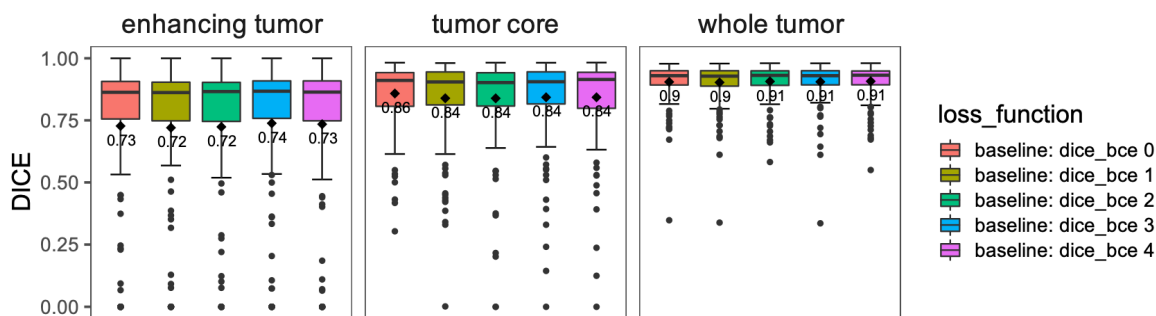


Figure 15: Illustration of training variance. Diamonds indicate mean scores. To get an estimate of the randomness involved in our training process, we trained our *nnU-net* implementation five times with *DICE+BCE* loss. We observe a dice performance between *0.72* and *0.74*, which is lower than the performance of our loss candidates.




Figure 16: Scree plot. Multiple tests (see top right corner) indicate that two components are sufficient for the PCA.




Figure 17: Three segmentation examples with varying expert assessment. T1, T1c, T2, and FLAIR images are displayed from left to right. The segmentations are overlaid in *red* for *necrosis*, *yellow* for *enhancing tumor* and *green* for *edema*. The 15 expert radiologists rated the top example *superior* with a mean of 5.73 stars, the middle example *mediocre* with a mean of 3.33 stars and the bottom example *inferior* with a mean of 1.93 stars.

### Appendix C. Bias correction




Figure 18: Raw (non-averaged) expert assessment in star rating across exams for the different experimental conditions. Diamonds indicate mean scores.




Figure 19: Mean star ratings per participant and method for MR segmentation rating Experiment 1. Some participants reveal a strong / positive biases compared to the mean ratings, compare Figure 23. However, when correcting for these biases participants are remarkably consistent in their ratings, compare Figure 3.




Figure 20: Mean star ratings per participant and method for ultrasound segmentation rating experiment 2. Similar to the MR experiment, some participants reveal a strong / positive biases compared to the mean ratings. However, when correcting for these biases participants are remarkably consistent in their ratings, compare Figure 7.




Figure 21: Raw (non-averaged) expert assessment in star rating across exams in the ultrasound segmentation rating experiment 2. Diamonds indicate mean scores. Expert radiologists rated the *nnu* (inspired by (Isensee et al., 2021)) and *dpm* (inspired by (Kamnitsas et al., 2017, 2015)) candidate algorithms much higher than the control condition (*con*) consisting out of purposely wrongly manufactured segmentations.



ARE WE USING APPROPRIATE SEGMENTATION METRICS?




Figure 22: Mean star ratings per participant and condition for MR segmentation rating Experiment 6. Some participants reveal strong biases in relation to the mean ratings. However, when correcting for these biases, participants are remarkably consistent in their ratings, compare Figure 12.




Figure 23: Individual bias per participant in Experiment 1. Diamonds indicate the mean. The boxplots illustrate the difference between the participants individual ratings and the mean rating per exam across all participants in Experiment 1. While some participants are very congruent with the average, e.g. participant *ablaze\_gull*, we observe strong inter-individual differences for others like *miserly\_willet*. One can obtain bias corrected ratings by subtracting the respective bias from the participants individual ratings resulting, compare Figure 3.

## Appendix D. Linear mixed models

Following, we illustrate the model outputs and model diagnostics for the linear mixed models derived in Section 3. The model diagnostics are generated with *performance* (Lüdtke et al., 2021). The model coefficients serve as weights for our loss functions. For the *channel-wise\_weighted* loss functions the more clinically relevant *tumor core* and *enhancing tumor* channel are weighted with a factor of *five* over the *whole tumor* channel for the channel wise aggregation.

Note: in the model outputs the *segmentation method* is referred to as *condition*. Further, the *exam* is encoded as *patient* (in our dataset there is only one exam per patient so the variables encode the same concept).

D.1 *gdice\_bce* lossListing 1: R model output for *gdice\_bce* loss

---

```

Linear mixed model fit by REML.
t-tests use Satterthwaite's method ['lmerModLmerTest']
Formula: stars ~ BCE + GDICE_W + (1 | patient) + (1 | condition)
  Data: mm_losses_df

REML criterion at convergence: 106.9

Scaled residuals:
   Min       1Q   Median       3Q      Max
-1.7200 -0.3908  0.1478  0.4664  2.1067

Random effects:
   Groups      Name      Variance Std.Dev.
condition (Intercept) 0.11107  0.3333
patient   (Intercept) 0.21882  0.4678
Residual                    0.09391  0.3065
Number of obs: 72, groups:  condition, 26; patient, 24

Fixed effects:
              Estimate Std. Error    df t value Pr(>|t|)
(Intercept)    4.0671      0.3146 65.7925  12.927 < 2e-16 ***
BCE             -0.4624      0.1024 50.4863  -4.514 3.83e-05 ***
GDICE_W        -0.7462      0.2896 58.8458  -2.577  0.0125 *
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
      (Intr) BCE
BCE    -0.642
GDICE_W 0.814 -0.306
              R2m      R2c
[1,] 0.3911532 0.8650808

```

---

ARE WE USING APPROPRIATE SEGMENTATION METRICS?

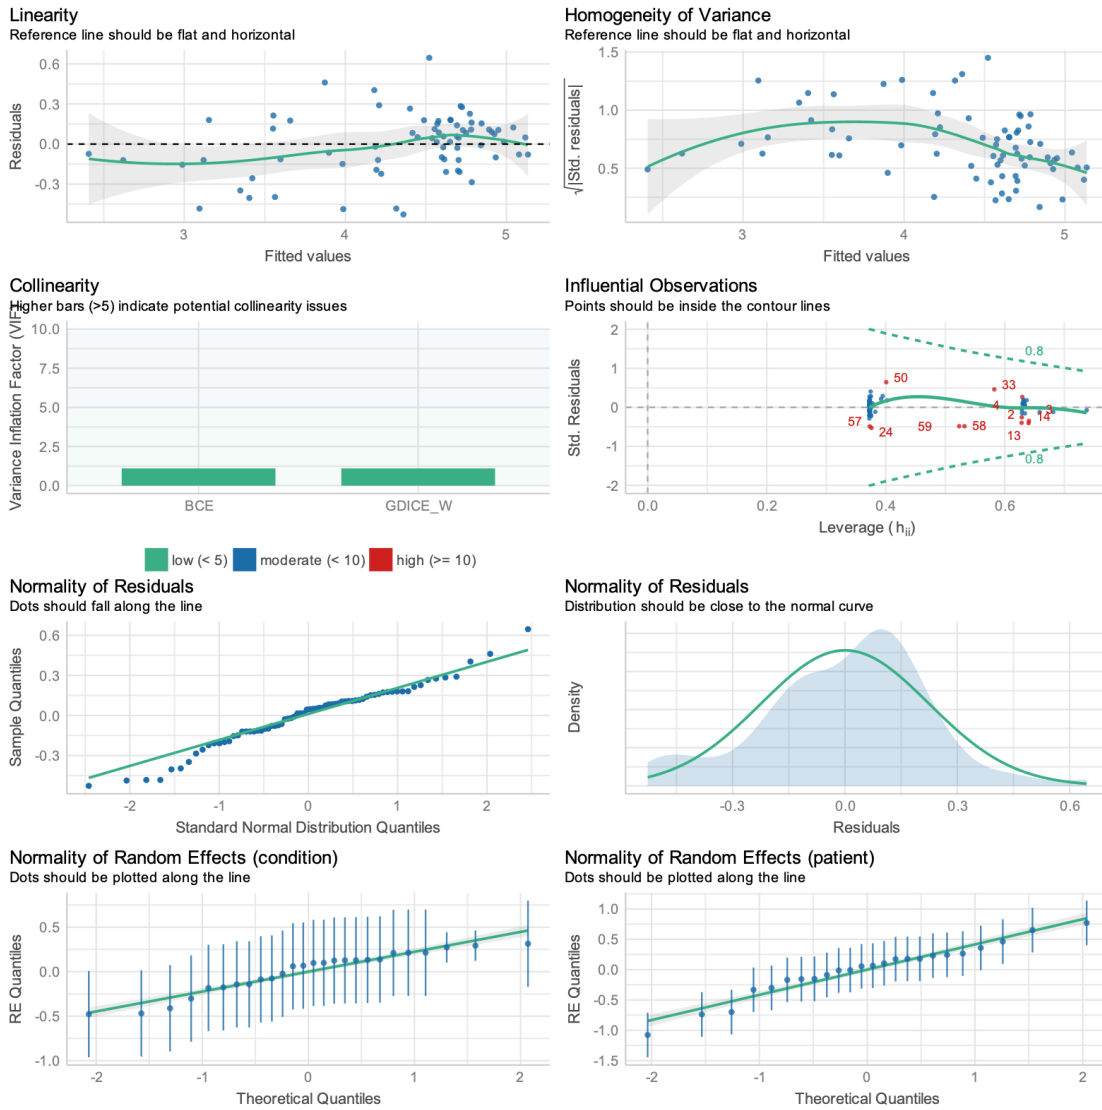


Figure 24: Model diagnostics for the *gdice\_bce* loss. The loss function combined *BCE* and *GDICE.W* with weights  $0.4624$  and  $0.7462$  obtained from the respective coefficients in the linear mixed model. The model achieved a *Pseudo R2* of  $0.391$ . )

D.2 *gdice\_bce\_ss* lossListing 2: R model output for *gdice\_bce\_ss* loss

---

```

Linear mixed model fit by REML.
t-tests use Satterthwaite's method [
lmerModLmerTest]
Formula: stars ~ BCE + GDICE_W + SS + (1 | patient) + (1 | condition)
Data: mm_losses_df

REML criterion at convergence: 98.3

Scaled residuals:
    Min       1Q   Median       3Q      Max
-1.9685 -0.4159  0.1278  0.4612  2.1035

Random effects:
   Groups      Name      Variance Std.Dev.
condition (Intercept) 0.11433  0.3381
patient   (Intercept) 0.25798  0.5079
Residual                    0.08433  0.2904
Number of obs: 72, groups:  condition, 26; patient, 24

Fixed effects:
              Estimate Std. Error      df t value Pr(>|t|)
(Intercept)   4.4306      0.3829   63.7986  11.570 <2e-16 ***
BCE            -0.3267      0.1693   42.6522  -1.930  0.0603 .
GDICE_W       -0.4570      0.3439   50.1523  -1.329  0.1899
SS            -18.2016     13.9181   54.4509  -1.308  0.1964
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
      (Intr) BCE      GDICE_
BCE      0.125
GDICE_W  0.863  0.289
SS      -0.569 -0.788 -0.556
              R2m      R2c
[1,] 0.40454 0.8900317

```

---

## ARE WE USING APPROPRIATE SEGMENTATION METRICS?

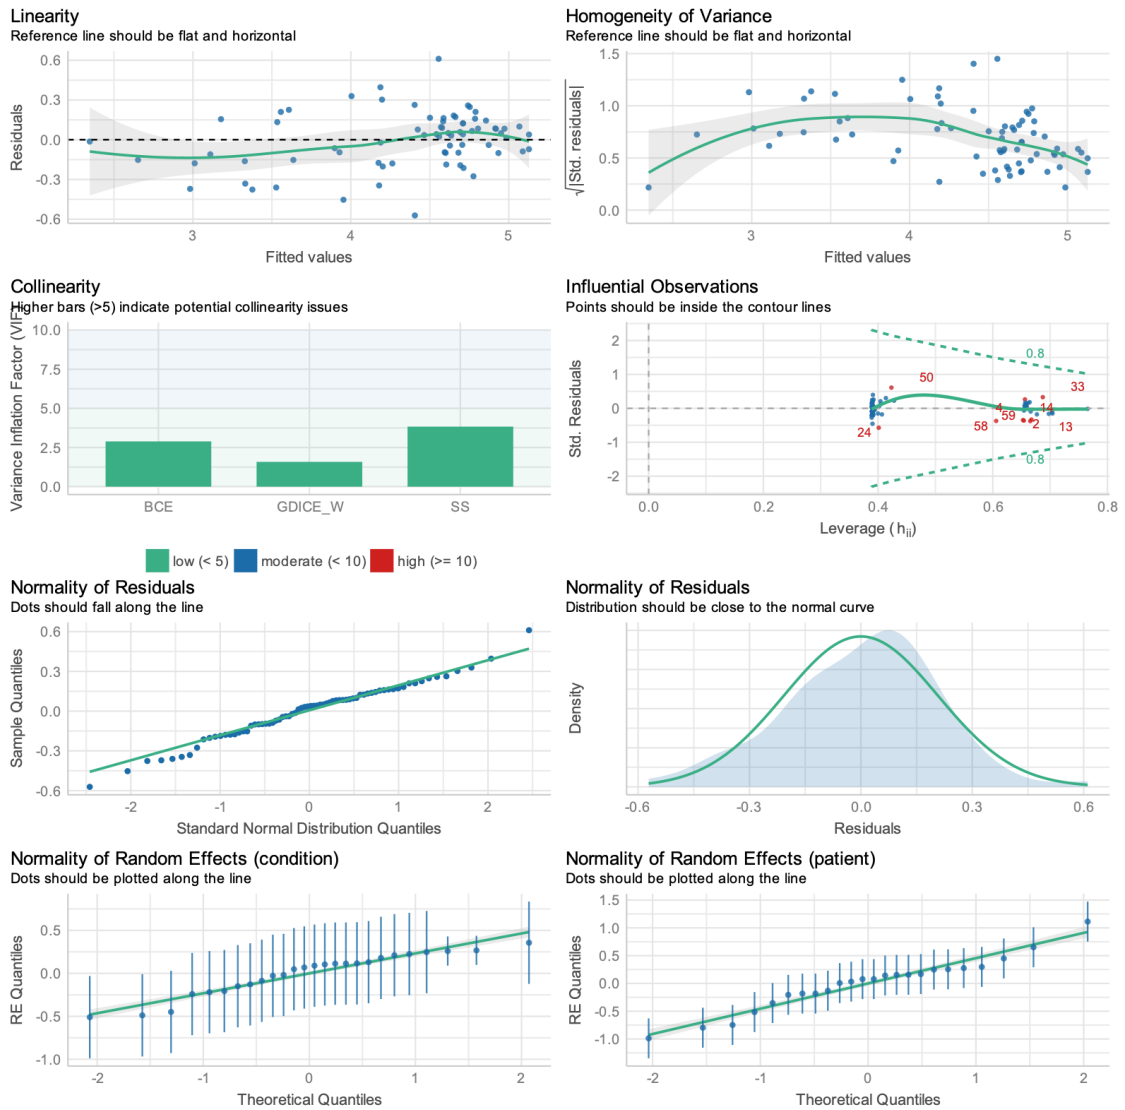


Figure 25: Model diagnostics for the *gdice\_bce\_ss* loss. The loss function combined *BCE*, *GDICE.W* and *SS* with weights 0.3267, 0.4570 and 18.2016 obtained from the respective coefficients in the linear mixed model. The model achieved a *Pseudo R2* of 0.405.

**D.3 channel wise losses: whole tumor channel**

Listing 3: R model output for whole tumor channel

---

```

Linear mixed model fit by REML. t-tests use Satterthwaite's method ['
  lmerModLmerTest']
Formula: stars ~ GDICE_W + SS + BCE + (1 | patient) + (1 | condition)
Data: wt

REML criterion at convergence: 315.2

Scaled residuals:
    Min       1Q   Median       3Q      Max
-2.21560 -0.57743  0.00251  0.61567  2.83251

Random effects:
 Groups   Name                Variance Std.Dev.
condition (Intercept) 0.2637    0.5135
patient  (Intercept) 0.1590    0.3988
Residual                    0.1630    0.4037
Number of obs: 216, groups:  condition, 26; patient, 24

Fixed effects:
              Estimate Std. Error    df t value Pr(>|t|)
(Intercept)   3.1165      0.9290 207.7590   3.355 0.000944 ***
GDICE_W       -1.5876      0.9146 204.4111  -1.736 0.084114 .
SS            -4.0027     10.2444 196.2805  -0.391 0.696425
BCE           -0.3039      0.0785 158.8889  -3.872 0.000157 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
      (Intr) GDICE_ SS
GDICE_W  0.985
SS       -0.893 -0.899
BCE      0.492  0.545 -0.729

```

---



## ARE WE USING APPROPRIATE SEGMENTATION METRICS?

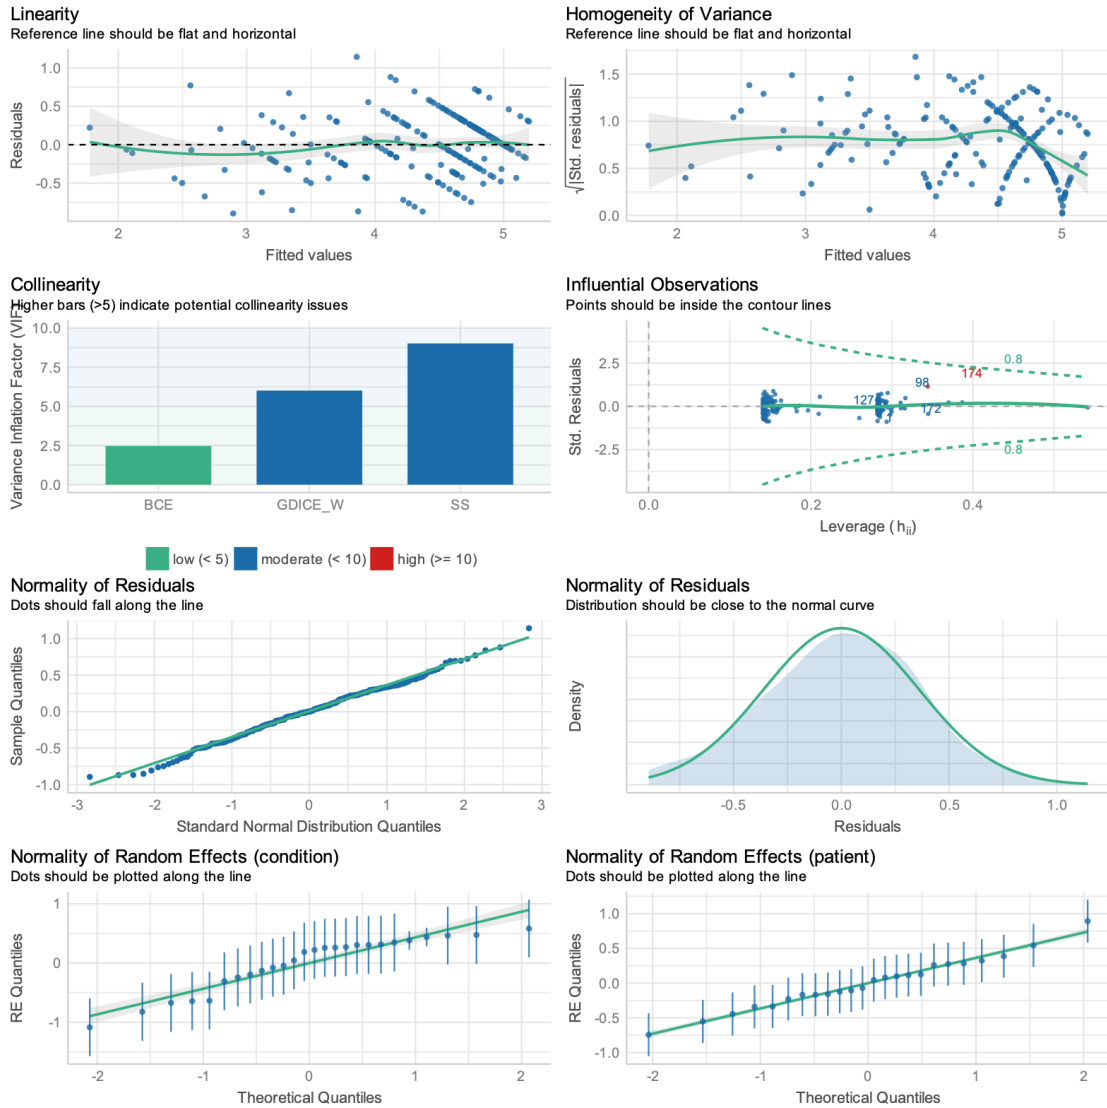


Figure 26: Model diagnostics for the *whole tumor* channels of the *channel\_wise* and *channel\_wise\_weighted* loss functions. The loss functions combined *GDICE\_W*, *SS* and *BCE* with weights  $1.5876$ ,  $4.0027$  and  $0.3039$  obtained from the respective coefficients in the linear mixed model. Notably, we observed higher variation inflation than for *gdice\_bce\_ss* loss combining the same components over all channels. The model achieved a *Pseudo R2* of  $0.270$ .

## D.4 channel wise losses: tumor core channel

Listing 4: R model output for tumor core channel

---

```

Linear mixed model fit by REML.
t-tests use Satterthwaite's method ['lmerModLmerTest']
Formula: stars ~ GDICE_W + BCE + (1 | patient) + (1 | condition)
Data: tc

```

```
REML criterion at convergence: 366.6
```

```
Scaled residuals:
```

```

      Min       1Q   Median       3Q      Max
-2.8922 -0.6517  0.0618  0.6587  3.1648

```

```
Random effects:
```

```

Groups   Name             Variance Std.Dev.
condition (Intercept) 0.1892   0.4349
patient  (Intercept) 0.1852   0.4303
Residual                    0.2112   0.4596

```

```
Number of obs: 216, groups: condition, 26; patient, 24
```

```
Fixed effects:
```

```

              Estimate Std. Error    df t value Pr(>|t|)
(Intercept)  3.76303     0.25975 19.38353  14.487 7.44e-12 ***
GDICE_W      -0.77646     0.26230 12.01047  -2.960  0.0119 *
BCE          -0.21026     0.04496 40.97802  -4.676 3.16e-05 ***
---

```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Correlation of Fixed Effects:
```

```

      (Intr) GDICE_
GDICE_W  0.817
BCE      -0.492 -0.282

```

---

## ARE WE USING APPROPRIATE SEGMENTATION METRICS?

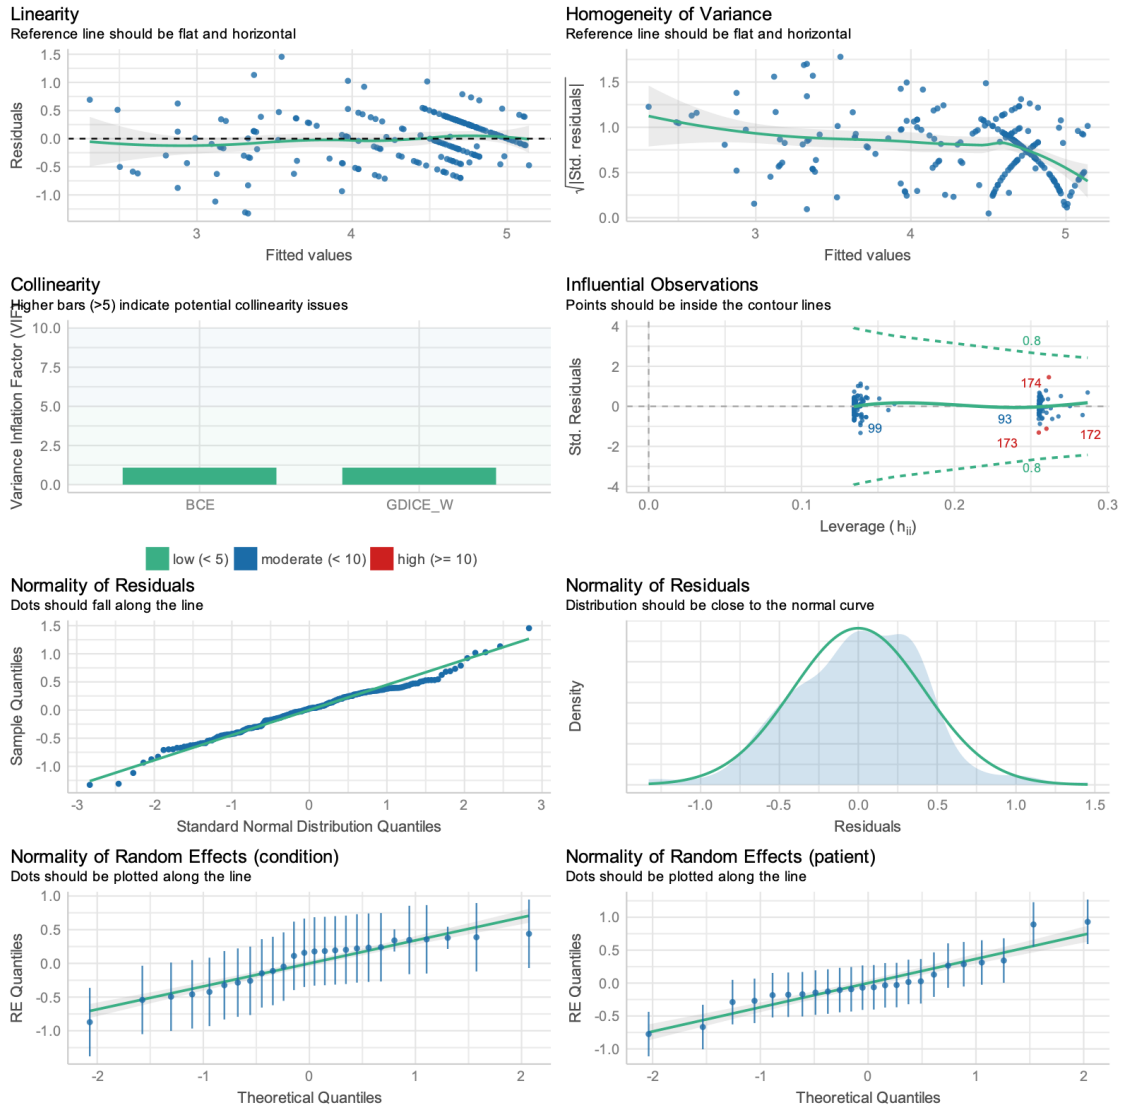


Figure 27: Model diagnostics for the *tumor core* channels of the *channel\_wise* and *channel\_wise\_weighted* loss functions. The loss functions combined *GDICE.W* and *BCE* with weights *0.77646* and *21026* obtained from the respective coefficients in the linear mixed model. The model achieved a *Pseudo R2* of *0.347*.

**D.5 channel wise losses: enhancing tumor channel**

Listing 5: R model output for enhancing tumor channel

---

```

Linear mixed model fit by REML.
t-tests use Satterthwaite's method ['lmerModLmerTest']
Formula: stars ~ GDICE_W + (1 | patient) + (1 | condition)
  Data: et

REML criterion at convergence: 383.5

Scaled residuals:
   Min       1Q   Median       3Q      Max
-2.8596 -0.6098  0.0881  0.6515  3.6149

Random effects:
   Groups      Name          Variance Std.Dev.
condition (Intercept) 0.3716    0.6096
patient   (Intercept) 0.1203    0.3469
Residual                    0.2308    0.4805
Number of obs: 216, groups:  condition, 26; patient, 24

Fixed effects:
              Estimate Std. Error      df t value Pr(>|t|)
(Intercept)   3.3635      0.2169  42.8141  15.509 < 2e-16 ***
GDICE_W       -0.9002      0.2318  26.7340  -3.884 0.000609 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Correlation of Fixed Effects:
      (Intr)
GDICE_W 0.729

```

---

## ARE WE USING APPROPRIATE SEGMENTATION METRICS?

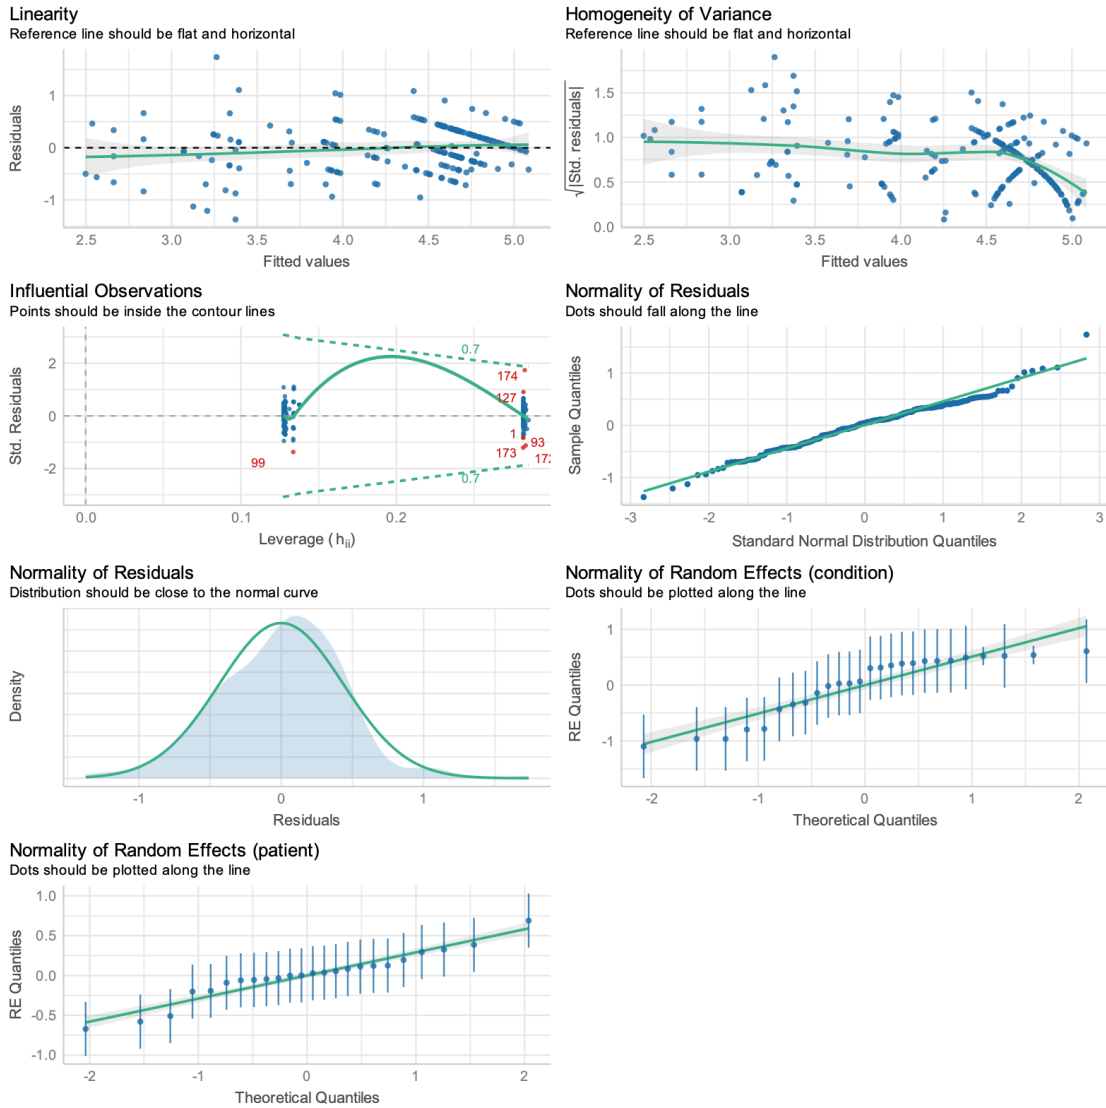


Figure 28: Model diagnostics for the *enhancing tumor* channels of the *channel\_wise* and *channel\_wise\_weighted* loss functions. The loss functions combined *GDICE.W* and *BCE* with weights *0.77646* and *21026* obtained from the respective coefficients in the linear mixed model. The model achieved a *Pseudo R2* of *0.112*, which is not surprising as our previous analysis showed that the clinically relevant *enhancing tumor* label is very difficult to capture.